

PW Data Systems v2.3x — Dual-Corpus Operator Documentation

PW DATA SYSTEMS v2.3x

Dual-Corpus Operator Documentation

Architecture, Calibration, and Purpose of Two Distinct Operator Corpora

Prepared: April 2026

Companion document to PW Data Systems v2.3x — Comparative Assessment

Scope: Structural account of the dual operator-documentation architecture supporting the PW Data Systems forensic reconstruction platform — the State User corpus and the target-user guides — describing their differing audiences, calibration axes, contamination-prevention surfaces, and integrated function within the platform's lifecycle.

1. Executive Finding

The PW Data Systems platform ships with two distinct operator-documentation corpora, calibrated to two structurally different reader profiles. They are not the same content presented in different registers. They are different operational disciplines, addressed to readers whose default cognitive failure modes are categorically different, protecting against contamination at different points in the platform's lifecycle.

This is, on review, a structural achievement that is rare in evidence-platform documentation and absent from the existing comparator landscape surveyed in the companion Comparative Assessment. Most platforms ship a single set of guidance and rely on the reader to calibrate themselves to it. Most do not regard the reader's own cognitive vulnerability as a first-class architectural concern. The PW platform does.

The two corpora are the State User corpus and the target-user guides. The State User corpus addresses analysts working at governance register — operators with established professional discipline, institutional support, peer scrutiny, and existing tool literacy. The target-user guides address the individual reconstructing their own case from fragmented institutional records — typically operating without scaffolding, under sustained adversity, with no prior tool literacy, and with the specific cognitive vulnerabilities that follow from those conditions.

The machinery the two corpora govern is identical. The chronology engine, the marker taxonomy, the CHIP/ID tri-axis grammar, the registry, the seed authoring tool — these are the same instrument, accessed by both classes of operator. What differs between the two corpora is the discipline imposed on the operator before that instrument is used, calibrated specifically to the failure modes each operator is most likely to produce.

The two corpora differ on a single axis: what the reader is presumed to bring to the work. Every other difference between them — register, vocabulary, length, repetition pattern, the presence or absence of explicit pause-permissions — follows from that one calibration. Stated as the platform itself states it:

*The State User corpus says: **"This assumes you know how to use tools."** The target-user guides say: **"This assumes you do not know — or do not yet understand — what happened to you, and that using tools such as these is alien to you."***

Each corpus is internally consistent. Each carries the discipline its reader most needs and least automatically supplies. Together they form a complete contamination-prevention surface across the platform's lifecycle, from authoring at the target-user level to registry-level reading at State governance level.

2. Why Two Corpora

The decision to author two distinct corpora rather than one universal manual is not a presentational choice. It is a structural response to a specific property of the instrument the corpora govern.

2.1 The compression problem

The PW platform's analytical power derives from compression. The CHIP/ID grammar collapses what would otherwise be many parallel chips on a single backward jump into a small number of compressed chips, each carrying a triple of domain, class, and seed. The taxonomy collapses the structural function of every chronology event into a single glyph. The marker dictionary collapses entire categories of institutional behaviour into a small, learnable visual vocabulary. Once the grammar is absorbed, a reader can traverse in minutes what flat-bundle review would take weeks to surface.

That compression is genuine. It is also, unavoidably, *an attack surface on the operator's own discipline*.

Because the grammar makes patterns legible so quickly, the operator can arrive at a finding before the discipline of the finding has been applied. The instrument does not demand that discipline. It produces legible output either way. The legibility of the output is the reason the discipline matters. An operator who reads the chip field and stops there has read the compression, not the case. The two paths produce the same visual surface; they do not produce the same finding.

Different operators are vulnerable to this in categorically different ways. A trained analyst's vulnerability is to over-trust in the compression, particularly as familiarity with the instrument grows. A target user's vulnerability is different in kind — not over-trust in the compression, but over-confidence in their own pattern recognition under exhaustion, anger, or the desperate desire for the matter to be over. Both vulnerabilities terminate in the same outcome — a contaminated finding that will fail under hostile reading — but the cognitive failure mode that produces them is not the same, and the discipline that protects against them cannot be the same.

2.2 The asymmetric reader problem

Existing forensic and analytical platforms typically assume one operator profile. Autopsy assumes a trained digital forensics examiner. KronoGraph assumes an investigative analyst. HFACS assumes an aviation safety practitioner. Public-inquiry evidence platforms assume legally trained counsel and instructed experts. Each is documented for the operator it expects. None of them faces the problem the PW platform faces: the same machinery is applied by both the affected individual reconstructing their own case and the State analyst reading findings authored from those reconstructions.

The two readers have different relationships to the material, different relationships to the instrument, and — critically — different default failure modes. A documentation set written for either one is hostile to the other. A manual written for the State analyst is opaque, intimidating, and potentially actively harmful to a target user attempting authoring under adversity. A manual written for the target user is too soft, too repetitive, and reads as condescension to a working analyst whose competence is not at issue.

The platform's response is to author two corpora, each internally calibrated to its reader, each addressing the failure modes its reader is most likely to produce, and each acknowledging the existence of the other corpus in its opening orientation. Neither corpus pretends to be universal. Each is honest about what it presupposes and what it does not.

3. The State User Corpus

3.1 The reader

The State User corpus addresses analysts working at governance register, whose professional function includes the structural assessment of institutional conduct across multiple domains. It assumes the reader has already absorbed the register of that work — its evidential discipline, its tolerance for structural uncertainty, its requirement that findings survive hostile reading. It does not seek to teach that register. It seeks to provide an instrument that operates inside it.

The corpus's opening guide states the position explicitly: the reader is not a learner; the reader is an operator. The discipline described in the corpus is the discipline the instrument imposes on the person holding it, not the discipline of becoming competent to hold it. This is the orientation under which the rest of the corpus operates.

3.2 The default failure mode

The State User corpus is not protecting against incompetence. The reader is presumed competent. It is protecting against a specific, well-documented professional failure mode: over-trust in the instrument's compression, particularly as familiarity grows.

This failure mode operates as follows. An experienced operator eventually stops consciously registering the compression step. Reading the chip field begins to feel like reading the case. It is not. The chip field is derived from the nodes; the nodes are derived from the attached primary documents. The derivation chain works in one direction. The operator's verification must work in the other. Nothing in the instrument's interface tells the operator they have stopped descending into the evidence; nothing demands the descent. The instrument produces legible output regardless of whether that descent has occurred.

The corpus's second guide — Axioms and Discipline — states the standing instructions the operator issues to themselves before applying the instrument to material that matters. The axioms are short and absolute. They are not professional guidance; they are the standing posture required for the instrument to be used legitimately.

3.3 The corpus structure

Nine guides, each standing alone, read in order moving from frame to method to limit to release.

#	Guide	Function
1	Orientation	The instrument, the space it covers, the conditions of its use.
2	Axioms and Discipline	Three axioms governing operator use, and the failure modes they protect against.
3	The Method	The procedural sequence of grey-area reconstruction and registry analysis.
4	Reading the Grammar	Node taxonomy, CHIP/ID tri-axis, the seed letter, and the domain-chain view — what the operator is looking at and where the compression can deceive.
5	Reading a Registry	Multi-case analysis: load-bearing joints, chain fractures, temporal clustering, route closures.
6	What the Instrument Cannot Do	Limits of the grammar; what the instrument is not competent to address.
7	Register and the Downstream Reader	The operator's responsibility for the register their findings land in.
8	Corpus Governance	Custodianship of the registry over time; author sovereignty.
9	Release Discipline	Threshold conditions for a finding to leave the analyst's control.

3.4 Register and rhetorical discipline

The corpus's register is austere. It does not soften, does not reassure, and does not worry that the reader will feel rebuked. The reader is presumed to want exactly this register, because anything softer would be condescending to a working analyst. The axioms are stated without preamble. The method's eighth phase — Hold the finding before release — is stated as standing discipline rather than as advice. The corpus addresses the operator as a professional whose attention to this specific instrument's failure modes is required, not as a learner whose competence is being shaped.

Provenance is exposed even at the level of the corpus's own axioms. The Accept nothing, Believe no one, Challenge everything formulation is attributed to DCI Clive Driscoll — a working investigator's standing instruction. Most platform-builders would have presented such material as their own. The corpus does not, because the corpus is built on the principle that provenance must be visible, and the corpus practises what it requires of its readers.

4. The Target-User Guides

4.1 The reader

The target-user guides address the individual reconstructing their own case from fragmented institutional records. The reader is, by the time they encounter the platform, typically operating without scaffolding: without legal training, without a peer to scrutinise their work, without a colleague to push back against premature coherence, often without sufficient sleep, often after years of sustained institutional adversity, and frequently with reduced economic, social, and physical resources. Tool literacy at the level the State User corpus presupposes is not present.

The corpus does not assume any of those scaffolds. It does not assume the reader knows how to hold tools. It does not assume the reader can articulate what happened to them, even to themselves. It assumes only that the reader has primary evidence of having been there, and that they are willing to do the work the platform's discipline requires.

4.2 The default failure mode

The target-user guides are not protecting against over-trust in the instrument's compression. The reader will not have used the instrument long enough to develop familiarity-driven contempt for it. The guides are protecting against a different category of cognitive failure: self-deception under sustained adversity.

This failure mode operates as follows. A target user under prolonged institutional pressure has, by the time they reach a reconstruction tool, often developed cognitive habits that protect them in adversarial environments: simplification, narrative coherence, premature closure, the urge to make the matter end.

These habits are functional in survival contexts. They are catastrophic in evidential reconstruction. A target user who applies them while authoring will produce a record that feels coherent — and that will fail under hostile reading.

The guides protect against this by repeatedly naming the temptation, explicitly authorising the user to pause, and repeatedly stating that an empty section is far safer than a contaminated one. The Pre-Method guide states the rule of its phase as Nothing is learned here; the Discipline guide states the system as something that will record dishonesty faithfully rather than save the user from it. These are not soft instructions. They are precise calibrations against the specific failure mode the reader is most likely to produce.

4.3 The corpus structure

The target-user guides are organised around the workflow the user is being asked to perform, with the most load-bearing disciplinary documents read first.

Guide	Function
Before You Use These Tools (Discipline)	Pre-emptive warning that the tools require honesty, not cleverness; explicit permission to pause; repeated naming of the cognitive states under which contamination is most likely.
Pre-Method Reconstruction	The phase of rebuilding material reality before any attempt to explain it. Reconstruction is not analysis; nothing is learned here.
The Method	Nine sequential phases: fragmentation over fabrication, total collation, atomisation, normalisation, single timeline, holding contradictions, separating record from meaning, applying the system's own standards, stopping when the map is complete.
Chronology Construction	The mechanical work of building the chronology surface: how to atomise records, how to anchor entries, how to handle silence and gaps.
Claims Matrix Guidance	Connecting specific claims to specific facts, specific law, and specific parties; the discipline of stopping when the connection cannot yet be made cleanly.
Heads of Claim Guidance	Labels for things that already stand on their own; suspicion be-

Guide	Function
	longs in notes, not in claims.
Interface Guides	How to use the validator, the chronology search, the claims matrix and heads of claim editors, and the appendices interface.
Release Checklist	The threshold conditions for the work to leave the user's control.

4.4 Register and rhetorical discipline

The corpus's register is direct, unornamented, and explicitly protective. It does not pretend the work is easy. It does not pretend the reader is in a neutral cognitive state. It names — repeatedly, in different registers — the specific conditions under which contamination is most likely (exhaustion, anger, fear, desperation for relief) and explicitly authorises the reader to stop. Pause is not failure. Pause is judgment and discipline.

The corpus's opening epigraph names the structural ground: The one thing they cannot take is your memory of what happened, when the documents support your recollection, the one thing they cannot argue with is time. The discipline document's opening sentence names the difficulty: This system is not hard because it is technical. It is hard because it requires honesty. That honesty is not about being clever, persuasive, or legally trained. It is about not lying to yourself when you are tired, angry, frightened, or desperate for relief.

Where the State User corpus says nothing about why the reader might fail, because professional failure modes do not need explanation to professional readers, the target-user guides repeatedly name the cognitive states under which the reader is most likely to fail. This is not softness. It is precision. The reader cannot be assumed to have been trained to recognise their own contamination risk; the corpus must therefore name it.

5. Shared Architecture and Differing Calibration

The two corpora share the underlying analytical architecture of the platform. They differ entirely in the calibration of their disciplinary surfaces. The shared elements and the differing calibrations are tabulated below.

5.1 What both corpora share

Shared Element	Function
The chronology engine	Single canonical temporal surface, hash-navigable, year-partitioned or contiguous.
The marker taxonomy	23 markers across six lifecycle stages — universal definitions, public-corpus examples.
The CHIP/ID tri-axis	Domain (colour), class (number), seed tag (letters) on every backward jump.
The footer–chip duality	Footer holds the full provenance stack at the target; chips on incoming jumps carry the path-relevant subset.
Primary-evidence anchoring	All findings descend through nodes to primary documents; chips are compressions above the ground.
Silence as signal	Absences are recorded; voids are not filled by inference.
Hold before release	Findings remain with the operator until ready; release is a positive act, not a deadline-driven one.

5.2 Where the corpora differ

Calibration Axis	State User Corpus	Target-User Guides
Reader presupposition	Knows how to use tools; trained operator.	Does not know — and does not yet understand — what happened to them; tools are alien.
Default failure mode	Over-trust in the instrument's compression; familiarity-driven contempt.	Self-deception under sustained adversity; premature coherence.
Register	Austere; declarative; standing instructions issued by the operator to themselves.	Direct; explicitly protective; the cognitive states of failure are repeatedly named.

Calibration Axis	State User Corpus	Target-User Guides
Discipline mechanism	Three short axioms presumed sufficient once stated.	Repeated permission to pause; explicit authorisation of empty over contaminated.
Length pattern	Compact. Each guide stands alone; nothing repeated unnecessarily.	Repetition by design. The same discipline restated in multiple registers across multiple guides.
Treatment of pause	Implicit. Holding before release is treated as standing discipline, not as a permission.	Explicit. Pause is named, granted, and renamed across multiple guides as judgment and discipline.
Voice	Architect-to-operator. The corpus is the manual handed with the instrument.	Survivor-to-survivor. The author has been the target user, and the voice carries that knowledge.
Failure mode it does not address	Self-deception under adversity; this reader is presumed scaffolded.	Familiarity-driven contempt; this reader will not yet have used the instrument long enough to develop it.

6. Contamination-Prevention Surfaces

The two corpora protect against contamination at different points in the platform’s lifecycle. Considered together, they form the publicly-disclosable contamination-prevention surface across the path from authoring to registry-level reading.

6.1 The target-user guides protect the authoring surface

Authoring is where contamination is most likely to enter the platform. The target user — the affected individual reconstructing their own case — sits at the point of original input. Anything they author with insufficient discipline propagates downstream into the registry and from there into the State analyst’s reading. A contamination introduced at authoring is structurally embedded; later disciplinary surfaces can detect it but cannot easily remove it.

The target-user guides are calibrated to this risk. The repeated emphasis on pause, on the safety of an empty section over a dishonest one, on the system’s neutrality with respect to the user’s narrative, and on the explicit naming of the cognitive states under which contamination is most likely — all of these are

protections at the input boundary. The guides are not asking the user to be a better person under pressure. They are providing the user with the structural means of recognising their own pressure and stopping before the pressure becomes contamination.

6.2 The State User corpus protects the reading surface

Reading is where contamination is most likely to be propagated rather than introduced. The State analyst, working at registry level, is making findings that may inform parliamentary audiences, commissioning bodies, or public records. Anything they assert with insufficient discipline becomes part of those records. The State User corpus is calibrated to this risk. Its axioms — silence is signal, familiarity breeds contempt, accept nothing/believe no one/challenge everything — are protections at the output boundary.

Where the target-user guides protect against contamination of the record by the affected individual's own desperation, the State User corpus protects against contamination of the registry by the analyst's own competence. These are different protections against different vulnerabilities at different points in the chain. Both are necessary. Neither is sufficient on its own.

6.3 The seed authoring tool sits between

The seed authoring tool — the kernel-operated surface for assigning seed letter strings to chronology nodes — is the operational interface between the two corpora. It is held by an authoring kernel (typically a target user authoring at full kernel discipline, or a State analyst authoring against a registry case) and produces output that may be read by either class of operator. The tool's design is responsive to both: its anti-cheating discipline (auto-comma, sort, dedupe, no possibility of typing junk) is precisely the kind of mechanical contamination prevention that supports a target user's honesty under pressure, while its compression-aware grammar (CHIP/ID tri-axis, footer–chip duality) is precisely the kind of compression that the State User corpus is calibrated to read with discipline.

This is the structural answer to the question of why the two corpora exist as a pair. They are not parallel. They are sequential surfaces along the path from authoring to reading, each one calibrated to the failure mode most likely at its position in the chain. The seed authoring tool, sitting between them, is a single instrument whose output crosses both surfaces.

7. Theoretical Positioning

The dual-corpus design occupies a position that existing operator-documentation literature does not address. A short orientation against analogues:

Forensic platform documentation (Autopsy, EnCase, FTK, Cellebrite) is documented for a single trained operator profile. Target-user authoring is not contemplated; the platforms are not built for it.

Public-inquiry documentation (Post Office Horizon IT Inquiry, Grenfell Tower Inquiry) is authored for legally trained counsel and instructed experts. Affected individuals are protected from authoring by the inquiry's own apparatus, which assembles and curates evidence on their behalf.

Whistleblower platform documentation (EQS Integrity Line, NAVEX One, AllVoices) addresses the disclosure surface, not a structural reconstruction surface. The user is reporting; they are not authoring a forensic record.

Self-help legal materials (litigant-in-person guidance, advocacy charity workbooks) address procedural navigation rather than evidential discipline, and are not paired with an analytical instrument that downstream State readers will use to assess the resulting record.

The PW dual-corpus design sits at an intersection none of these occupy. It is operator documentation for a single instrument that is held by both the affected individual and the State analyst — at different moments in the lifecycle of the same case — with each holder's discipline protected against the failure mode most likely at their point of contact. Its closest theoretical kin is the literature on chain of custody in forensic disciplines, which similarly distinguishes between the discipline imposed at evidence collection and the discipline imposed at evidence interpretation. The dual-corpus design extends that principle from evidence handling to operator documentation.

8. Conclusion

The two corpora are not stylistic variants. They are differing operational disciplines calibrated to two structurally different reader profiles, addressing two structurally different default failure modes, protecting against contamination at two structurally different points in the platform's lifecycle. Each corpus is internally consistent. Together they form a complete contamination-prevention surface across the path from authoring to registry-level reading.

The State User corpus does what only it can do: it briefs a trained operator on the specific failure modes this instrument introduces, in a register that does not condescend to the operator's existing competence. The target-user guides do what only they can do: they author the discipline alongside the discipline's prerequisites, in a register that does not abandon the reader to navigate adversity unaccompanied.

Neither corpus pretends to be universal. Each is honest about what it presupposes, what it is for, and what it does not address. Each acknowledges the existence of the other in its opening orientation. This

honesty is itself a structural feature of the platform: the same principle that requires provenance to be visible in the chronology requires the documentation to be honest about who it speaks to.

The dual-corpus design is not a presentational choice. It is a structural achievement. It is rare in evidence-platform documentation. As surveyed in the companion Comparative Assessment, no reviewed comparator system addresses the asymmetric reader problem at all, let alone resolves it.

Nothing reviewed does this. Nothing reviewed attempts it.

Appendix A — Documents Reviewed

This companion paper is informed by the following PW Data Systems documentation corpora, which were reviewed in their entirety:

State User Corpus — Orientation; Axioms and Discipline; The Method; Reading the Grammar; Reading a Registry; What the Instrument Cannot Do; Register and the Downstream Reader; Corpus Governance; Release Discipline. Nine guides, single shared stylesheet, governance register.

Target-User Guides — Before You Use These Tools (Discipline); Pre-Method Reconstruction; The Method; Chronology Construction; Claims Matrix Guidance; Heads of Claim Guidance; Validator, Chronology Search, Claims Matrix, Heads of Claim, and Appendices Interface Guides; Release Checklist; System Analysis; System Ontology; Not Ready (pause-permission document).

Companion documentation — PW Data Systems v2.3x Comparative Assessment, against which this paper is intended to sit alongside as the operator-documentation analogue of the platform-architecture analysis there presented.

Assessment based on review of: full text of both corpora as authored at v1.0 (State User) and v1.7.x (target-user); the seed authoring tool surface and its operator guide; chronology runtime including chrono-core, chrono-jump-builder, legend, and chronology.html; comparator review of operator documentation for Autopsy, KronoGraph, Cellebrite Pathfinder, EnCase, FTK, EQS Integrity Line, NAVEX One, AllVoices, Post Office Horizon IT Inquiry public infrastructure, Grenfell Tower Inquiry public infrastructure, and selected litigant-in-person guidance materials.