

Method, Redaction, and Chain of Custody

This document exists so that anyone, including the City, Flock, or a critic, can check what was done to these records between the City's server and this download.

1. Chain of custody

Requester	Magnetic Norse
Agency	City of Kalamazoo Department of Public Safety
Request number	P022224-061526
Request filed	June 2026. The Mayor's determination dates the submission June 12; the records portal logs the request as created June 15
Partial denial	July 8, 2026
Appeal filed	July 8, 2026
Records delivered	July 22, 2026, OneDrive transfer in 2 parts, plus 1 email attachment
Mayor's written determination	July 23, 2026, appeal denied
Delivered master file	Audit Logs - Complete.csv, 38,728,527 bytes, 140,561 data rows
SHA-256 of that master file	980608204b2b81c5243354f03c2d482b3c6da47ede2d3521e3c9bae5535db4b2

On that hash. The master file is not in this pack, because both audit CSVs here are derived from it and publishing a third copy adds nothing. The hash is here so that if the master is ever handed to somebody who wants to check this work, they can confirm it is the same file these numbers came from. It is a promise I can be held to. It is not proof the City sent that file, because only the City can attest to that.

The appeal did not win the logs. The City's July 8 determination denied the audit logs on the grounds that they are "not held in the regular course of business" and that the City was not obligated to create a new public record, a compilation, or a summary. After the appeal was filed, the City supplemented its production on its own initiative and delivered the logs on July 22 with the license plate and case number columns redacted.

The determination of July 23 reached 3 appealed items. It upheld 2 of them and ruled the third moot.

Part 2, the data-sharing agreements, went moot, on the finding that the supplemental production had "effectively granted" that part in full.

On part 3, the camera inventory, the Mayor **upheld** the withholding of camera locations under MCL 15.243(1)(n) and 15.243(1)(y). His reasoning, in his words, is that disclosing exact camera sites "would allow individuals to map covered corridors, identify coverage

gaps, and actively evade law enforcement detection." Worth noting that the City's own July 8 denial had withheld this item under different subsections, (1)(s) and (1)(u).

Part 5, the audit logs, also drew an **upheld**, on the plate and case number redactions. The exemptions the City cited and the Mayor found permissible are MCL 15.243(1)(a), (1)(b)(i), (1)(b)(iii), (1)(b)(v) and (1)(d) for the plates, and (1)(b)(i) and (1)(b)(v) for the case numbers. Those are the subsections that cover this pack.

The determination is in the pack so you can read it yourself.

The City also took the position that the request was not a valid FOIA request, because a request made by a limited liability company must name an individual agent under MCL 15.233. The City processed it anyway. That position is in the determination and is reproduced without comment.

The delivered files have not been modified. They are retained unaltered as the master copies. Every step below was performed on a copy.

2. What the City produced

Flock generates 2 standing monthly reports per agency. The City pulled 6 months of both, so 12 files covering January through June 2026, and combined them into a single CSV using Power Query. The `Source.Name` column stamps every row with the file it came from.

Series	Rows	Officer name field	Contents
<month>-Kalamazoo MI PD-Network-Audit.csv	115,606	*** on every row	every agency that searched Kalamazoo's network, Kalamazoo PD included
<month>-Kalamazoo MI PD-Audit.csv	24,955	populated on every row	Kalamazoo PD only

On the *.** My reading is that this is Flock's own export format for the network report, and that nobody at the City applied it. The evidence is that it appears on 100% of the rows in one report series and on 0% of the rows in the other, consistently across all 6 months, which is the shape of a format difference. I have not seen Flock documentation saying so, and no one at the City or at Flock has confirmed it. Treat it as an inference with the evidence attached, not as an established fact.

14 columns: `Source.Name`, ID, Name, Org Name, Total Networks Searched, Time Frame, License Plate, Reason, Case #, Filters, Search Time, Search Type, Text Prompt, Moderation.

The camera deployment tracker arrived the same day, July 22, 2026, as an attachment to an email rather than through the OneDrive transfer, which is why it sat apart from the rest of the production and was missing from the first version of this pack. It is

Flock Camera Deployment Tracker 2026-07-22.xlsx in 03_Source_Documents, byte-identical to what the City sent, and it hashes to 53b428b6fa84e26b888d10f850cee526302a20f15d2e55bbf25696ea609c96d1. 42 device rows, 10 columns, with Location Name, Latitude and Longitude reading REDACTED on every row. A cleaned copy with the install year pulled out of the scheduling column is in 02_Analysis/10_camera_inventory.csv. The spreadsheet carries the name of whoever last saved it in its document properties, as any spreadsheet does; it was left there because re-saving the file would break the hash above.

3. Deduplication, and the check that nothing was altered

140,561 rows contain 116,017 distinct values in Flock's ID column. 24,542 IDs appear in both report series. 1 ID appears 4 times, twice in each series, which is where the 2 extra rows below come from.

Before treating those as duplicates, every one was compared field by field across both copies. **Across all 24,542 cross-series IDs there are zero conflicts** on Org Name, Time Frame, Search Time, Reason, Total Networks Searched, Search Type, License Plate, Case #, Text Prompt or Moderation.

The Filters field does differ, on 8,206 of those IDs. The network report reads Redacted where the local report carries the actual filter selections, for example SUV, Chevrolet, black, Redacted. That is a difference in how Flock generates the 2 exports, and it says nothing about what the search was. The deduplicated file keeps the local copy, which is the more complete of the 2, so no filter detail is lost. It is called out here because "field by field with zero conflicts" would otherwise be an overstatement.

2 separate searches would carry 2 different Flock IDs. These carry one. What looks like duplication is a single search listed in 2 reports.

The arithmetic closes exactly:

```
140,561 raw rows
- 24,542 cross-series duplicate IDs
-      2 byte-identical rows
= 116,017 distinct searches
```

Neither report contains the other. The network report covers all 10 agencies, and the local report covers only Kalamazoo PD, but 412 Kalamazoo PD searches sit in the local report and not in the network one. Every Kalamazoo PD search that is in the network report is also in the local report, so the gap runs one way.

Most of that gap has an explanation, and it is an inference rather than something Flock or the City has confirmed. The network report appears to log searches that queried Kalamazoo's network; the local report logs searches run by a Kalamazoo PD account whatever network it hit. 388 of the 412, or 94.2%, touched one network or none. All 180

apilookup searches in the file touched zero networks and all 180 are in this set. A search that reached no network has nothing to record in a per-network audit. Of the local-only searches that touched exactly one network, the filters name other states; of the searches present in both reports that touched exactly one network, they do not.

A further 20 touched between 2 and 5 networks, which is close enough to the same reading to sit under it. That accounts for 408 of the 412.

4 rows resist it: 3 mobile searches logged within 3 minutes of each other on January 9 at 896 networks, and 1 lookup at 148. Those are still unexplained.

4. What the City redacted before release

2 columns arrived already redacted. Neither was touched further.

- **License Plate** reads REDACTED on 98,737 rows. The other 17,280 rows are blank. Those break down as area search (13,452), search - Mobile (3,330), visual (389), freeform (52), multiGeo (41), searchSummary - Mobile (14), freeform - Mobile (2). That is 17,280 exactly. Those are search types that normally take no plate input, but the word normally is doing work: 6 search rows, 5 search - Mobile rows and 9 multiGeo rows do carry a plate. 20 rows in 116,017. Nothing in the file explains them.
- **Case #** reads REDACTED on 49,617 rows and is blank on 66,400. The file shows whether a case number existed. It never shows what the number was.

The consequence is that nothing in this release links a plate to a vehicle or a person. Methods used elsewhere that consolidate repeated partial-plate searches into single investigative episodes cannot be run on this production.

5. What was removed before publication

The notes column has been reviewed 2 times for redaction. The first pass, in July 2026, looked for names, dates of birth, phone numbers, addresses, ages and case numbers. **A second pass on August 6, 2026 looked for the categories the first pass had not searched for, and found license plates, state identification numbers, internal report numbers and a physical description.** Both passes are logged, and REDACTION_LOG.csv marks which pass each entry came from. A reader checking this file against an earlier copy needs to know why the counts moved, so the second pass is logged as its own pass.

A third review, also on August 6, 2026, re-derived every figure in this pack from the CSVs and re-ran every scan. It removed nothing further from the data. It did correct 3 errors in the redaction log, listed in section 5f.

5a. Officer names

The `Name` column carried 177 distinct Kalamazoo PD account names across 24,954 searches. Each was replaced with a pseudonym, `Officer 001` through `Officer 177`.

Pseudonyms were assigned by a random permutation. The order carries no information about volume, seniority, alphabetical position, or anything else, and cannot be derived from this data or from any public record. The mapping is one-to-one and stable, so the same account always receives the same pseudonym and per-account analysis stays intact. `02_Analysis/09_kalamazoo_pd_searches_per_officer_pseudonymized.csv` reports search counts and documentation rates per account.

*** rows were left as ***.

The name-to-pseudonym map has not been published and exists only in a private working file. The random seed used to generate it was not recorded.

Why. These are line officers rather than policymakers, and naming them would expose individuals to consequences this record cannot justify. What the record supports is a statement about a department. Any journalist who wants the names can file the same request and receive them, the same way I did.

What a determined reader can still do. Every row in the audit CSVs carries the pseudonym and a timestamp, so anyone can rebuild an account's shift pattern, day-of-week pattern, gaps, first date and last date. That is unavoidable if per-account analysis is going to survive at all.

5b. Free-text notes

The `Reason` column is free text and the City did not redact it. It is structured as `Category - operator note`. 115,515 rows carry that structure. A further 502 rows have no category separator at all; those were reviewed as notes in their own right, which is how 2 identification numbers sitting alone in the field were caught on the second pass.

How the notes were reviewed. Not by sampling. Take the `Reason` column, split each value on the first occurrence of " - ", and keep the part after it; where there is no separator, keep the whole value. On the published file that yields 91,109 rows carrying a note and **2,718 distinct notes**, in which **192 words appear on 60 or more rows**. Those frequent ones are police vocabulary (`fugitive`, `narcotics`, `udaa`, `warrant`, `stolen`, and so on) with no personal names among them. **Every note containing any less common word was read individually**, which came to **1,611 distinct notes**. Separately, every note was cross-checked against the 177-name officer roster. That second check caught one entry the manual read had missed.

Those counts are computed on the deduplicated file in this pack, using the split rule above, so they reproduce from what you have. Counts taken before redaction differ slightly, because replacing a case number or a plate with a label collapses some strings together.

A word-frequency method has a blind spot by construction. It never reads a note whose every word is common. So a second, frequency-blind sweep was run over all 2,718 distinct notes, looking for 2 capitalized words in a row anywhere in a note regardless of how common either word is, which is the shape of a personal name. It returned roughly 200 hits and none of them is a person. The closest is Carter's RF, which is a retailer and a retail fraud.

Pattern scans were run for phone numbers, Social Security numbers, email addresses, dates of birth, street addresses, ages, Michigan license plate formats, and state or offender identification number formats. Results: no Social Security numbers, no email addresses, 1 phone number, 2 dates of birth, 5 address references, 16 license plates, 32 identification numbers.

What was found and removed:

Type	Distinct entries	Rows
Named private individual	33	144
State or offender identification number	32	167
License plate	16	24
Internal report number, including 2 structured serials	12	56
Named officer or federal agent	6	18
Street address	1	4
Physical description of a private individual	1	4
Specific incident location	1	4
Out-of-state county naming a warrant's origin	1	6
Total	103	417 of 140,561, or 0.30%

10 rows carried 2 of these, so the entry rows sum to 427 while the distinct row count is 417.

2 of the name entries also carried a date of birth, 1 carried an age, and 1 carried a personal phone number.

On the license plates. The City redacted the License Plate column. It did not redact the notes column, and 16 operators typed a plate in there instead of, or as well as, the plate field. In one case the same plate appears in a note on one row and in the redacted plate field on the row logged 27 seconds later. Those 16 are the reason an earlier draft of this document was wrong to say the release contained no plates anywhere. It does not now.

On the identification numbers. 32 entries carried what appear to be Michigan state identification or offender numbers, some of them explicitly labeled SID or MDOC in the

note. A number of that kind identifies one specific person and keys directly into a criminal history record, which makes it a stronger identifier than a name. All 32 are gone.

On the report numbers. The second pass looked specifically for identifiers that were neither a name nor a case number, and found 12: 10 incident report numbers written as IR # nnnn, and 2 structured serials. They sit in the same class as the case numbers the City redacted in its own column, so they came out too, across 56 rows. They were replaced with a flat [REPORT REDACTED] rather than a stable label, because none of them carries enough rows for a per-report count to mean anything.

Replacements are explicit markers: [NAME REDACTED], [DOB REDACTED], [PHONE REDACTED], [AGE REDACTED], [ADDRESS REDACTED], [OFFICER REDACTED], [PLATE REDACTED], [ID REDACTED], [REPORT REDACTED], [DESCRIPTION REDACTED], [LOCATION REDACTED], [COUNTY REDACTED]. A reader can always see that something came out and what kind of thing it was. The rest of each note is intact.

2 entries were generalized a second time. In both, the name was already gone, but what was left named the incident narrowly enough to find the person anyway: one gave the exact location of a distinctive theft, the other gave the out-of-state county a warrant came from alongside a date of birth marker. Those 2 details came out of the data and out of the log.

Where a call was close, the entry was redacted. 10 entries could plausibly be a business, a street, a county, or a legal term instead of a person. All 10 were removed anyway. Losing 10 ambiguous entries out of 103 costs the record almost nothing, and it costs the searches nothing at all.

5c. Case numbers typed into the notes column

The City redacted the Case # column. Some operators also typed case numbers into the free text of the Reason column, and the City did not redact those.

225 case numbers appear that way, across 3,306 rows. Each was replaced with a stable label, [CASE-A] through [CASE-HQ], ordered by how many rows carry it within each review pass. The same case number always receives the same label.

That keeps every finding intact. [CASE-A] appears on 980 searches by the Michigan State Police between March 2 and June 29, all filed under Drugs/Narcotics, which is one investigation consuming just under 1% of 6 months of searches across an entire camera network. The label carries that. The number itself would add nothing and would mean publishing something the City chose to withhold.

The first review pass labeled 112 of these. The second pass labeled the remaining 113, including team-prefixed serials such as the ones the Michigan State Police narcotics and fugitive teams use. An earlier draft of this document implied all case numbers had been handled after the first pass. They had not been.

Note on counting: near-identical tokens differing only in zero-padding, spacing or capitalization were treated as the same case, since within a single agency and category

they are the same number typed several ways. Where an operator's typo produced a genuinely different string, it received its own label. Counts by label are therefore a floor.

Michigan offense codes such as 3500, 35001, 35 and 0900, and the 5-digit family (09001, 13002, 22001, 24002, 54001 and so on), radio codes such as 10-10 and 10-22, and team placeholders of the form Wmt-tbd-26 all appear in the notes column. Those are codes, not case numbers, and they were left alone. A Michigan incident report number carries a year prefix and runs 8 to 10 characters; none of these does. They should not be counted as investigations. 608 searches were filed under a case reference that literally reads "to be determined."

5d. What was deliberately left in

- **Category labels on sensitive calls.** Searches labeled Suicidal, Welfare Check, juvenile runaway, CSAM, missing elderly are in the file in full. None of them contains a name. They are bare category text. They are also the only way to measure how often this system gets used for the purposes it is defended with, so removing them would damage the record without protecting anybody. Section 6 covers the risk that comes with that choice.
- **Incident locations** such as intersections and street names, apart from the one specific street address. These are blotter-level detail.
- **All 54 AI freeform prompts.** Reviewed individually. Every one describes a vehicle. None describes a person or carries identifying information.
- **The Filters column.** Every distinct value was tokenized and read. It is built from 135 distinct tokens: vehicle make, body type, color, a US state or territory, Canada, Mexico, the word Redacted, and on 2 rows the word false. No free text, and nothing personal.

5e. Verification

5 automated passes were run on the finished files:

1. A residual scan confirming that no removed name, date of birth, phone number, address, license plate or identification number survives anywhere in any published file, in any capitalization. 450 removed strings against every file in the pack, zero survivals.
2. A structural scan for anything still matching a case-number, plate, report-number or identification-number shape. What it returns now is offense codes, radio codes, team placeholders, and a university commencement event tag, all of which are meant to be there and are listed in 5c.
3. A cross-column scan. None of the removed values leaked into another field.
4. A column-by-column comparison against the pre-redaction file, showing that ID, Org Name, Total Networks Searched, Time Frame, Search Time, Search Type, License Plate, Case #, Filters, Text Prompt, Moderation and Source.Name are byte-identical, that no row was added or dropped, and that per-account search counts are unchanged. Only Reason and Name differ, which are the 2 columns the redaction touches.

5. A byte-level identity scan of every file in the pack, covering raw bytes, UTF-16 in both byte orders, decompressed PDF object streams, and extracted PDF text, for anything identifying the person who made this pack or the software used to make it. The generated PDFs carry no metadata dictionary, no XMP packet and no annotations. The source documents keep the originating agency's and vendor's metadata, which is what makes them verifiable.

The published `REDACTION_LOG.csv` carries an accounting of every change: the category it sat under, what kind of information came out, what the text reads now, how many rows it touched, and which review pass caught it.

On the row counts in that log. For case labels the count is exact, because the label is already visible in the file and counting it adds nothing. For everything else the count is a band (1, 2-5, 6-20, 21 or more). The reason is that an exact per-entry count on a generic marker tells a reader how many separate people sit behind a block of otherwise identical rows, and how their searches divide, which is close to the thing redacting the identifier was meant to prevent. The bands still let you audit the work. If you want the exact total, count the rows carrying a marker in the CSV itself; it comes to 417.

The log does not carry the original text, because a log that reprints what it removed would defeat the redaction. 3 entries in it were reworded because the surrounding detail was specific enough to identify the incident even with the name gone.

5f. What the August 6 review changed

The third review corrected 3 things in `REDACTION_LOG.csv` and nothing in the data.

1. 6 log entries had a stray trailing character left on the end of the replacement text, which was the last character of the identification number that had been removed. None of those strings existed in the released CSV, which is how they were caught. They are gone.
2. 1 case label was listed twice, once under each review pass, so the log's case-number rows summed to 3,339 instead of 3,306. The duplicate is gone.
3. 3 entries described the surviving text as more generalized than it actually is. The log now matches the file exactly. Every replacement string in the log can be found in the audit CSVs; that is now a check, and it passes on all 361 rows of the log, which is 258 case entries plus the 103 personal-identifier entries in the table above.

6. Known limitations

- The audit log records that a search was run. It does not record what the search returned or what was done with the result.
- Officer names exist only for Kalamazoo PD. Every outside-agency row shows `***`, so per-officer analysis is impossible for 78.5% of the searches.

- 412 IDs appear in the local report but not the network report. Section 3 accounts for 408 of them as an inference. 4 are still unexplained.
- 20 rows carry a license plate on a search type that normally takes none. Section 4 has the breakdown. Nothing in the file explains them.
- 162 rows in the deduplicated file, 205 in the as-produced file, carry search timestamps on July 1, 2026 UTC. Those are June 30 local-time searches crossing the UTC midnight boundary, and they appear as a 2026-07 row in 02_Analysis/03_monthly_volume_by_agency.csv, which is annotated. The period covered is January through June. The first search in the file is 00:48 Eastern on January 1 and the last is 23:58 Eastern on June 30.
- **A search can be joined to a public police alert.** Sensitive category labels were kept, and every row carries a timestamp, so a reader with a local news archive can sometimes match a search to an incident. In most cases that adds nothing, because the incident and the people in it are already public. In one class it adds something real. Agencies publish missing-person alerts with a name and a photo but do not usually publish that the person was suicidal, and this file does carry that word on a date. 167 rows carry a mental-health term in the operator's note, across 44 combinations of date and agency, 11 of them at departments small enough that one day means one incident. Generalizing those notes was considered and rejected. It would take the count of searches mentioning suicide or a suicidal subject from 147 to 2, and destroy the only measurement of how often this system gets used for the purpose it is most often defended with.
- The MOU in this pack is signed only by Flock, on June 11, 2026; the signature and date fields for both Kalamazoo County Consolidated Dispatch and Kalamazoo Public Safety are empty, and the document says it commences once fully executed. Its numbered sections run 8, 10, 11, with no section 9. It is reproduced exactly as produced.
- Grand Rapids PD appears only between January 1 and February 18, 2026, then not at all. All 4,428 of its searches fall in that window. Nothing in the file explains the stop.
- Reason text is whatever an operator typed. It is inconsistent, frequently misspelled, and often a single word or an internal code. Keyword counts in the findings sheet are counts of what got typed, not of what happened, and the pattern behind each count is published in 02_Analysis/08_reason_keyword_counts.csv so you can disagree with the definition and recompute.

7. Reproducing this

Load either CSV with all columns as strings and no automatic null conversion. The `Time Frame` column contains embedded newlines and will break a strict CSV parser, so use a tolerant one. Deduplicate on `ID` before counting anything.

```
import pandas as pd
df = pd.read_csv(path, dtype=str, keep_default_na=False, engine='python')
```

```
df.columns = [c.strip() for c in df.columns]
searches = df.drop_duplicates(subset='ID') # 116,017
```

Any count taken from the as-produced file without deduplicating will overstate Kalamazoo PD's activity by roughly 24,500 searches.

8. Corrections

If a figure here is wrong, the underlying data is in this pack. Check it and say so. Corrections get posted publicly.

mags@magneticnorse.com

Records lawfully obtained under the Michigan Freedom of Information Act, MCL 15.231 et seq. Magnetic Norse is a project of Central Ventures LLC, a for-profit business. Not affiliated with Flock Safety, Flock Group Inc., or any government agency. This document reports what the records show. It does not establish intent.