Juniper Contamination: Investigation Postmortem

Juniper Contamination Incident
Seekers Spirits | Cambodia
July–August 2022

Incident type	Supplier ingredient contamination - juniper berries
Discovery method	Accidental: routine tasting during batch comparison
Affected batches	Batches 19, 20, 21, 22 (Gold Gin, Dry Gin, Jason Kong Butterfly Pea Gin)
Contaminated lot	PDI-JNP-220620: Supplier A, received 2022-06-20
Discovery date	2022-08-15
Bottles quarantined	715 bottles across four batches, three SKUs (Dry Gin appears in batches 20 and 22)
Root cause confirmed	Individual ingredient: musty off-character isolated to juniper lot
Resolution	Full batch quarantine, supplier quality gate introduced, lot-level testing mandated

Table of Contents:

Section 1: Systems Context

1.1 Operations Overview

Seekers Spirits is a craft distillery based in Phnom Penh, Cambodia, producing premium spirits for domestic on-trade, off-trade, and export markets.

At the time, production ran on a single 300L still. All botanicals were sourced from third-party suppliers, with juniper berries forming the backbone across three core SKUs: Seekers Mekong Dry Gin, Seekers Mekong Gold Gin, and Jason Kong Butterfly Pea Gin.

Quality control (QC) was my responsibility end-to-end, from tasting through to final release decision. For each batch, I recorded ABV, clarity, and a qualitative sensory assessment, and made the final go/no-go call before any stock left the warehouse.

The production flow followed a simple sequence: ingredients were received and visually inspected, batches were produced on the still, and finished product was assessed through QC before release. In practice, there were no formal quality gates upstream of production, no automated monitoring between stages, no lab-based analysis, and no lot-level intake checks beyond visual inspection. Any failure introduced at the ingredient stage had an unobstructed path through production before it could be surfaced.

1.2 Ingredient Supply Chain

Juniper berries were sourced from a Vietnam-based botanical supplier (Supplier A). Lot PDI-JNP-220620 (50kg, received 20 June 2022) was used across four consecutive batches between July and August 2022. At receipt, only a visual inspection was performed. There was no sensory check, no comparison against previous deliveries, and no established quality baseline for the ingredient.

As a result, the lot entered production without being verified against any sensory standard. A problem introduced at receipt would only surface at final QC, after the ingredient had already been used across multiple batches.

1.3 Observability Gaps

The following gaps in observability were present prior to the incident and contributed to its late detection:

No defined standard for incoming inspection
Whether an ingredient was accepted depended on who inspected it. Ingredients were checked at receipt for visible issues such as bruising or mold, and for obvious smell defects, but there was no defined standard to assess them against. As a result, quality was not checked consistently before ingredients entered production.

No flavour baseline for incoming raw ingredients
There was no defined reference for how the raw ingredient should taste. Without that, a new lot could be accepted even if its flavour was different from earlier lots. The variation was not identified before production and only became clear later through batch comparison.

No fixed reference point for batch assessment
Each finished batch was assessed against a different reference bottle chosen at the time, rather than against one fixed comparison batch. Because the reference point kept changing, it was harder to tell whether the new batch had changed or whether the comparison sample was simply different.

No cross-batch visibility
Even when batches were compared against a reference bottle, each assessment was treated separately. There was no way to track flavour changes in the same product over time, so a repeated issue could be missed even if each batch appeared acceptable on its own.

Detection happened at the end of the process, not upstream where problems entered. There were no checks at the ingredient stage and no way to spot problems developing as production moved forward, so issues could pass through the full process and only be found at final quality control.

1.4 System Flow and Detection Architecture

The diagram below shows the process flow, where contamination entered, and where it was detected. Without checks earlier in the process, the issue moved through production and was only found at final QC.

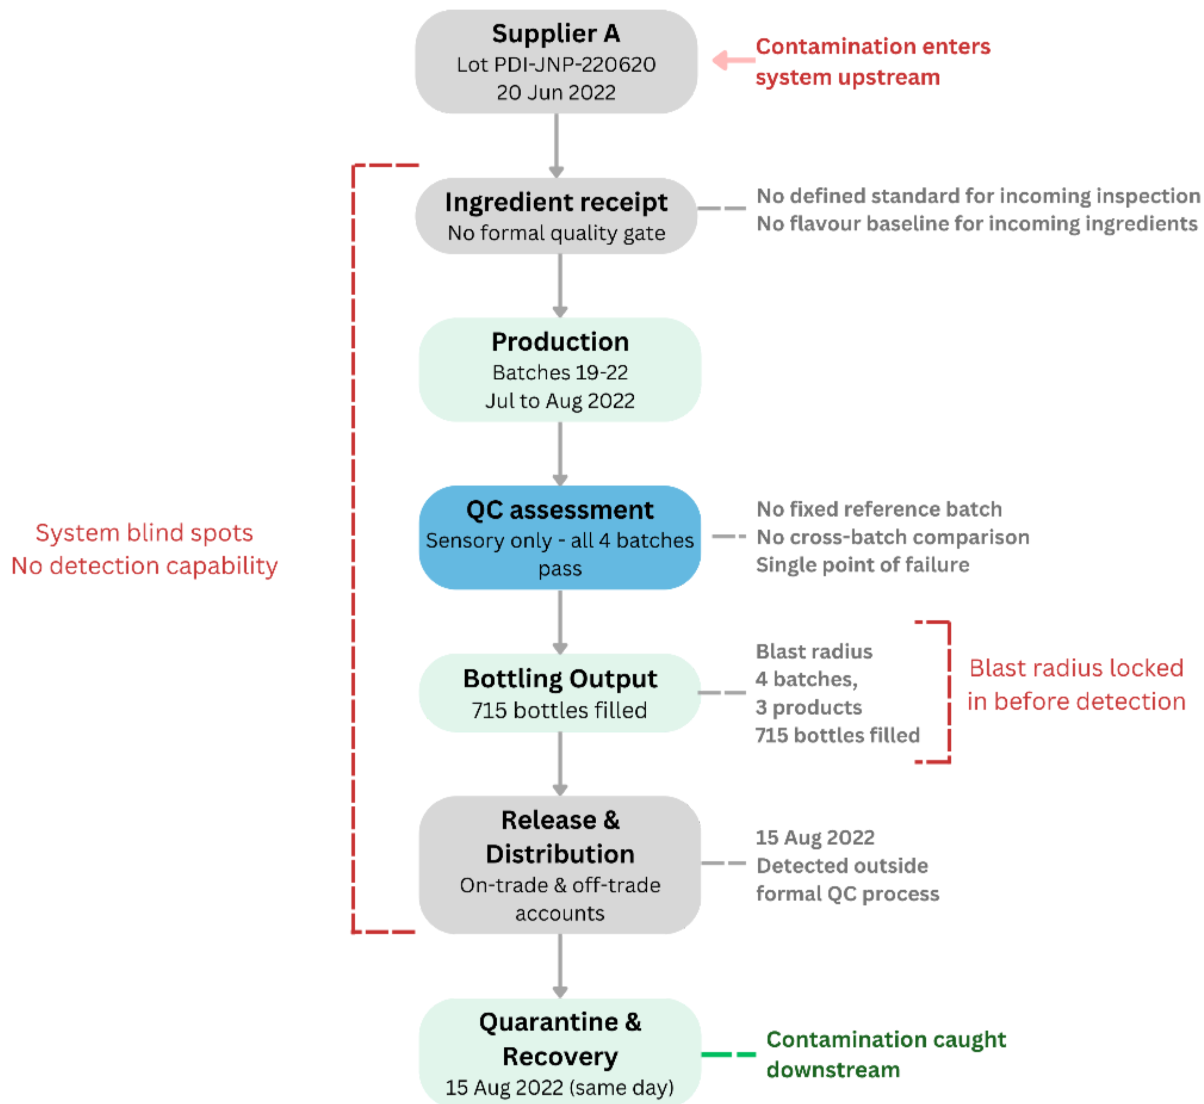

By the time the failure was detected, it had already passed through every stage of production. The impact was driven not just by the failure itself, but by how long it remained undetected. In this case, that gap lasted eight weeks.

Section 2: Discovery Timeline

2.1 Pre-Discovery: Batches Pass Initial QC

Between 11 July and 8 August 2022, four batches were produced using juniper lot PDI-JNP-220620: batch 19 (Gold Gin), batch 20 (Dry Gin), batch 21 (Jason Kong), and batch 22 (Dry Gin). All four passed quality control and were approved for bottling. Sensory notes described them as acceptable, with no clear deviation from the expected profile.

Critical gap

Quality control included side by side comparison against existing finished stock, but the reference bottle was chosen informally rather than from one designated baseline batch. With no fixed reference, comparisons relied partly on memory of previous assessments. As a result, small changes across consecutive batches were harder to catch.

2.2 Discovery: Accidental Detection During Batch Comparison

On 15 August 2022, during a routine quality check of a later batch, I selected a bottled sample as the reference for comparison. The bottle I picked was from one of the contaminated batches. When I tasted it, I noticed a clear difference from the expected profile: a musty undertone and an unusual bitterness.

The issue was found outside the formal QC process. It was only identified because an affected bottle was used as the reference in an unrelated tasting.

2.3 Initial Hypothesis

When I first noticed the off-notes, I thought it might be palate fatigue or a problem with the sample. I retasted it several times and checked the original QC records, which showed that the batch had previously passed. At that stage, it was not clear whether the issue was with my perception or with the product.

To rule out a sampling issue, I pulled a second bottle from the same batch. It showed the same flavour problem. The issue was batch-wide. I then pulled batches before and after it and assessed them in sequence to find where the problem started and stopped.

At that point, I had confirmed the issue was real and affected the whole batch. My working hypothesis was that it was limited to that one batch, most likely related to recipe, process, or storage. But as more batches showed the same problem across different products, batch-level causes became less likely. The pattern pointed to a shared ingredient as the likely source.

At that point, I shifted my mental model from a batch-level failure to a shared input failure. This reframed the investigation from diagnosing an isolated issue to identifying common variables across affected batches.

2.4 Investigation: Ingredient Isolation

Once a shared ingredient became the most likely cause, I reviewed the production records and listed the botanicals used in the affected batches. After confirming through tasting that the issue was limited to batches 19 to 22, I queried the database to identify the ingredients shared across all four:

```
SELECT i.ingredient_name, COUNT(DISTINCT bi.batch_id) AS batch_count
FROM batch_ingredients bi
JOIN stock_receipts sr ON bi.receipt_id = sr.receipt_id
JOIN ingredients i ON sr.ingredient_id = i.ingredient_id
WHERE bi.batch_id IN (19, 20, 21, 22)
GROUP BY i.ingredient_name
HAVING COUNT(DISTINCT bi.batch_id) = 4
ORDER BY i.ingredient_name;
```

This returned five shared ingredients. After removing the base ingredients common to all batches (Neutral Grain Spirit, RO Water, and Sugar), two botanical candidates remained: Juniper Berries and Fresh Lime Peel.

I reviewed supplier lots and delivery records for both ingredients and checked them by smell and taste. Neither showed any obvious issue. This was the first dead end. Direct sensory checks of the raw materials were not enough to identify the source.

To isolate the source, I moved to single botanical distillation. This let me test one input at a time and confirm the source directly. I started with juniper because it made up the largest share of botanical weight across the recipes and was the most likely to affect flavour. The juniper distillation reproduced the same off-notes: a musty undertone and unusual bitterness. That confirmed juniper as the source of the deviation.

2.5 Root Cause Confirmation: Lot-Level Traceability

After juniper was identified as the source of the deviation, I traced it back through supplier and inventory records. All affected batches (19, 20, 21, and 22) were linked to the same juniper lot: PDI-JNP-220620. No other juniper lots were used in those batches.

I then reviewed batches produced outside that lot window and confirmed they were clean. The issue was limited to juniper lot PDI-JNP-220620. No other batches were affected.

2.6 Investigation Timeline (table)

Date	Event	Detail	Confidence	Basis
Jun 20, 2022	**Lot received**	Juniper lot PDI-JNP-220620 received from Supplier A (Vietnam-based botanical supplier). Visual inspection only. No baseline recorded.	-	Pre-incident. No concern raised.
Jul 11 – Aug 8	**Production**	Four batches produced using contaminated lot: B19 (Gold Gin), B20 (Dry Gin), B21 (Jason Kong), B22 (Dry Gin).	-	Pre-incident. No concern raised.
Jul–Aug 2022	**Initial QC passes**	All four batches pass quality control. No deviation detected. Approved for bottling.	-	Pre-incident. No concern raised.
Aug 15, 2022	**Accidental discovery**	Contaminated bottle used as reference comparison. Musty undertone and an atypical bitterness detected. Deviation identified.	~50%	Single bottle, single taster. Could be sample error, palate fatigue, or storage effect. Real but not confirmed.
Aug 15 (Step 1)	**Batch confirmation**	Second bottle pulled from same batch. Deviation confirmed. The issue is batch-wide, not sample-specific.	~75%	Sample error and storage ruled out. Cause still unknown. Could be recipe, process, equipment, or ingredient.
Aug 15 (Step 2)	**QC record review**	Historical QC records were reviewed. No earlier deviation had been flagged.	~80%	Contamination present across the full batch window. Process and assessor error less likely.
Aug 15 (Step 3)	**Blast radius scoping**	Testing conducted across recent batches using the same juniper lot. Batches 19–22 confirmed affected. No impact outside this lot window.	~85%	Four batches affected, all tied to a single lot window. Pattern strongly points to ingredients.

Aug 15 (Step 4)	**Ingredient isolation**	Single botanical distillation conducted. Off-notes isolated to juniper lot PDI-JNP-220620.	~100%	All other ingredients are clean. Root cause confirmed with certainty.
Aug 15	**Quarantine**	All four batches quarantined. Stock recovery initiated. Supplier A notified.	-	Action taken on confirmed finding.
Post-Aug 15	**Remediation**	Mandatory lot-level testing introduced for all future juniper receipts. Quality gate updated.	-	Structural fix implemented.

Decision tradeoff:

Customer-visible actions were held until the evidence was clear. Acting too early risked unnecessary disruption. Acting too late risked more contaminated products reaching customers. Quarantine was triggered only once the source had been confirmed. The investigation moved in clear steps, ruling out other explanations one by one. Each check was designed to give a clear yes or no answer, so the investigation could move quickly without compromising the decision.

Section 3: Root Cause Analysis

3.1 Root Cause

The root cause was a contaminated juniper lot, PDI-JNP-220620, from a supplier in Vietnam.

The exact nature of the contamination was not identified. The off-notes, a musty undertone and unusual bitterness, were consistent with either microbial activity or oxidation during storage or transit, but no lab analysis was carried out to confirm this.

The investigation confirmed that the problem came from an upstream raw material, not from the production process, and identified which lot was responsible. What happened to the ingredient lot prior to receipt was not established.

3.2 Contributing Factors

No defined standard for incoming inspection
Botanical ingredients were accepted based on visual checks only. There was no flavour check at receipt and no structured comparison against earlier deliveries. That meant the contamination could pass into production without being identified.

No flavour baseline for incoming ingredients
There was no reference for how juniper from this supplier should taste. Without that, lot-to-lot variation was harder to identify before production and was only recognised later.

No cross-batch visibility
Each batch was compared against a reference bottle, but those comparisons were not tracked across runs. That made a small repeated change across multiple batches easier to miss, especially if the new profile had started to feel normal by the second or third assessment.

Contamination was too subtle to trigger detection
The off-notes were present but not strong. Batches 20 and 22 (Dry Gin) showed the clearest deviation, batch 21 (Jason Kong) showed a moderate deviation, and batch 19 (Gold Gin) showed the weakest. Because the reference point changed with each assessment, there was no fixed baseline. That made the gradual change harder to spot in any single QC check.

3.3 Five Whys Analysis

Problem: Contaminated juniper was used across four consecutive production batches. The issue was not identified until 15 August 2022, after four batches had passed QC, been bottled, and entered distribution.

1. Why did contaminated juniper pass QC across four consecutive batches?
The contamination was subtle. Each batch was compared against a reference bottle chosen at the time rather than one fixed baseline. Because the reference point kept changing, small changes across consecutive batches were harder to spot.

2. Why was there no fixed reference point for batch assessment?
There was a defined quality control process with structured qualitative assessment, but it did not specify a fixed reference bottle for comparison. As a result, the comparison point could change from one assessment to the next.

3. Why did the QC process not specify a fixed comparison bottle?
The QC process was designed to assess each batch against the expected standard, not to detect gradual change caused by variation in incoming ingredients. A fixed reference batch would only have mattered if that risk had been recognised. It was not.

4. Why was ingredient variation not recognised as a risk?
Supplier consistency was assumed rather than verified. Because this supplier had not caused a quality issue before, lot-to-lot variation was not treated as a risk. There was no intake check to test that assumption.

5. Why was there no intake check to test that assumption?
Supplier selection was based on price and availability. Quality and consistency were not formally assessed, and there was no requirement to verify them at intake or during the relationship. As a result, no detection step was ever built into the intake process.

Structural root cause:

Supplier quality and consistency were not treated as reliability requirements. The system relied on an assumption of consistency that was never tested, so no detection step was built into intake. As a result, final QC became the only point where problems could be found.

Section 4: Impact Assessment

4.1 Affected Stock

batch_id	product_name	batch_date	Severity
19	Seekers Mekong Gold Gin	2022-07-11	Mild
20	Seekers Mekong Dry Gin	2022-07-18	Severe
21	Jason Kong Butterfly Pea Gin	2022-07-25	Moderate
22	Seekers Mekong Dry Gin	2022-08-08	Severe

Over 56 days, the issue passed through four production runs and led to a blast radius of 715 bottles across three products.

4.2 Channel Exposure

By the time the issue was found, affected bottles had already been distributed to on-trade accounts such as bars, restaurants, and hotels, and to off-trade retail accounts including supermarkets and specialty retailers. Once stock had left the warehouse, customer exposure could not be fully avoided. Recovery was prioritised by account volume, with the highest-volume accounts contacted first to recover as much stock as possible.

Around 70% of the affected stock was recovered. The remaining 30% was assumed to have been sold or consumed before collection.

The issue was detectable, but the QC process was not set up to catch it reliably. In a side by side comparison, the difference was clear. But in routine single-batch assessment, it was easier to miss. The lack of customer complaints fits that. The contamination was noticeable, but not strong enough for most customers to report without something to compare it against.

No stock had been shipped internationally. The issue was contained within the domestic Cambodia market. This mattered because a Mongolia market launch was scheduled for mid 2023. If the issue had not been found and resolved in August 2022, the affected juniper lot could have stayed in use and entered production again during the export window.

Section 5: Response and Recovery

5.1 Immediate Actions

- All four affected batches quarantined upon confirmation
- Quarantined stock destroyed following confirmation of contamination
- Stock recovery initiated across on-trade and off-trade accounts
- Supplier A notified of the contaminated lot
- Remaining stock of lot PDI-JNP-220620 removed from production
- New juniper lot sourced, validated, and approved before use

5.2 Production Recovery

Production resumed on the same day the contamination was confirmed. There were no active runs using the affected juniper lot, and a replacement lot was already in stock and cleared before production began.

The first gin batch after the contamination used the new juniper lot from Supplier A. Before production, the lot was tested through single botanical distillation against the documented reference batch, using the same method that had confirmed the contamination. The profile was clean and matched expectations.

Quality control records rated the batch as "Very Good," with juniper character and intensity restored to the expected level. Later batches were also consistent. This confirmed that the issue was limited to lot PDI-JNP-220620 and that neither the equipment nor the process contributed to the failure.

Section 6: Structural Improvements

6.1 Changes Introduced

Mandatory lot-level testing for juniper
Every incoming juniper delivery is now checked before entering production. A small sample is distilled and tasted against the documented reference batch. Lots that do not match the baseline are rejected.

Flavour baseline documented for all approved suppliers
A reference flavour profile is now documented for each approved supplier. When a new ingredient lot arrives, it is tasted against that reference before entering production. This replaces memory-based assessment with a consistent point of comparison and makes ingredient lot-to-lot differences easier to spot at receipt.

Fixed reference baseline for QC
Quality control now uses one fixed reference batch for side by side tasting. Previously, different historical batches were used for comparison, so the reference point kept changing. Using one consistent baseline makes small changes easier to spot.

Lot-level traceability added to production records.
Each botanical delivery is now recorded by lot number, and each production batch is linked to the ingredient lot it used. This makes it possible to scope the blast radius of future incidents quickly. What previously took hours of manual record tracing can now be done in seconds with a single database query.

See example below.

6.2 Blast Radius Query Example

Operational question

"We have a contaminated lot. What did it touch, and how fast can we scope it?"

Query:

```
SELECT b.batch_id, b.product_name, b.batch_date, sr.supplier_lot_reference
FROM batch_ingredients bi
JOIN stock_receipts sr ON bi.receipt_id = sr.receipt_id
JOIN batches b ON bi.batch_id = b.batch_id
WHERE sr.supplier_lot_reference = 'PDI-JNP-220620'
ORDER BY b.batch_date;
```

Output:

batch_id	product_name	batch_date	supplier_lot_reference
19	Seekers Mekong Gold Gin	2022-07-11	PDI-JNP-220620
20	Seekers Mekong Dry Gin	2022-07-18	PDI-JNP-220620
21	Jason Kong Butterfly Pea Gin	2022-07-25	PDI-JNP-220620
22	Seekers Mekong Dry Gin	2022-08-08	PDI-JNP-220620

Four batches, one lot. Blast radius confirmed in under a second.

The lot was received on 20 June 2022. Detection came on 15 August, 56 days later, after all four batches had been produced and approved for bottling. By the time the query was run, 715 bottles across three products had already been filled and distributed. The query did not reduce the blast radius. It showed it immediately, which made targeted recovery possible.

What this demonstrates

This is the first step in blast radius scoping: identifying every production run linked to a contaminated ingredient lot before taking any external action. In this incident, that meant hours of manual checking against paper records. The query does the same job in seconds, without the risk of transcription errors or missed records.

This works because batch_ingredients links stock_receipts to batches at the ingredient lot level. That schema decision, made at design time, is what makes rapid scoping possible under pressure. Without it, the blast radius query does not exist.

Full Blast Radius

See Appendix A1 for the extended traceability query, which adds *batch_outputs* to return bottle counts per SKU across all four batches, producing the full 715-bottle picture.

Section 7: Operating Playbook

This incident can be understood as a system failure where detection was not colocated with the point of failure. The investigation and response focused on reducing detection latency, limiting blast radius, and redesigning the system so that similar failures are surfaced earlier.

7.1. Investigation model used

This investigation followed a repeatable sequence:

1. Confirm the issue is real
2. Scope the blast radius before diagnosing the cause
3. Identify what the affected cases have in common
4. Isolate the most likely cause through controlled testing
5. Trace it back to the upstream source
6. Act only when the level of confidence matches the impact of the decision

On 15 August, I confirmed the issue with a second bottle pull, then tested batches forward and backward to find the boundary of the problem before I knew the cause. Once the affected batch window was clear, I queried the shared ingredients and moved into diagnosis. Root cause confirmation came last, through single botanical distillation, before any external action was taken.

This order matters. By scoping first, recovery could be targeted as soon as the cause was confirmed. Each step narrowed the range of possible explanations, so that by the time an irreversible action was needed, the decision was supported by evidence.

7.2 Core Principles

Detection latency is a system design problem

Accidental detection is not a success. It shows there is a gap in the system. In this incident, the contaminated ingredient entered the process at receipt but was not found until final quality control, after production was complete. Checks were concentrated at the end of the process rather than where the problem entered.

Systems are more reliable when checks sit as close as possible to where problems are introduced. When detection comes late, even small issues can move through multiple stages before they are found. The fix was not to rely more heavily on final QC. It was to move detection upstream.

Blast radius scoping before escalation

Before any external communication, the investigation established which batches were affected and which were not. This meant confirming which batches used the same juniper lot and which used a different one. That allowed recovery actions to be targeted and avoided unnecessary escalation.

Confirmed finding before irreversible action

Quarantining batches and starting stock recovery were irreversible, customer-visible actions. They were only taken once the source had been confirmed. Single botanical distillation provided that confirmation.

Structural fix over process fix

The response did not rely on reminders or asking people to be more careful. It changed the system itself through mandatory lot testing, documented baselines, and cross-batch comparison. Each change closed a structural gap so the same kind of failure would be found earlier by design.

7.3 What I Would Do Differently

Treat supplier onboarding as a reliability problem

Supplier selection was treated mainly as a commercial decision: price, availability, and relationship. It did not include how a quality failure would be detected or how quickly its impact could be traced. Those are reliability questions and should be addressed before the first delivery, not after the first failure.

Instrument traceability from day one

At the time of the incident, batch tracking existed but was incomplete. A full traceability system had been delayed by other priorities and was treated more as admin than as a reliability requirement. That was a mistake. Without an ingredient lot to batch link, there is no quick way to see which runs are affected, and the investigation falls back on manual record tracing.

This was still manageable at the volume we were operating at. At higher volume, or across multiple markets, it would not have been. Traceability is not something to build once you need it. By the time you need it, the blast radius is already running.

Appendix: Database Evidence

Note on database evidence

The database used to model this incident was built retrospectively as part of this portfolio. At the time, the investigation relied on paper records and direct sensory assessment. Building the schema afterwards made it clear which design decisions would have changed investigation speed in real time. The most important was modelling batch_ingredients as a junction table linking stock_receipts to batches at the lot level. Without that decision, the blast radius query does not exist.

The queries shown here follow the same investigative steps used during the incident. What originally took hours of manual record tracing can now be done in seconds through structured data.

A1. Contamination Traceability Query

The key question during the investigation was simple: which batches used this contaminated ingredient lot, which SKUs were produced from those batches, and how many bottles of each had already been filled? Answering it manually meant cross-checking stock receipts, batch ingredient logs, production records, and bottling records.

The traceability chain runs through four tables:

- **stock_receipts:** identifies the contaminated lot by supplier_lot_reference
- **batch_ingredients:** links stock receipt lots to production batches
- **batches:** records each production run, including date and product
- **batch_outputs:** records bottles filled per SKU from each batch

Traceability Query

The query below traces the contaminated lot through to its final bottled output.

```
SELECT sr.supplier_lot_reference, b.batch_id, b.batch_date,
       b.product_name, bo.sku, bo.bottles_filled
FROM stock_receipts sr
JOIN batch_ingredients bi ON sr.receipt_id = bi.receipt_id
JOIN batches b ON bi.batch_id = b.batch_id
JOIN batch_outputs bo ON b.batch_id = bo.batch_id
WHERE sr.supplier_lot_reference = 'PDI-JNP-220620'
ORDER BY b.batch_id, bo.sku;
```

A2. Query Output

Running this query against the Seekers database returns the following.

supplier_lot_reference	batch_id	batch_date	product_name	sku	bottles_filled
PDI-JNP-220620	19	2022-07-11	Seekers Mekong Gold Gin	SMGG200	60
PDI-JNP-220620	19	2022-07-11	Seekers Mekong Gold Gin	SMGG700	100
PDI-JNP-220620	20	2022-07-18	Seekers Mekong Dry Gin	SMDG200	55
PDI-JNP-220620	20	2022-07-18	Seekers Mekong Dry Gin	SMDG700	90
PDI-JNP-220620	21	2022-07-25	Jason Kong Butterfly Pea Gin	JKBG200	50
PDI-JNP-220620	21	2022-07-25	Jason Kong Butterfly Pea Gin	JKBG700	105
PDI-JNP-220620	22	2022-08-08	Seekers Mekong Dry Gin	SMDG050	100
PDI-JNP-220620	22	2022-08-08	Seekers Mekong Dry Gin	SMDG200	60
PDI-JNP-220620	22	2022-08-08	Seekers Mekong Dry Gin	SMDG700	95

Total: 4 batches, 3 products, 9 bottling runs, 715 bottles quarantined.

A3. What This Demonstrates

This query traces a real incident that led to the quarantine of 715 bottles across three product lines. The key point is the design choice that made the query possible. The traceability chain only exists because batch_ingredients was built as a junction table linking stock_receipts to batches at the ingredient lot level (see A5. Ingredient Traceability Schema). Without it, the same question has to be answered by manually cross-checking stock receipts, batch logs, production records, and bottling records.

This incident was also a data visibility problem. The contamination was real, and the investigation method was sound, but the speed and confidence of the response depended on whether the data structure could show the full affected range quickly. With the right schema, that answer is available in seconds. Without it, time is spent pulling records while the impact is still growing.

System design is not separate from incident response. It determines how quickly you can act when something goes wrong.

A4. QC Records: Initial and Retrospective Assessments

The QC records below show both the original passing assessments and the retrospective findings recorded on 15 August 2022. Together, they show that the contamination was present but not identified at the time. The query retrieves records for batches 19 to 22, including approval status and notes, for both assessment types.

The issue was not individual error. The contamination was subtle, and the QC process was not set up to catch this kind of issue reliably at the time.

Query

```
SELECT
    batch_id,
    qc_reference,
    qc_date,
    approved,
    notes
FROM qc_records
WHERE batch_id IN (19, 20, 21, 22)
ORDER BY batch_id, qc_date;
```

QC Records - Database Output

batch_id	qc_reference	qc_date	approved	Notes (excerpt)
19	QC-SMGG-220711-1	2022-07-18	t	Batch within acceptable parameters. Botanical profile consistent with house standard. No faults detected. Approved for bottling.
19	QC-SMGG-220711-2	2022-08-15	f	Retrospective investigation assessment. Musty undertone and atypical bitterness. Linked to PDI-JNP-220620.
20	QC-SMDG-220718-1	2022-07-25	t	Batch within acceptable parameters. Botanical profile consistent with house standard. Approved for bottling.
20	QC-SMDG-220718-2	2022-08-15	f	Retrospective investigation assessment. Flavour degradation confirmed. Linked to juniper lot PDI-JNP-220620. Batch quarantined.
21	QC-JKBP-220725-1	2022-08-01	t	Batch within acceptable parameters. No faults detected on initial assessment. Approved for bottling.
21	QC-JKBP-220725-2	2022-08-15	f	Retrospective investigation assessment. Strongest off-character expression: musty undertone and pronounced atypical bitterness. Linked to juniper lot PDI-JNP-220620. Batch quarantined.
22	QC-SMDG-220808-1	2022-08-15	t	Batch within acceptable parameters. No faults detected on initial assessment. Approved for bottling.
22	QC-SMDG-220808-2	2022-08-15	f	Retrospective investigation assessment. Flavour degradation confirmed. Linked to juniper lot PDI-JNP-220620. Batch quarantined.

A5. Ingredient Traceability Schema

The ERD below shows the PostgreSQL schema used for the traceability queries in this appendix. The batch_ingredients table links stock_receipts to batches at the ingredient lot level, which is what makes blast radius scoping possible. This database was built retrospectively. See the note at the start of the appendix.