Cooling System Failure:

Investigation Postmortem

Seekers Spirits | Cambodia
May – June 2023

Incident type	Cooling system capacity limit during sustained production
Affected batch	Batch 50 - Seekers Mekong Dry Gin
Production date	29 May 2023
Discovery date	12 June 2023
Root cause confirmed	Post-August 2023
Batches at risk	5 (full Mongolia shipment requirement)
Batches affected	1 (containment at earliest viable intervention point)
Resolution	Cooling system redesigned. Four structural improvements implemented Q3 2023.

Table of Contents:

Section 1: Systems Context

1.1 The Operation

Seekers Spirits is a craft distillery in Phnom Penh, Cambodia. At the time of the incident, production ran on two stills connected to a coolant reservoir that removed heat during distillation. If the coolant became too warm, condensation became less reliable and spirit quality could drop.

The main still was Mae, a 300L system with a stable record. Under sustained load, coolant reservoir temperature consistently stayed within 43 to 50°C. Mae was the baseline for process and product quality. No prior failure had been attributed to Mae.

The second was a 2000L still added to increase capacity for the Mongolia market launch. Before the cooling failure, it had completed only two runs: Orange Liqueur batches 48 and 49, both lasting 7.5 hours. Batch 50, the first gin run on the system, ran for 12 hours.

Production was run by a small team. Final authority over batch approval, shipment release, and production sequencing sat with me. All stop, continue, and escalation decisions were made unilaterally.

Quality control took place after production, following a mandatory resting period. Batches were assessed by sensory evaluation and marked go or no go. For the 2000L still, there were no automated thresholds, instrument-based criteria, or fixed operating limits for long runs. Decisions were made from experience rather than defined checks.

1.2 The Constraints

The failure happened against a fixed shipment deadline. The Mongolia market entry represented over thirty percent of projected near-term revenue. Freight and permits had already been paid for. A delay would have meant writing off those costs, missing the shipping window, and putting the distributor relationship at serious risk.

By the time the failure appeared, there was no flexibility left in the schedule. The period between May and the August shipment left no margin for failure. Each production run used time and material that could not be recovered, and the resting period could not be shortened without risking quality. No runway remained.

Quality was also non-negotiable. This was the product's first commercial release in that market. A delay risked losing the market entry. Shipping compromised products risked losing the relationship. Each outcome was unacceptable, but they pushed in opposite directions. One pushed for speed. The other limited it. Every recovery decision had to sit between those two constraints.

The failure happened in that setting: at scale, under deadline pressure, and in higher ambient temperatures than the system had been tested in.

1.3 Observability Gaps

Gap 1: Display-Only Observability With No Retention

The system displayed distillate output temperature in real time during each run, but nothing was written to storage. During the gin run on 29 May 2023 (batch 50), temperature was not observed because condensation was taken as confirmation that the system was working.

The only recorded temperature for batch 50 is 61.8°C, noted informally after the run when the output felt unusually hot. By the time it was recorded, the batch had already completed. The reading was high. No action was taken. It later became the only thermal data point available for the investigation.

In practice, data was generated but not stored. When the investigation began, no production-time data remained.

Gap 2: Coolant Reservoir Not Instrumented During Production

Coolant reservoir temperature was never measured during production. The 1 to 2°C per hour increase was only identified after shipment, when the reservoir was instrumented during controlled testing and the result confirmed by the manufacturer.

During batch 50's 12-hour run and during all earlier runs, that increase was not visible. There was nothing to monitor, no defined threshold, and no point where intervention could happen.

The investigation was limited by two gaps: data that was not retained, and data that was never captured during production. Without either, there was no record of how the system behaved. The failure could only be reconstructed after the fact.

Section 2: Discovery Timeline

2.1 Pre-Incident Context (May – Early June 2023)

The first production runs on the 2000L still were two Orange Liqueur batches, batch 48 (8 May) and batch 49 (15 May). Each ran for 7.5 hours, passed quality control (QC), and matched the expected product profile. These runs were the only reference for how the 2000L still behaved during production prior to the incident.

During these runs, I observed that the 2000L still ran hotter than the 300L still, with distillate output moving from cool to warm to hot based on touch. No temperature data or coolant reservoir readings were recorded. When both Orange Liqueur batches passed QC weeks later, I treated the temperature pattern as normal. When the same pattern appeared during batch 50, I applied the same reasoning without re-examining the assumption.

Batch 50, the first gin run on the 2000L still, began on 29 May.

2.2 Discovery and Containment - 12 June 2023

9:00 AM: Initial Detection

Signal	During QC tasting of batch 50, I noticed the flavour was off. The gin was muted, with a flat finish that did not match the usual profile.
Interpretation	Possible causes included palate bias, sample contamination, not enough resting time, a mistake in the recipe or weighing error, or variation during the run. Nothing clearly pointed to one cause.
Confidence	~50%. Based on a single tasting and a single sample. The issue seemed real but not yet confirmed across the batch.
Action	I asked the production team to taste independently.

10:00 AM: Signal Confirmation

Signal	The production team confirmed the issue. Batch 50 was then tasted alongside a batch from the 300L still. The difference was clear side by side.
Interpretation	Palate bias, ingredient variance, and storage effects were largely ruled out. The issue appeared to come from the production system, most likely the 2000L still.
Confidence	~70%. The issue was real and affected the batch, but the root cause was still unknown.
Action	Prepared to contain the batch.

10:30 AM: Containment Decision

Signal	One batch had completed distillation and was resting (20% of required Mongolia volume). Batch 2 preparation was underway with botanicals already weighed. No bottling or labelling had occurred.
Interpretation	The issue was likely coming from the production system. The root cause remained unknown.
Confidence	~70%. Enough to act.
Action	Batch 50 tank locked. Batch 2 preparation halted and fresh botanicals discarded. Exposure limited to one batch.

Afternoon: Hypothesis Elimination

Signal	Batch-level causes were checked one by one. Independent tasting ruled out palate bias. Sample and resting effects were checked. Recipe and weighing were reviewed. Logs and system data showed no errors or alarms.
Interpretation	No operator error was found. There were no alarms, procedural violations, or clear deviations. Nothing at the batch level explained the issue. Attention shifted to the 2000L still as the likely source.
Confidence	~75–80% system-level issue. Each check reduced the likelihood of a batch-level cause. The 2000L still was now the most likely source.
Action	The investigation shifted from a batch issue to a system issue. A diagnostic run on the 2000L still was considered and rejected because it would use material without improving the decision in the time remaining.

Evening: Recovery Planning and Deliberate Escalation Delay

Signal	By evening, continuing production on the 2000L still within the shipment window was no longer viable. The focus shifted to recovery.
Interpretation	Escalating without a clear plan would lead to open-ended diagnosis under time pressure. It was better to escalate with defined options and a recommended path.
Confidence	~80%. Enough to commit to a recovery path, but not enough to identify the root cause.
Action	Escalation was delayed until the following morning. The evening was used to confirm recovery on the 300L still, calculate the number of runs needed with a 20% buffer, and define the recovery plan.

2.3 Escalation and Recovery: 13 June – late June 2023

On the morning of 13 June, I brought department heads together with a recovery plan, not an open problem. The root cause was still unknown. All Mongolia production moved to the 300L still with immediate effect, and the 2000L still remained paused until the cause was identified.

The escalation was used to commit to a plan, not reopen diagnosis. Reopening the problem would have cost time and broken alignment. Three actions were set: production was prioritised around the Mongolia shipment over all other products, additional botanicals were purchased to support the recovery plan, and production frequency increased.

Production on the 300L still moved into sustained operation, with two distillations per day across the recovery window. The tradeoff was clear: lower output in exchange for using a system known to produce product that passed quality control. With a fixed shipment deadline and non-negotiable quality requirement, the priority became using the system that was known to work.

A 20% buffer was built into the production plan. This allowed for failed batches or delays without forcing rushed runs. One batch failed because of an incorrect recipe input, but it was identified immediately and scrapped with no further impact.

As recovery continued, the main risk shifted from the system to the team. When signs of fatigue appeared, including changes in body language, less precise log entries, or skipped SOP steps, production was slowed and the buffer absorbed the shortfall instead of maintaining pace. Production finished by late June. After the mandatory resting period, inventory was finalised and held pending outbound logistics.

2.4 Post-August 2023: Shipment and Root Cause Identification

The Mongolia shipment departed in August 2023. Once it had left, I went back to investigate the issue.

The investigation found a gradual coolant temperature increase of 1 to 2°C per hour, reaching approximately 24°C over a 12-hour run. I shared this with the manufacturer, and we confirmed that the cooling reservoir was too small for sustained operation.

The issue had not appeared in earlier testing because the testing conditions did not match production. European test environments ran at 15 to 20°C, compared with 30 to 35°C in Cambodia, and test runs lasted 6 to 8 hours rather than 12 to 24 hours. The pre-cooler also depended on city water temperature, which was much higher in Cambodia.

This changed a core assumption. Visible condensation had been treated as enough evidence that the system was working properly. It was not. The incident showed that visible condensation did not guarantee quality. From that point on, distillate output temperature was tracked.

Full volume was delivered on schedule and met quality requirements. No customer commitments were broken.

2.5 Decision Checkpoints and Confidence Progression

Checkpoint	Timestamp	Confidence	Basis	Action Taken
Pre-incident thermal observation	May–early June	~10%	Two Orange Liqueur runs completed at 7.5 hours with no quality impact. Temperature pattern noted but not investigated.	Documented. Continued operating.
Initial QC deviation detected	12 Jun, 9:00 AM	~50%	Single taster, no cross-check. Palate bias, sample error, or insufficient resting time could not be ruled out.	Independent tasting requested.
Deviation confirmed	12 Jun, 10:00 AM	~70%	Multi-taster confirmation. Persists under direct comparison. Ingredient and storage variance reduced.	Containment decision initiated.
Containment - production paused	12 Jun, 10:30 AM	~70%	Stopping was recoverable. Continuing without knowing the cause was not.	Batch 50 tank locked. Batch 2 preparation halted and botanicals discarded.
Batch-level hypothesis eliminated	12 Jun, afternoon	~75%	Logs and system data reviewed. No procedural violations or operator errors found.	Investigation shifted from batch-level cause to system-level cause.
System-level framing confirmed	12 Jun, late afternoon	~80%	All batch-level causes eliminated. The 2000L still identified as the most likely source.	Recovery options modelled. Diagnostic run rejected as it would consume material without improving the decision.
Recovery path bounded	12 Jun, evening	~80%	300L still confirmed viable for recovery. Required run count calculated with buffer.	Escalation deferred to morning to present a defined plan rather than an open problem.
Escalation and alignment	13 Jun, morning	~80%	Recovery path presented with bounded options. Uncertainty acknowledged explicitly.	Production restarted on the 300L still. Botanical procurement secured.
Root cause confirmed	Post-August 2023	100%	1–2°C/hr drift confirmed by post-incident sensor measurement and manufacturer validation.	Cooling system redesigned. Structural improvements implemented to prevent recurrence.

Section 3: Root Cause Analysis

3.1 Mechanism of Failure

The 2000L still was connected to a cooling reservoir designed for the 300L legacy still. During longer runs, coolant returned to the reservoir faster than it could cool, so reservoir temperature rose by 1 to 2°C per hour from a starting point of about 30°C. Over a 12-hour gin run, that added roughly 24°C.

The same pattern occurred during the Orange Liqueur batches, but those runs stopped at 7.5 hours, before temperature rose far enough to affect flavour. The gin batch ran for 12 hours, so the temperature kept rising.

As coolant temperature increased, condensation became less reliable. Vapour was not fully condensed, distillate output temperature rose, and spirit quality dropped.

Coolant temperature progression: batch 50 vs control

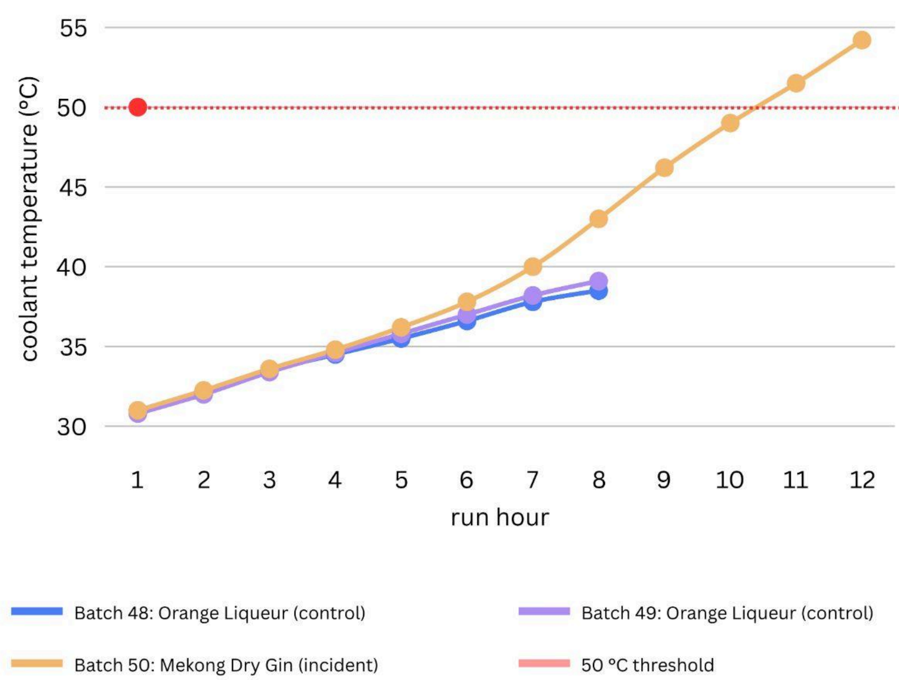

3.2 Failed Assumptions

Three assumptions, each reasonable in isolation, led to the failure and delayed detection. Each was based on limited evidence from small-scale, stable conditions and did not hold under production conditions at scale.

Assumption 1: Manufacturer configuration was sufficient across environments

Belief: The manufacturer's recommended setup, including the pre-cooler, was expected to work under local conditions.

Why it seemed reasonable: The manufacturer had asked detailed questions about the operating environment and recommended the pre-cooler specifically for Cambodia's climate. This suggested that the setup would work under local conditions.

Why it failed: The pre-cooler was connected to city water, which in Cambodia was significantly warmer than in European installations. Its effectiveness was assumed rather than tested under local conditions. As a result, it could not remove enough heat during long runs.

Assumption 2: Short-run behaviour was representative of long-run behaviour

Belief: If the system worked during Orange Liqueur production, it would work across runs on the 2000L still.

Why it seemed reasonable: The system appeared stable in production. Short runs finished without affecting flavour, and that pattern had held in both recipe development and scale-up.

Why it failed: The Orange Liqueur runs ended before temperature rose far enough to affect flavour. The behaviour only held within a 7.5-hour window. The assumption was applied beyond the conditions where it had been observed.

Assumption 3: Scale would not introduce capacity limits

Belief: Scaling from the 300L legacy still to the 2000L still in Cambodia would not introduce capacity limits.

Why it seemed reasonable: The Orange Liqueur runs on the 2000L still completed without issue, and the manufacturer had confirmed the setup for those conditions. There was no local precedent for cooling failure, and no comparable system at this scale.

Why it failed: Cooling capacity did not increase with system size. During long runs, the 2000L still generated more heat than the reservoir could remove. The reservoir had been sized for the 300L system. Scale exposed a limit that smaller runs and earlier production had not shown.

3.3 Why Testing Did Not Surface the Failure

Pre-deployment testing did not match the conditions in which the failure occurred. Runs were shorter, environmental conditions were different, and no instrumentation was in place to track coolant reservoir temperature during production. The system appeared stable until full-scale production showed where it failed.

Section 4: Impact Assessment

4.1 Affected Batch

batch_ID	product_name	batch_date	still_used	status	Disposition
50	Seekers Mekong Dry Gin	2023-05-29	2000L still	Quarantined	Failed QC. Temperature increased during the run. Batch quarantined. Not shipped.

4.2 Exposure Progression

At the point of containment, one batch had been produced against a total of five. If production had continued, batches 2 through 5 would have run under the same conditions. The blast radius was reduced from five batches to one.

4.3 Reversibility at Point of Containment

All downstream steps were still reversible at the point of containment. At 10:30 AM on 12 June, batch 50 was resting in tanks, and botanicals for batch 2 had been prepared but not yet added to the still. No bottling, labelling, or outbound movement had occurred. Customer exposure was zero.

Batch 50 tank locked; batch 2 preparation halted; botanicals discarded due to short usability window. Exposure was limited to one batch.

Section 5: Structural Improvements

Detection

Coolant reservoir temperature instrumented and monitored
A 50°C threshold has been set for coolant reservoir temperature across all batches. This threshold was taken from the highest recorded temperature on the 300L still, which consistently produces spirit that passes quality control. Above this level, cooling becomes less effective and product quality begins to drop.

During Batch 50, distillate output temperature was visible during the run and acted as an indicator of the cooling reservoir's performance. It was not recorded or acted on. Coolant reservoir temperature was not measured at all. Any deviation now requires immediate investigation while the batch is still active. This includes checking reservoir temperature and confirming that the chiller is operating as expected.

Decision Constraints

Hard stop on weak quality signals before raw material commitment
Any weak signal, sensory or operational, is now enough to halt production before raw materials enter the still. Once added, the batch is committed and cannot be recovered. The stop is a fixed rule, not a judgment call. This delays irreversible commitments until the issue is understood, while action is still possible.

Clear decision authority with clear escalation trigger
One person holds final authority to halt production, but the trigger for review is now fixed instead of being decided in the moment. Any threshold breach requires a documented review before the run continues. When one person both sees the signal and makes the decision, weak signals are easier to dismiss. The documented review removes that reliance.

Execution Controls

SOP versioning and update loop
When a batch requires a workaround or exposes an issue, the relevant SOP is updated before the next batch, and the earlier version is kept. During the incident, the assumption that visible condensation was enough to confirm quality was never written into the process. Versioned SOPs record changes directly, so the process does not depend on who is on shift.

Section 6: Operating Playbook

6.1 Investigation Model

The investigation follows a repeatable sequence used when the cause is unknown and the impact is still growing:

1. **Confirm the signal is real.**
   Rule out perception error, sampling error, and one-off noise before acting. At minimum, the signal should be confirmed independently.

2. **Scope before diagnosing.**
   Establish how much output is affected before identifying the cause. Scope first, then diagnose. Otherwise, containment may address only part of the problem.

3. **Eliminate highest-probability causes first.**
   Work through the possible causes in order of likelihood and ease of checking. Each one ruled out narrows the problem and increases confidence in what remains.

4. **Act before root cause is confirmed.**
   Containment should not wait for certainty when each additional cycle adds more at-risk output. Act when pausing is clearly safer than continuing, not when the cause is fully known.

5. **Escalate with a clear plan, not an open problem.**
   Escalate with options, not just a problem. If escalation happens before a recovery path is defined, it slows the decision. If it happens with clear options and a recommended path, a decision can be made immediately.

6. **Confirm root cause after stabilisation, not before.**
   When stabilisation is the main constraint, root cause investigation can wait, as long as that decision is explicit and a review is scheduled once the system is stable.

6.2 Core Principles

Principle 1: Repeated unexplained behaviour should be investigated, not ignored

A pattern that repeats under similar conditions without explanation is not noise. It should be checked before the batch continues. When the cause is unclear and no baseline exists, pause and investigate. Waiting for a threshold does not work when no threshold has been defined. Without a clear limit, the issue may not be noticed during the run. By the time it is obvious, it is too late to act.

Principle 2: Reversibility should guide decisions

When confidence is incomplete, choose the option that can be reversed at the lowest cost. In practice, this means choosing the option with limited downside over one where downside grows with each step. Do not wait for certainty. Delay irreversible commitments until confidence is higher. When reversibility is ignored, decisions become commitments before they are properly understood.

Principle 3: Decision runway must be actively preserved, not passively assumed

Under deadline pressure, options narrow over time. Preserving decision runway requires action: contain early, delay escalation until there is a clear plan, and pause long enough to understand the decision being made. Escalating too early adds coordination without improving the decision. Escalating too late removes options. Escalation should be triggered by a predefined condition or a fixed time, not by judgment under pressure. When runway is not preserved, decisions get made by the deadline instead of by choice.

Principle 4: Commitment points require explicit recognition before they are reached

A commitment point is any step where correction becomes costly or difficult. As it approaches, slow down, restate the assumptions, and check whether to continue. Moving forward on momentum instead of stopping to review is what turns a recoverable situation into an irreversible one.

6.3 Transferability Across Systems

The principles in this section come from operating a physical production system under real constraints: fixed deadlines, limited instrumentation, no automated alerts, and a failure that developed without crossing a defined threshold. They apply to software reliability for the same reason: the decisions are the same.

A gradual rise in temperature without a defined threshold is similar to latency rising under load without triggering an alert. In both cases, the change is visible before any limit is crossed, but without a baseline or an explicit check, it is easy to miss. What matters is deviation from the baseline, not whether a threshold has been breached. Monitoring shows the result, not the conditions that caused it.

Containment before root cause maps to halting a deployment while error rates are rising, even before a defined limit is crossed. Pausing is reversible. Continuing under uncertainty is not.

Reverting to the 300L system rather than continuing to debug the 2000L system under deadline pressure is the same as rolling back to a known stable version rather than attempting a live fix. When time is the constraint, predictability matters more. A known baseline with lower capacity is better than an unknown system with higher capacity when recovery is uncertain.