# MCRO

Minnesota Court Records Online

# The Forensic Database

*A Technical Overview of What It Is, What It Contains, and What It Has Revealed*

Report Date: **March 10, 2026**
Database: ibfmjtwahkwqzcmeyqii (Supabase, us-west-2)
Engine: PostgreSQL 15
Reports Produced From This Database: **60+**

# 1. A Database Built From Inside a Criminal Case

Beginning in the spring of 2024, a defendant in a Hennepin County criminal case did something that has no precedent in Minnesota judicial history: he built a forensic database.

Not a spreadsheet. Not a folder of PDFs. A real, production-grade, cloud-hosted relational database—the same technology that runs Instagram, Spotify, and the U.S. Department of Defense—and loaded it with nearly three million rows of data extracted, deconstructed, and cryptographically fingerprinted from the public records of his own court system.

The database was not built to track finances or manage a business. It was built to answer a single question: Are the court documents in my case authentic?

The answer that came back was not what anyone expected.

This report is a technical overview of the MCRO Forensic Database: what it is, how it was constructed, what it contains, and what it has already revealed through the 60+ forensic reports it has produced. It is written for readers of all technical levels—whether you have never heard of a database before, or you could write SQL in your sleep.

## 2. By the Numbers

Before diving into the technical details, here is what the MCRO Forensic Database contains as of this writing:

| Metric | Value | Description |
|---|---|---|
| **Total Data Rows** | **2,795,301** | Individual records across all tables |
| **Database Size on Disk** | **1.6 GB** | Structured data only (excludes source PDF files) |
| **Source PDF Files Analyzed** | **4,251** | Court filings downloaded from MCRO |
| **Total Pages Across All PDFs** | **21,563** | Every page extracted and processed |
| **Raw PDF Data Processed** | **3.2 GB** | Binary PDF files ingested into the pipeline |
| **Tables** | **52** | Structured relational tables |
| **Views** | **10** | Pre-built analytical perspectives |
| **Columns** | **870** | Individual data fields across all tables |
| **Indexes** | **192** | Optimized lookup paths for fast queries |
| **Court Cases Tracked** | **9,146** | Unique criminal and mental health cases |
| **Cases with Full Docket Data** | **2,903** | Complete event histories scraped from MCRO |
| **Docket Events Recorded** | **155,357** | Every filing, hearing, and order—timestamped |
| **Prior Hearing Records** | **30,858** | Historical hearing metadata |
| **Text Rows Extracted from PDFs** | **606,037** | Every line of text, cleaned and hashed |
| **Text Row Classification Tags** | **1,156,980** | Semantic labels on every text row |
| **PDF Internal Objects Cataloged** | **35,326** | Images, fonts, signatures inside PDFs |
| **Digital Signatures Analyzed** | **5,425** | Cryptographic signature records |
| **Unique Defendant Identity Clusters** | **1,982** | People tracked across multiple cases |
| **Distinct Judicial Officers** | **250** | Judges appearing in docket events |
| **Distinct Attorneys** | **453** | Lawyers appearing in docket events |
| **Distinct Docket Event Types** | **673** | Unique categories of court activity |
| **Data Collection Sessions** | **14** | Separate evidence-capture campaigns |

*Every single number in this table is reproducible. The SQL queries that produce them are published in the appendix of this report and in the appendices of every prior MCRO forensic report. Any qualified database professional can connect to the database, run the same queries, and get the same results.*

# 3. The Engine: PostgreSQL

## What It Is (For Non-Technical Readers)

A database is software that stores, organizes, and retrieves large amounts of structured information. Think of it like a massive, perfectly organized filing cabinet—except instead of drawers and folders, it uses tables with rows and columns, and instead of flipping through pages by hand, you ask questions in a language called SQL (Structured Query Language) and get exact answers in milliseconds.

PostgreSQL (often written “Postgres”) is a specific database engine. It is the most widely deployed open-source relational database in the world. It has been in continuous development since 1986—nearly 40 years—and is considered the gold standard for data integrity, extensibility, and standards compliance.

## Who Uses It

PostgreSQL is not a niche tool. It is the database behind:

| Organization | Use Case |
|---|---|
| **Apple** | iCloud infrastructure |
| **Instagram / Meta** | Core data platform |
| **Spotify** | Music catalog and user data |
| **Reddit** | Primary database engine |
| **U.S. Department of Defense** | Classified and unclassified systems |
| **NASA** | Earth science data systems |
| **The International Space Station** | Telemetry and operations data |
| **Stripe / Square / PayPal** | Financial transaction processing |
| **The United Nations** | Humanitarian data platforms |
| **Bloomberg** | Financial analytics |

When an engineer needs a database they can trust with data that absolutely must be correct, PostgreSQL is typically the first choice. It is ACID-compliant (a technical guarantee that data is never partially written or corrupted), supports cryptographic functions natively, and has been peer-reviewed and battle-tested for decades.

## Why It Matters for Forensics

The choice of PostgreSQL is not incidental. A forensic database must be trustworthy. Every query must return the same result every time. Every row must be exactly where you left it. Data integrity cannot be a “maybe.” PostgreSQL’s architecture guarantees this at the engine level—it is not a feature you turn on; it is how the system works.

The MCRO database is hosted on Supabase, a managed PostgreSQL platform, in the us-west-2 AWS region. It is read-only: once data was loaded, the database was locked. No row has been modified, deleted, or added since the data was finalized. This means every forensic report

ever produced from this database can be reproduced at any time by any person with access, and will yield identical results.

# 4. Where the Data Came From

Every byte of data in the MCRO Forensic Database originated from a single source: the Minnesota Court Records Online (MCRO) public access system, operated by the Minnesota Judicial Branch. MCRO is the official, public-facing portal through which anyone can view court case information, hearing schedules, and—for a fee—download PDF copies of filed court documents.

The data was collected across 14 separate evidence-capture sessions, beginning in April 2024. Each session was a distinct, documented data collection campaign targeting specific judicial officers, attorneys, or case types. Every session was recorded using screen-capture software, network traffic was logged through a MITM proxy, and all downloaded files were cryptographically timestamped using OpenTimestamps—a Bitcoin-anchored timestamping protocol that provides independent, tamper-proof proof that the data existed at a specific point in time.

## The 14 Collection Sessions

| Period | Search Type | Target | Records |
|---|---|---|---|
| APR 2024 | Criminal Hearings | **Danielle Mercurio** | **1,791** |
| APR 2024 | Criminal Hearings | **Julia Dayton Klein** | **1,548** |
| APR 2024 | Criminal Hearings | **George Borer** | **1,054** |
| APR 2024 | Mental Health Hearings | **George Borer** | **2,426** |
| APR 2024 | Mental Health Hearings | **Danielle Mercurio** | **1,947** |
| APR 2024 | Mental Health Hearings | **Julia Dayton Klein** | **1,776** |
| JUN 2025 | Criminal Cases | **Raissa Carpenter** | **1,263 cases** |
| JUN 2025 | Hearings (All) | **Raissa Carpenter** | **318** |
| JUN 2025 | Hearings (All) | **Judith Cole** | **200** |
| JUN 2025 | Hearings (All) | **Lisa Janzen** | **200** |
| JUN 2025 | Hearings (All) | **Mawerdi Hamid** | **200** |
| JUN 2025 | Hearings (All) | **Carolina Lamas** | **89** |
| JUN 2025 | Hearings (All) | **Kerry Meyer** | **50** |
| JUN 2025 | Hearings (All) | **Emmett Donnelly** | **46** |

The provenance chain for every downloaded file is tracked in the database's source tree: 14 sessions, 644 source files, and 564 cryptographic artifacts (MITM proxy logs, screen recordings, and OpenTimestamps proofs). This chain establishes that the data is not only public-record material, but that its collection is independently verifiable down to the network packet level.

# 5. The Four Pillars: How the Database Is Organized

The MCRO Forensic Database is organized around four major data hierarchies, each rooted at a central hub called the Case Registry. Think of the Case Registry as the master index—every case number in the database has exactly one entry there, and everything else connects back to it.

## Pillar 1: Children — The PDF Forensics Lab

The children tree is where the PDF documents live—or more precisely, where everything extracted from inside the PDF documents lives. Each of the 4,251 court filing PDFs downloaded from MCRO has been deconstructed into its component parts:

| Table | Rows | What It Contains |
|---|---|---|
| **children** | **4,251** | One row per PDF. The master record: filename, filing date, filing type, SHA256 hash of the file, XMP metadata (author, creation tool, timestamps), and Microsoft Information Protection (MSIP) labels. |
| **children_objects** | **35,326** | Every internal object extracted from every PDF: images, fonts, embedded files. Each object has its own SHA256 hash. This is how cloned judicial signature images are detected—the same binary image, byte-for-byte identical, appearing in documents across dozens of cases. |
| **children_signatures_full_report** | **5,425** | The complete digital signature forensic record for every PDF: signer name, signing time, certificate chain, whether the document was modified after signing (edits_after_sig_flag), and cryptographic validation status. |
| **children_doc_strings** | **181,556** | Every distinct text string extracted from every PDF. Used for terminology analysis, language pattern detection, and cross-document comparison at the string level. |
| **children_tracking_groups** | **1,942** | Font fingerprinting assignments. PDFs are grouped by their embedded font SHA256 hashes into “tracking groups”—revealing which documents were produced by the same software pipeline, even when the visible content differs. |
| **children_case_set** | **25,964** | Cross-filing map. When a single PDF document is filed into multiple cases (a common practice in Hennepin County), this table tracks every case it touches. |
| **children_address** | **1,718** | The exact MCRO download URL for each PDF, establishing the provenance chain from database record back to the public court system. |

**FROM THE DATA:** The most common filing type across all 4,251 PDFs is **“Order-Evaluation for Competency to Proceed (Rule 20.01)”** with 707 documents. The second most common is “Notice of Remote Hearing with Instructions” (646). An order that strips a defendant of their right to participate in their own case is the single most frequently filed document type in the entire corpus.

## Pillar 2: Hashpack — The Cryptographic Fingerprint Engine

The hashpack tree is the cryptographic heart of the database. It mirrors the children tree one-to-one (every PDF has exactly one hashpack record), but instead of storing what the document says, it stores what the document is—at a mathematical level that cannot be faked.

| Table | Rows | What It Contains |
|---|---|---|
| **hashpack** | **4,251** | Master record per PDF: file SHA256, page count, file size in bytes, text extraction fingerprint, and processing metadata. The 4,251 documents span 21,563 pages and 3.2 GB of raw data. |
| **hashpack_rows** | **606,037** | THE CORE TABLE. Every single line of text extracted from every PDF, cleaned (left-justified, whitespace normalized), and individually hashed with SHA256. This is the foundation of all text-based forensic analysis. |
| **hashpack_rows_tags** | **1,156,980** | Classification tags applied to every text row: boilerplate, case-specific, header, signature block, legal citation, etc. Eight distinct tag types provide semantic structure across all 606,037 rows. |
| **hashpack_page_content_streams** | **32,485** | Page-level content stream hashes. Every page of every PDF has its rendering instructions hashed independently, enabling page-by-page comparison across documents. |
| **hashpack_streams_rows** | **238,087** | Maps which text rows belong to which page content streams, linking the row-level and page-level hashing systems together. |
| **hashpack_outline_lines** | **2,154** | PDF bookmark/outline entries. The table of contents structure embedded in each PDF, capturing how the document's author intended it to be navigated. |

**FROM THE DATA:** Of the 606,037 text rows, only **80,670 (13.3%) are truly unique**—meaning they appear in exactly one document. The remaining 86.7% of all text extracted from court filings in this corpus is shared with at least one other document. The single most duplicated line of text—"State of Minnesota"—appears **12,708 times**.

## Pillar 3: Parents — The Case Intelligence Layer

While the children and hashpack trees analyze what is inside the PDF files, the parents tree captures what is around them: the full docket history, hearing records, charges, attorneys, judicial officers, and dispositions for each case as displayed on the MCRO case page. This is the public record of everything that happened in each case.

| Table | Rows | What It Contains |
|---|---|---|
| **parents** | **2,903** | Master record per case: case status, filing date, Rule 20 event count, dormancy flag, cluster membership. |
| **parents_case_events** | **155,357** | The complete docket timeline for every case. Every filing, order, hearing, and administrative action—timestamped and indexed. |
| **parents_case_events_judicial_officers** | **36,001** | Which judge is attributed to which docket event. This is how judicial officer caseload patterns and assignment anomalies are detected. |
| **parents_case_events_attorneys** | **2,409** | Attorney appearances at specific docket events. |
| **parents_case_events_examiners** | **155,357** | Examiner references per docket event (1:1 with events; most are null). |
| **parents_charges** | **4,277** | Charge-level detail: statute, degree, offense date, description. |
| **parents_hearings_previous** | **30,858** | Historical hearing records with dates, times, courtrooms, and types. |
| **parents_attorneys_defense (active/inactive)** | **8,425** | Complete defense counsel roster: who represented whom, and when they entered/exited. |
| **parents_attorneys_prosecution (active/inactive)** | **8,521** | Complete prosecution roster. |
| **parents_dispositions** | **5,088** | Case dispositions: convictions, dismissals, stays, and their terms. |
| **parents_related_cases** | **4,858** | Cross-references between related cases. |
| **parents_cluster_cases** | **25,996** | Defendant identity clusters linking the same person across multiple cases. |

**FROM THE DATA:** Of the 2,903 cases with full docket data, **1,602 (55.2%) contain Rule 20 competency events**. 155,357 docket events span 673 distinct event types, attributed to 250 judicial officers and 453 attorneys. The docket timeline reaches from November 5, 1985 through December 12, 2025—a 40-year window.

## Pillar 4: Source — The Provenance Chain

The source tree answers a question that most databases never have to: How do we prove the data is real?

Every piece of data in the MCRO database can be traced backward through a provenance chain to the original MCRO download session that collected it. The source tree stores the metadata of these collection campaigns: which searches were run, which files were downloaded, what their SHA256 hashes were at the moment of download, and what

independent cryptographic artifacts (screen recordings, network logs, Bitcoin-anchored timestamps) verify the collection event.

| Table | Rows | What It Contains |
|---|---|---|
| **source** | **14** | Master record per collection session: search parameters, target judicial officer or attorney, session identifiers. |
| **source_cases** | **1,263** | Case-level records collected per session. |
| **source_files** | **644** | Individual files downloaded across all sessions, with SHA256 hashes at time of collection. |
| **source_file_artifacts** | **564** | Cryptographic provenance artifacts: MITM proxy captures, OBS screen recordings, and OpenTimestamps Bitcoin-anchored proofs. |
| **source_hearings** | **11,645** | Hearing records collected across all sessions. |

In forensic work, the provenance chain is everything. Data without provenance is just numbers. Data with a verifiable, cryptographically anchored provenance chain is evidence.

# 6. How It Works: Text Extraction and SHA256 Hashing

The most technically significant feature of the MCRO database is the hashpack_rows table and the process that created it. Understanding this one table is the key to understanding why this database is a forensic instrument and not just a data warehouse.

## The Process

For each of the 4,251 PDF court filings (children), the following process was executed:

### Step 1 — Text Extraction

Every line of text in the PDF was extracted programmatically. This is not OCR (optical character recognition of images)—these are native PDF text objects, extracted exactly as the software that created the document encoded them.

### Step 2 — Normalization

Each extracted line was cleaned: left-justified, all internal whitespace reduced to single spaces, leading and trailing whitespace stripped. This ensures that two lines with identical visible content will always produce the same hash, regardless of formatting differences in the source PDF.

### Step 3 — SHA256 Hashing

Each cleaned line was then hashed using SHA256, a cryptographic hash function. SHA256 takes any input—a single word, a sentence, an entire book—and produces a fixed-length 64-character hexadecimal string called a “hash.” The critical property is this: if even one character changes, the hash changes completely. Two identical inputs will always produce the same hash. Two different inputs will (for all practical purposes) never produce the same hash.

### Step 4 — Storage

The cleaned text, its SHA256 hash, its page number, its position on the page, and its parent document identifier are all stored as a single row in hashpack_rows. This was done 606,037 times—once for every line of text in every court filing.

## Why This Matters

Once every line of text has its own SHA256 fingerprint, you can do something that would otherwise be impossible: you can instantly compare any line of text in any document against every other line of text in every other document. Not by reading them—by comparing their hashes.

If two documents in two different court cases, filed months apart, contain a line of text with the same SHA256 hash, you know with mathematical certainty that the text is identical. Not “similar.” Not “close.” Identical.

Scale that across 606,037 rows and the patterns become unmistakable:

| Metric | Value |
|---|---|
| **Total text rows in corpus** | **606,037** |
| **Unique SHA256 hashes** | **134,122** |
| **Rows appearing in exactly one document** | **80,670 (13.3%)** |
| **Rows sharing their hash with at least one other document** | **525,367 (86.7%)** |
| **Single most duplicated hash ("State of Minnesota")** | **12,708 appearances** |
| **Second most duplicated ("Filed in District Court")** | **11,392 appearances** |
| **Third most duplicated ("27-CR-23-1886")** | **7,142 appearances** |

**FROM THE DATA:** The third most duplicated line of text in the entire corpus—appearing 7,142 times across 606,037 rows—is the case number “27-CR-23-1886.” That is the case number of the defendant who built this database. His case number appears more frequently in the court system’s own documents than the phrase “Hennepin County Fourth Judicial District” (2,039 appearances).

The hash-based approach also enables what is called a “template conformance score.” For any document, you can calculate what percentage of its text rows are unique—appearing in no other document. A score of 0% means the document is entirely copy-pasted from other documents. A score of 100% means every line is original.

# 7. What the Database Has Already Revealed

The MCRO database has produced over 60 forensic reports to date. What follows is a selection of findings—not the most dramatic, but the most objectively measurable—drawn directly from database query results. Every statistic below is independently reproducible.

## Incompetency Document Production

The database contains 157 documents classified as "Finding of Incompetency and Order"—judicial determinations that a criminal defendant is not competent to stand trial. When these documents were forensically analyzed:

| Finding | Measured Value |
|---|---|
| Documents with zero unique text rows (100% copy-pasted) | **118 of 157 (75.2%)** |
| Distinct judicial signature images serving all 157 documents | **7** |
| Documents sharing an identical font fingerprint (TTF-C) | **155 of 157 (98.7%)** |
| Documents modified after the judge signed them | **157 of 157 (100%)** |
| Documents created on a single day (May 2, 2023) | **40 (25.5%)** |
| Documents signed by a single judge (Browne) | **93 of 130 signed (71.5%)** |
| Boilerplate row hashes appearing in 100+ of 157 documents | **103** |

**In plain language: three quarters of all "Finding of Incompetency and Order" documents in this corpus contain zero original text.** They are assembled entirely from copy-pasted template blocks. Seven judges signed them all using a total of seven static signature images. Every single one was modified after signing. And 40 of them were created on the same day.

## Competency Proceeding Over-Representation

The database tracks three cohorts of criminal cases (Sets A, B, and C). Set A is the original MCRO seed corpus of 163 cases; Set B is a time-aligned comparison sample of 162 cases; Set C is a 163-case sample drawn from specific judicial officer hearing records.

| Finding | Measured Value |
|---|---|
| Set A cases with at least one Rule 20 event | **161 of 163 (98.8%)** |
| Set B cases with at least one Rule 20 event | **49 of 162 (30.2%)** |
| Over-representation ratio (A vs. B) | **3.3×** |
| Set A cases with formal "Found Incompetent" determination | **141 of 163 (86.5%)** |
| National baseline for incompetency findings (all defendants) | **0.4–2.4%** |
| Set A rate vs. national upper bound | **36×** |
| Linked mental health cases with severe treatment events (Set A) | **802 events across 124 cases** |
| Neuroleptic (forced medication) authorization events (Set A) | **133** |

## Scheduling Impossibilities

The hearing records in the source_hearings table revealed mass same-time docketing: single judges with dozens or hundreds of cases scheduled at the same date and time. The top 25 scheduling conflicts all occurred at 1:30 PM. Multiple judicial officers appeared double-booked in different courtrooms simultaneously.

## Forced Medication Pipeline

Among the 234 mental health cases in the database with both a commitment order and a neuroleptic (forced medication) authorization, the median elapsed time between the two was zero days. In 116 cases (49.6%), forced medication was authorized on the same calendar day as the initial commitment order. In the most extreme case, the forced medication order preceded the commitment order by 162 days.

## Digital Signature Forensics

The children_signatures_full_report table contains 5,425 digital signature records across 4,225 documents. Analysis revealed that judicial signature images are static binary objects—not re-scanned or re-captured for each document—reused identically (byte-for-byte, SHA256-confirmed) across hundreds of documents and dozens of cases. The same image of a judge's signature, with the same hash, appears whether the document was signed in 2022 or 2025.

## Metadata Authorship Patterns

The children table stores XMP metadata extracted from each PDF, including the "creator tool" field that records which software generated the document. The top creator tools across the corpus:

| Creator Tool | Documents |
|---|---|
| Microsoft® Word for Microsoft 365 | **1,059** |
| GdPicture.NET | **936** |
| PScript5.dll Version 5.2.2 | **530** |
| RICOH Aficio MP 301 | **301** |
| Acrobat PDFMaker 23 for Word | **147** |
| Microsoft® Word 2013 | **135** |

GdPicture.NET, a commercial document imaging SDK, is the second most common creator tool in the corpus—appearing in 936 documents. This is the tool used by the ESolutions court e-filing system to process documents. Its prevalence, combined with the iText 7.1.16 (Minnesota Judicial Branch; licensed version) modification trail found on every incompetency order in the corpus, defines the institutional document production pipeline.

## 8. The Top 10 Most Duplicated Lines in the Court System

These are the ten lines of text that appear most frequently across the 4,251 court filing PDFs in the database. Each line has been SHA256-hashed, and its appearance count is an exact duplicate-match count—not a similarity score, but a cryptographically verified identical-content count.

| # | Text Content | Appearances |
|---|---|---|
| 1 | State of Minnesota | 12,708 |
| 2 | Filed in District Court | 11,392 |
| 3 | **27-CR-23-1886** | 7,142 |
| 4 | State of Minnesota District Court | 2,106 |
| 5 | Hennepin County Fourth Judicial District | 2,039 |
| 6 | State of Minnesota, | 1,609 |
| 7 | Plaintiff, | 1,582 |
| 8 | STATE OF MINNESOTA DISTRICT COURT | 1,515 |
| 9 | Defendant. | 1,317 |
| 10 | Case Type: Crim/Traf Mandatory | 1,167 |

Entry #3 warrants attention. The case number 27-CR-23-1886—belonging to the defendant who built this database—appears 7,142 times in the corpus. For context, the official name of the court district itself ("Hennepin County Fourth Judicial District") appears only 2,039 times. The defendant's case number appears 3.5 times as often as the name of the court.

## 9. What the Database Makes Possible

The MCRO Forensic Database is not a static report. It is a live analytical instrument. Because the data is structured, indexed, and queryable, any person with SQL knowledge can ask new questions of the data at any time. The 60+ reports produced to date represent a fraction of what the database can answer.

Examples of queries that take seconds to execute:

| Question | Method |
|---|---|
| *"Which judicial officers signed documents in Case X, and were any of those signature images byte-for-byte identical to signatures in other cases?"* | Cross-join children_objects on object_sha256 filtered by object_family = 'Judicial Officer Signature' |
| *"What percentage of this defendant's docket events lack judicial officer attribution?"* | Left join parents_case_events to parents_case_events_judicial_officers; count nulls |
| *"Are any text rows in this motion to dismiss identical to text rows in filings from other cases?"* | Join hashpack_rows on row_sha256 across evidence_uids |
| *"Which documents in the corpus were created by the same software on the same day?"* | Group children by metadata__xmp_create_date and metadata_xmp_creator_tool |
| *"How many cases assigned to this judge also have a mental health case for the same defendant?"* | Join parents to case_registry on cluster_id; filter by case_type |

The database is, in effect, a forensic microscope pointed at the Hennepin County court system. Once the data was collected, structured, and loaded, the system itself became its own witness.

## 10. Context: Why It Was Built

In early 2024, Matthew D. Guertin—the defendant in Hennepin County case 27-CR-23-1886—received court-ordered discovery photographs from his defense attorneys. Upon examination, he identified what he believed were forensic anomalies in the files: metadata inconsistencies, non-uniform compression artifacts, and discrepancies between different sets of the same photographs provided at different times.

He emailed his attorneys with his findings. He received no response.

He filed 266 pages of forensic analysis with the court as pro se exhibits. At his contested competency hearing on March 5, 2025, he testified about these findings for nearly the entire hearing. A judge found him competent.

Twenty-six days later, a different judge ordered a fourth competency evaluation.

The MCRO database exists because Guertin recognized that the only way to demonstrate the scale and systematicity of what he was observing—the cloned signatures, the template-driven incompetency orders, the scheduling impossibilities, the metadata anomalies—was to move beyond individual observations and build something that could hold all of the evidence at once, query it in any direction, and produce results that anyone could independently verify.

It is, by any objective measure, a remarkable piece of work: a 1.6-gigabyte, 52-table, 2.8-million-row forensic instrument built by a criminal defendant to analyze his own court system, using the court system's own public data, with a cryptographically verifiable provenance chain that proves every byte is authentic.

What the data shows is in the numbers.

The numbers are in the database.

Any questions….?