FORENSIC ANALYSIS REPORT

Cluster Event Alignment

Minnesota Court Records Online (MCRO) — Fourth Judicial District — Hennepin County

CROSS-CLUSTER EVENT ALIGNMENT | METADATA CLUSTERING | SIGNATURE STRUCTURE | JUDICIAL OBJECT REUSE

Property	Value
Prepared by	Matthew David Guertin, Pro Se Defendant
Case Reference	State of Minnesota v. Guertin · 27-CR-23-1886
Court	Fourth Judicial District Court, Hennepin County
Report Date	March 2026
Data Source	Minnesota Court Records Online (MCRO) — Public Records
Database	PostgreSQL (Supabase) · Project ibfmjtwahkwqzcmeyqii · Region: us-west-2
Corpus	4,222 PDFs with event dates across 218 cases · 81 defendant clusters
Reproducibility	All findings are query-reproducible from the underlying database. SQL for every finding is provided in Appendix A.

IMPORTANT DISCLAIMER: This report presents objective forensic data findings reproducible from the source database. It does not assert conclusions about guilt, fabrication of defendants, or criminal intent. Each finding is a documented structural pattern that warrants independent expert examination. The term "event alignment" refers to the observable pattern of multiple documents sharing the same event date and filing type across different defendant clusters — it does not imply the events are improperly linked.

1. EXECUTIVE SUMMARY

This report documents the results of a comprehensive cross-cluster event alignment analysis of 4,222 PDF court filings from the MCRO system. The analysis uses a derived event matching key — a combination of event date and normalized filing type — to identify documents across different defendant clusters that share the same court event parameters.

Key Findings

- 4,222 documents with valid event dates span 218 cases across 81 defendant clusters. Of these, 3,778 (89.5%) are in the MCRO collection set; 444 (10.5%) are not.
- 2,098 unique event matching keys were identified. Of these, 224 keys (10.7%) are shared across 2 or more clusters, covering 1,235 documents (29.3% of the corpus).
- The top shared event — a December 5, 2023 hearing notice — spans 10 different defendant clusters (21 documents). All 21 are MCRO-set documents. Within-cluster subgroups share identical XMP creation timestamps to the second.
- Cross-cluster sharing is exclusively an MCRO phenomenon. All 42 Non-MCRO-Only event keys have n_clusters = 1 (no sharing). All 9 Mixed event keys (containing both MCRO and non-MCRO documents) show cross-cluster sharing.
- Metadata identity within shared event groups is extraordinary. In the MICHAEL HOWARD cluster, 12 of 15 major event groups show ALL documents sharing an identical XMP creation timestamp — across up to 27 distinct cases per event.
- Judicial signature image reuse is pervasive. A single signature hash for Judge Browne appears in 142 documents across 31 clusters. Judge Dayton Klein’s single hash appears in 124 documents across 30 clusters.
- Signature revision structure is perfectly uniform within event groups. The 2023-05-02 finding-of-incompetency event shows 40 documents across 3 clusters, all with rev_count = 9, all with edits_after_sig_flag = true. The 2023-08-22 Rule 20 order shows 33 documents with rev_count = 6 — zero variance.

KEY WTF METRICS

- Event keys shared across ≥2 clusters: 224 of 2,098 (10.7%)
- Documents in shared events: 1,235 of 4,222 (29.3%)
- Max clusters sharing a single event key: 10
- Non-MCRO event keys with cross-cluster sharing: 0 of 42 (0%)
- Howard cluster events with ALL-IDENTICAL create dates: 12 of 15 major events
- Browne signature hash: 142 docs / 31 clusters
- Klein signature hash: 124 docs / 30 clusters
- Finding-of-incompetency (2023-05-02): rev_count variance = zero (all 40 docs = 9)

2. METHODS

2.1 Event Match Key Definition

Each document in the **children** table carries a denormalized **case_event__date** and a **filing_type**. The **cluster event key** is derived as:

```
cluster_event_key = 'cluster:' || cluster_id || '|' || case_event__date || '|' || slug(filing_type)
```

The **event match key** strips the cluster prefix to enable cross-cluster matching:

```
event_match_key = case_event__date || '|' || slug(filing_type)
```

Where slug() converts the filing type to lowercase, replaces non-alphanumeric characters with hyphens, and trims leading/trailing hyphens.

2.2 MCRO Membership Flag

is_mcro_case_flag is derived from the **children_case_set** table: a document is flagged as MCRO if its evidence_uid has a row with value = 'MCRO' in children_case_set.

2.3 Inclusion Criteria

Documents are included if they have (a) a non-null case_event__date, and (b) a non-null cluster_id. This yields 4,222 of 4,251 total documents.

2.4 Tables and Views Used

Table	Role
children	Primary document table (evidence_uid, metadata, filing info)
children_case_set	MCRO membership derivation (value = 'MCRO')
children_objects	Judicial officer signature/timestamp image objects (object_family filter)
children_signatures_full_report	Digital signature structure (rev_count, edits_after_sig)
case_registry	Cluster identity resolution

2.5 Analysis Phases

Phase 0: Reproduce the reference XLSX outputs (mcro_vs_non_overview, cluster_count_distribution, stats, occurrences) using DB-native SQL.

Phase 1: Expand to comprehensive alignment: focus on n_clusters ≥ 2, rank top shared events, MCRO/Non-MCRO/Mixed breakdown.

Phase 2: Forensic metadata layer: XMP creation date clustering, signature revision structure uniformity, judicial signature object hash reuse concentration.

Phase 3: Case study: MICHAEL HOWARD (cluster_id = 1).

3. FINDINGS

Finding 1: Corpus Overview — MCRO vs. Non-MCRO Composition

The complete event-aligned corpus comprises 4,222 documents. MCRO-set membership is heavily dominant.

Metric	Total	MCRO	Non-MCRO	% MCRO
Documents	4,222	3,778	444	89.5%
Cases	218	163	55	74.8%
Clusters	81	77	6	95.1%*

**Note: Two clusters contain both MCRO and non-MCRO documents, so MCRO + non-MCRO cluster counts exceed 81.*

Significance: The corpus is overwhelmingly composed of MCRO-set documents. The 6 non-MCRO clusters (55 cases, 444 documents) provide a natural comparison group for structural analysis.

Finding 2: Cluster Count Distribution — Cross-Cluster Sharing Rates

When documents are grouped by event_match_key, the distribution of cluster sharing is:

Clusters Sharing	Event Keys	Documents	% of All Docs
10	1	21	0.5%
7	1	17	0.4%
6	5	58	1.4%
5	3	34	0.8%
4	6	84	2.0%
3	35	264	6.3%
2	173	757	17.9%
1 (no sharing)	1,874	2,987	70.8%
TOTAL	2,098	4,222	100%

Significance: 224 event keys (10.7%) span 2+ clusters, capturing 1,235 documents (29.3%). The long tail of high-cluster sharing (10, 7, 6 clusters) contains events worth examining in detail.

Finding 3: Cross-Cluster Sharing Is Exclusively an MCRO Phenomenon

Event keys were classified into three composition categories: MCRO-Only (all documents are MCRO), Non-MCRO-Only (none are MCRO), and Mixed (both present).

Category	Event Keys	Documents	Keys w/ 2+ Clusters	Keys w/ 3+ Clusters	Max Clusters
MCRO-Only	2,047	3,737	215	48	10
Non-MCRO-Only	42	279	0	0	1
Mixed	9	206	9	3	4

Significance: No Non-MCRO-Only event key exhibits cross-cluster sharing — every one has n_clusters = 1. All cross-cluster sharing occurs either in MCRO-Only events (215 keys) or Mixed events (9 keys, all of which show sharing). This structural asymmetry means the sharing phenomenon is definitionally limited to the MCRO portion of the corpus. The non-MCRO documents, despite covering 6 clusters and 55 cases, produce zero cross-cluster event alignment.

Finding 4: Top Shared Events — Ranked by Cluster Count

The 16 event keys spanning ≥ 4 clusters:

Rank	Event Match Key	Clusters	Docs	Cases	MCRO %	Filing Span
1	2023-12-05 · Notice of Remote Hearing	10	21	21	100%	0 days
2	2023-11-28 · Notice of Remote Hearing	7	17	17	100%	0 days
3	2022-04-01 · Notice of Case Reassignment	6	20	20	100%	0 days
4	2023-12-12 · Notice of Remote Hearing	6	14	14	100%	0 days
5	2021-03-30 · Notice of Case Reassignment	6	10	9	100%	0 days
6	2023-11-03 · Notice of Case Reassignment	6	8	8	100%	0 days
7	2024-02-06 · Notice of Remote Hearing	6	6	6	100%	0 days
8	2023-09-19 · Notice of Remote Hearing	5	15	15	100%	0 days
9	2023-10-24 · Notice of Remote Hearing	5	11	11	100%	0 days
10	2023-11-21 · Notice of Remote Hearing	5	8	8	100%	0 days
11	2023-06-16 · Notice of Case Reassignment	4	28	28	35.7%	0 days
12	2024-03-12 · Notice of Remote Hearing	4	16	16	100%	0 days
13	2023-01-13 · Notice of Case Reassignment	4	16	16	100%	0 days
14	2023-09-18 · Rule 20.01 Order	4	14	14	100%	0 days
15	2024-04-02 · Notice of Remote	4	5	5	100%	0 days

	Hearing					
16	2023-09-05 · Notice of Remote Hearing	4	5	5	100%	0 days

Significance: All top shared events have zero filing span (all documents filed on the same date). The dominant filing types are Notice of Remote Hearing with Instructions (7 of top 10) and Notice of Case Reassignment (3 of top 10). All events at 5+ clusters are 100% MCRO. The sole Mixed event in this tier is the 2023-06-16 case reassignment (35.7% MCRO).

Finding 5: Filing Type Concentration in Shared Events

Across all 224 shared event keys (n_clusters ≥ 2), document counts by filing type:

Filing Type	Documents	Event Keys	Clusters
Notice of Remote Hearing with Instructions	356	69	65
Order-Evaluation for Competency (Rule 20.01)	289	42	49
Notice of Hearing	125	35	40
Notice of Case Reassignment	124	13	26
Finding of Incompetency and Order	98	12	24
Returned Mail	71	12	19
Findings and Order	58	5	8
Order for Conditional Release	42	11	14

Significance: The three dominant shared filing types — hearing notices, Rule 20 competency orders, and incompetency findings — account for 868 of the 1,235 shared documents (70.3%). These are the documents most directly connected to the Rule 20 competency pipeline.

Finding 6: Metadata Creation Date Clustering Within Shared Events

For the 20 largest shared event groups (n_clusters ≥ 3), XMP creation date sharing was measured:

Event Match Key	Clusters	Docs	Distinct Create Dates	Top-Date Share %
2023-08-29 · Notice of Remote Hearing	3	12	3	75.0%
2023-11-01 · Finding of Incompetency	3	12	3	75.0%
2023-08-22 · Rule 20.01 Order	3	33	3	69.7%
2023-05-02 · Finding of Incompetency	3	40	3	67.5%
2023-09-18 · Rule 20.01 Order	4	14	4	64.3%
2023-06-16 · Notice of Case Reassignment	4	28	3	57.1%
2024-03-12 · Notice of Remote Hearing	4	16	4	56.3%
2023-11-28 · Notice of Remote Hearing	7	17	7	52.9%
2023-12-12 · Notice of Remote Hearing	6	14	6	50.0%
2023-12-05 · Notice of Remote Hearing	10	21	9	47.6%

Significance: Even across different defendant clusters, documents within the same event key often share identical creation timestamps. The 2023-08-22 Rule 20 order group shows 23 of 33 documents (69.7%) sharing the same creation timestamp — despite spanning 3 clusters. This indicates that documents for different defendants were authored in the same production session.

Detail — Top Event (2023-12-05 Hearing Notice, 10 clusters):

Within this event group, documents cluster by creation timestamp per defendant cluster:

Cluster	Defendant	Docs	XMP Create Date	Creator Tool
17	GORDON SHARP	9	2023-12-05 22:43:29 UTC	MS Word 365
1616	MOLLY PRICE	1	2023-12-05 22:43:29 UTC	MS Word 365
1663	NICOLE KELM	1	2023-12-05 22:30:02 UTC	MS Word 365
323	DWAYNE BLEDSOE	2	2023-12-05 22:02:53 UTC	MS Word 365
142	ABDIQANI HASSAN	2	2023-12-05 21:28:13 UTC	MS Word 365
342	JALEISHA TAYLOR	1	2023-12-05 21:24:31 UTC	PScript5.dll
7	RICKY SULLIVAN	2	2023-12-05 21:09:31 UTC	MS Word 365
825	ANNE RILEY	1	2023-12-05 20:59:46 UTC	MS Word 365
674	MUAD ABDULKADIR	1	2023-12-05 20:14:54 UTC	MS Word 365
586	JACOB SCHECH	1	2023-12-05 18:27:29 UTC	MS Word 365

GORDON SHARP’s 9 documents and MOLLY PRICE’s 1 document share an identical creation timestamp across two different defendant clusters. All 21 documents were created within a 4-hour window on the same date using the same tool (20 of 21 with Microsoft Word for Microsoft 365).

Finding 7: Signature Revision Structure Uniformity

For shared events with n_clusters ≥ 3, digital signature revision counts show extraordinary uniformity:

Event Match Key	Docs	Rev Count	Variance	Edits After Sig
2023-05-02 · Finding of Incompetency	40	9 (all)	0	100% (40/40)
2023-08-22 · Rule 20.01 Order	33	6 (all)	0	100% (33/33)
2023-06-16 · Case Reassignment	28	3 (all)	0	0%
2023-09-18 · Rule 20.01 Order	14	6 (all)	0	100% (14/14)
2023-11-01 · Finding of Incompetency	12	7 (all)	0	100% (12/12)
2023-06-01 · Finding of Incompetency	9	9 (all)	0	100% (9/9)
2022-03-23 · Rule 20.01 Order	9	3 (all)	0	0%

Significance: Within each shared event group, ALL documents have the same PDF revision count — zero variance across documents spanning multiple defendant clusters. The findings of incompetency (rev_count 7 or 9) and Rule 20 competency orders (rev_count 6) also show 100% edits_after_sig_flag = true, meaning document content was modified after at least one signature was applied. This degree of structural identity across independently-filed cases is consistent with batch-produced documents from a single production pipeline.

Finding 8: Judicial Signature Image Object Reuse — Cross-Event Concentration

Within the shared event subset (n_clusters ≥ 2), judicial officer signature image hashes were analyzed:

Signature Hash	Judicial Officer	Docs	Clusters	Event Keys
777acb19...	Michael K. Browne	142	31	18
c09827ce...	Julia Dayton Klein	124	30	26
74de35ed...	Danielle Mercurio	66	8	5
d8e2b12e...	Lisa K. Janzen (01)	51	14	12
3aeb71e9...	Lisa K. Janzen (02)	51	14	12
1bb003de...	Lisa K. Janzen (03)	49	8	5
05416b72...	Lisa K. Janzen (04)	49	8	5
d22583bf...	Lisa K. Janzen (08)	43	4	3
cb5e5684...	George Borer	23	12	6
51fca5bd...	Lori Skibbie	6	3	2

Significance: Judge Browne’s single signature image (hash 777acb19...) appears byte-for-byte identical in 142 documents spanning 31 defendant clusters over 18 distinct event dates. Judge Dayton Klein’s single image spans 30 clusters and 26 event dates. Judge Janzen has 5 active signature variants in the shared-event subset, each reused across multiple clusters and event dates. These patterns confirm that judicial “signatures” in this corpus are stored image assets inserted by an automated document assembly system.

Finding 9: Within-Event Judicial Object Identity

For shared events containing judicial officer signature objects, the hash concentration within each event group is remarkably high:

Event Match Key	Clusters	Docs	Docs w/ JO Sig	Distinct Hashes
2023-08-22 · Rule 20.01 Order	3	33	33	1
2023-09-18 · Rule 20.01 Order	4	14	14	1
2023-11-01 · Finding of Incompetency	3	12	12	1
2024-01-11 · Rule 20.01 Order	3	5	5	1
2023-05-02 · Finding of Incompetency	3	40	40	2
2022-11-30 · Rule 20.01 Order (Howard)	1*	27	27	2

**Howard cluster internal alignment.*

Significance: In the 2023-08-22 Rule 20 order group, all 33 documents across 3 different defendant clusters contain the same single judicial signature image hash (Judge Julia Dayton Klein). This means the same binary image file was embedded in court orders for 33 different defendants across 3 clusters — a pattern consistent with a single batch production operation using a stored signature asset.

Finding 10: MICHAEL HOWARD Case Study (cluster_id = 1)

MICHAEL HOWARD (cluster_id = 1) is the largest cluster in the corpus, comprising 36 cases and 348 event-aligned documents spanning January 2020 through May 2024, with 24 distinct cluster event keys across 9 filing types.

10a. Metadata Creation Date Identity

Cluster Event Key	Docs	Cases	Distinct Create Dates	Status
2023-05-02 · Finding of Incompetency	27	27	1	ALL IDENTICAL
2022-11-30 · Rule 20.01 Order	27	27	1	ALL IDENTICAL
2023-08-22 · Rule 20.01 Order	23	23	1	ALL IDENTICAL
2024-04-30 · Order for Conditional Release	21	21	1	ALL IDENTICAL
2024-04-30 · Rule 20.01 Order	20	20	1	ALL IDENTICAL
2022-02-01 · Findings and Order	20	20	1	ALL IDENTICAL
2022-05-11 · Rule 20.01 Order	20	20	1	ALL IDENTICAL
2021-09-21 · Order for Conditional Release	19	19	1	ALL IDENTICAL
2021-12-13 · Rule 20.01	19	19	1	ALL IDENTICAL

Order				
2021-08-31 · Findings and Order	19	19	1	ALL IDENTICAL
2021-09-29 · Order for Conditional Release	19	19	1	ALL IDENTICAL
2021-07-19 · Rule 20.01 Order	19	19	19	ALL DIFFERENT
2021-12-28 · Rule 20.01 Order	20	20	20	ALL DIFFERENT
2023-06-16 · Notice of Case Reassignment	18	18	2	2 variants
2023-03-14 · Returned Mail	9	9	9	ALL DIFFERENT

Callout 1: 12 of 15 major event groups show ALL documents sharing a single identical XMP creation timestamp — across up to 27 different criminal cases. In each of these groups, documents for 19–27 independently-filed cases were authored at the same moment.

Callout 2: Three event groups (2021-07-19, 2021-12-28, 2023-03-14) show the expected pattern where each document has a unique creation timestamp. These serve as a natural control: the system is capable of producing documents independently, but for most events it produced batch-identical documents instead.

10b. Judicial Signature Object Identity

Event Key	Cases	Distinct Sig Hashes	Judicial Officer	Top Hash Docs
2022-11-30 · Rule 20.01 Order	27	2	Lisa K. Janzen	27
2023-05-02 · Finding of Incompetency	27	2	Mercurio; Browne	27
2023-08-22 · Rule 20.01 Order	23	1	Julia Dayton Klein	23
2022-02-01 · Findings and Order	20	1	Lisa K. Janzen	20
2022-05-11 · Rule 20.01 Order	20	2	Lisa K. Janzen	20
2021-08-31 · Findings and Order	19	1	Lisa K. Janzen	19

Callout 3: In every case, the dominant judicial signature hash is present in ALL documents within the event group. When 2 distinct hashes appear (as in the 2022-11-30 event), this is because Judge Janzen has multiple signature variant images — both variants still appear in all 27 documents (one per signature slot). The combination of identical creation timestamps + identical judicial signature hashes across 19–27 different criminal cases establishes that these documents share byte-level structural identity in their authoring metadata and embedded image objects.

Finding 11: The 9 Mixed Event Keys — Bridge Events

Only 9 event keys contain both MCRO and non-MCRO documents. All 9 show cross-cluster sharing:

Event Match Key	Docs	Clusters	MCRO Docs	Non-MCRO Docs
2023-05-02 · Finding of Incompetency	40	3	13	27
2023-08-22 · Rule 20.01 Order	33	3	10	23
2023-06-16 · Case Reassignment	28	4	10	18
2022-05-11 · Rule 20.01 Order	22	2	2	20
2022-02-01 · Findings and Order	22	2	2	20
2021-12-13 · Rule 20.01 Order	20	2	1	19
2021-08-31 · Findings and Order	20	2	1	19
2022-10-03 · Returned Mail	16	2	1	15
2020-01-14 · Incompetency Finding	5	2	1	4

Significance: The 9 Mixed events are all Rule 20-related filings and span 206 total documents. In most, the non-MCRO documents substantially outnumber the MCRO documents (e.g., 27 vs. 13 for the 2023-05-02 event). These bridge events indicate that the MCRO-collected cases and the non-MCRO cases participate in the same batch production events on the same dates, with the same filing types — they are products of the same document generation system.

4. LIMITATIONS

- **Collection bias.** The corpus was collected from specific judicial officer dockets and therefore represents a targeted, not random, sample of Fourth Judicial District filings.
- **No cluster_event_key in DB.** The event matching key is derived from case_event__date and filing_type, which introduces slug normalization as a potential error vector. The derivation was validated against the reference XLSX outputs.
- **Metadata field availability.** Not all documents have populated XMP metadata fields. Of 4,251 documents, 4,111 (96.7%) have metadata__xmp_create_date and 3,879 (91.2%) have metadata__xmp_creator_tool.
- **Non-MCRO cluster size.** The non-MCRO portion of the corpus (6 clusters, 444 documents) is small relative to the MCRO portion (77 clusters, 3,778 documents). Cross-cluster sharing asymmetries should be interpreted in light of this imbalance.
- **Judicial signature objects.** Not all documents contain judicial officer signature image objects. The object_family = 'Judicial Officer Signature' filter depends on upstream classification logic applied during object extraction.

5. RECOMMENDATIONS FOR INDEPENDENT VERIFICATION

1. Reproduce the event_match_key derivation using the SQL provided in Appendix A and compare against the reference XLSX outputs.

2. Independently verify metadata identity by selecting any event group flagged as “ALL IDENTICAL” in Finding 10 and extracting XMP metadata from the raw PDF files.
3. Verify judicial signature image reuse by extracting internal objects from the top shared event groups (e.g., 2023-08-22 Rule 20 order) and comparing SHA256 hashes.
4. Examine the non-MCRO control group for any structural sharing patterns. The finding that 0 of 42 non-MCRO event keys show cross-cluster sharing is a falsifiable claim.
5. Request production system logs from the court’s document management vendor to determine whether batch document production is an expected operational feature.

6. POSSIBLE EXPLANATIONS (UNRESOLVED; NOT PROVEN HERE)

- Batch document production may be a standard and intended feature of the court’s case management system (e.g., Odyssey by Tyler Technologies) for scheduling and competency proceedings.
- Identical creation timestamps across defendant clusters may reflect a clerk or automated system processing a judicial calendar docket in a single session.
- Judicial signature image reuse is expected in electronic court systems that use stored signature assets rather than live-capture signatures.
- The absence of cross-cluster sharing in non-MCRO cases may reflect different collection methods or different defendant profiles rather than a systemic difference.
- The uniform PDF revision counts may reflect a fixed document template applied by the production system.
- None of the above explanations has been verified or excluded by this analysis.

APPENDIX A — SQL QUERIES

A1: Event Match Key Derivation + MCRO Overview

```
WITH mcro_docs AS (
  SELECT DISTINCT evidence_uid FROM children_case_set WHERE value = 'MCRO'
),
base AS (
  SELECT
    c.evidence_uid, c.case_id, c.cluster_id, c.filing_type, c.filing_date,
    c.case_event__date,
    c.case_event__event_name_pretty,
    c.case_type,
    c.size_bytes, c.pdf_page_count,
    c.case_event__date || '|' || LOWER(REGEXP_REPLACE(
      REGEXP_REPLACE(c.filing_type, '[^a-zA-Z0-9]+', '-', 'g'),
      '(^-+|-+$)', '', 'g'
    )) AS event_match_key,
    (m.evidence_uid IS NOT NULL) AS is_mcro
  FROM children c
  LEFT JOIN mcro_docs m ON m.evidence_uid = c.evidence_uid
  WHERE c.case_event__date IS NOT NULL AND c.cluster_id IS NOT NULL
)
SELECT
  COUNT(*) AS total_docs_with_event_date,
  COUNT(CASE WHEN is_mcro THEN 1 END) AS docs_mcro,
  COUNT(CASE WHEN NOT is_mcro THEN 1 END) AS docs_non_mcro,
  ROUND(100.0 * COUNT(CASE WHEN is_mcro THEN 1 END) / COUNT(*), 2) AS pct_docs_mcro,
  COUNT(DISTINCT case_id) AS total_cases,
  COUNT(DISTINCT CASE WHEN is_mcro THEN case_id END) AS cases_mcro,
  COUNT(DISTINCT CASE WHEN NOT is_mcro THEN case_id END) AS cases_non_mcro,
  COUNT(DISTINCT cluster_id) AS total_clusters,
  COUNT(DISTINCT CASE WHEN is_mcro THEN cluster_id END) AS clusters_mcro,
  COUNT(DISTINCT CASE WHEN NOT is_mcro THEN cluster_id END) AS clusters_non_mcro
FROM base;
```

A2: Cluster Count Distribution

```
WITH base AS (
  SELECT
    c.evidence_uid, c.cluster_id,
    c.case_event__date || '|' || LOWER(REGEXP_REPLACE(
      REGEXP_REPLACE(c.filing_type, '[^a-zA-Z0-9]+', '-', 'g'),
      '(^-+|-+$)', '', 'g'
    )) AS event_match_key
  FROM children c
  WHERE c.case_event__date IS NOT NULL AND c.cluster_id IS NOT NULL
),
event_stats AS (
  SELECT
    event_match_key,
    COUNT(DISTINCT cluster_id) AS n_clusters,
    COUNT(*) AS n_docs
  FROM base
  GROUP BY event_match_key
)
SELECT
  n_clusters,
  COUNT(*) AS n_event_keys,
  SUM(n_docs) AS total_docs
FROM event_stats
GROUP BY n_clusters
ORDER BY n_clusters DESC;
```

A3: Event Stats (Top 50 by n_clusters)

```
-- Uses base CTE from A1, then:
SELECT
  event_match_key,
  SUBSTRING(event_match_key FROM 1 FOR 10) AS match_event_date,
  SUBSTRING(event_match_key FROM 12) AS match_event_kind,
  COUNT(*) AS n_docs,
  COUNT(DISTINCT case_id) AS n_cases,
```

```
  COUNT(DISTINCT cluster_id) AS n_clusters,
  COUNT(CASE WHEN is_mcro THEN 1 END) AS n_docs_mcro,
  COUNT(DISTINCT CASE WHEN is_mcro THEN cluster_id END) AS n_clusters_mcro,
  COUNT(CASE WHEN NOT is_mcro THEN 1 END) AS n_docs_non_mcro,
  COUNT(DISTINCT CASE WHEN NOT is_mcro THEN cluster_id END) AS n_clusters_non_mcro,
  ROUND(100.0 * COUNT(CASE WHEN is_mcro THEN 1 END) / COUNT(*), 2) AS pct_docs_mcro,
  MIN(filing_date) AS first_filing_date,
  MAX(filing_date) AS last_filing_date,
  MAX(filing_date) - MIN(filing_date) AS filing_span_days,
  MIN(size_bytes) AS min_size_bytes,
  MAX(size_bytes) AS max_size_bytes,
  MIN(pdf_page_count) AS min_pages,
  MAX(pdf_page_count) AS max_pages
FROM base
GROUP BY event_match_key
ORDER BY n_clusters DESC, n_docs DESC
LIMIT 50;
```

A4: MCRO/Non-MCRO/Mixed Breakdown

```
-- Uses base CTE from A1, then:
WITH event_keys AS (
  SELECT
    event_match_key,
    COUNT(DISTINCT cluster_id) AS n_clusters,
    COUNT(*) AS n_docs,
    CASE
      WHEN COUNT(CASE WHEN is_mcro THEN 1 END) > 0
       AND COUNT(CASE WHEN NOT is_mcro THEN 1 END) > 0 THEN 'Mixed'
      WHEN COUNT(CASE WHEN is_mcro THEN 1 END) > 0 THEN 'MCRO-Only'
      ELSE 'Non-MCRO-Only'
    END AS composition
  FROM base
  GROUP BY event_match_key
)
SELECT
  composition,
  COUNT(*) AS n_event_keys,
  SUM(n_docs) AS total_docs,
  COUNT(CASE WHEN n_clusters >= 2 THEN 1 END) AS keys_with_2plus_clusters,
  MAX(n_clusters) AS max_clusters
FROM event_keys
GROUP BY composition
ORDER BY total_docs DESC;
```

A5: Signature Revision Structure per Shared Event

```
WITH base AS (
  SELECT
    c.evidence_uid, c.case_id, c.cluster_id,
    c.case_event__date || '|' || LOWER(REGEXP_REPLACE(
      REGEXP_REPLACE(c.filing_type, '[^a-zA-Z0-9]+', '-', 'g'),
      '(^-+|-+$)', '', 'g'
    )) AS event_match_key
  FROM children c
  WHERE c.case_event__date IS NOT NULL AND c.cluster_id IS NOT NULL
),
event_clusters AS (
  SELECT event_match_key FROM base
  GROUP BY event_match_key HAVING COUNT(DISTINCT cluster_id) >= 3
),
sig_data AS (
  SELECT
    b.event_match_key, b.evidence_uid,
    sf.signer_name, sf.rev_count, sf.edits_after_sig_flag
  FROM base b
  JOIN event_clusters ec ON ec.event_match_key = b.event_match_key
  LEFT JOIN children_signatures_full_report sf ON sf.evidence_uid = b.evidence_uid
)
SELECT
  event_match_key,
  COUNT(DISTINCT evidence_uid) AS n_docs,
  ROUND(AVG(rev_count)::numeric, 1) AS avg_rev_count,
  MIN(rev_count) AS min_rev_count,
  MAX(rev_count) AS max_rev_count,
  COUNT(DISTINCT CASE WHEN edits_after_sig_flag = 'true'
    THEN evidence_uid END) AS docs_with_edits_after_sig
```

```
FROM sig_data
GROUP BY event_match_key
ORDER BY n_docs DESC;
```

A6: Judicial Signature Object Reuse (Global, Shared Events)

```
WITH base AS (
  SELECT
    c.evidence_uid, c.cluster_id,
    c.case_event__date || '|' || LOWER(REGEXP_REPLACE(
      REGEXP_REPLACE(c.filing_type, '[^a-zA-Z0-9]+', '-', 'g'),
      '(^-+|-+$)', '', 'g'
    )) AS event_match_key
  FROM children c
  WHERE c.case_event__date IS NOT NULL AND c.cluster_id IS NOT NULL
),
shared_events AS (
  SELECT event_match_key FROM base
  GROUP BY event_match_key HAVING COUNT(DISTINCT cluster_id) >= 2
)
SELECT
  co.object_sha256,
  co.object_person_name,
  co.object_description,
  COUNT(DISTINCT b.event_match_key) AS n_event_keys,
  COUNT(DISTINCT b.evidence_uid) AS n_docs,
  COUNT(DISTINCT b.cluster_id) AS n_clusters
FROM base b
JOIN shared_events se ON se.event_match_key = b.event_match_key
JOIN children_objects co ON co.evidence_uid = b.evidence_uid
WHERE co.object_family = 'Judicial Officer Signature'
GROUP BY co.object_sha256, co.object_person_name, co.object_description
ORDER BY n_docs DESC;
```

A7: MICHAEL HOWARD — Metadata Creation Date Identity

```
WITH howard AS (
  SELECT
    c.evidence_uid, c.case_id,
    c.filing_type, c.filing_date,
    c.metadata__xmp_create_date,
    'cluster:1|' || c.case_event__date || '|' || LOWER(REGEXP_REPLACE(
      REGEXP_REPLACE(c.filing_type, '[^a-zA-Z0-9]+', '-', 'g'),
      '(^-+|-+$)', '', 'g'
    )) AS cluster_event_key
  FROM children c
  WHERE c.cluster_id = 1 AND c.case_event__date IS NOT NULL
)
SELECT
  cluster_event_key,
  COUNT(*) AS n_docs,
  COUNT(DISTINCT case_id) AS n_cases,
  COUNT(DISTINCT metadata__xmp_create_date) AS distinct_create_dates,
  CASE WHEN COUNT(DISTINCT metadata__xmp_create_date) = 1
    THEN 'ALL IDENTICAL'
    ELSE COUNT(DISTINCT metadata__xmp_create_date)::text || ' variants'
  END AS create_date_sharing
FROM howard
GROUP BY cluster_event_key
ORDER BY n_docs DESC;
```

A8: MICHAEL HOWARD — Judicial Signature Object Reuse

```
WITH howard AS (
  SELECT
    c.evidence_uid, c.case_id,
    c.filing_type,
    'cluster:1|' || c.case_event__date || '|' || LOWER(REGEXP_REPLACE(
      REGEXP_REPLACE(c.filing_type, '[^a-zA-Z0-9]+', '-', 'g'),
      '(^-+|-+$)', '', 'g'
    )) AS cluster_event_key
  FROM children c
  WHERE c.cluster_id = 1 AND c.case_event__date IS NOT NULL
)
SELECT
```

```
  h.cluster_event_key,
  COUNT(DISTINCT h.case_id) AS n_cases,
  COUNT(DISTINCT co.object_sha256) AS distinct_sig_hashes,
  string_agg(DISTINCT co.object_person_name, '; '
    ORDER BY co.object_person_name) AS judges
FROM howard h
JOIN children_objects co ON co.evidence_uid = h.evidence_uid
WHERE co.object_family = 'Judicial Officer Signature'
GROUP BY h.cluster_event_key
ORDER BY n_cases DESC;
```