SAMPLE 2 · A ROLLOUT THAT HAS NOT PAID OFF

A technician app that has not lifted first-time fix

This is a fictional case. The organisation, people and figures are invented. The manager and colleagues were simulated from a written case file, so this shows how the Snapshot works, not a customer result. Everything from "Decision" onwards is the readout exactly as the Snapshot produced it, laid out for print; nothing in it was edited. The figures in it were calculated in code from the numbers supplied.

THE SITUATION

A fictional field-service company, ten months after rolling out a mobile app to its technicians. The vendor was offering a US$65,000-a-year AI dispatch module as the fix.

WHAT THE MANAGER TYPED TO START

> *We rolled out a mobile app for our technicians about ten months ago — job details, parts lookup, photos, customer sign-off. The case was first-time fix rate going from 70% to 85%. The KPI pack says we're at 72%. Honestly I don't know why it hasn't moved; the techs seem to like the app. The vendor is offering an AI dispatch module for $65,000 a year that they say will fix it. The Operations Director wants a recommendation. What do you suggest?*

HOW THE INVESTIGATION RAN

12 evidence requests, one at a time	**3** colleagues invited: field technician, call centre team leader, parts manager	**34** evidence items recorded	**14** explanations tested: 3 supported, 0 weakened or ruled out

References in square brackets, such as [E6], point to the evidence record at the back. The questions asked are listed there too.

DECISION

Correct the work

Do not buy AI dispatch now. Correct the existing information defects and compare repair outcomes consistently before funding another module.

1. **Proposed: remove fault-description truncation from the technician's job information.**
   Who: Integration technical lead implements; · Review: Proposed initial review: 2026-11-29, checking approval, execution and technical acceptanc… · Needs agreement
2. **Proposed: shorten stock-refresh delay and explicitly resource confirmation where a dispatch decision depends on a part being aboard.**
   Who: Integration lead changes and checks the refresh; · Review: Proposed initial review: 2026-11-29, checking the actual setting, execution, stock agreem… · Needs agreement
3. **Proposed: correct the basis used to judge whether the app delivered its benefit.**
   Who: Operations reporting lead reconstructs the comparable measure; · Review: Proposed review: 2026-11-29, before any purchase recommendation relies on the reported sh… · Needs agreement

Stop if: At the proposed purchase review, if there is still no checkable link between the module's capabilities and this operation's failed first visits, and no supported value case covering its full costs, close the current pur…

Why

The record identifies concrete problems in the current setup, while the vendor's claimed benefit has not been demonstrated for your work. Fixing these problems is justified; promising that they will deliver the first-time-fix target is not.

Next steps

1. Proposed: remove fault-description truncation from the technician's job information.

Needs agreement from CRM-to-app service owner, identity unconfirmed, with the Operations Director sponsoring the change

What changes	Change the integration so technicians can read the complete CRM fault description before attending the job. If the present app field cannot hold it, provide a supported, job-linked full-text view rather than silently discarding the remainder. Compare long source descriptions with the technician-visible version before release. This changes a documented design limitation; no conflicting written rule has been supplied.
Who does it (proposed)	Integration technical lead implements; call-centre lead and a field technician check usability and completeness. Identities and capacity require confirmation.
Measure	Completeness of source descriptions in the technician-accessible job view, clarification calls caused by missing text, and subsequent repeat visits recorded as insufficient job information. Distinguish descriptions exceeding the current limit from shorter descriptions.
Starting point	The specification truncates text beyond 255 characters (E16). The technician estimates frequent clarification calls, but neither their frequency nor associated repeat visits has been measured (E5, E8).
Review (proposed)	Proposed initial review: 2026-11-29, checking approval, execution and technical acceptance. Review repair outcomes after corrected visits have completed the current 30-day repeat-visit window; that outcome review necessarily follows the initial review.
Would show this is wrong	If complete descriptions reach technicians but missing-information calls and related repeat visits persist, truncation is not a sufficient explanation for those outcomes. That would not negate the demonstrated information loss in the old mapping.

Watch-outs

- Needs Integration engineering and release capacity: not confirmed.
- Needs Approval to change the mapping or provide a supported full-text job view: needs approval.
- Needs Call-centre and technician time for source-to-app acceptance checks: not confirmed.
- If descriptions become longer or an integration release changes field handling: partly works.
- Extra work falls on: Engineering and support initially, with acceptance-check work for the call centre and technicians.
- If it goes wrong, nobody owns the recovery work yet.
- Safeguard: Propose source-to-app completeness checks for releases and a recoverable deployment. The service owner must authorise support and rollback responsibility before release; capacity is unknown. These targeted checks protect against renewed silent information loss without a wider system replacement. No disruption is established by the record.

2. Proposed: shorten stock-refresh delay and explicitly resource confirmation where a dispatch decision depends on a part being aboard.

Needs agreement from Integration service owner for the refresh setting, and parts/dispatch process owner for stock-confirmation responsibilities

What changes	Have the integration owner validate the reported setting and change the reported weekly refresh to nightly, checking that consumption transactions reach the app. Nightly data can still be wrong after intraday consumption. For jobs whose scheduling depends on a specific part being aboard, explicitly propose that the parts coordinator obtains technician confirmation before dispatch promises availability. The record does not establish existing reservation or physical-check controls: these safeguards are new proposals requiring an owner, authorisation and capacity. Until that capacity is approved, do not treat the display alone as confirmed availability.
Who does it (proposed)	Integration lead changes and checks the refresh; parts coordinator obtains required confirmations; technician confirms physical availability; dispatch uses the confirmation. These assignments are not agreed.
Measure	Agreement between consumption records, displayed stock and confirmed physical availability, especially later in the week; separately track repeat visits recorded as part not on van. Record confirmation effort so the correction does not hide additional work.
Starting point	The parts manager reports Monday syncing and stale availability later in the week (E27). A technician reports a displayed contactor that was absent (E9). There is no measured mismatch rate or verified link to scheduling. Part not on van accounts for 31% of recorded repeat-visit reasons in the supplied report, not a proven share caused by stale data (E22).
Review (proposed)	Proposed initial review: 2026-11-29, checking the actual setting, execution, stock agreement and whether confirmation work is staffed. Review first-time-fix effects only after affected visits complete the 30-day repeat-visit window.
Would show this is wrong	If nightly refresh improves stock accuracy but parts-related returns do not change, stale display data is not a sufficient explanation for those returns. If accuracy does not improve, examine the consumption feed rather than assuming refresh frequency was the whole problem.

Watch-outs

- Needs Integration capacity to validate, configure and support the change: not confirmed.
- Needs Authorisation for the configuration change: needs approval.
- Needs Parts-coordinator and technician time for proposed pre-dispatch confirmation of required stock: needs approval.
- If faster intraday parts consumption or reduced parts-coordinator availability makes overnight stock data less dependable: partly works.
- Extra work falls on: Integration support, the parts coordinator, technicians and dispatch. The proposed confirmation safeguard adds work; it is not an established existing control.
- If it goes wrong, nobody owns the recovery work yet.
- Safeguard: Before activation, propose an authorised integration-support role for failed-sync recovery and a parts-process owner for handling unresolved stock discrepancies. Their capacity and authority remain unconfirmed. Target physical confirmation to stock-dependent scheduling decisions rather than checking every item on every van. The effort may prevent avoidable visits and enable reliable reassignment, but its cost is unmeasured. Do not claim the display is live stock or remove any subsequently verified controls.

3. Proposed: correct the basis used to judge whether the app delivered its benefit.

Needs agreement from Operations performance-measure owner, with the Operations Director deciding how the original target will be assessed

What changes	Keep the current 30-day repair-stays-fixed measure. Recalculate baseline and post-app performance using that same rule and consistent eligibility, showing underlying job counts. If historical records cannot support reconstruction, explicitly withdraw the like-for-like benefit comparison and establish a consistent current baseline. Do not lower the target or attribute the apparent gap to the app merely because the definition changed.
Who does it (proposed)	Operations reporting lead reconstructs the comparable measure; the performance-measure owner approves it; the Operations Director decides the target basis. Ownership is not agreed.
Measure	Whether baseline and current results use the same return window, eligibility and job-count basis; the resulting comparable first-time-fix rate, if reconstruction is possible.
Starting point	The original 70% baseline used a 14-day repeat window. From month 4 the pack uses 30 days. Months 1–2 were 71% and 72% before that change; later months are reported between 71% and 73% (E3, E4).
Review (proposed)	Proposed review: 2026-11-29, before any purchase recommendation relies on the reported shortfall.
Would show this is wrong	A consistent reconstruction showing similarly limited improvement would weaken measurement change as an explanation for the apparent shortfall. It would not remove the documented comparability defect in the supplied pack.

Watch-outs

- Needs Historical job dates, same-fault return links and eligible-job records: not confirmed.
- Needs Reporting analyst capacity: not confirmed.
- Needs Approval of a consistent outcome definition and target basis: needs approval.
- If job mix changes or historical records prove insufficient for reconstruction: partly works.
- Extra work falls on: The reporting role and measure owner.
- If it goes wrong, nobody owns the recovery work yet.
- Safeguard: Agree the definition and record requirements, and retain underlying counts so later changes can be checked. Analyst capacity and authority remain unconfirmed. If reconstruction is infeasible, a consistent prospective baseline is preferable to an expensive attempt to manufacture historical precision. This does not establish whether the original target was realistic.

Why not a test instead: The narrow information-loss, stock-display and measurement-comparability problems are findings with identifiable corrections. Correct them through the appropriate owners and verify implementation; a test is not needed to establish that losing text or displaying stale stock is undesirable. Their contribution to failed first visits remains uncertain and must be measured separately.

Why not correct it directly: Corrections are recommended only within the supported findings. There is not enough evidence to replace dispatch, assign blame for parts preparation, reset the 85% target or claim that these corrections will deliver it.

Stop condition: At the proposed purchase review, if there is still no checkable link between the module's capabilities and this operation's failed first visits, and no supported value case covering its full costs, close the current purchase request without an order. This means stop this procurement proposal, not switch off the existing app or conclude that AI dispatch could never help. — by 2026-11-29, or before signing if the purchase decision is earlier.

What is known

- About ten months after rollout, the manager reports 72% first-time fixes against an original expectation of moving from 70% to 85%. The date by which the target was due is unknown. [E1, E23]
- From month 4, the repeat-visit window widened from 14 to 30 days. Later reported rates are therefore not directly comparable with the original baseline. The first two post-app months were also below target before that change. [E3, E4]
- The integration specification discards fault-description text beyond 255 characters. A technician reports incomplete notes requiring clarification calls. [E16, E5]
- The stock display can show a consumed part as available. The parts manager reports weekly refreshing and an available nightly setting, but no implemented correction has been demonstrated. [E27, E28, E9]
- The supplied repeat-visit report records part not on van at 31%, insufficient job information at 18%, specialist needed at 22%, customer access at 14% and other at 15%. These are reason categories, not established underlying causes or preventable shares. [E22]
- The manager recalls a vendor offer of about $65,000 annually and up to roughly a 10-point improvement, supported by other-industry cases rather than analysis of this operation's data. [E30, E31]

The explanations

Errors in app-supported job information create corrective work that prevents first-visit completion. Supported enough for the next decision

Accounts for: The documented mapping explains why long fault descriptions lose text and is consistent with clarification calls.
Does not account for: It does not establish which failed first visits were affected or how much repair performance would improve after correction.

Weekly van-stock synchronisation leaves consumed parts displayed as available, potentially misleading pre-visit parts preparation. Supported enough for the next decision

Accounts for: Stale availability is supported by the parts manager's account and the technician's stock discrepancy.
Does not account for: Its influence on dispatch and repeat visits is unverified. Nightly refreshing would still leave intraday uncertainty.

A change in first-time-fix measurement makes the reported 72% non-comparable with the original 70% baseline. Supported enough for the next decision

Accounts for: The changed repeat window causes a break in comparability between the original baseline and later KPI figures.
Does not account for: It does not establish the size of that effect or explain the entire shortfall; early post-app results were also below target.

Parts unavailable on technicians' vans prevent first-visit completion, independently of whether preparation has an assigned owner. Open – not yet tested

Accounts for: Actual missing parts are a distinct possibility supported by the recorded part-not-on-van category.
Does not account for: The report does not distinguish stock-display errors, unavailable supply, preparation failures or faults only identified during the visit.

Specialist capability is not available at the first visit for some faults, requiring another visit. Open – not yet tested

Accounts for: The specialist-needed category makes first-visit capability or assignment a relevant alternative.
Does not account for: There is no job-level evidence showing whether specialist need was predictable or would be addressed by the proposed module.

Customer-access constraints prevent completion at the first visit. Open – not yet tested

Accounts for: Customer access is a recorded reason for returns and may constrain results independently of app

information.
Does not account for: Its preventability and relationship to dispatch have not been established.

Necessary pre-visit parts preparation has not been assigned to an accountable owner. Open – not yet tested
Accounts for: An accountability gap in pre-visit preparation could contribute to missing parts.
Does not account for: Preparation responsibilities and histories were unavailable. That does not establish unowned work or existing reservation controls.

The original business case overestimated how strongly the app's feature set could influence first-time fixes. Open – not yet tested
Accounts for: An overoptimistic original benefit expectation could explain why available app features have not delivered the target.
Does not account for: The original assumptions and supporting evidence were not supplied, so this remains untested.

Increased incoming job demand is straining the capacity needed to complete repairs on the first visit. Open – not yet tested
Accounts for: Higher demand could strain first-visit capacity.
Does not account for: No demand or capacity comparison was supplied.

Administrative work has shifted from office staff to technicians, reducing time available for diagnosis and repair during the first visit. Open – not yet tested
Accounts for: Administrative work transferred to technicians could reduce repair time.
Does not account for: No before-and-after task evidence establishes such a transfer.

A review queue for app-submitted job evidence holds up resolution beyond the first visit. Open – not yet tested
Accounts for: An evidence-review queue could delay resolution.
Does not account for: No submission, review or release history establishes a queue.

The previous recording process is still running alongside the app, consuming technician time through duplicate entry. Open – not yet tested
Accounts for: Duplicate recording could consume technician time.
Does not account for: The CRM feeding the app is not itself evidence of unnecessary duplicate entry.

Technician performance measures reward completing visits rather than achieving a durable first-time fix. Open – not yet tested
Accounts for: Conflicting performance measures could encourage visit completion over durable repair.
Does not account for: No technician targets or reward rules were supplied.

The tool does not fit the real work Open – not yet tested
Accounts for: Poor fit with real work could extend beyond the specific information defects already found.
Does not account for: The record does not support a broader conclusion that the app is unsuitable; technician approval does not establish effectiveness either.

Figures

No figures computed yet.

- The proposed annual fee is approximately $65,000 according to the manager's recollection. Charges beyond that fee and implementation requirements are unknown, so it is not a complete investment cost. [E30, E32]
- The claimed improvement of up to roughly 10 points is a vendor upper-bound claim, not a measured local outcome. It cannot support a savings or payback estimate here. [E30, E31]
- Clarification calls and stock discrepancies indicate operational friction, but labor cost, avoidable return cost and recoverable capacity are unquantified. The recorded reason percentages must not be treated as achievable savings. [E5, E9, E22]
- The corrections also consume engineering, reporting and operational time. Obtain approval for that work and the proposed stock-confirmation safeguard rather than treating internal effort as free. [E19, E26, E28]

Solved or only promised

- A working correction that preserves complete fault descriptions has not been demonstrated.: **unknown**
- Nightly stock refresh is reported as an option, not an implemented or promised fix.: **unknown**
- A comparable historical benefit measure has not been produced.: **unknown**
- The dispatch uplift is a vendor promise, not a demonstrated result for this operation.: **promised**

Not established

- That dispatch is the principal constraint, or that AI dispatch would be ineffective.
- That the identified defects were introduced by the app or explain a quantified share of failed first visits.
- That 85% is attainable through the app, remains the right target, or was defined on the current counting basis.
- That parts preparation is unowned, or that reservation and physical-confirmation controls currently operate.
- That repeated accounts of the same technician incident constitute independent incidents.
- That corrective owners, authority, staffing or review dates have been agreed.

Missing evidence

- Comparable eligible-job and first-time-fix counts with consistent repeat-visit rules.
- Job histories linking information, parts, specialist or access problems to visit outcomes.
- Confirmed technical and operational decision owners, implementation capacity and correction costs.
- Verified module pricing, implementation requirements, full costs and evidence of fit with this operation's failure causes.
- The original benefit assumptions and the economic value of avoided return visits.

What could change this conclusion

- A consistent outcome series could show more or less improvement than the current pack suggests.
- Successful technical corrections followed by unchanged repair outcomes would redirect attention to the remaining repeat-visit causes.
- Operation-specific evidence linking dispatch capability to preventable failures, with supported value exceeding full cost, could justify reconsidering the module.
- If proposed corrections lack authority, feasible implementation or staffed safeguards, narrow or defer those changes rather than transferring unsupported work to technicians or parts staff.

Appendix A: the questions asked

Each round asked for one thing someone already had. A question could go to the manager or, by email link, to a colleague who sees only that question.

	Answered by	Request
C1	—	Send the existing monthly first-time-fix report covering the three months used for the original baseline and every month since the app rollout. Include four fields: reporting month, eligible jobs, jobs fixed on the first visit, and the counting rule used, including any changes. Raw counts are enough; there is no need to calculate percentages. If you don't have exactly this, send the closest thing you have and say how it differs, or tell us what you have seen first-hand.
C2	—	From the KPI pack you already have, send the first-time-fix figures for months 4–10 by whichever breakdown it already supports: site or location, shift, job type or category, or day of week. Use one existing breakdown only. Include up to four fields: month, group, reported first-time-fix figure, and number of jobs covered if shown. Send raw figures as displayed; do not calculate anything or rebuild missing counts. If you don't have exactly this, send the closest thing you have and say how it differs, or tell us what you have seen first-hand.
C3	field technician	For the last five same-fault repeat jobs you personally attended in the past month, share the existing job-history entries showing four things: job reference, what prevented completion at the first visit, what parts preparation occurred and who did it, and any app information or function that needed correcting or could not support the job. Please distinguish recorded reasons from what you directly observed. If you don't have exactly this, send the closest thing you have and say how it differs, or tell us what you have seen first-hand.
C4	field technician	For your last three completed jobs this week, describe the recording and approval work you actually did, using existing job entries if available. Give four details for each task: what the task was, who does it now versus before the app, where it is entered including any duplicate record, and any reviewer involved with the time submitted and time released if visible. Include any paperwork or older system still used alongside the app. If you don't have exactly this, send the closest thing you have and say how it differs, or tell us what you have seen first-hand.
C5	—	Send the technician performance scorecard or targets you already use for the latest completed month. Include up to four fields: measure, target, result, and any linked reward or consequence stated in the record. We need to distinguish targets for visits completed from targets for repairs that stay fixed. If you don't have exactly this, send the closest thing you have and say how it differs, or tell us what you have seen first-hand.
C6	call centre team leader	For the last three technician calls you personally handled about job notes stopping mid-sentence, share what the existing call or job entries show: job reference and date, original call-centre note, note visible to the technician if available, and what information you supplied during the call. If no entries are available, describe the most recent such call you handled firsthand, including whether you could see both versions. If you don't have exactly this, send the closest thing you have and say how it differs, or tell us what you have seen first-hand.
C7	field technician	Send the latest existing defect ticket or support thread about CRM fault descriptions being cut to 255 characters in the technician app. Include four details if shown: named technical owner, current status, proposed or implemented remedy, and evidence that the remedy works on a real long fault description. A promised fix is different from a working one. Look in the integration support queue or vendor correspondence. If you don't have exactly this, send the closest thing you have and say how it differs, or tell us what you have seen first-hand.
C8	—	For the last three same-fault repeat visits you personally scheduled in the past month, share the existing dispatch entries or your firsthand account of four details: job reference and date, why another visit was needed, how the first technician was selected, and what changed to make the return visit possible. Use only details already visible to you; distinguish recorded reasons from your own observations. If you don't have exactly this, send the closest thing you have and say how it differs, or tell us what you have seen first-hand.

	Answered by	Request
C9	—	Send the section of the original app approval paper or benefits slide that explains the increase from 70% to 85% first-time fixes. Include four details if recorded: repeat-visit causes the app was expected to remove, evidence supporting that expectation, date the benefit was due, and expected annual value. Send the existing figures and assumptions without calculating anything new. If you don't have exactly this, send the closest thing you have and say how it differs, or tell us what you have seen first-hand.
C10	parts manager	For the most recent three jobs you personally handled in the past month where a return visit was needed because a part was not on the van, share the existing parts entries or your firsthand account. Include four details: job reference, when the required part became known, whether it was available in stores before the first visit, and who was responsible for getting it onto the van under the rule used. Look in the parts reservation or issue history; no new report is needed. If you don't have exactly this, send the closest thing you have and say how it differs, or tell us what you have seen first-hand.
C11	—	Send the existing AI dispatch proposal or vendor email behind the $65,000 annual offer. Include four details if stated: which repeat-visit causes it addresses, evidence from actual repair outcomes supporting the claim, implementation requirements, and total charges beyond the annual fee. Send what is already in the proposal or correspondence; do not seek anything new from the vendor. If you don't have exactly this, send the closest thing you have and say how it differs, or tell us what you have seen first-hand.
C12	field technician	For the last three visits you personally scheduled in the past week using the app's van-stock display, share the existing dispatch entries or what you saw firsthand. Include four details: job reference and date, part availability shown in the app, actual availability if subsequently confirmed, and what happened to the visit because of any difference. We need to distinguish stale stock that actually influences scheduling from a display problem that dispatch already works around. Look at the dispatch screen or existing job messages; no new report is needed. If you don't have exactly this, send the closest thing you have and say how it differs, or tell us what you have seen first-hand.

Appendix B: the evidence record

Each item is frozen when it is recorded and graded as a fact (from a record) or an observation (an account). The summaries are the Snapshot's own.

	Type	Source	Summary
E1	Observation	Manager opening statement	Opening statement
E2	Observation	Manager via C1: Manager, C1 response	The supplied pack lacks the job counts needed to check the percentages and uses an annual baseline average rather than the requested baseline months.
E3	Fact (record)	Manager via C1: Monthly KPI pack reproduced by Manager, C1	The pack records first-time-fix rates near the baseline, including the first two months before the counting-rule change. Month 3 is unavailable and later months are supplied only as a range.
E4	Fact (record)	Manager via C1: Monthly KPI pack, page 2 note reproduced by Manager, C1	The documented repeat-visit window widened from 14 to 30 days beginning in month 4, changing the first-time-fix measure.
E5	Observation	field technician via C3: Field technician, C3 firsthand account	The technician reports incomplete job notes that require a call to the call centre. This establishes a reported information gap and additional clarification work, not its cause or effect on first-time fixes.
E6	Observation	field technician via C3: Field technician, C3 firsthand account	The technician reports one discrepancy between app-displayed van stock and physical stock. The account does not establish why the discrepancy occurred or whether it caused a repeat visit.
E7	Observation	field technician via C3: Field technician, C3 account of evidence availability	The requested job-level records and parts-preparation details were unavailable from this contributor. This limits linkage to repeat visits; it does not establish that preparation was absent or unassigned.

	Type	Source	Summary
E8	Observation	field technician via C4: Field technician, C4 firsthand account	The technician reports incomplete job notes that require a call to the call centre. The frequency is an estimate, and the account does not establish that these jobs required repeat visits.
E9	Observation	field technician via C4: Field technician, C4 firsthand account	The technician reports one discrepancy between app-displayed van stock and physical stock. Its origin and effect on first-visit completion are unknown.
E10	Observation	field technician via C4: Field technician, C4 statement of evidence limits	The response does not supply the requested evidence about work transfers, duplicate recording or approval queues. These explanations remain untested, not disproved.
E11	Observation	Manager via C5: Manager, C5	The manager relays technicians' favorable opinion. This does not establish first-time-fix improvement or how performance incentives affect their work.
E12	Observation	Manager via C5: Manager, C5	The requested technician targets and incentive rules were unavailable to the manager for this answer. Their absence from this response does not establish that such rules do not exist.
E13	Fact (record)	Manager via C5: Monthly KPI pack reproduced by Manager, C5	The manager reproduces the same KPI figures already recorded in E3. This is repeated evidence, not an independent confirmation or new performance measurement.
E14	Fact (record)	Manager via C5: Monthly KPI pack, page 2 note reproduced by Manager, C5	The reproduced counting-rule note repeats E4 and supports the existing finding that later rates are not directly comparable with the baseline. It does not explain the entire performance shortfall.
E15	Observation	call centre team leader via C6: Call centre team leader, C6	The call-centre team leader reports that fault descriptions originate in the CRM and reach the technician app through a sync. This identifies the information path, not unnecessary duplicate recording.
E16	Fact (record)	call centre team leader via C6: CRM-to-app sync specification reproduced by call centre team leader, C6	The reproduced specification explicitly truncates fault descriptions beyond 255 characters when mapping them to app job notes. This establishes a documented information-loss mechanism, but does not show which actual jobs were affected. No conflicting written requirement has been supplied.
E17	Observation	field technician via C7: Field technician, C7 firsthand account	The technician again reports incomplete notes requiring clarification calls. This repeats the account already recorded; it is not independent corroboration, a measured frequency, or evidence linking the information loss to repeat visits.
E18	Observation	field technician via C7: Field technician, C7 evidence availability	The respondent does not have the requested technical records. This does not establish that no ticket or vendor correspondence exists.
E19	Observation	field technician via C7: Field technician, C7 evidence limits	The technician cannot verify technical ownership, remediation status or successful correction. This limitation does not contradict the sync specification or establish an ownership gap.
E20	Observation	Manager (relaying for Dispatcher who personally scheduled recent same-fault return visits) via C8: Manager relaying for dispatcher, C8	The manager relays complaints and a claim that these issues predate the app. This is hearsay, not a verified baseline comparison or evidence that the app introduced these problems.
E21	Observation	Manager (relaying for Dispatcher who personally scheduled recent same-fault return visits) via C8: Manager relaying for dispatcher, C8	The answer provides no job-level scheduling evidence linking repeat visits to technician selection, preparation or changes before the return. These mechanisms remain unresolved.

	Type	Source	Summary
E22	Fact (record)	Manager (relaying for Dispatcher who personally scheduled recent same-fault return visits) via C8: App report reproduced by manager, C8	The reproduced report identifies recorded reasons for repeat visits. Parts availability, information, specialist requirements and customer access are distinct explanations worth separating. These categories do not establish underlying causes, preventability, app attribution or whether AI dispatch would address them.
E23	Observation	Manager via C9: Manager, C9	The manager cannot supply the original benefit assumptions, supporting evidence, delivery date or annual value. This limits evaluation of the business case; it does not establish that the original assumptions were unrealistic or undocumented.
E24	Fact (record)	Manager via C9: Monthly KPI pack reproduced by Manager, C9	The manager reproduces the figures already recorded in E3 and E13. These are not new measurements or independent corroboration, and they do not establish the original benefit assumptions.
E25	Fact (record)	Manager via C9: Monthly KPI pack, page 2 note reproduced by Manager, C9	The repeated counting-rule note is consistent with the existing measurement-comparability finding. It provides no new evidence about how much the changed window affects the reported rate.
E26	Observation	parts manager via C10: Parts manager, C10	The parts manager cannot supply the requested job-level records. This leaves parts preparation, availability before the visit and responsibility unverified; it does not establish an ownership gap.
E27	Observation	parts manager via C10: Parts manager, C10	The parts manager reports weekly stock synchronisation and stale displayed availability later in the week. This is consistent with the technician's stock discrepancy, but neither the setting nor its effect on repeat visits has been verified.
E28	Observation	parts manager via C10: Parts manager, C10	The parts manager reports an available configuration alternative that has not been requested. This is not evidence of an implemented change, a promised fix or departure from a written rule.
E29	Observation	parts manager via C10: Parts manager, C10	The contributor explicitly limits the account: it does not explain the requested return visits or establish when their required parts became known.
E30	Observation	Manager via C11: Manager, C11: recollection of vendor proposal	The manager recalls an approximate annual fee and an upper-bound benefit claim. Neither is reproduced from the proposal, and the claimed improvement is not a measured result for this operation.
E31	Observation	Manager via C11: Manager, C11: recollection of proposal support	The manager recalls evidence from other industries but no analysis of this operation's data. This leaves the module's fit with the recorded repeat-visit causes unestablished; it does not prove the module would be ineffective.
E32	Observation	Manager via C11: Manager, C11: statement of available information	The answer supplies neither a cause-to-capability link nor implementation requirements or full costs. This is a limit of the available answer, not proof that the proposal omits those details.
E33	Observation	field technician via C12: Field technician, C12	The technician repeats the stock-display discrepancy already reported in E6 and E9. This is a specific firsthand account, not another independent incident or confirmation that scheduling was affected.
E34	Observation	field technician via C12: Field technician, C12	The technician cannot link the discrepancy to a dispatch decision or visit outcome. The requested dispatch evidence remains unavailable; this does not establish that dispatch ignores the display.