LEARNING DESIGN DOCUMENT

First 15

A Scenario-Based Onboarding Experience for Last-Mile Delivery Professionals

CLIENT	MileOne Logistics (fictional)
DESIGNER	Kwame Yeboah · Learning Experience Designer
MODALITY	Blended self-paced · Rise 360 + Storyline 360
SEAT TIME	37–48 minutes across three LMS items
STANDARDS	WCAG 2.1 AA · SCORM 1.2 and 2004 4th Ed
VERSION	1.0 · September 2026

"The first fifteen minutes of a shift are the cheapest place to solve a problem, and the most expensive place to miss one."

MileOne Logistics is a realistic fictional organisation created for this portfolio project. No proprietary information or client data is used or implied. Every figure in this document is a design target, a seat-time estimate, or a proposed measure — never a reported result.

nimfah.com/work/first-15-last-mile-onboarding

Contents

HOW TO READ THIS DOCUMENT

Sections 01 to 05 are the design case: the problem, the model, the audience, the gap, and the objectives. Sections 06 to 08 are the build: architecture, branching, and the verified scoring model. Sections 09 to 13 are the production apparatus a client would hold me to.

01

The business problem

A last-mile delivery operation loses more service and safety ground in the first fifteen minutes of a shift than in the eight hours that follow. That window is where an associate finds the tire-pressure alert, the label they cannot read, the medical delivery with a receiving window, and the dispatch note that today's volume is heavier than planned. It is also the window where a new hire has the least context and the most pressure to simply leave.

Conventional onboarding is badly matched to this. It front-loads policy — handbooks, compliance modules, a checklist to acknowledge — then sends a nine-day-old associate into a morning that demands judgment, prioritisation, and the confidence to interrupt a dispatcher. The gap is not knowledge. Most new associates can recite the pre-shift checklist. The gap is decision-making under competing pressure, which is the one thing a checklist cannot teach and a knowledge check cannot measure.

THE DESIGN PROBLEM, STATED ONCE

How do you let someone practise the first fifteen minutes of a bad morning forty times before it counts?

Three costs in scope

- **Preventable roadside and safety events** traceable to a pre-shift condition that was noticed and not resolved.
- **Avoidable first-attempt failures** on deliveries where the information needed to succeed was available at 7:10 a.m. and not acted on — unreadable addresses, unconfirmed receiving windows, access instructions that cannot be followed as written.
- **Escalation friction in both directions:** real risks raised too late to be managed, and low-stakes items raised as emergencies, which consumes coordinator capacity and teaches associates that escalating is punished.

Organisational context

DIMENSION	CONTEXT
Operation	Regional last-mile carrier. Contracted parcel volume plus a growing medical-courier line with receiving-window and cold-chain requirements. Mixed cargo-van fleet across metro stations.
Hiring	High-volume, seasonally spiked. Compressed classroom window before an associate is on a route. Significant first-90-day attrition, so training cost per retained associate is sensitive to seat time.
Dispatch	Coordinators frequently promoted from the associate seat with operational knowledge but no training in coordination, prioritisation, or escalation intake.
Existing training	Policy and compliance modules, a printed pre-shift checklist, and shadowing with a senior associate. Nothing that rehearses decisions.
Culture constraint	Stations run on time pressure. Any training that reads as “slow down and be careful” will be discarded on the floor. The design has to make good judgment faster, not slower.

The framing that makes this fundable: MileOne is not buying a course. It is buying a shared decision standard that two roles can both be held to.

02

The decision model

Everything downstream — every lesson, scene, variable, and job aid — descends from three ideas. A reviewer should be able to recite them after five minutes. That is the test of a good instructional model.

One · Service, Safety, and Time

Three forces compete at the start of every shift. Learners resolve them in a fixed order of precedence, because "use your judgment" is not a teachable instruction and "safety first" alone does not tell anyone what to do second.

DIMENSION	THE RULE LEARNERS LEARN	WHY THE ORDER MATTERS
SAFETY	Resolve before the van moves. A gate, not a trade-off.	Anything that puts the associate, the public, or the load at risk is not negotiable against time. The only non-negotiable dimension.
SERVICE	Protect the promise you can keep. Renegotiate the one you cannot — early.	Service failures are usually caused by silence, not delay. A delay flagged at 7:12 is an operational adjustment; the same delay found at 11:40 is a broken promise.
TIME	The variable you manage, not the one you protect.	New associates protect the clock, because the clock is what visibly judges them. The model reframes time as the resource spent to buy safety and service.

Two · The 90-second triage

Every pre-shift discovery gets one of three dispositions, decided in under ninety seconds. This is the behavioural heart of the course and the thing learners carry into the van on a card.

	DISPOSITION
FIX	I can resolve this myself, inside the pre-shift window, safely, and within policy.
FLAG	This needs someone else's authority, information, or resources. Escalate now, with a complete message, and keep working while it is handled.
NOTE	This does not change today's plan, but it must be on the record.

> **THE INSIGHT THIS RESTS ON**
>
> New hires do not fail because they cannot tell a big problem from a small one. They fail because they treat every problem as binary — deal with it alone, or stop everything — and so they under-escalate real risk and over-escalate noise. Three dispositions beat two.

Three · The four escalation triggers

"Escalate when appropriate" is not criteria. These are. Any one of them turns a discovery into a Flag.

TRIGGER	DEFINITION
RISK	Roadworthiness, load security, or anyone's physical safety — the associate's, the public's, a customer's.
CUSTODY	Chain-of-custody or compliance conditions: cold chain, signature required, controlled or regulated goods, anything sealed or carrying a temperature indicator.

TRIGGER	DEFINITION
PROMISE	A committed window or service level now at risk. The trigger is the risk to the promise, not the confirmed failure.
AMBIGUITY	Information that cannot be resolved with what is available, where being wrong has a cost. An address you are 80% sure of is ambiguity.

THE TWO RULES THE SCENARIO TEACHES THROUGH CONSEQUENCE

A slow leak is a Flag. A cold-morning pressure variance is a Note.

Escalating with a complete message costs 40 seconds. Escalating with "there's a problem" costs twenty minutes.

03

Learner personas

Both personas are fictional composites created for this project.

Jordan Reyes

Delivery Associate · day 9 on route · primary audience

FIELD	DETAIL
AGE / TENURE	26 · second week solo
BACKGROUND	Two years warehouse fulfilment, six months food delivery. First job with a company vehicle and a route of their own.
DEVICE	Reads on a phone. Will not open a desktop-only resource. Skims; reads carefully only once something has gone wrong.
MOTIVATION	Wants to be good at this quickly, and wants to be seen as someone who does not need help.

> *"The tire light was on yesterday too and it was fine. Dispatch already said we're heavy today, so I'm not going to be the person who calls in about a label at seven in the morning."*

WHAT IS ACTUALLY IN THE WAY

- No model for the middle ground — every discovery feels like handle it quietly or make a scene, and quietly always feels safer socially.
- Cannot yet tell a degrading condition from a stable one. The tire light yesterday and the tire light today look identical.
- Believes escalating reflects on their competence. Nobody has told them that escalating early is what competent associates do.
- Time is the only feedback they get. Nobody has ever said "you rolled out four minutes late and that was the right call."

Marcus Boateng

Dispatch Coordinator · six weeks promoted · secondary audience

FIELD	DETAIL
AGE / TENURE	31 · three years as an associate before promotion
BACKGROUND	One of the station's strongest associates. Promoted for reliability, not coordination skill. Received a systems walkthrough and a login.
DEVICE	Two monitors and a headset. Will read a well-structured page at the desk. No time for anything in the 6:45–7:30 window.
MOTIVATION	Does not want to be the coordinator blamed for a preventable incident, and does not want to be the one whose routes roll out late.

> *“What I need at 7:10 is a sentence with four things in it. What I get is ‘hey, quick question’ and then four minutes of context. And the ones I actually need to hear about — I find out at eleven o'clock.”*

WHAT IS ACTUALLY IN THE WAY

- Escalation intake is unstructured, so triage happens conversationally and consumes the window.
- No language for coaching an associate in ten seconds without sounding like a reprimand — so he either fixes it himself or says nothing.
- Was never taught the model he is now expected to enforce, which makes enforcement feel arbitrary.
- Learned as an associate that escalating was discouraged, and is now the person who needs it to happen.

04

Performance gap

	CURRENT	DESIRED	GAP TYPE
Pre-shift discovery	Associates notice issues and proceed, or stop entirely and wait for instruction	Reach a Fix / Flag / Note disposition within 90 seconds and act on it	**Skill + mental model**
Escalation decision	Governed by personal comfort with interrupting a coordinator	Governed by four stated triggers, applied consistently	**Criteria absent**
Escalation quality	“There's a problem with one of my packages”	Four fields: condition, identifier, request, deadline	**Skill + no model given**
Route sequencing	Optimiser sequence accepted as given; custody items absorbed into volume	Time-sensitive and custody items protected before volume is absorbed	**Skill + authority ambiguity**
Documentation	Decisions documented inconsistently and after the fact	Exceptions scanned and noted at the point of decision	**Process + habit**

The gap is overwhelmingly one of skill and mental model, with one genuine knowledge component (the four triggers) and one process component (documentation). This is the finding that drives the modality split: a small amount of content to know, and a large amount of judgment to practise.

Design principles

Written as decision rules, so they can settle arguments during production.

- **Practise the decision, not the policy.** If a screen can be answered by remembering a sentence, it belongs in Rise as a knowledge check, not in the scenario.
- **Consequence before correction.** Show what the choice caused, then name the principle. Never the reverse.
- **Every distractor is someone's reasonable Tuesday.** No joke options. Each wrong choice must be one a competent, rushed person would defend.
- **Three dispositions, never two.** Any framework that forces a binary produces over-escalation.
- **Indigo means you can touch it.** One reserved interaction colour, never used for emphasis.
- **Accessible by construction.** If a design needs a retrofit to be accessible, it is the wrong design. Decide it at the storyboard, not at QA.

05

Learning objectives

TERMINAL OBJECTIVE

By the end of this experience, given a pre-shift review containing competing safety, custody, service, and volume pressures, learners will be able to triage each discovery into Fix, Flag, or Note within ninety seconds, justify the disposition using MileOne's four escalation triggers, and roll out with all custody-controlled and time-sensitive commitments protected.

Objective-to-evidence mapping

#	OBJECTIVE	BLOOM	INTERACTION	ASSESSMENT EVIDENCE	WORKPLACE BEHAVIOUR
LO1	Classify five mixed discoveries as Fix / Flag / Note and state the deciding trigger	**Apply**	Rise L2 Process block, then an 8-item Sorting activity; Rise L4 knowledge check	Sorting pattern; KC4 score; dashboard inspection completeness	Reaches a disposition on every discovery instead of proceeding by default
LO2	Distinguish a degrading vehicle condition from a stable variance	**Analyse**	Rise L3 Labelled graphic and Tabs; Storyline Decision Point 1	DP1 choice and points; KC2; remediation flag	Airs up and logs a cold variance; flags a 17-psi overnight drop as a suspected leak
LO3	Select and justify the safest next action under competing priorities	**Evaluate**	Storyline DP3 with consequence layers; Rise L5 scenario block	DP3 choice; Safety and Service deltas; debrief justification	Protects the medical receiving window under volume pressure
LO4	Compose a four-field escalation message	**Create**	Rise L5 Accordion and scenario block; Storyline DP2	DP2 message variant scored against the four-field rubric	Escalations arrive complete and actionable; no follow-up call needed
LO5	Sequence a route protecting custody and window commitments	**Apply**	Rise L3 Flashcards; Storyline manifest card and DP3	DP3 sequencing choice; Time budget consumed; Service score	Medical tote sequenced first; volume absorbed after commitments
LO6	Appraise a customer access instruction against safety and authorisation policy	**Evaluate**	Rise L4 Accordion; Storyline DP4	DP4 choice and Safety delta; documentation sub-choice	Declines to enter a dwelling; delivers to the safest compliant point and documents

DELIBERATELY NOT OBJECTIVES

- **Operating the handheld** — assumed prerequisite; teaching it would double seat time.
- **Reciting the pre-shift checklist** — the checklist exists and is not the gap.
- **Vehicle maintenance diagnosis** — associates are not mechanics. The objective is to distinguish degrading from stable well enough to route the decision.
- **De-escalating an angry customer** — real, adjacent, and a separate module. Including it would blur the shift-start frame that makes this teachable in forty-five minutes.

06

Course architecture

ITEM	MODULE	CARRIES	TIME	COMPLETION
1	**Rise 360 Prepare for the Shift**	Lessons 1–5. Concepts, criteria, worked examples, four knowledge checks. Mobile-first.	19–24 min	All lessons viewed; knowledge checks attempted
2	**Storyline 360 The First 15**	Branching scenario. Hub dashboard, four scored decision points, remediation, three-dimension debrief.	8–12 min	Debrief reached. Success at score ≥ 80%
3	**Rise 360 Field Guide + Toolkit**	Job aid, escalation card, reflection, manager coaching guide. Built to reopen on a phone.	10–12 min	All sections viewed; reflection submitted

Total seat time 37–48 minutes; a typical learner completes in about 45.

What is built where, and why

This is the section a senior reviewer reads first, because tool choice is where most blended designs are decided by habit.

WHY RISE FOR THE CONCEPT LAYER

Rise reflows. A learner opening Lesson 3 on a phone in a van gets a single readable column; the same content in Storyline would be a scaled-down 1280×720 stage with 14px type. Rise also lets a learner leave and return mid-lesson without suspend-data anxiety, which matters when training happens in fifteen-minute gaps. And its block library fits the actual need — a labelled graphic is the correct form for spatial inspection knowledge, and a sorting activity is the correct form for applying a three-way classification to ambiguous cases.

WHY STORYLINE FOR THE SCENARIO

Branching with persistent state is what Storyline is for. The design needs numeric variables that accumulate across scenes, conditional layers that fire on a threshold, a re-decide loop that preserves a remediation flag, and a debrief that reads the whole path back. Rise's scenario block is genuinely good, but it is a self-contained interaction — it cannot carry a score across four decision points, gate a control on visited states, or drive a conditional debrief.

THE ALTERNATIVE CONSIDERED AND REJECTED

Embedding the Storyline scenario inside Rise as a Storyline block produces a single seamless package with one completion record — genuinely attractive for learner experience. Rejected for the primary configuration because it collapses the assessment's score reporting into the parent Rise course's completion, which costs the decision-level reporting the measurement plan depends on. Where a client's LMS cannot support learning paths, the embedded configuration is the right fallback and the design supports it without change.

LMS tracking

SETTING	VALUE	RATIONALE
Storyline publish profile	SCORM 2004 4th Ed primary; SCORM 1.2 fallback also delivered	2004 separates completion_status from success_status, so “finished the scenario” and “met the standard” are distinct records — exactly the distinction a remediation path creates. It also lifts 1.2's 4096-character suspend-data ceiling.
Passing score	80%	Reachable only by resolving the Safety gate and getting at least three of four decisions at or near optimal.

SETTING	VALUE	RATIONALE
Scoring construction	Graded Pick One per decision, points assigned per choice (Feedback by Choice)	A binary right/wrong model would misrepresent decisions that are defensible but costly.
Completion trigger	Set complete when the learner reaches the action-plan slide	Completion should require the debrief, because the debrief is where the learning consolidates.
xAPI	Optional. Statements per decision carrying choice, dimension deltas, and the cumulative path string	Enables decision-path analytics where an LRS exists. Designed as an enhancement — the SCORM package is fully functional without it.

07

The branching scenario

“The First 15: Choose the Next Best Move” · 8–12 minutes · 38 slides across 7 scenes · 4 scored decision points

NARRATIVE SETUP

It is 7:10 a.m. at Station DC-3. Ninety-six stops, second week solo on Route 214. The pre-shift review surfaces five discoveries: a rear-right tire at 61 psi against a spec of 80 — and 78 yesterday; a signature-required shipment whose suite number is obscured; a refrigerated medical tote for a clinic that receives between 8:00 and 9:30, sequenced by the handheld as stop 14; a dispatch message that volume is 18% over forecast with two associates out; and a fourteen-month-old customer note offering a gate code and asking the associate to step inside an unlocked porch door.

Three rules that hold it together

- **No real-time clock.** Time is consumed by choices, never by the wall clock. A learner can sit on a decision for four minutes and lose nothing. This is an accessibility requirement (WCAG 2.1 SC 2.2.1) and an instructional one — a countdown produces test anxiety, not judgment.
- **Consequences arrive where they would really arrive.** The cost of a bad pre-shift decision shows up at stop 61 or 11:40 a.m., not immediately. Delayed causal feedback is what changes a mental model.
- **Every distractor is defensible.** Each wrong option is one a competent, rushed associate would make and could explain, so a learner cannot pass by elimination.

Branching map

```
START
 ├─ Title  ·  motion + audio preferences  ·  cinematic cold open
 ├─ Mission brief + confidence check (5 sliders)
 │
 ├─ SHIFT-START DASHBOARD  ◄── hub ──────────────────────┐
 │    ├─ Van 4412 ──────────► back to hub ──────────────┤
 │    ├─ Manifest ──────────► back to hub ──────────────┤
 │    ├─ Dispatch message ──► back to hub ──────────────┤
 │    └─ Customer notes ────► back to hub ──────────────┘
 │    [ Roll out disabled until all four are opened ]
 │
 ├─ DP1  VEHICLE SAFETY ALERT
 │    A  air up + log + flag ──► consequence ──► DP2
 │    B  air up and roll ──────► DELAYED coaching, stop 61 ──► DP2
 │    C  drive and monitor ────► CRITICAL ──► micro-lesson ──► re-decide
 │    D  park and wait ────────► over-escalation coaching ──► DP2
 │
 ├─ DP2  PACKAGE-LABEL EXCEPTION   A / B / C / D ──► DP3
 ├─ DP3  PRIORITY DELIVERY CONFLICT   A / B / C / D ──► DP4
 │
 ├─ DP4  CUSTOMER ACCESS / SAFETY NOTE
 │    A  plan non-entry + flag the note ──► DEBRIEF
 │    B  follow the note ──► CRITICAL ──► micro-lesson ──► re-decide
 │    C  plan undeliverable ──► DEBRIEF
 │    D  decide at the gate ──► coaching ──► DEBRIEF
 │
 ├─ PERFORMANCE DEBRIEF   three dimensions, threshold markers
 ├─ PATH REVIEW           your four choices and what each cost
 ├─ PATTERN FEEDBACK      1 of 5 conditional layers
 ├─ TRANSFER ACTION PLAN  text entry + post-confidence   [COMPLETE]
 └─ RESULTS               score, pass, retry             [SUCCESS ≥ 80%]
```

The decisions

DP	SITUATION	PREFERRED ACTION	WHAT THE DISTRACTORS TEACH
1	Rear right at 61 psi. Spec 80. Yesterday 78. Station compressor working.	Air it up, log it, and flag the overnight drop to fleet — Fix and Flag.	Air-and-roll misses that a moving condition is a Flag. Drive-and-monitor is a category violation. Park-and-wait costs 51 minutes for something resolvable in two.
2	Signature-required shipment; barcode scans, suite number unreadable. Still loading.	Scan as an exception, photograph, and message dispatch with the tracking number while loading.	“Fairly sure” works nine times in ten, which is what makes it a habit. A voice call without the tracking number costs four minutes instead of forty seconds.
3	Clinic window 08:00–09:30, sequenced as stop 14 (arrives 10:20). Volume +18%.	Resequence the clinic first, confirm the window, and flag a number: 85 of 96, need 11 covered.	Absorbing volume treats a planning problem as an effort problem. “Volume's too high” hands over a feeling instead of a constraint.
4	Fourteen-month-old note: gate code, friendly dog, step inside the unlocked porch door.	Plan a non-entry delivery now, and flag the note so the account gets corrected.	Following the note crosses two lines at once. Deferring to the gate chooses to decide later with less time and more pressure — and if it works, that is worse.

08

Scoring model

Three tracked dimensions. Safety and Service are 0–100 scores starting at 70; Time is a fifteen-minute budget that choices consume and that may go negative. Modelling time as a budget rather than a score teaches the precedence order structurally, without a line of narration.

Full matrix

DP	OPTION	PTS	SAFETY	SERVICE	TIME	OUTCOME
DP1	**A Air up + log + flag**	25	+15	0	−3	Preferred
	B Air up and roll	15	−10	−5	−2	Delayed coaching, stop 61
	C Drive and monitor	0	−20	0	0	CRITICAL → remediation
	D Park and wait for fleet	10	+10	−15	−9	Over-escalation coaching
DP2	**A Exception + photo + flag**	25	0	+15	−1	Preferred
	B Load it, sort it later	5	0	−15	0 (−4 at 9:47)	Delayed coaching
	C Set aside undeliverable	10	0	−10	−2	Coaching
	D Call, no tracking number	15	0	+5	−4	Four-field coaching
DP3	**A Resequence + flag a number**	25	0	+15	−3	Preferred
	B Absorb and push	0	0	−20	−2	Delayed coaching, 9:41
	C “Volume too high”, wait	10	0	−5	−6	Coaching
	D Resequence, absorb silently	15	0	+5	−2	Half-right coaching
DP4	**A Plan non-entry + flag note**	25	+5	+10	−2	Preferred
	B Follow the note	0	−25	0	0	CRITICAL → remediation
	C Plan undeliverable	10	+5	−10	−1	Coaching
	D Decide at the gate	5	−10	−5	0	Coaching

VERIFIED, NOT ASSERTED

The matrix was checked by enumerating all 256 paths. Exactly 13 reach the 80% pass mark without triggering remediation — 5.1%, the correct by-chance figure for a four-item, four-option assessment at an 80% cut score. The enumeration also found two real bugs: a learner with a high Safety score could choose to enter a dwelling and never trip the safety interrupt, and the score floor could re-trap a learner in an infinite remediation loop. Both are fixed below.

Thresholds and remediation

GATE	CONDITION	EFFECT
Categorical critical choice	DP1-C or DP4-B	Critical layer fires unconditionally. These are category violations — driving on a 19-psi-low rear tire, entering a dwelling — not score events.
Critical safety floor	Safety ≤ 55 and no prior remediation	Catches accumulation (DP1-B then DP4-D lands at 50). Fires at most once per attempt, which is what prevents the re-trap.

GATE	CONDITION	EFFECT
Remediation	Either trigger	Micro-lesson of two screens, ~40s. Safety restored to 60 — above the floor, well below the 80 threshold. The attempted option is disabled with a rewritten label. Re-decide caps at 15 points.
In-scenario success	Safety ≥ 80, Service ≥ 70, Time ≥ 0	Decisive-and-safe debrief
LMS success	Score ≥ 80	success_status = passed
LMS completion	Action-plan slide reached	completion_status = completed, independent of pass

Pattern feedback

Feedback is selected by pattern, not by total, because two learners on the same score can need opposite advice. Giving them the same paragraph is the most common failure in scenario debriefs.

LAYER	SIGNATURE	WHAT THIS LEARNER NEEDS TO HEAR
Recovered well	Remediated, then scored ≥ 80	Name the recovery and do not congratulate past it. The second decision counts; the first is on the record.
Decisive and safe	Safety ≥ 80, Service ≥ 70, Time ≥ 0	Name the mechanism that worked — proportionate escalation, and naming a number — so it is repeatable.
Careful but costly	Safety ≥ 80, Service < 70 or Time < 0	Permission to resolve things themselves. The problem is authority, not judgment.
Fast but exposed	Safety < 80, Time ≥ 3	The precedence rule, stated flatly. Time saved cannot buy safety back. The only pattern where the tone is direct rather than warm.
Needs another run	Safety < 80 and Service < 70	Two takeaways maximum, no inventory of errors, and an explicit invitation to re-run.

09

Visual system

Deep charcoal foundation, warm ivory content surfaces, one burnished gold accent used structurally, and one restrained indigo reserved exclusively for things the learner can act on. In this system indigo means “you can touch this” and gold means “this is structural or evaluative” — which is why a learner never has to hunt for what is clickable and a keyboard user never has to guess where focus is.

Palette, with measured ratios

Ink	Slate	Graphite	Ivory	Paper	Bone
#14161A	#1F232A	#2A2F37	#F4F1EA	#FBFAF7	#E7E2D7

Gold	Gold light	Gold deep	Indigo	Teal	Crimson
#B8873B	#D9A84E	#775418	#2F4B8F	#2F6F5B	#9E2B2B

PAIR	RATIO	VERDICT	PERMITTED USE
Ivory on Ink	16.06:1	AAA	Body copy on dark
Ink on Ivory	16.06:1	AAA	Body copy on light
Gold on Ink	5.66:1	AA	Rules, meters, markers, large type on dark
Gold light on Ink	8.33:1	AAA	Small labels on dark
Gold on Ivory	2.83:1	**FAILS**	**Never. The most attractive pairing in the system and unusable for text at any size.**
Gold deep on Ivory	6.07:1	AA	The only gold permitted as text on a light surface
Indigo on Ivory	7.38:1	AAA	Interactive text and slider tracks
Ivory on Indigo	7.38:1	AAA	Button labels
Indigo light on Ink	7.75:1	AAA	Links and interactive text on dark
Teal on Ivory	5.25:1	AA	Confirm / safe, with an icon and a word
Amber deep on Ivory	5.45:1	AA	Caution, with an icon and a word
Crimson on Ivory	6.57:1	AA	Risk, with an icon and a word
Secondary on Ivory	5.29:1	AA	Meta text (#5F646A, replacing #6B7076 at 4.43:1)
Secondary on Ink	8.10:1	AAA	Meta text on dark

> **THREE TOKENS FAILED THE FIRST PASS AND WERE REPLACED**
>
> Gold #B8873B as text on ivory measured 2.83:1. Caution #A8641C on ivory measured 4.13:1 — fine for large text, failing for body. Secondary text #6B7076 measured 4.43:1, a near miss that would have failed an audit. A designer working by eye would have shipped all three, because they look right.

Typography and treatment

ROLE	FACE	FALLBACK	USAGE
Display	Bodoni Moda	Georgia	Scene titles, cold open, debrief headers. Never below 28px — high-contrast didone serifs lose their thin strokes at small sizes, which is legibility, not taste.
Interface	Inter	Arial / Helvetica	Everything a learner reads to make a decision. 18px minimum.
Data	IBM Plex Mono	Consolas	Pressure readings, times, tracking numbers, scores. Tabular figures matter — 61 and 80 must occupy the same width or the gauge readout jitters.

IMAGERY, EXCLUDED ON PURPOSE

- Cartoon office illustration — wrong register, and reads as a template.
- Stock “diverse team smiling at a laptop” — costs the credibility the documentary grade buys.
- Gradient meshes and glassmorphism — dated, and they reduce text contrast.
- Drop shadows for depth — the system uses surface value steps and hairlines, which survive scaling.

10

Accessibility

WCAG 2.1 Level AA for both Rise and Storyline output, decided at the storyboard rather than audited at the end. Where that produced a design consequence rather than a checkbox, the consequence is named.

REQUIREMENT	HOW IT IS MET
Contrast	Every pair calculated, not judged. Ratios published. Gold is never text on a light surface; Gold-deep exists for that case.
Keyboard	Every interaction operable. Sliders on arrow keys with a numeric readout. Modal layers set focus on show, with Continue last in focus order.
Focus states	3px ivory outer ring with a 1px ink inner stroke, built once on the slide master. 7.38:1 on indigo, 16.06:1 on charcoal, and visible on ivory — a single-colour ring fails on at least one of the three surfaces.
Focus order	Set explicitly on all 38 slides. The rule throughout: data before choices, so a screen-reader user never backtracks to find out what 61 psi means.
Button labels	Full meaning, never position. Disabled controls state their reason: “Roll out — unavailable. Open all four review items first.”
Alt text	Three categories: decorative for atmosphere; descriptive for decision-relevant imagery; and never an image for content in disguise — the manifest, message thread, note cards, and gauge readings are all live text.
No colour alone	Every state carries an icon or a word: and “LOW”, a dot and “BELOW”, a lock glyph and a rewritten label.
Captions	Written in Storyline's caption editor, never burned in. Transcript downloadable. Audio off by default and never the sole carrier — the module is complete muted.
Plain language	Grade 8 target, Grade 7 for Lesson 1. Terminology defined at first use. One exception: character dialogue is written as people speak.
Reduced motion	Learner-set before any content. Implemented as duplicate static slides rather than conditional animation. Exceeds AA (SC 2.3.3 is AAA).
No timed interaction	There is no timer anywhere. Time is consumed by choices, never by the clock (SC 2.2.1).
Representation	Cast spans age, gender, ethnicity, body type, with visible assistive equipment among background associates. Competence never correlates with identity.

THE CONSTRAINT THAT IMPROVED THE DESIGN

The first version had a visible countdown, because pressure felt authentic. Removing it for SC 2.2.1 forced the design to create pressure through consequence instead — and consequence is what teaches. Accessibility is usually framed as a cost; here it made the instruction better.

KNOWN RISK, NAMED RATHER THAN HIDDEN

Rise's sorting activity is drag-based. The keyboard alternative is verified in the published output at QA, and if it proves keyboard-hostile the activity is replaced with a matching-style knowledge check carrying the same eight items. Naming the risk and its fallback is more useful than asserting the interaction is fine.

11

Measurement plan

STATED ONCE, PLAINLY

Everything in this section is a proposed framework. This is a portfolio prototype with no live learner population — there are no results here, no completion rates, no performance data, and no operational outcomes. Only the measures I would instrument and the decision each one informs. In a live implementation I would validate the baseline across at least one full hiring cohort before attributing any change to the training.

METRIC	WHY IT MATTERS	SOURCE	TIMING	DECISION IT INFORMS
Path completion	Whether the design fits the pre-route window at all	LMS path	Weekly	Whether to split modules or shorten seat time
Module A → B drop-off	Isolates a failing transition from failing content	LMS item completion	Weekly	Whether the closing button and expectation copy need rewriting
Score distribution	Tests whether the cut score discriminates	Storyline Results	Per cohort	Whether to adjust 80% or revise a distractor
Per-decision choice distribution	Locates the exact misconception	xAPI or SCORM interaction data	Per cohort	Which single Rise screen to rebuild
Decision-path strings	Surfaces error combinations a per-question view hides	Path string variable	Per cohort	Whether two lessons are interacting badly
Remediation rate	How often a categorical safety violation is chosen first	Remediation flag	Per cohort	Whether the Safety gate needs teaching earlier
Time in module	Seat-time honesty; station scheduling depends on it	LMS	Monthly	Whether the published 37–48 minutes is accurate
Confidence shift (paired)	Catches the confident-but-incapable case	Pre/post sliders	Per cohort	Whether a criterion is taught as a rule without application
Field-guide re-access	The cheapest transfer signal available	LMS access log	Monthly	Whether the job aid is the right artefact
Manager observation	Level 3. Whether the behaviour appears on a real route	Ride-along checklist	Days 10–30	Whether coaching or content is the gap
Escalation completeness	The most direct behavioural proxy available	Sampled dispatch messages	Monthly	Whether Lesson 5 needs a rebuild or managers need the intake standard
Pre-shift window duration	Tests the claim that good judgment is faster	Operational timestamps	Baseline + 30/60/90	Whether the design's core promise holds
Avoidable first-attempt failures	The operational number a director cares about	Exception codes	Baseline + 30/60/90	Whether to extend the pattern to other shift-start contexts

Kirkpatrick mapping

LEVEL	INSTRUMENTED BY	MEASURES	STATUS
1 Reaction	The prototype	Completion, drop-off, time in module, two-item relevance rating	**BUILT**
2 Learning	The prototype	Score, per-decision distribution, path strings, remediation rate, paired confidence, knowledge checks	**BUILT**

LEVEL	INSTRUMENTED BY	MEASURES	STATUS
3 Behaviour	A client	Manager observation, escalation completeness sampling, field-guide re-access	**PROPOSED**
4 Results	A client	Pre-shift window duration, avoidable first-attempt failures, preventable roadside events	**PROPOSED**

HOW THE CHEAPEST INSTRUMENT IMPROVES THE COURSE

The scenario records every learner's choice at every decision point, so the modal wrong answer is a diagnosis of a specific screen rather than a general observation about learners. If the cohort's modal error is “air it up and roll”, Lesson 3 is not landing the degrading-versus-stable distinction, and the fix is one Rise screen — half a day. A course that can be improved in half-day increments off real data is a different asset from one that gets rewritten every two years.

12

Constraints, assumptions, and risks

Constraints

CONSTRAINT	DESIGN CONSEQUENCE
Seat time ≤ 50 minutes	Three modules, none longer than 24 minutes, each independently completable in a sitting
Mobile access is unavoidable	All reading and reference content in Rise, which reflows. Storyline targeted at desktop and tablet with a mobile-tolerant layout
No proprietary content or branding	Fictional carrier, invented characters, illustrative figures, labelled as such on every deliverable
Single designer-developer, no voice or video budget	Text-first narration with optional synthesis; documentary photography over filmed video; audio never the carrier of essential information
Must publish to a standard LMS	SCORM 1.2 and 2004 4th Ed profiles specified; xAPI treated as an enhancement, not a dependency
Literacy and language range	Plain-language target of Grade 8 or lower; no idiom in instructions; terminology defined at first use

Assumptions

- Associates have completed basic equipment and policy onboarding; the design does not teach the handheld.
- A pre-shift window of roughly fifteen minutes is the operational norm, with associate discretion inside it.
- Escalation to a coordinator is culturally permitted and technically available at shift start.
- Station leads or coordinators can run a short ride-along conversation in the first two weeks, which the manager guide depends on.
- The LMS reports SCORM score and completion, and learners have individual accounts.
- **Content accuracy on vehicle specifics, cold-chain handling, and access policy would be validated by a subject-matter expert before release. In a live implementation this is a gating step, not a formality — the pressure thresholds and custody rules in this prototype are plausible, not authoritative.**

Content risks

RISK	MITIGATION
Operational inaccuracy — a threshold or policy that does not match a real operation	Every operational specific is labelled illustrative; a live build gates release on SME validation; the model is designed to survive a change in the specifics
The scenario reads as a trick	Two of four decisions have a defensible non-optimal choice with a survivable cost; feedback names the reasoning, not the rule; scoring rewards pattern, not single answers
“Escalate everything” as the learned lesson	One decision point explicitly penalises over-escalation on Time and Service; the Note disposition exists to absorb it
Policy drift	Criteria centralised in Lesson 2 and the job aid; the Storyline content layer is separated from the interaction shell so text updates do not require rebuilding logic
Tone misread as blame	Feedback addresses the decision, never the learner; no character is humiliated; remediation is framed as a second pass, not a failure
Fictional client mistaken for real experience	Disclosure on the case study page, this document's cover, the job aid footer, and the Storyline title slide

13

Production plan

PHASE	WORK	EFFORT
1 Discovery and research	Operational context, failure-mode framing, business problem, constraints	1.5 d
2 Content outline	Decision model, criteria, the five discoveries, worked examples and non-examples	1.5 d
3 Learning design	Personas, objectives and mapping, modality strategy, architecture, branching, scoring with verification	3 d
4 Visual design	Palette with measured contrast, type pairing, icon set, masters, key screen designs	3 d
5 Rise development	Five lessons plus Module C, all blocks, knowledge checks, attachments, mobile pass	4 d
6 Storyline development	38 slides, 25 variables, trigger logic, layers, states, remediation, debrief conditionals	7 d
7 Audio and media	Photography or licensing and grading, UI renders, captions, optional narration	3 d
8 Accessibility review	Focus order, alt text, keyboard, screen reader, zoom	2 d
9 QA	Full checklist, path spot-checks, both critical routes, both remediation loops	2 d
10 LMS testing	Both SCORM packages: completion, score, pass, fail, resume, retry	1 d
11 Portfolio documentation	Case study copy, this document, storyboard, visual assets, silent walkthrough	3 d
12 Publication	Build the page, SEO, accessibility statement, prototype hosting, cross-browser check	1 d
	Total	**32 d**

Prioritised MVP

The smallest thing that makes a hiring manager want to talk to you. Everything else is polish on an asset that already worked.

- **Playable Storyline prototype — two decision points, dashboard, debrief.** The only irreplaceable asset. A case study without a playable artefact is a claim about competence.
- **Published case study page with the full narrative.** The prototype proves you can build; the page proves you can think.
- **The one-page job aid.** Cheapest credibility in the set. It signals transfer thinking, which most portfolios omit entirely.
- **One fully built Rise lesson.** Proves purposeful block selection. One lesson does this as well as five.
- **This design document.** The artefact that survives being forwarded to someone who will not click a link.

If only one addition to version two

A think-aloud test with three delivery drivers, written up honestly. A portfolio piece that includes real user findings — especially findings that contradicted the designer — is in a different category from one that does not.

First 15

A scenario-based onboarding experience for last-mile delivery professionals. Designed, art-directed, and specified end to end by Kwame Yeboah.

CASE STUDY	nimfah.com/work/first-15-last-mile-onboarding
PROTOTYPE	Playable branching scenario · 8–12 minutes
STORYBOARD	Ten annotated production screens
CONTACT	kwame.nimfah@gmail.com

MileOne Logistics is a realistic fictional organisation created for this portfolio project. All scenarios, characters, and figures are illustrative. No proprietary or client information is represented, and no metric in this document is a reported outcome.