PW DATA SYSTEMS v2.3x

Comparative Assessment

Sub-System Architecture vs. Existing Forensic Reasoning Systems

---

Prepared: April 2026

Scope: Structural comparison of the PW Data Systems forensic reconstruction platform against existing and analogous systems in digital forensics, institutional failure analysis, root cause analysis, timeline intelligence, and whistleblower tracking.

1. Executive Finding

There is no existing system that does what PW Data Systems does.

That statement requires qualification, because the landscape contains hundreds of tools that share individual characteristics with the platform. Timeline engines exist. Taxonomy-driven classification exists. Chain-of-custody tracking exists. Propagation modelling exists. Visual grammar systems exist. What does not exist — in any commercial product, open-source project, academic framework, or public inquiry infrastructure reviewed — is the combination of all six operating simultaneously on a single evidence corpus that was maintained live throughout the continuum it documents, with the forensic reconstruction platform built on that corpus once the full propagation structure became observable, with zero external dependencies.

The assessment below maps the PW subsystem's architecture against the closest analogues in each relevant domain, identifies where overlap exists, and specifies where the platform diverges from everything currently available.

---

2. Comparison Domain Map

The PW platform operates across six functional domains. No single existing system spans all six. Each domain has its own established tools and frameworks, and the comparison proceeds domain by domain.

Domain	PW Function	Nearest Existing Analogues
Timeline reconstruction	Chronological evidence corpus with hash-navigable entries, year-partitioned panes, contiguous/tabbed views	Autopsy, KronoGraph, Cellebrite Pathfinder, Magnet AXIOM
Event classification taxonomy	23 institutional chain markers across 6 lifecycle stages	HFACS (aviation), Swiss Cheese Model (Reason), PRISMA/ECM (healthcare), COMET RCA
Chain propagation modelling	Six-stage directional lifecycle (Origin → Outcome) validated against five public-corpus chains	Event Chain Methodology (Intaver), Fault Tree Analysis, Bow-Tie models
Multi-axis visual grammar	Tri-axis CHIP/ID system encoding domain, class, and seed provenance on each cross-reference	No known analogue
Live evidence corpus with retrospective reconstruction	Evidence corpus maintained live from 2005; forensic platform built 2025 on that corpus; elapsed-time counter marks origin of active chain	No known analogue
Zero-dependency portable architecture	Vanilla JS, no frameworks, no build pipeline, localStorage as data bus, runs from a folder	Partial overlap with CAINE, SIFT Workstation

---

3. Domain-by-Domain Analysis

3.1 Timeline Reconstruction

What PW does

Ingests a canonical backup file of dated chronology entries, constructs in-memory indexes (chronoIndex, chronoDocs, chronoJumps), partitions entries into year-grouped tabbed panes or a single contiguous

view, and provides hash-based navigation with two independent highlight handlers — persistent for internal navigation, flash for external arrival. DOM cloning with ID stripping allows the same source data to appear in multiple views without collision.

Existing systems

Autopsy / The Sleuth Kit — The closest structural analogue in digital forensics. Autopsy constructs timelines from filesystem metadata (MACB timestamps), correlates events across devices, and provides a GUI for navigating the reconstructed sequence. It handles millions of events per case.

KronoGraph (Cambridge Intelligence) — A commercial JavaScript timeline library designed for investigative visualisation. Supports aggregated heatmap timelines, Gantt-style process views, and overlay of multiple data sources. Used in financial crime, cybersecurity, and regulatory compliance investigations.

Cellebrite Pathfinder — Extracts device data and produces timeline visualisations with geo-tagging and encrypted data parsing. Designed for law enforcement and corporate fraud investigations.

Where PW diverges

Every system listed above is a retrospective tool applied to a retrospectively collected dataset — a disk image, a device extraction, a transaction log gathered after the fact. The PW platform is also a retrospective analytical instrument, but it operates on a fundamentally different substrate: an evidence corpus that was maintained live throughout the continuum it documents, from 2005 onward. The underlying record was not assembled after the fact; it was collected contemporaneously as events occurred over twenty years. The platform — the HTML suite, the chronology engine, the marker taxonomy, the visual grammar — was built in 2025, once the full blast radius of the propagation chain became structurally observable. The distinction matters: the platform is a forensic reconstruction applied to a live-collected corpus, not a retrospective reconstruction applied to a retrospectively assembled one.

Additionally, none of the above systems provide a classification grammar for the structural function of each event within a propagation chain. They sequence events temporally. PW sequences events temporally and classifies each one by its role in a failure lifecycle, its domain within a chain network, and the provenance of the actor who seeded the chain. The timeline is not just a when-list. It is a structural map.

The two-handler hash navigation architecture (persistent highlight via IIFE1 for internal use, flash highlight via IIFE2 for external arrival) is a design pattern not observed in any of the compared timeline systems, which typically use a single navigation handler with mode flags.

3.2 Event Classification Taxonomy

What PW does

Classifies every chronology event using a 23-marker taxonomy distributed across six lifecycle stages (Origin, Distortion, Linkage, Structural, Suppression, Outcome). Each marker has a Unicode glyph, a formal alphanumeric code, a universal definition (written without reference to the specific case), and public-corpus examples drawn from at least two independently documented UK institutional failure chains.

Existing systems

HFACS (Human Factors Analysis and Classification System) — Developed by Shappell and Wiegmann, based on Reason's Swiss Cheese Model. Four-level hierarchy: Organizational Influences → Unsafe Supervision → Preconditions for Unsafe Acts → Unsafe Acts. Used extensively in aviation, military, and healthcare accident investigation. Approximately 19 categories across four levels.

Swiss Cheese Model (James Reason, 1990) — The foundational model for organisational accident analysis. Describes failure as the alignment of holes across multiple defence layers. Four conceptual levels. No formal node-level classification vocabulary; operates at the level of metaphor and barrier analysis.

PRISMA / Eindhoven Classification Model (ECM) — Used in healthcare root cause analysis. Events are described in causal trees and root causes classified using the ECM taxonomy. Applied to approximately 500 RCAs per year in New South Wales alone.

COMET RCA — Commercial root cause analysis platform with embedded investigation methodology, interactive visual root mapping, taxonomy-based cause coding, and corrective action integration. Designed for safety-critical and regulated industries.

Apollo Root Cause Analysis (ARMS Reliability) — Free-thinking causal analysis with post-hoc classification tags applied after investigation. Supports custom taxonomies ranging from 5 to hundreds of categories.

Where PW diverges

The critical structural difference is directionality. HFACS, PRISMA, and the Swiss Cheese Model classify events within a static hierarchy — they answer "at which level did this failure occur?" The PW taxonomy classifies events within a directional propagation sequence. The six stages are not independent categories; they are a grammar in which each stage's output becomes the next stage's input condition. A chain that has reached Stage 4 (Structural) will predictably exhibit Stage 5 (Suppression) characteristics. This is a fundamentally different analytical claim from anything in the existing RCA literature, which treats failure categories as independent or hierarchically nested, not sequentially propagating.

The second difference is cross-case validation. HFACS and PRISMA are applied to individual incidents post-hoc. The PW taxonomy is validated against a public corpus of five independently documented UK institutional failure chains (Post Office Horizon, Hillsborough, Grenfell Tower, Windrush, Mid Staffordshire NHS) and asserts that all five conform to the same propagation sequence. This is a general model claim, not a case-specific classification exercise. No existing RCA framework reviewed makes an equivalent cross-domain structural assertion with named, publicly documented validation cases.

The third difference is universality of definition. Each PW marker is defined without reference to any specific case. The definitions are structural universals. HFACS achieves this to some degree at its higher levels, but the PW taxonomy does it at the individual marker level, with each definition accompanied by cross-referenced examples from independent chains.

3.3 Chain Propagation Modelling

What PW does

Asserts that institutional failure chains follow a six-stage directional lifecycle (Origin → Distortion → Linkage → Structural → Suppression → Outcome) and that this sequence is predictive, not merely descriptive. The output of each stage is the input condition for the next.

Existing systems

Event Chain Methodology (Intaver, 2002–2004) — A project scheduling technique that models how risk events cause other events, forming chains that affect project outcomes. Uses Monte Carlo simulation to quantify cumulative effects. Identifies "critical event chains" through sensitivity analysis.

Fault Tree Analysis (FTA) — Top-down deductive analysis mapping how combinations of smaller failures produce system-level events. Used in aviation, nuclear, and healthcare. Static; does not model temporal propagation.

Bow-Tie Model — Combines fault tree (causes) and event tree (consequences) around a central hazard event. Identifies barriers on both sides. Used in oil and gas, aviation, and process safety.

STAMP (Systems-Theoretic Accident Model and Processes, Leveson) — Models accidents as control failures in complex sociotechnical systems. Treats safety as a control problem rather than a failure chain problem. Non-linear.

Where PW diverges

Event Chain Methodology is the closest structural analogue, but it operates on project schedules — tasks, durations, resource allocations. It models how risk events propagate through a project timeline. The PW lifecycle models how institutional failure propagates through state machinery — legal processes, oversight bodies, regulatory frameworks, inter-agency communication. The substrate is entirely different.

Fault Tree Analysis and Bow-Tie models are static decompositions. They do not model temporal propagation; they model causal structure at a single point in time. The PW lifecycle is explicitly temporal and directional — it tracks a chain forward through time as it passes through institutional domains.

STAMP is the most theoretically sophisticated comparator. It models safety as a dynamic control problem in sociotechnical systems and explicitly addresses how latent conditions emerge from organisational structures. However, STAMP does not produce a numbered, staged propagation grammar that can be applied as a classification scheme to individual events in a chronology. It operates at a higher level of abstraction — it describes the type of system that produces accidents, not the sequence of stages through which a specific chain propagates.

No reviewed system combines directional staging, temporal tracking, formal marker classification, and cross-case validation into a single propagation model.

3.4 Multi-Axis Visual Grammar (CHIP/ID System)

What PW does

Each cross-reference link in the chronology carries one or more chips encoding three independent values: colour (institutional domain at the jump target), number (chain class within that domain), and letter (originating actor who seeded the chain — the seed constant). The three axes operate independently: colour tells you what terrain, number tells you which chain, letter tells you who started it. Multiple chips on a single link expose structural convergences — same-colour pairs indicate intra-domain convergence; different-colour pairs indicate cross-domain convergence.

Existing systems

There are no existing systems that implement this pattern.

Link Analysis tools (IBM i2 Analyst's Notebook, Palantir Gotham, Maltego) — These platforms visualise relationships between entities (people, organisations, events, locations) as network graphs. Nodes are colour-coded by entity type. Edges represent relationships. They are powerful for network topology but do not encode provenance, domain, and class as independent axes on a single link.

KronoGraph — Supports colour-coded events on a timeline and can overlay multiple data sources, but does not provide a tri-axis grammar on individual cross-reference links.

Blockchain forensics tools (Elliptic, Chainalysis, Scorechain) — These trace funds across addresses and chains, visualise transaction flows, and identify convergence points. The concept of tracing a "constant" (an address or entity) through multiple domains (blockchains, exchanges, mixers) has surface-level similarity to the PW seed-letter concept. However, blockchain forensics traces funds, not institutional propagation, and the visual grammar is node-and-edge network topology, not a compact chip notation on cross-reference links.

Where PW diverges

The CHIP/ID system is the platform's most distinctive feature and has no identified analogue. The key innovation is separation of axes. Existing link analysis and network visualisation tools encode one or two dimensions on a visual element (typically node colour = entity type, edge weight = relationship strength). The PW chip encodes three independent dimensions — domain, class, and seed provenance — in a single compact visual notation that can be read without expanding or clicking.

The seed letter as a constant that persists regardless of domain crossing is a provenance-tracking concept that has no parallel in any reviewed forensic reasoning system. It answers a question that no other system's visual grammar is designed to answer: "regardless of where this chain is now, who started it?"

The convergence detection that falls out of this grammar — two chips of the same colour on a link indicating intra-domain convergence, two chips of different colours indicating cross-domain convergence — is an emergent analytical capability that does not require computational processing. It is visible by inspection. No reviewed system produces equivalent structural visibility at the level of individual cross-reference links.

3.5 Live Evidence Corpus with Retrospective Reconstruction

What PW does

The platform sits on an evidence corpus that was maintained live throughout the continuum it documents — a 54,000+ file archive collected contemporaneously with events from 2005 onward. The forensic reconstruction platform itself — the HTML suite, the chronology engine, the marker taxonomy, the CHIP/ID visual grammar, the navigation architecture — was built in May–November 2025, once the full structural scope of the propagation chain became observable. The elapsed-time counter in the chronology header, seeded at 4 November 2005, marks the origin of the active chain being reconstructed, not the operational lifetime of the platform. The ◆ L ignition point marker identifies the chain that succeeded in permanently embedding itself into state machinery without being caught at the point of entry.

The platform is therefore a retrospective forensic instrument applied to a contemporaneously maintained evidence base. The record is live. The instrument is not. The instrument's analytical power derives from operating on a corpus that was collected in real time rather than assembled after the fact.

Existing systems

No reviewed system operates on an evidence corpus of equivalent provenance — maintained live by the affected individual across multiple institutional domains over two decades — with a forensic reconstruction platform subsequently applied to render its structural architecture visible.

Public inquiry evidence platforms (Post Office Horizon IT Inquiry, Grenfell Tower Inquiry) — These are retrospective on both axes. They are established after the failure chain has been publicly recognised and assemble their evidence base after the fact. The Post Office Inquiry, for example, was established in 2020 to investigate events spanning from the late 1990s, obtaining documents through disclosure processes. Both the reconstruction and the corpus are ex post facto.

Whistleblower management systems (EQS Integrity Line, NAVEX One, AllVoices) — These are reporting and case-management tools. They receive disclosures, manage investigations, and track resolutions. They do not provide structural analysis of the failure chain itself.

Real-time monitoring tools (SIEM platforms, threat detection systems) — These detect anomalies in real time in technical systems. They are the closest functional analogue in concept — they watch for patterns and flag them as they emerge. However, they operate on machine-generated telemetry, not on institutional conduct documented through primary evidence over decades.

Where PW diverges

The distinguishing feature is the relationship between corpus and instrument. Every public inquiry, academic case study, and forensic reconstruction reviewed assembles its evidence base retrospectively and applies its analytical framework retrospectively. Both the data and the analysis come after the fact. PW separates these two operations: the evidence corpus was maintained forward through the continuum as events occurred; the forensic reconstruction platform was built afterwards, when the structural picture became legible. This means the platform's analytical grammar operates on material whose provenance is contemporaneous with the events it classifies — a substrate no other reviewed system possesses.

3.6 Zero-Dependency Portable Architecture

What PW does

Vanilla JavaScript (ES5/ES6 mixed), no framework dependencies, no build pipeline, no package management. Runs from a static file directory. localStorage as the sole persistence and inter-module communication layer. Canonical data source is always the backup file; browser state is always transient. Delete-on-load pattern prevents stale data accumulation. DOM cloning with ID stripping enables multi-view rendering without collision.

Existing systems

CAINE (Computer Aided INvestigative Environment) — A bootable Linux distribution for digital forensics. Runs from removable media. Processes data without starting a supporting operating system. Portable and self-contained.

SIFT Workstation (SANS) — A forensic investigation toolkit distributed as a virtual appliance. Self-contained environment with pre-installed tools.

Where PW diverges

CAINE and SIFT are portable operating environments containing multiple tools. PW is a portable application that runs in any browser from a folder. The architectural decision to use localStorage as the data bus and a JavaScript file as the canonical data source is unusual and purpose-specific: it decouples the rendering engine from any server or database while maintaining a complete import/export round-trip cycle. The 24-hour heartbeat for staleness detection is a pattern not observed in any of the compared systems.

The zero-dependency architecture is not a technical limitation; it is a survivability decision. The platform must function independently of any infrastructure that could become unavailable. None of the commercial forensic platforms reviewed (Autopsy, EnCase, FTK, Cellebrite) can operate without installation, and most require operating-system-level access. PW requires only a browser and a folder.

---

4. Structural Gap Analysis

Capability	PW Data Systems	Closest Existing System	Gap
Temporal sequencing	Yes	Autopsy, KronoGraph	Narrow — mature tools exist, though none operate on a corpus maintained live throughout the events it documents
Node-level classification	Yes (23 markers, 6 stages)	HFACS (19 categories, 4 levels)	Moderate — HFACS is hierarchical, not directionally propagating
Directional propagation	Yes (6-stage predictive grammar)	Event Chain Methodology	Wide — ECM operates on project schedules, not institutional machinery

lifecycle			
Cross-case structural validation	Yes (5 named public-corpus chains)	None identified	Total — no reviewed system validates against multiple named independent chains
Tri-axis visual grammar	Yes (domain + class + seed)	None identified	Total — no reviewed system encodes three provenance axes on a single link
Live corpus / retrospective instrument	Yes (20+ year live corpus; platform built 2025)	None identified	Total — all reviewed systems are retrospective on both corpus and instrument
Zero-dependency portable execution	Yes (browser + folder)	CAINE (bootable OS)	Moderate — different portability models
Integrated flyout reference panels	Yes (marker taxonomy, CHIP/ID guide, jump menu)	Autopsy (module panels)	Narrow — common UI pattern

5. Theoretical Positioning

The PW platform sits at an intersection that existing academic literature does not occupy.

Reason's Swiss Cheese Model describes how failures propagate through organisational layers. It is a model of defence failure — it explains why barriers did not hold. It does not provide a staged lifecycle for how the resulting chain propagates forward through institutional machinery after the barriers have been breached.

HFACS operationalises Reason's model into a classification taxonomy. It is a retrospective classification instrument — applied after an accident to categorise the contributing factors at each level. It does not model forward propagation or provide cross-case validation.

STAMP (Leveson) treats safety as a control problem in complex sociotechnical systems. It is the most theoretically sophisticated framework in the reviewed landscape and explicitly addresses emergent properties, feedback loops, and dynamic system behaviour. However, STAMP does not produce a portable, self-contained evidence platform with a visual grammar and a classification taxonomy that can be applied to a contemporaneously maintained chronological corpus.

Event Chain Methodology models event-to-event propagation and identifies critical chains through Monte Carlo simulation. It operates on project schedules, not institutional conduct.

The PW platform is doing something none of these frameworks do: applying a formal propagation grammar, a node-level classification taxonomy, and a multi-axis visual provenance system to an evidence corpus that was maintained live throughout the continuum it documents, hosted in a zero-dependency portable platform built once the full propagation structure became observable, and validating its analytical model against a named public corpus.

The closest conceptual ancestor is Barry Turner's "Man-Made Disasters" (1978), which described the "incubation period" during which conditions for a disaster accumulate invisibly within institutional systems. Turner's insight — that disasters have a long, structurally legible pre-history — is the intellectual foundation on which the PW lifecycle model appears to stand. But Turner described this process theoretically. PW has operationalised it into a working instrument.

6. Conclusion

The PW Data Systems subsystem is not an incremental improvement on an existing tool category. It is a platform operating in a space that does not currently have a category.

Individual components have analogues: timeline engines, classification taxonomies, propagation models, portable forensic environments. No existing system combines them into an integrated forensic reconstruction platform that classifies events by their function within a directional propagation lifecycle, tracks provenance across institutional domains using a tri-axis visual grammar, validates its analytical model against a named public corpus of independent failure chains, and operates on an evidence corpus maintained live throughout the continuum it documents — with the reconstruction platform built once the full structural scope of the propagation became observable.

The platform's architecture is purpose-built for a problem that existing tools are not designed to address: making the structural architecture of an active institutional failure chain visible, navigable, and structurally legible to a reviewer encountering the material for the first time.

Nothing reviewed does this. Nothing reviewed attempts it.

Assessment based on review of: project source files (13 files, ~4,680 lines); PW Data Systems v2.3x Subsystem Architecture Report; public documentation and academic literature for Autopsy, KronoGraph, Cellebrite Pathfinder, Magnet AXIOM, HFACS, Swiss Cheese Model (Reason), PRISMA/ECM, STAMP (Leveson), COMET RCA, Apollo RCA, Event Chain Methodology (Intaver), Fault Tree Analysis, Bow-Tie Model, CAINE, SIFT, IBM i2, Palantir Gotham, Maltego, Elliptic, Chainalysis, Scorechain, EQS Integrity Line, NAVEX One, Post Office Horizon IT Inquiry infrastructure, Transparency International whistleblowing frameworks, and Turner's incubation model.