Every Morning, 120 stores wait on one spreadsheet

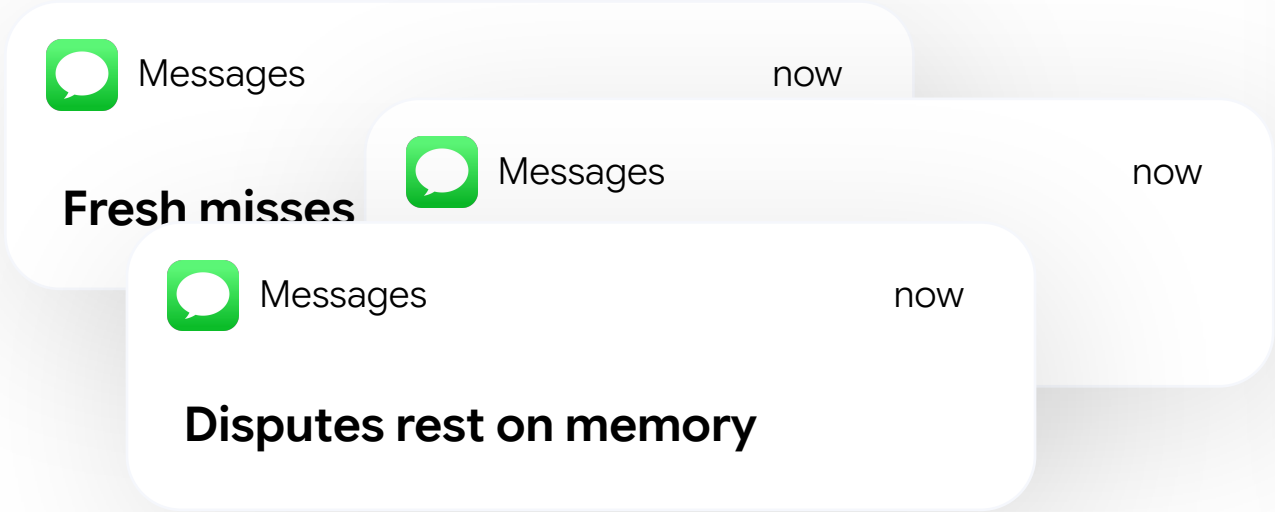
Fresh · 80 **Style · 25** **Tech · 15**

2 Depots · 60 vehicles, Only 16 can carry chilled

Every Morning, 120 stores wait on one spreadsheet

Fresh · 80 **Style · 25** **Tech · 15**

2 Depots · 60 vehicles, Only 16 can carry chilled



How a delivery runs today

Store Manager
Phones the order in

Dispatcher
Types it into a spreadsheet

Loader
Works from a printed list

Driver
Paper sheet · Invisible once on the road



Three problems we found on the **Ground Operations**

Vehicle Shortage

On peak days, there are more orders than vehicles

01

02

Dispatcher Stress

One dispatcher, up to 120 trips a day

Poor Coordination

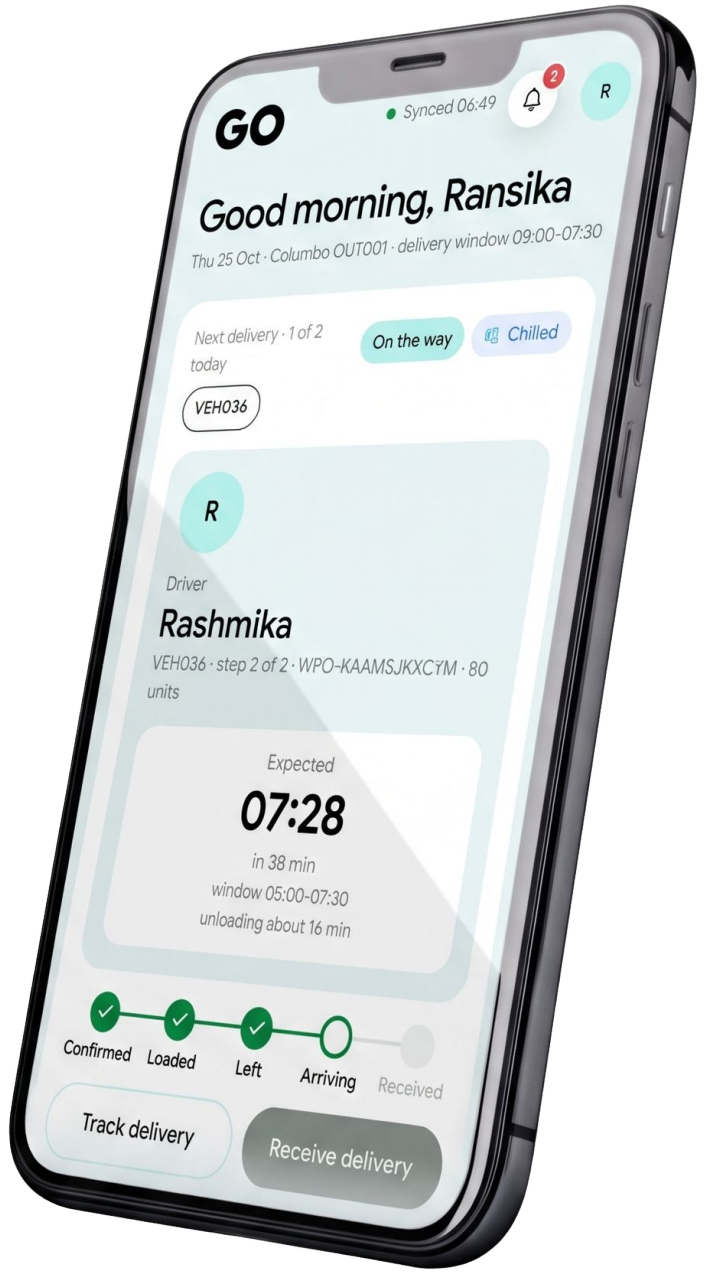
One late van starts a chain of phone calls

03

Ground Operations

4 People One app,
Every Handoff on record



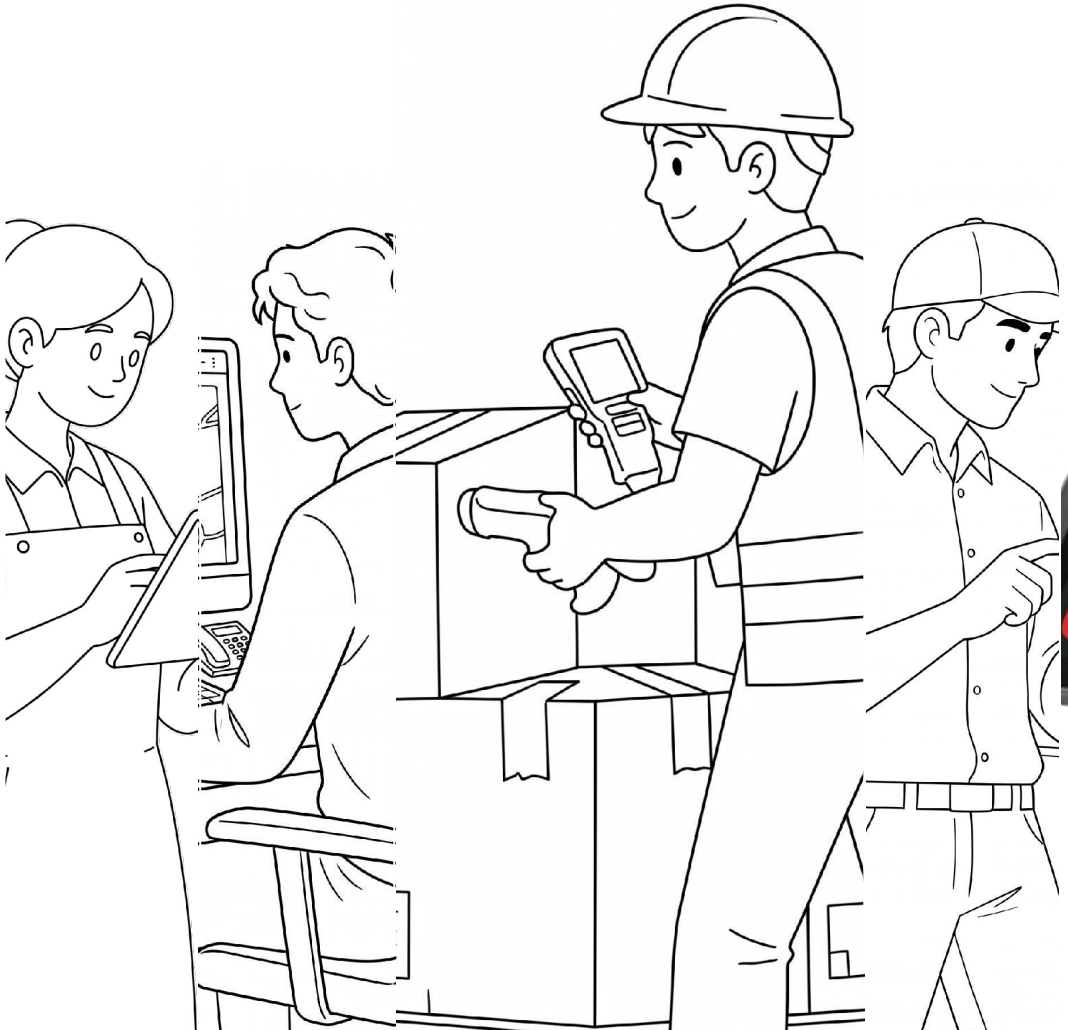


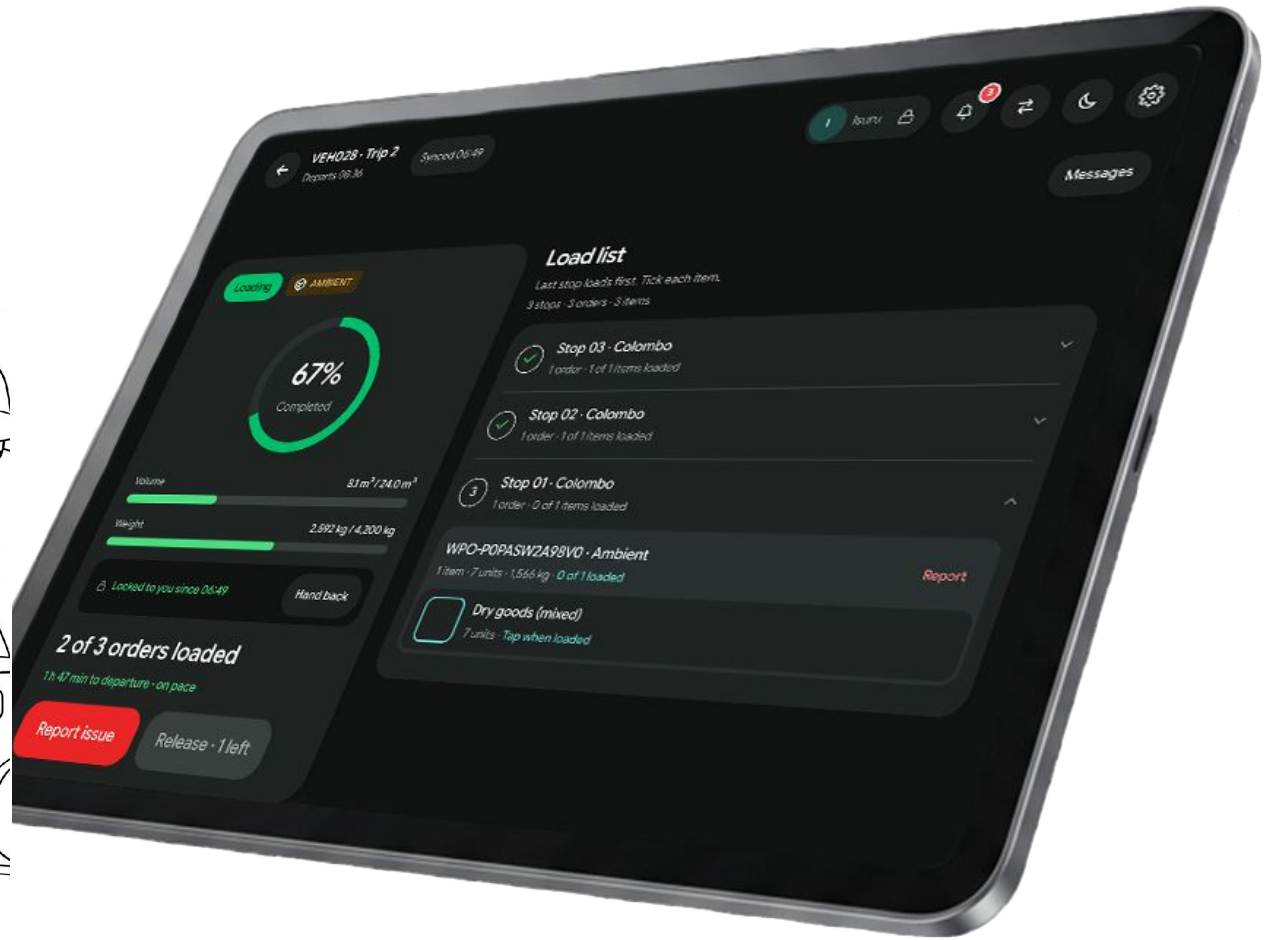
4 People One app,
Every Handoff on record





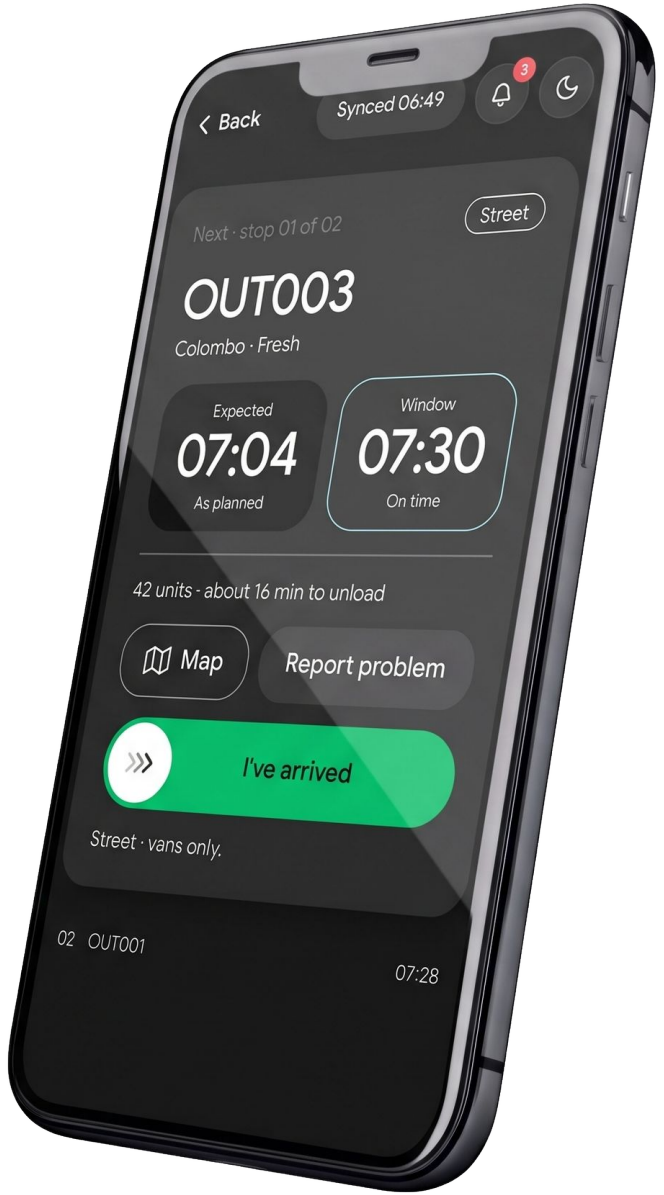
4 People One app,
Every Handoff on record





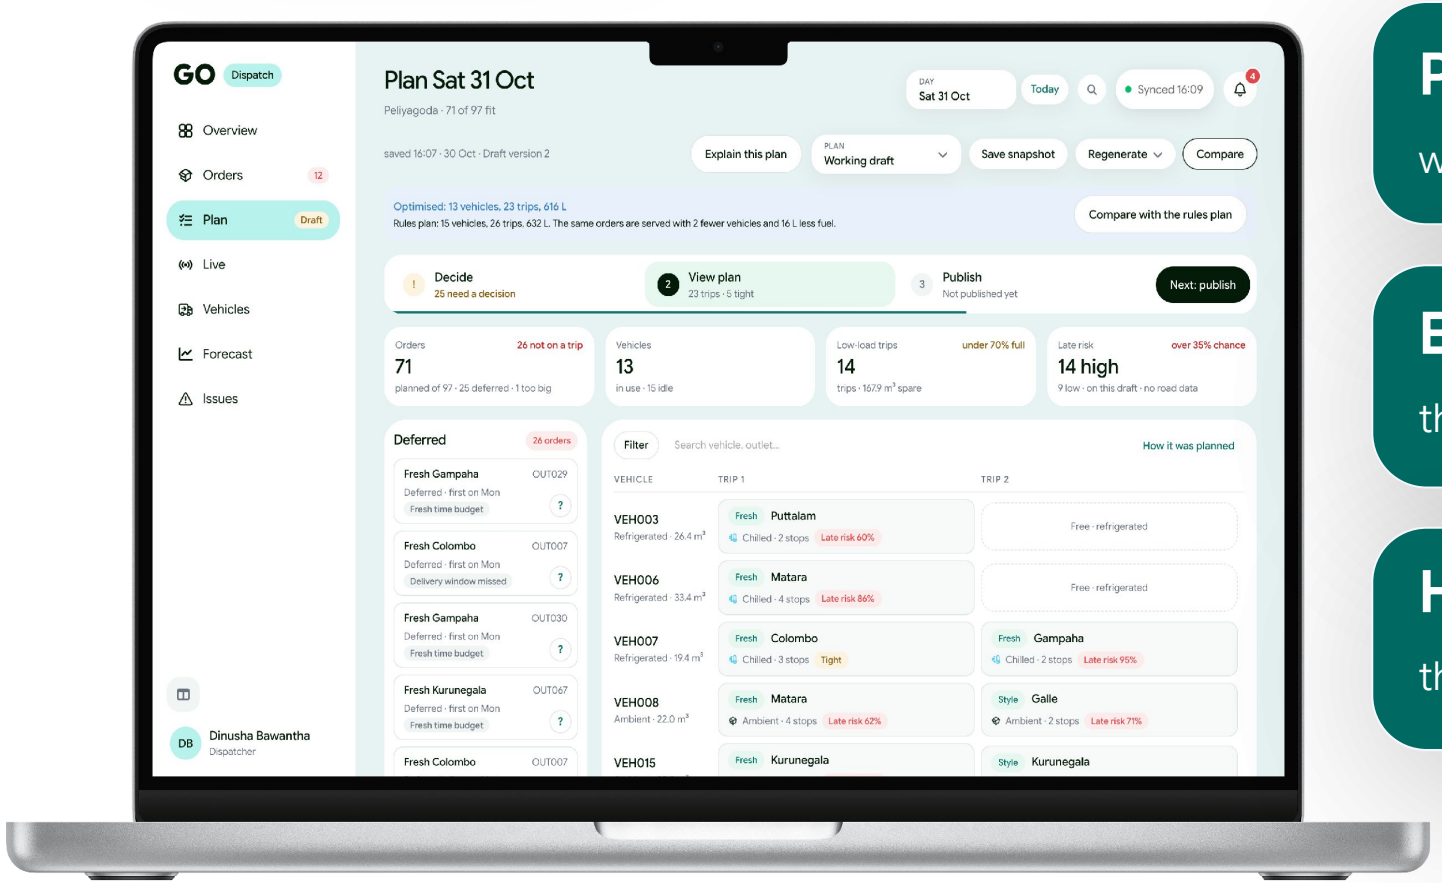
4 People One app,
Every Handoff on record





Problem 1 · Vehicle Shortage

When vehicles run short



Plan in seconds

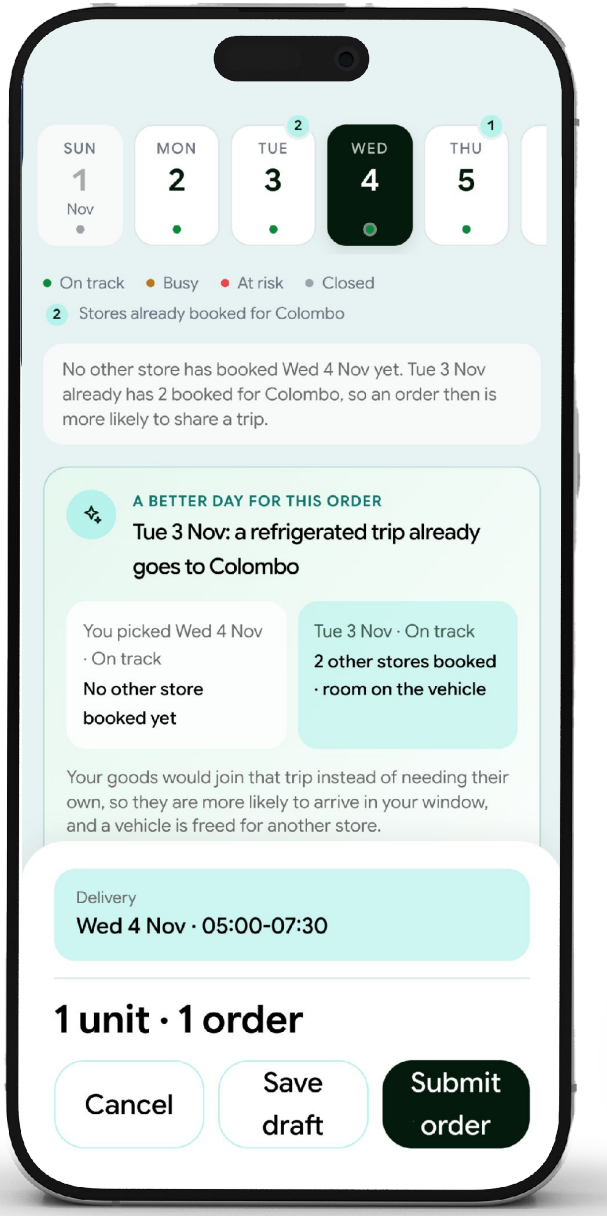
within every vehicle's limits

Every waiting order explained

the rule, and the smallest change that makes room

Hire ahead

the forecast shows the gap; the admin hires



Problem 1 · Vehicle Shortage

GO suggests a day when a vehicle already goes there

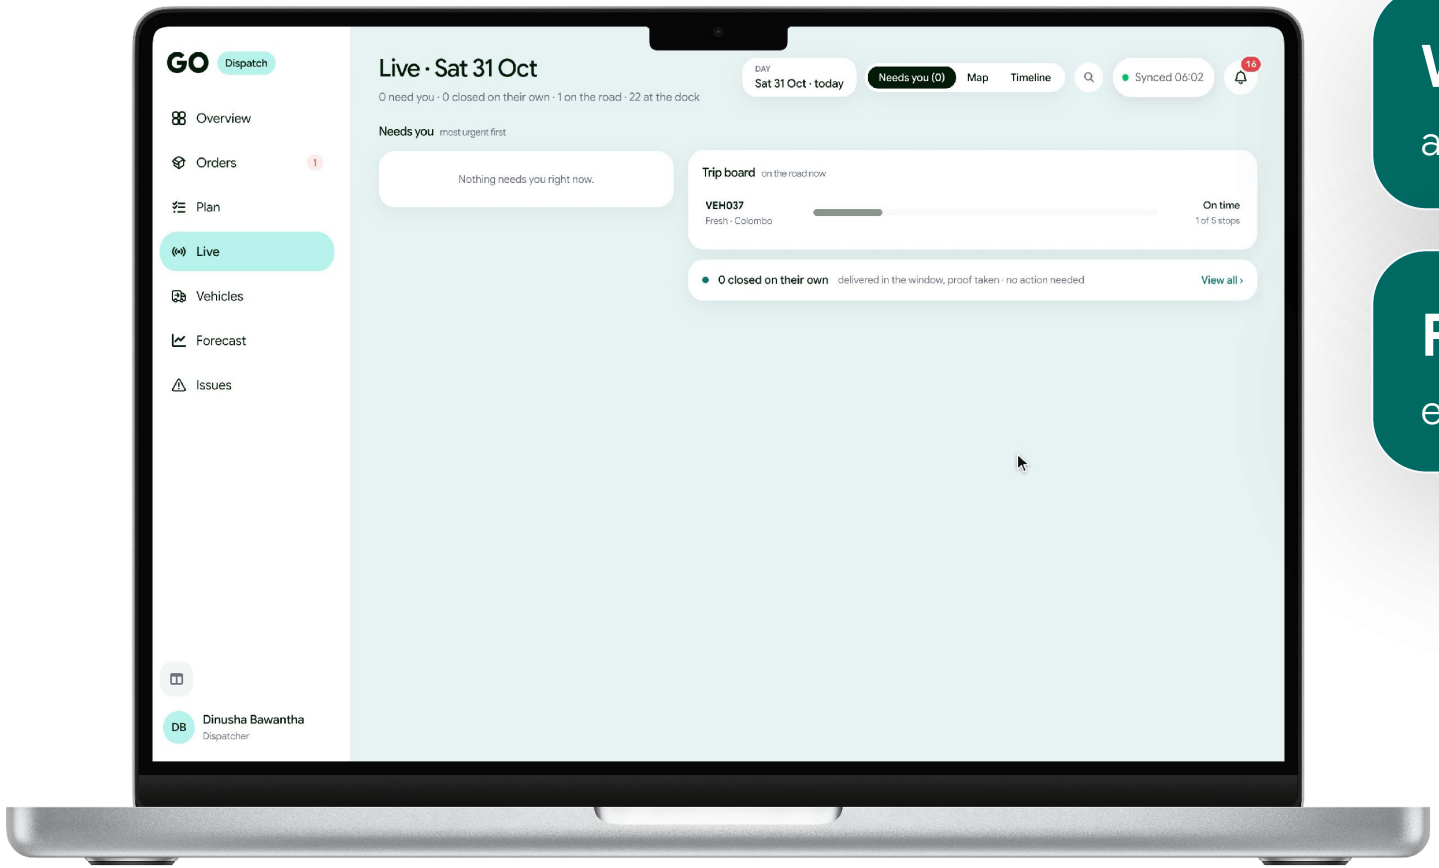
Mon	Tue	Wed	Thu	Fri	Sat
2 Nov	3 Nov	4 Nov	5 Nov	6 Nov	7 Nov
		your pick			

3 and 5 Nov: a vehicle already goes to your district

Only a suggestion. The store manager decides.

Problem 2 · Dispatcher Stress

GO watches with them



Watcher

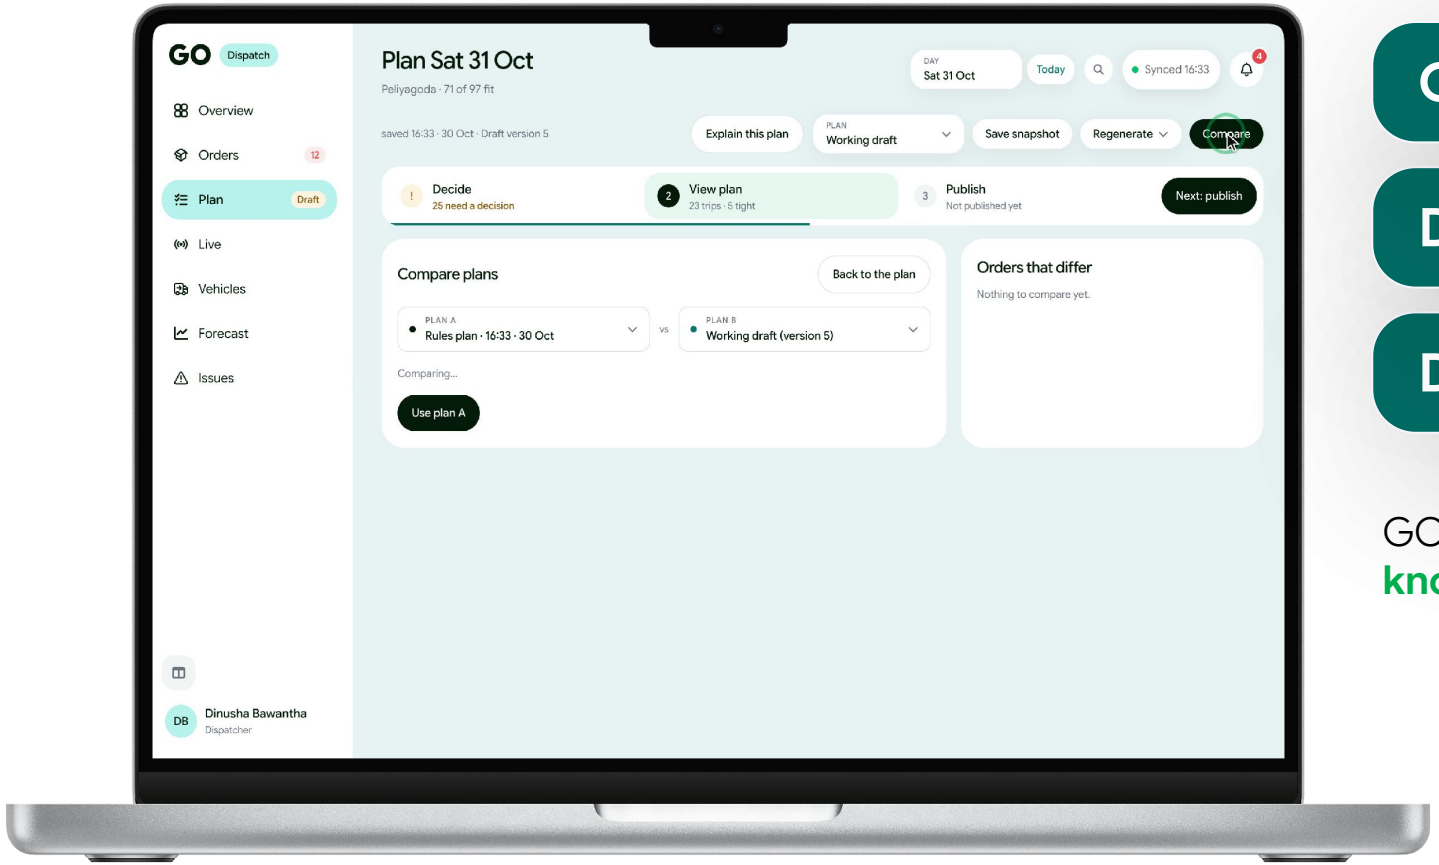
alerts only when a decision is needed

Focus page

everything that needs you, in one place

Problem 2 · Dispatcher Stress

More flexibility across tools



Compare plans

Drag & Drop

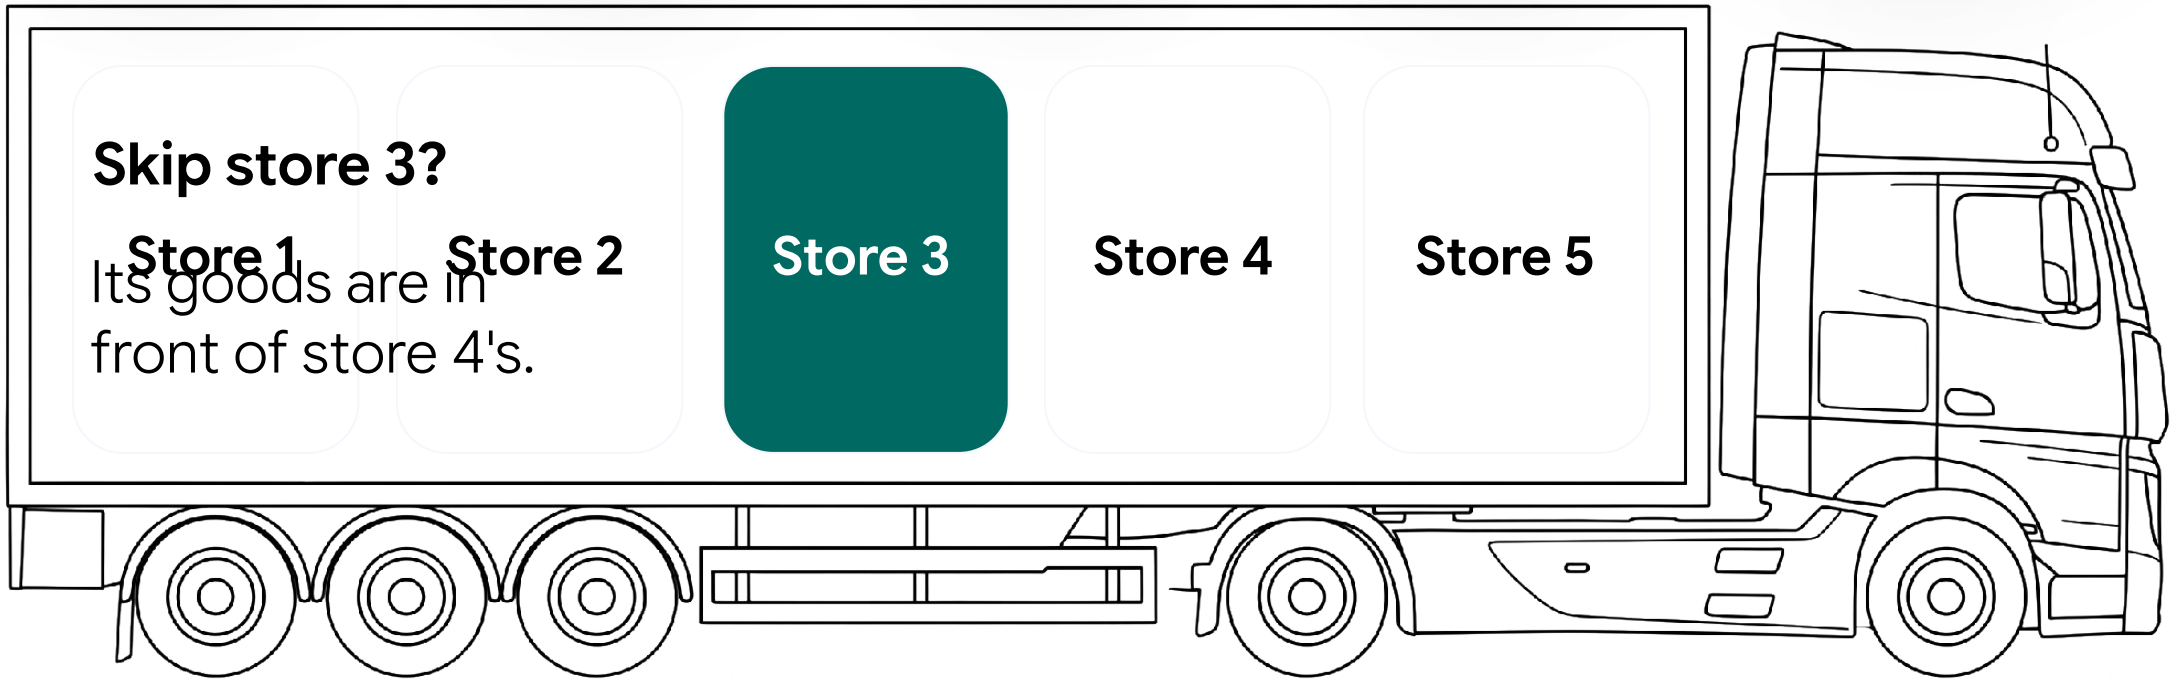
Defer

GO brings the tools. **The dispatcher brings the know how.**

Problem 3 · Poor Coordination

Now, a trip on the road

Store 1 — Store 2 — Store 3 — Store 4 — Store 5



Problem 3 · Poor Coordination

Ground operations break in ways no plan expects

On time, but the store closed early today

The store is open, the mall's loading bay is not

A van breaks down with chilled goods inside

Nobody at the store to receive

Loader: “all loaded.” Store: “three cartons missing.”

Delivered with no signal, while the dispatcher skipped that store

Problem 3 · Poor Coordination

Today, a chain of phone calls

07:12 → 07:31, one call at a time



One call at a time.
Nobody sees the whole picture.
Nothing on record.

Problem 3 · Poor Coordination

Three obvious answers, and why each breaks

1

Let AI decide everything

Who answers when it gets it wrong?

2

Buy a routing or tracking tool

It moves the van. It doesn't get store 4 to agree.

3

A WhatsApp group

Fast, but no structure. No record.

Problem 3 · Poor Coordination

Two kinds of problems

A Can be settled later

Missing goods · Damage · Shortages

Report with photos → an issue with an owner, tracked to the end

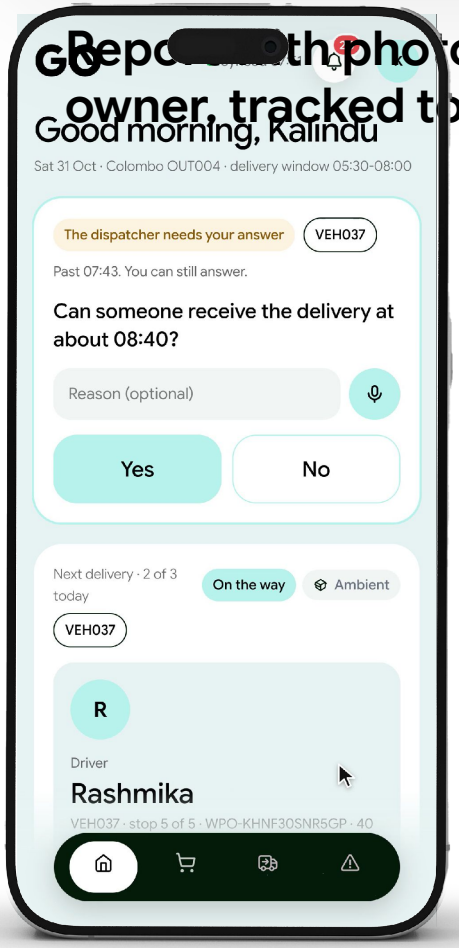
B Can't wait

Window closing · Van stuck · Store closed

A live decision, in minutes

A Can be settled later

Missing goods · Damage · Shortages



Report [redacted] photos → an issue with an owner, tracked to the end

Issue Reported, with photos

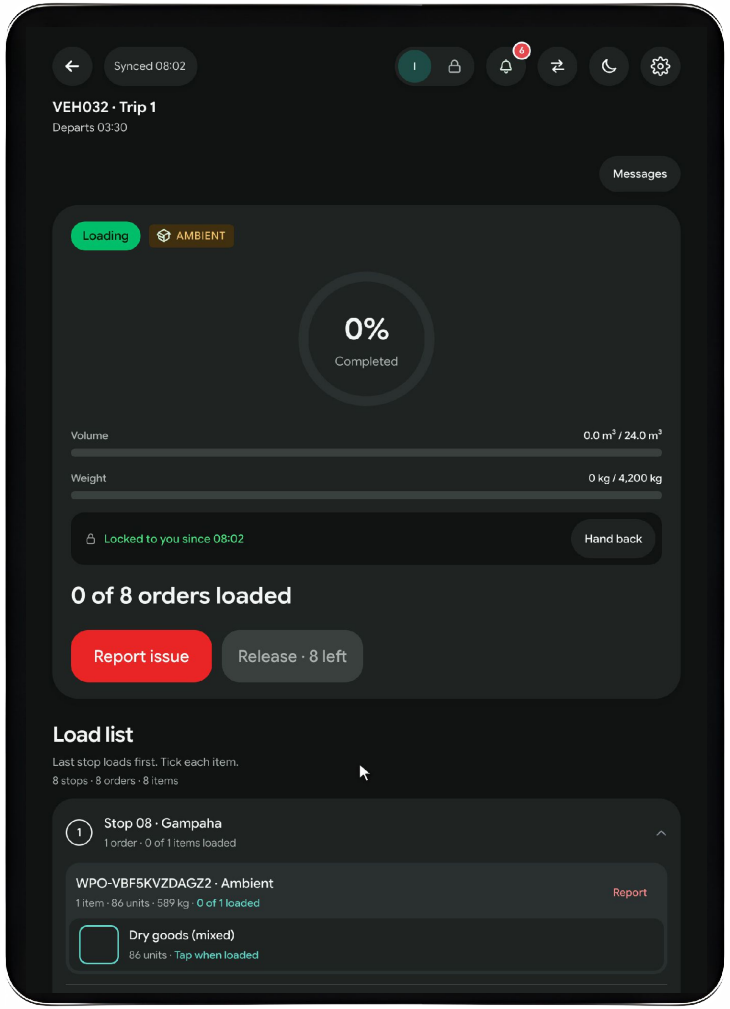
from the loader's tablet or the store's phone

One issue, one owner

tracked like a ticket; a loader's flag is linked, not repeated

Closed on record

who handled it, when and why



Problem 3 · One late van, a chain of calls

Every case, the same tools

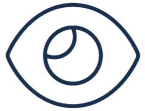
1

Report



one tap,
with photos

Watcher



spots it first

2

3

Ask



GO asks the
stores

Trip room



everyone, one
conversation

4

5

Needs you



a person
decides

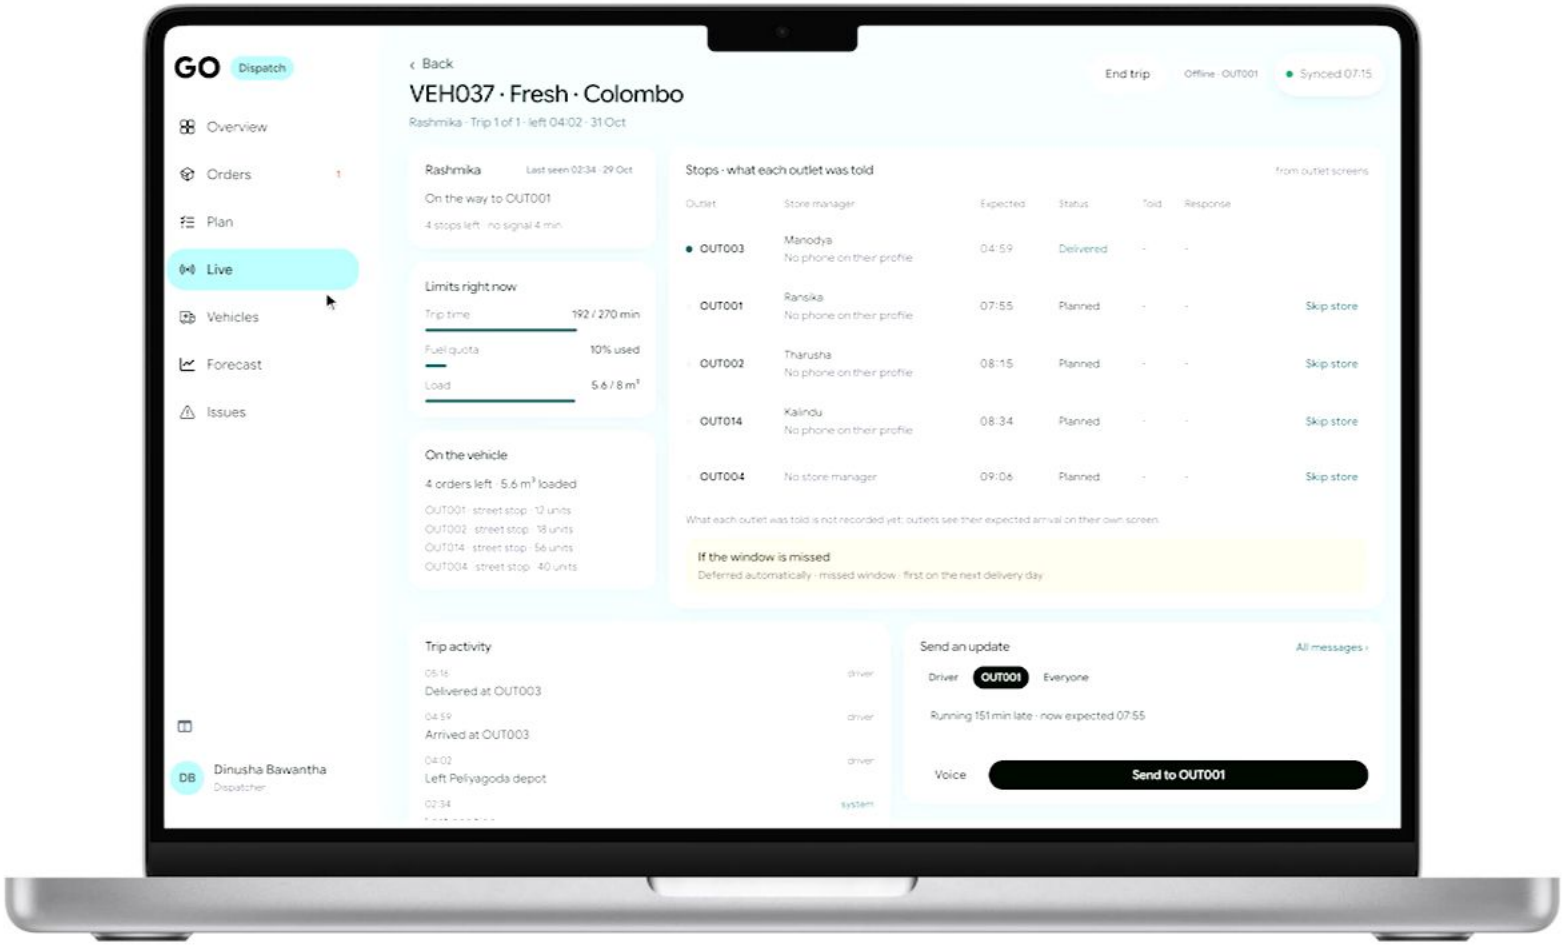
B One room for the whole trip

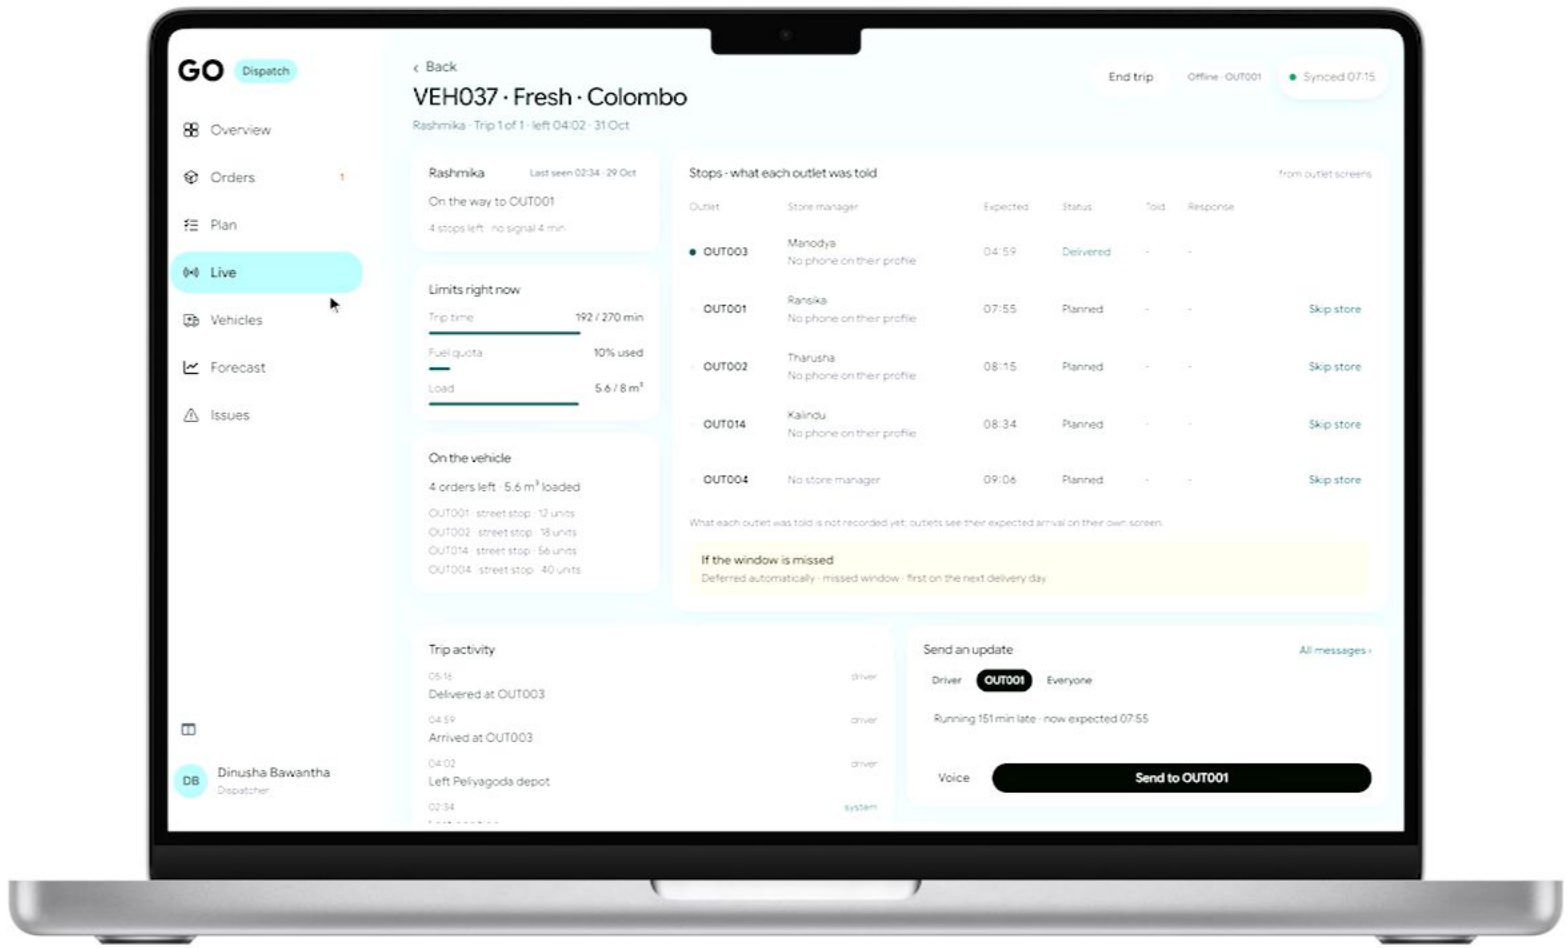
The dispatcher, the driver, every store on it

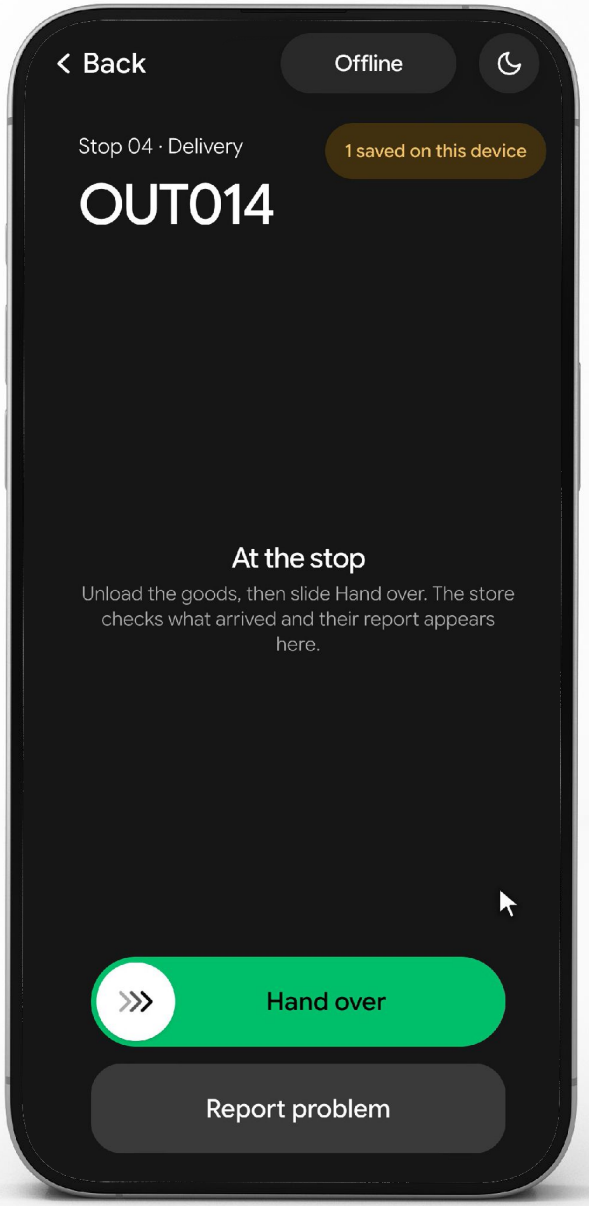
One picture, live, for everyone

GO asks, and suggests the next step

A person decides




GO
Dispatch
Overview
Orders
1
Plan
Live
Vehicles
Forecast
Issues
DB
Dinusha Bawantha
Dispatcher
‹ Back
VEH037 · Fresh · Colombo
Rashmika · Trip 1 of 1 · left 04:02 · 31 Oct
End trip
Offline · OUT001
Synced 07:15
Rashmika
Last seen 02:34 · 29 Oct
On the way to OUT001
4 stops left · no signal 4 min
Limits right now
Trip time
192 / 270 min
Fuel quota
10% used
Load
5.6 / 8 m³
On the vehicle
4 orders left · 5.6 m³ loaded
OUT001 · street stop · 12 units
OUT002 · street stop · 18 units
OUT014 · street stop · 56 units
OUT004 · street stop · 40 units
Stops · what each outlet was told
from outlet screens
Outlet
Store manager
Expected
Status
Told
Response
OUT003
Manodya
No phone on their profile
04:59
Delivered
OUT001
Ransika
No phone on their profile
07:55
Planned
Skip store
OUT002
Tharusha
No phone on their profile
08:15
Planned
Skip store
OUT014
Kalindu
No phone on their profile
08:34
Planned
Skip store
OUT004
No store manager
09:06
Planned
Skip store
What each outlet was told is not recorded yet; outlets see their expected arrival on their own screen.
If the window is missed
Deferred automatically · missed window · first on the next delivery day
Trip activity
05:16
driver
Delivered at OUT003
04:59
driver
Arrived at OUT003
04:02
driver
Left Peliyagoda depot
02:34
system
Send an update
All messages ›
Driver
OUT001
Everyone
Running 151 min late · now expected 07:55
Voice
Send to OUT001



Made for the driver

Works with no signal

the day's trip on the phone; proof sent when signal returns

Dark screen while driving

nothing to read, nothing to tap

Fuel QR, even offline

the fuel pass works at the pump with no signal

Next stop, one action

Voice notes to listen to, not read

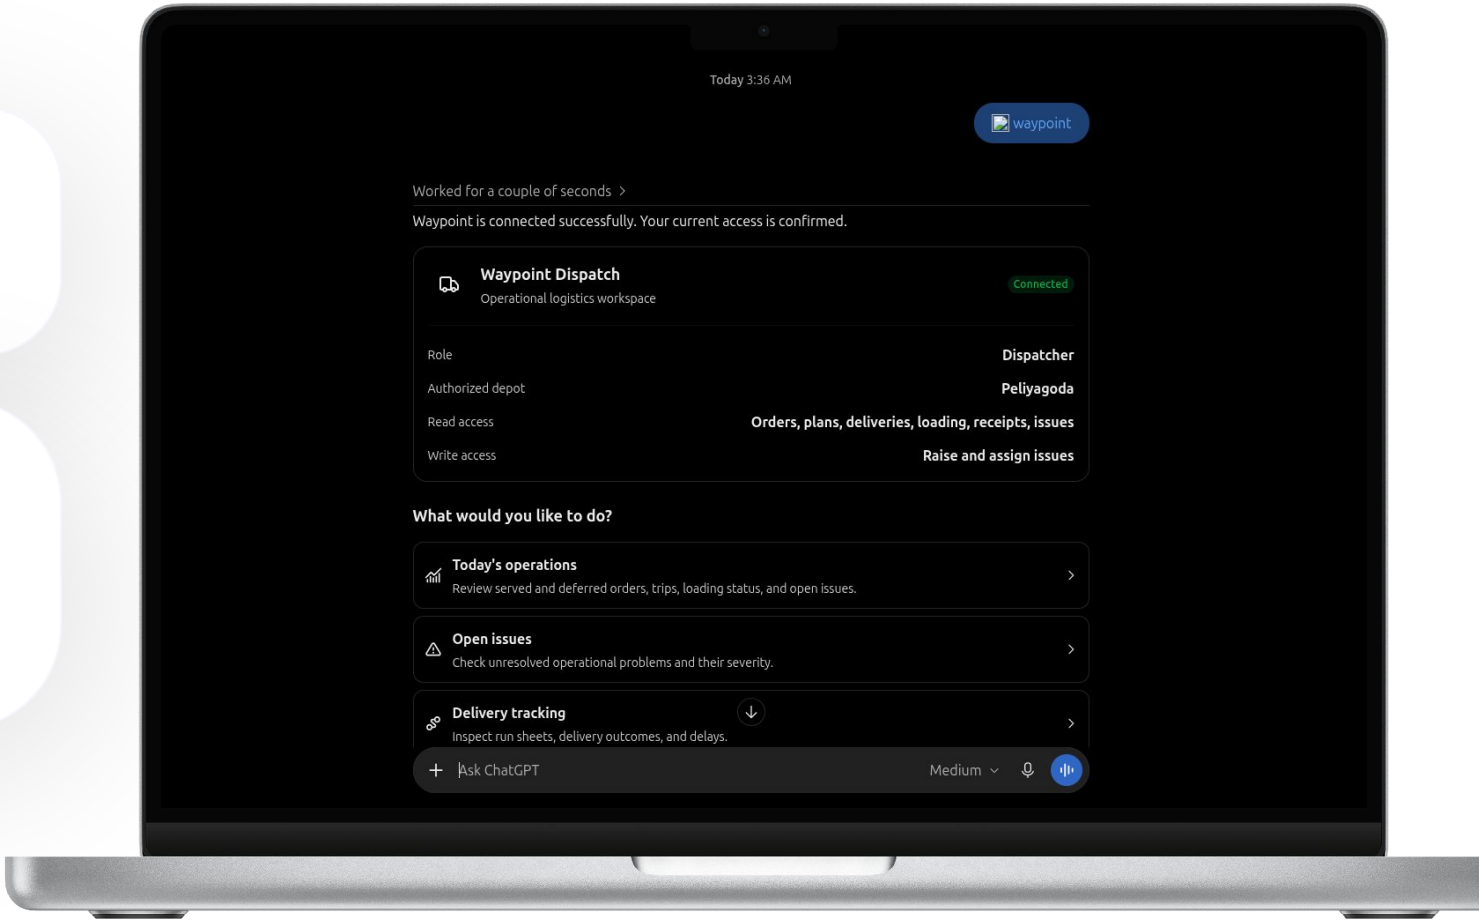
Ask GO with your own AI

MCP: connect your own AI

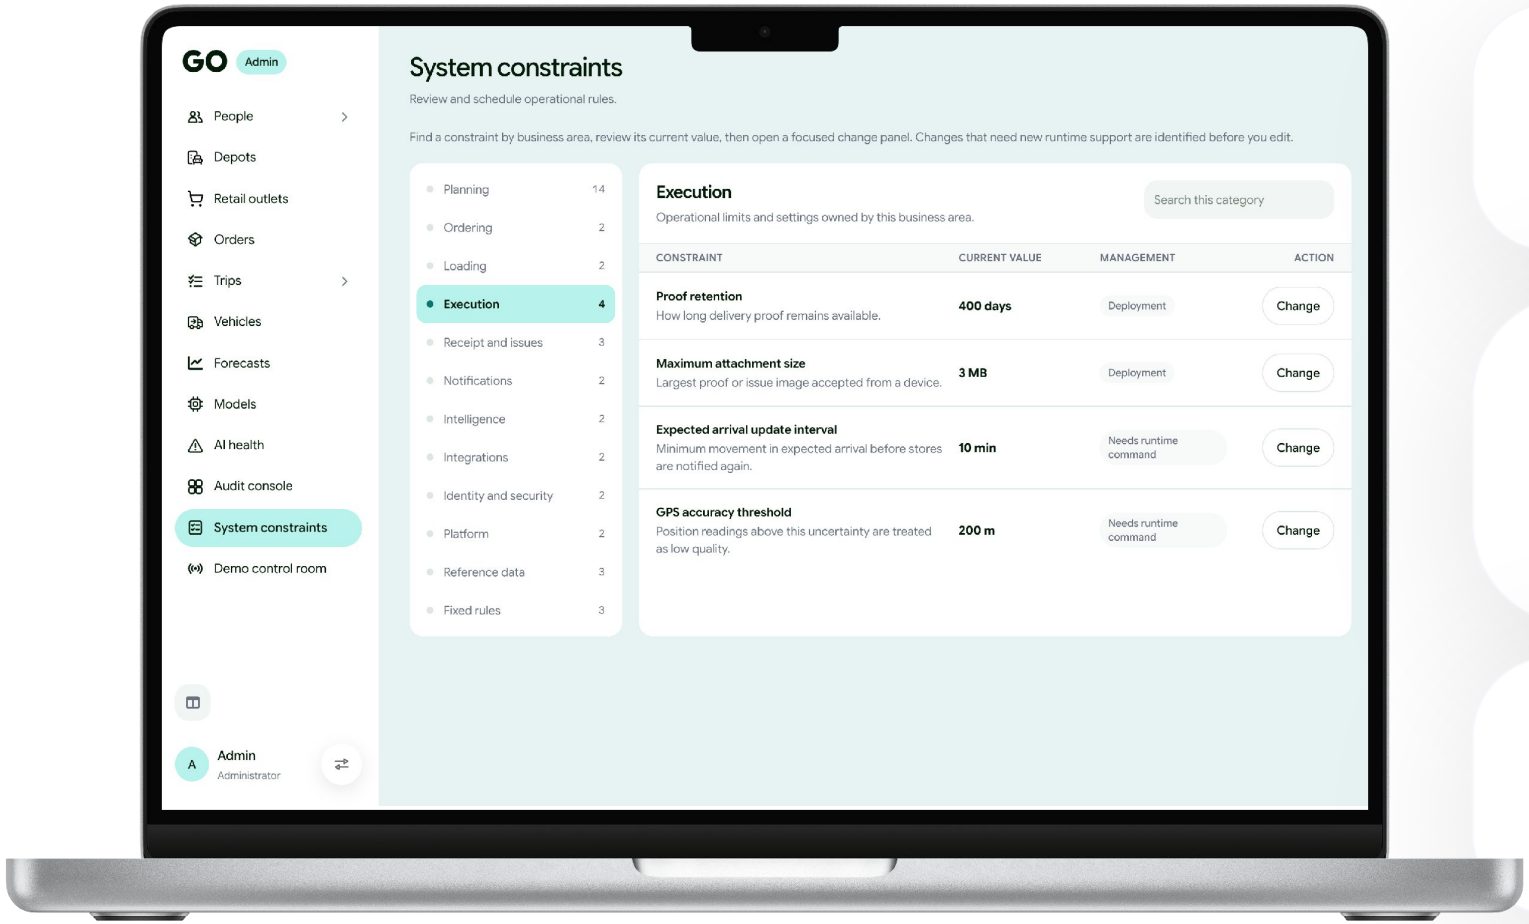
"Which trips are at risk this morning?"

Reads the day, within their permissions

trips, plans, issues, the audit trail



Run the rules, without a developer



Users, roles and policies

permissions written like a cloud console

Vehicles and planning parameters

changed without a new release

Every change audited

who, what, when and why

Designathon

Designed for people who **start at 3 AM**

Go Green
The next step to take · From our logo

Teal
Decisions and links · Calm & trusted

Night
Dark screens do for 3 AM · No glare in the cab



One action per screen

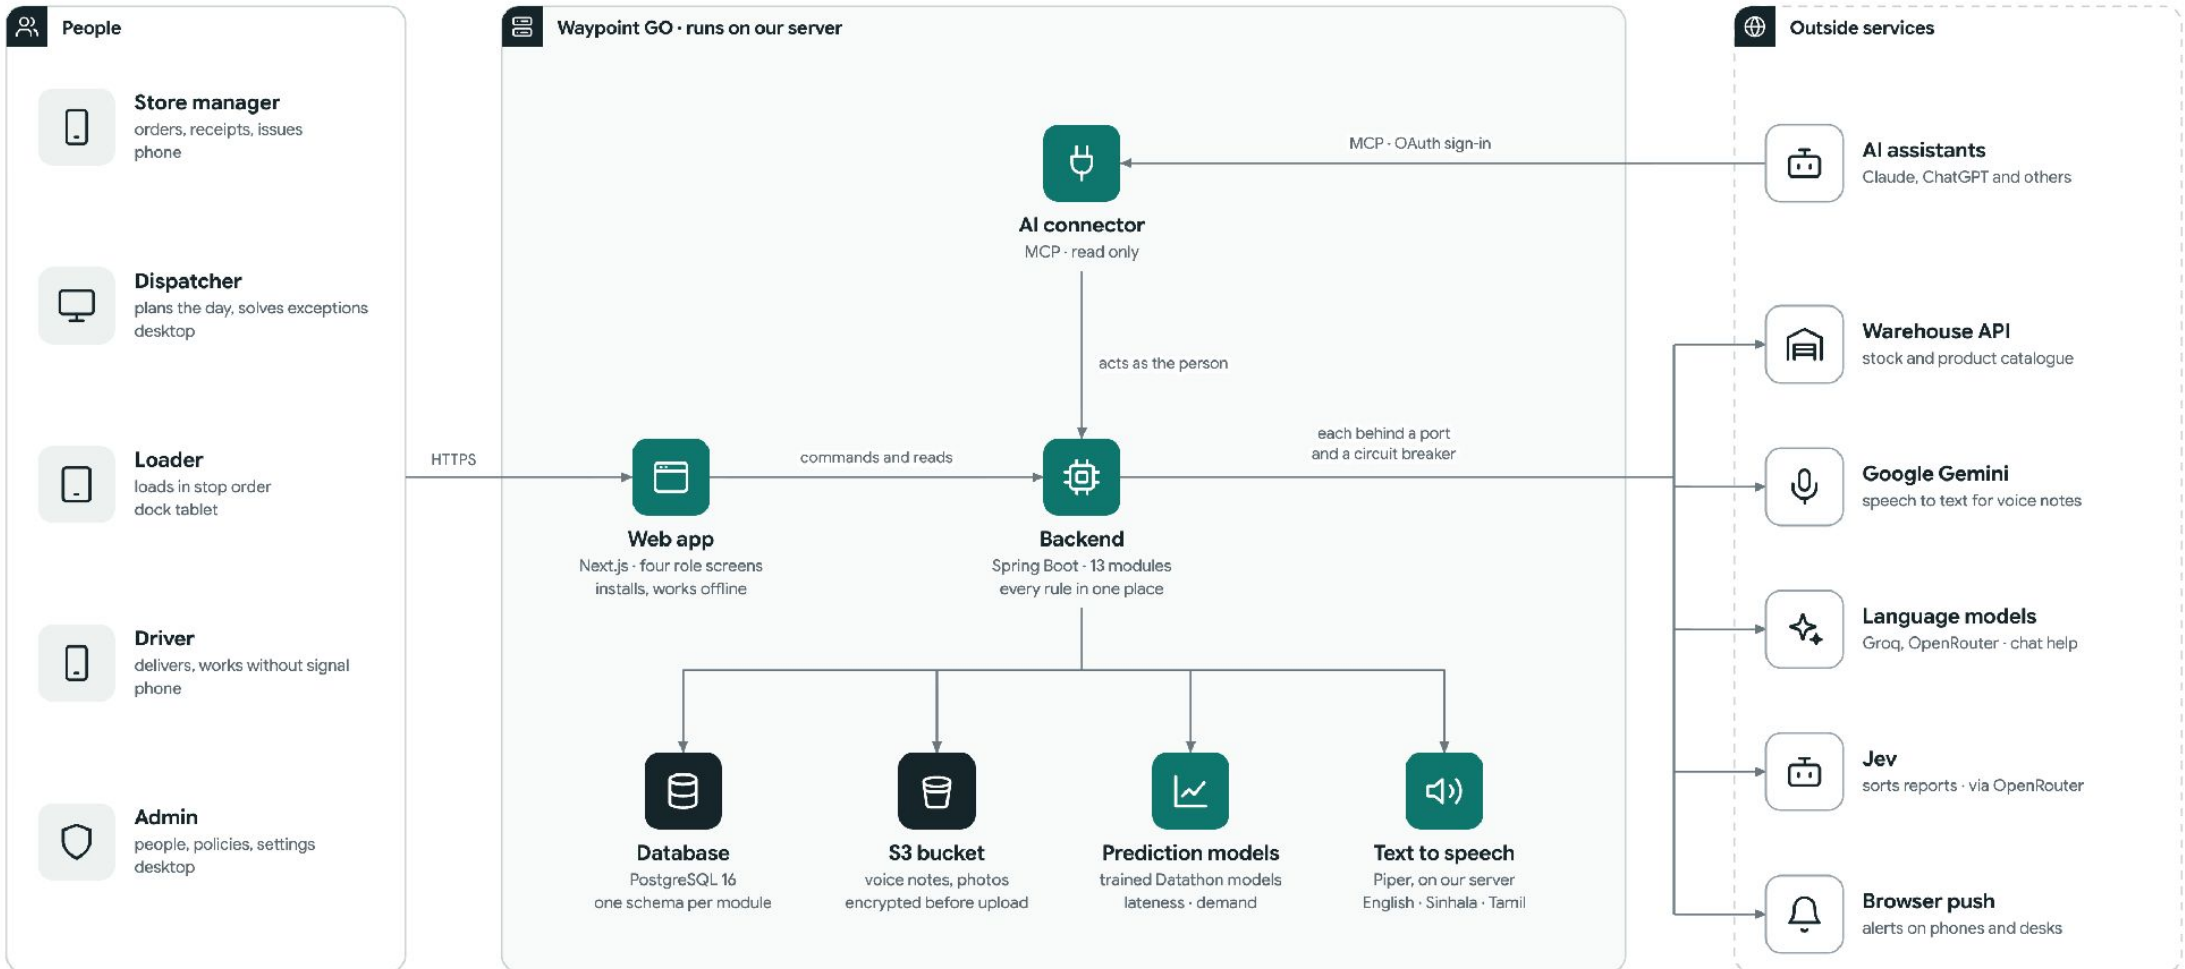
Slide to confirm

Sinhala · Tamil · English

Hackathon

Waypoint GO at a glance

Five kinds of people, one web app, one backend that holds every rule, and a few outside services behind safe connectors.



Part of GO · Keeps data · Outside service · Person · Request or call · Event, in the background · 1 Step in order

Waypoint GO · Team Xception

Tech stack



Spring Boot



Next.js



Tailwind CSS



PostgreSQL



FastAPI



scikit-learn



Docker



Cloudflare



Google Gemini



Jev by TypeSafe

Datathon

Three kinds of intelligence, under one rule

Task 1

Predict

Which stops will be late, and how long unloading takes

Task 2A

Foresee

Demand ten weeks ahead, as trucks, refrigerated trucks and drivers

Task 2B + Live planner

Decide

The best plan when trucks run short, and why

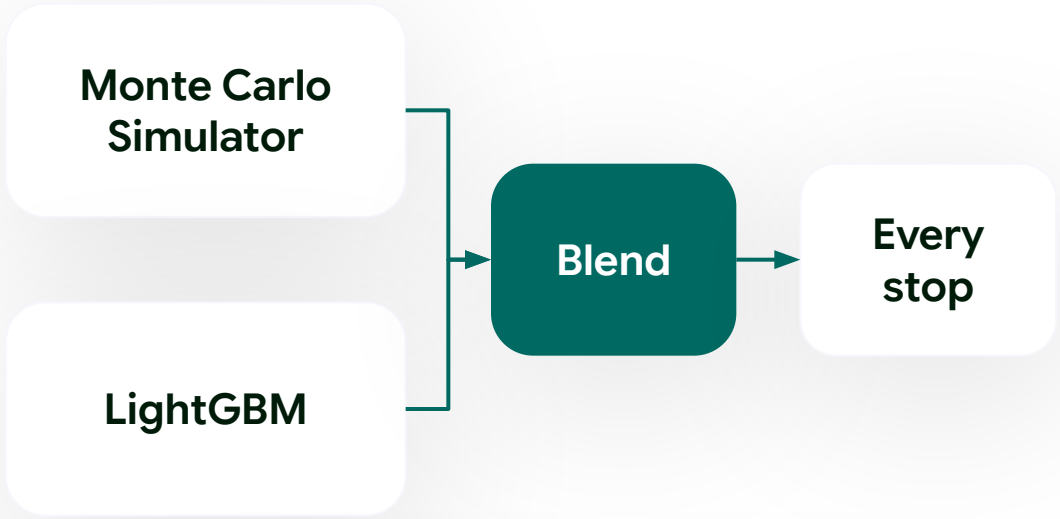
Rules First > **Our own models** > **Language AI, wording only** > **A person decides**

If a model is ever down, GO falls back to a safe estimate and says so on screen

Task 1 · Predict

Late risk and unloading time, predicted before the truck leaves

How it predicts



Why it matters: delay builds along the route



Lateness probability

56

out of every 100

late deliveries **SAVED**

Service time prediction

< 4 min

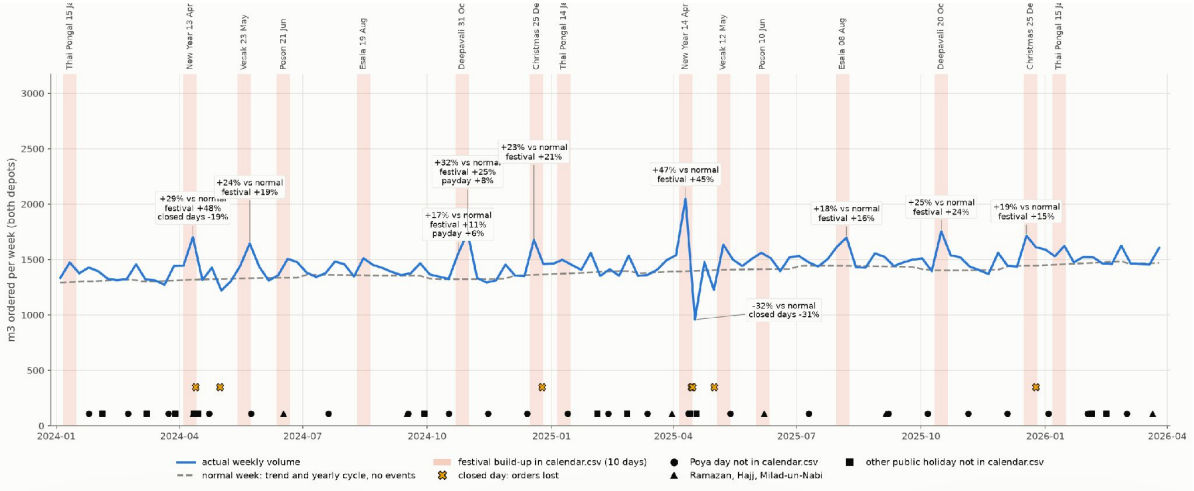
unloading-time error

- **Staff ready**
- **Dock ready**
- **Gate ready**

Task 2A · Foresee

Every demand peak has a cause

Weekly demand over two years, every peak labelled

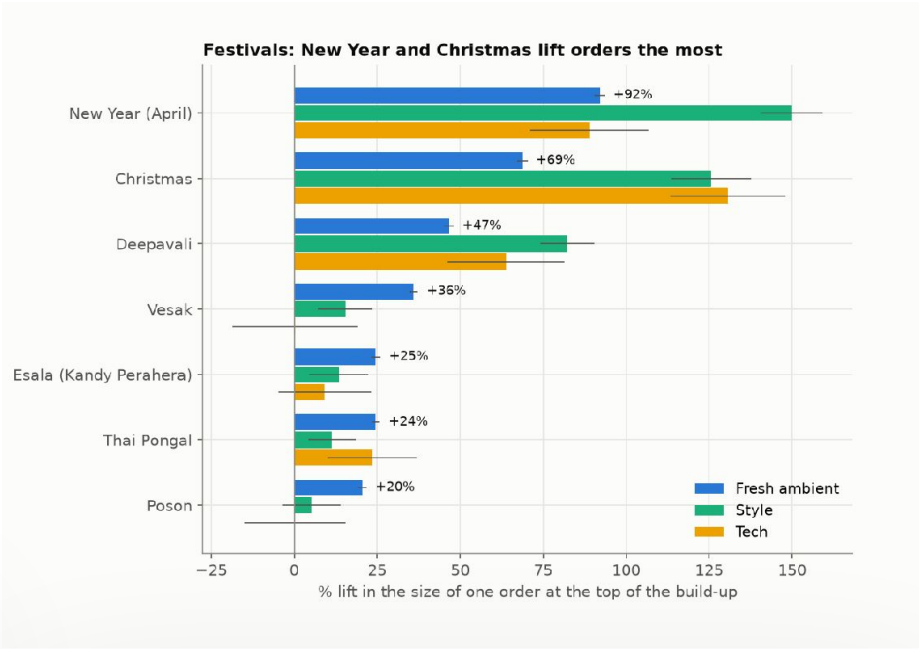


How it forecasts

Orders per store × **Size per order** = **Weekly forecast**

Orders: each store's ordering days, from its history
Size: one model per load: Fresh ambient, Fresh chilled, Style, Tech

Festival lift, measured per brand

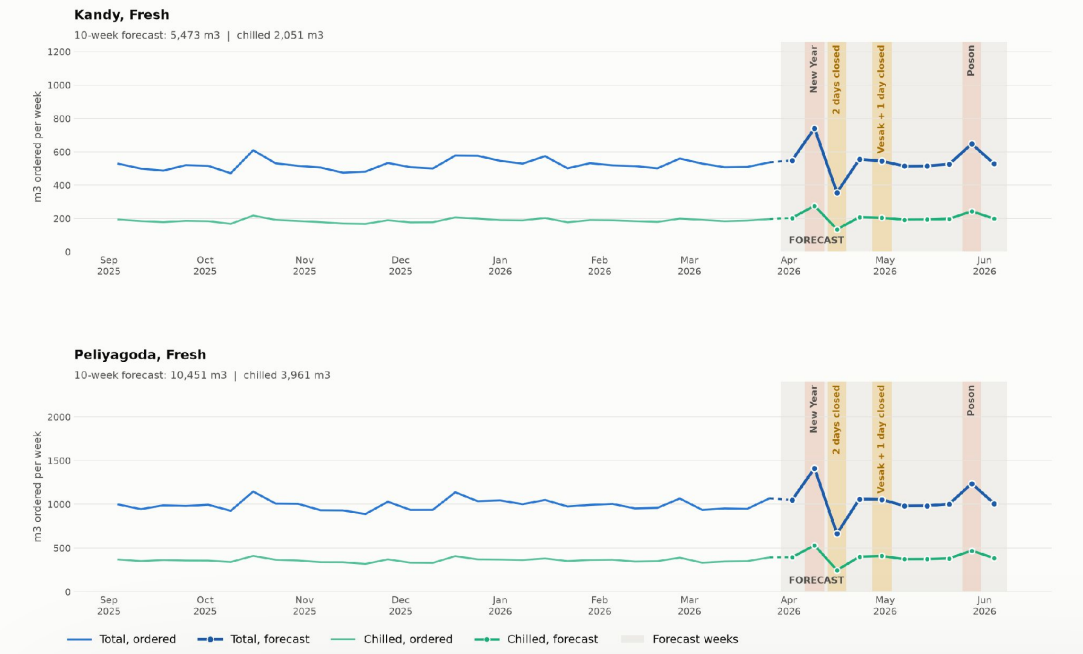


Each cause becomes a forecast input
Festival ramps, payday, monsoon and season. Holidays with no lift are left out.

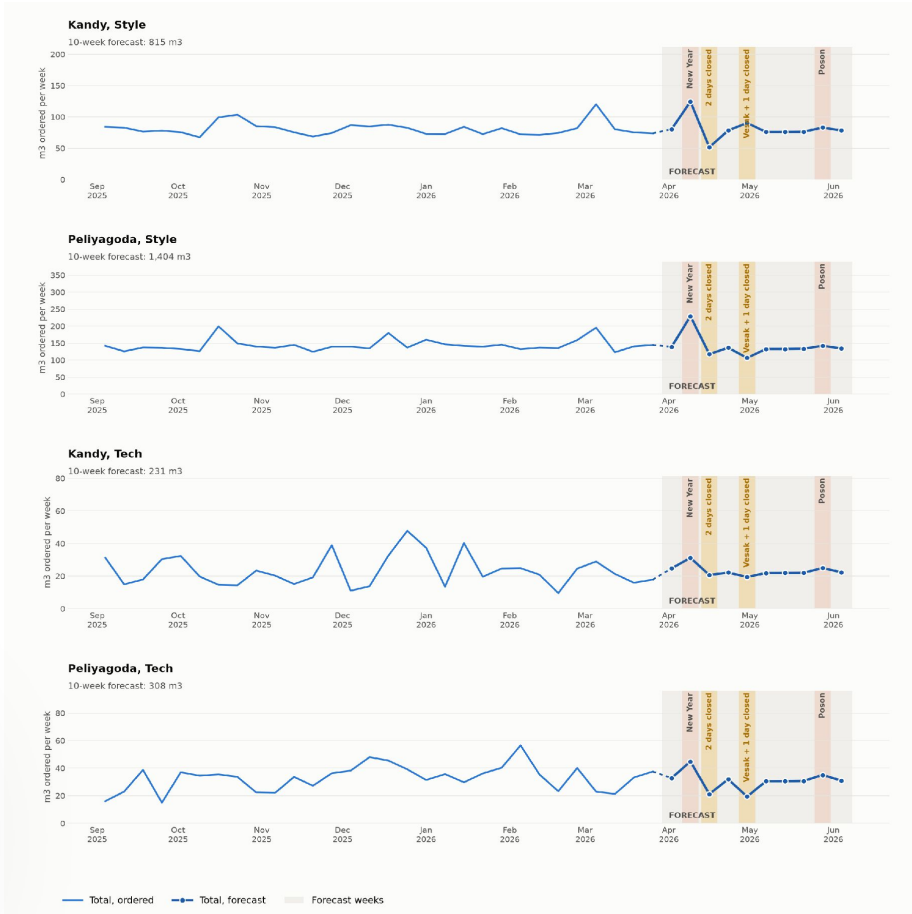
Task 2A · Foresee

Off by about 3%, ten weeks ahead

Fresh, by depot: total and chilled, m³ a week



Style and Tech, by depot: m³ a week



3%

Forecast error, ten weeks ahead

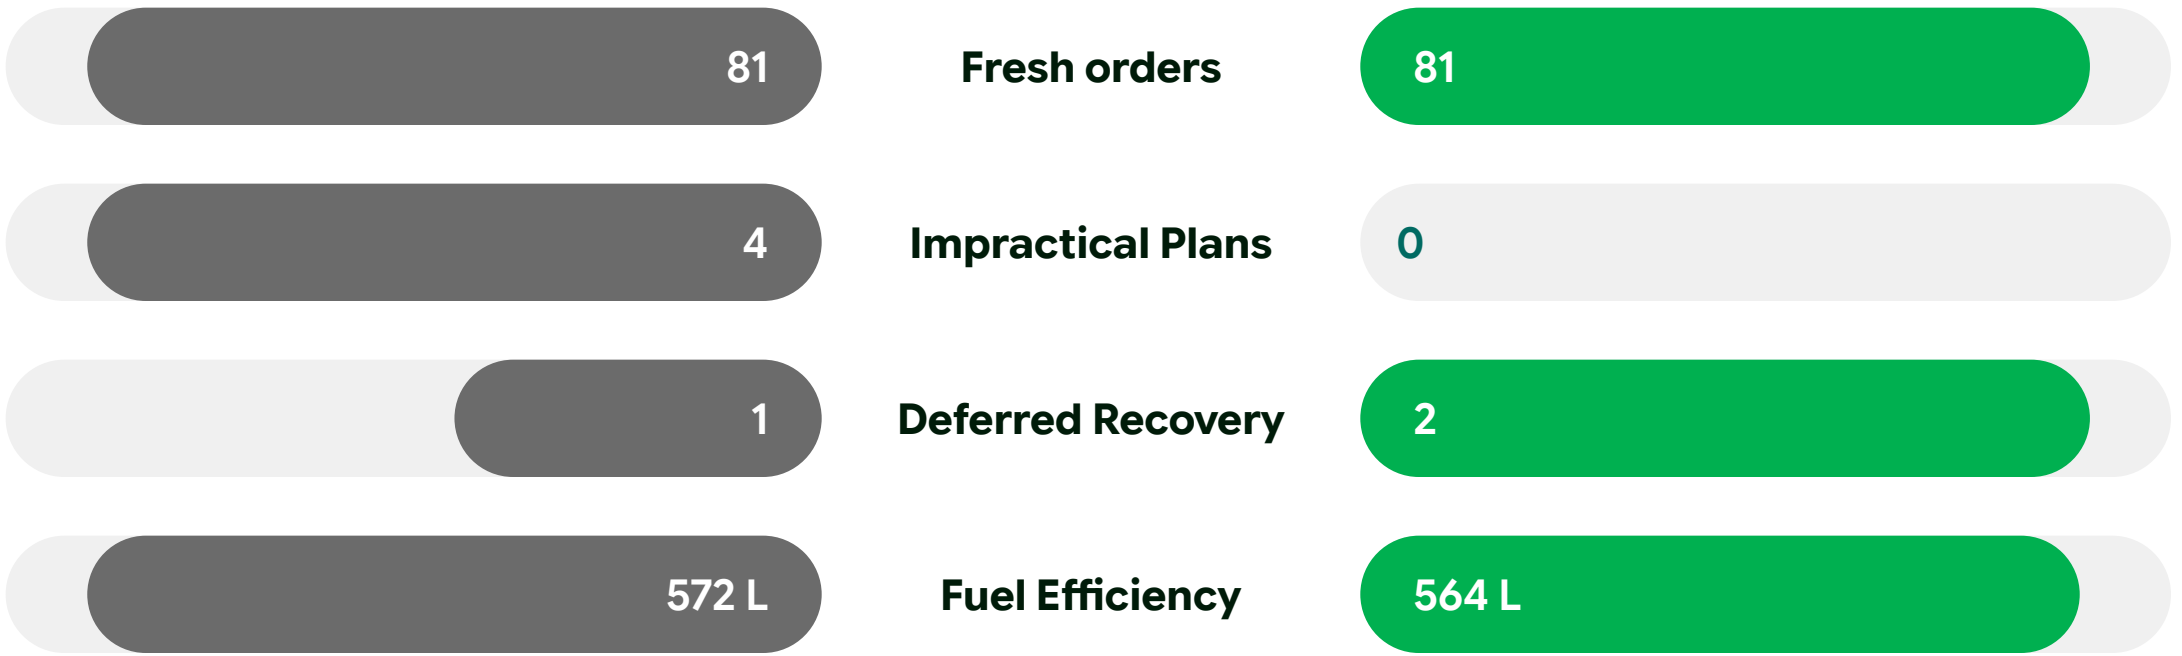
WAPE

Weighted Absolute Percentage Error

Task 2B · Decide

Priority Orders, Practical Plans

September 2025 simulation . 26 Operating Days . Daily averages



MCP · Your own AI

Ask GO with your own AI

Connect Claude or ChatGPT through MCP and ask GO in plain words

Claude · ChatGPT

“Which trips are at risk this morning?”

GO answers from the day's records

Read-only

Your AI can look, never change.

Trips, plans, issues and the audit trail, only within the permissions you already have

Your permissions only

it sees only what that person can already see

Every call audited

who asked, what was read and when

By default, every AI connection is read-only.

Business value

What changes for Waypoint

Four problems people live with today, and what GO changes for each

Today

One spreadsheet, one person

With GO

More orders, same fleet

every order checked against every rule; fewer trips and litres

Today

Fresh misses 08:00

With GO

Fresh on shelves by 08:00

late risk seen before the truck leaves

Today

The peak is a surprise

With GO

Hire before the peak

the forecast turns weeks into vehicles, reefers and drivers

Today

Disputes rest on memory

With GO

Nothing lost, nothing argued

every handoff on record, even with no signal

Four people, One app. Every handoff on record.

Any Questions?

WayPoint | GO

Thank you!

Team | **Xception**

Appendix

Click a topic to jump to it · every backup slide has an Index button to come back

12 How GO is built

13 The planning engine

14 Task 1 · Late risk and unloading time

15 Task 2A · Demand forecast

16 Task 2B · The peak-day plan

17 Trust, safety and privacy

18 Works offline, fails gracefully

19 Business case

20 Rollout plan and risks

21 Beyond Waypoint: market and revenue

22 From design to product, and our honest limits

23 Engineering quality

24 Numbers at a glance

Use these slides during Q&A only

How GO is built

One web app for every role, a modular backend, and the models served beside it

People

Store manager · desktop or phone

Dispatcher · desktop

Loader · shared tablet, PIN

Driver · phone, offline

AI assistants · Claude, ChatGPT via MCP

Web app (Next.js PWA)

Four role workspaces

Offline queue, replayed exactly once

Service worker and caches

Sinhala, Tamil, English

Backend (Spring Boot)

14 modules, from ordering to intelligence

One command path: permission check + audit

Events and a live change stream

Data and models

PostgreSQL 16 · row-level security · 188 migrations

Model service (FastAPI): late risk + demand

Checksums and file scan; fallback estimate if down

Language AI gateway: off by default, masked, capped

Runs on one server (≈ LKR 19,000 a month) with Docker Compose · Prometheus, Grafana, Loki and alerts · tests on every pull request

The planning engine

Every order checked against every rule; people make the final call, with a reason

Step	What happens
1 · Place	Orders placed in priority order; earlier deferrals first (9 rules)
2 · Re-plan reefers	Refrigerated vehicles planned again as a whole; kept only if better
3 · Optimise	Same orders on fewer vehicles, trips and litres (1 vehicle, 3 trips, 15 L saved)
4 · Decide	Dispatcher places by hand, swaps or keeps deferred, always with a reason
5 · Publish	Every rule checked again; loaders, drivers and stores are told

Check	13 checks per order
Vehicle	available, home depot, refrigerated, van-only
Trip	one brand and district, one temperature
Capacity	weight and volume, no split orders
Time	delivery window, 2 trips a day, Fresh 270 min
Fuel	weekly quota per vehicle

Explanations come from the engine, not guesswork. Each waiting order shows the rule that stopped it, what was fine, where it could go, the closest it came and the smallest change. A what-if re-runs the day with one more vehicle without changing the plan.

Why 72 in the app and 77 in the Datathon? The app also enforces every store's delivery window and fuel quota; the Datathon plan used the booklet's time budgets, where 3 stops arrived 16 to 87 minutes late.

Task 1 · Late risk and unloading time

Labels built from the route records; a simulator and LightGBM, blended

6 unseen weeks	Statistical	LightGBM	Ensemble	Selected
Unloading error (min)	5.24	3.75	3.66	**3.72**
Late log-loss	0.302	0.133	0.125	**0.131**
Correct at 0.5	85.8%	94.5%	94.6%	**94.5%**
False alarms	368	115	89	**95**
Training time	0.01 s	18 s	16 min	**1.5 min**
Per order / size	–	–	2 s · 33 MB	**48 ms · 9.7 MB**

Service time = leave − later of (arrival, window open)

Late = arrival after the window closes

Simulator: hundreds of runs per route, delays carried stop to stop

Blend in log-odds: 50/50 for lateness, 55/45 for minutes

Caught 488 of 678 late; 84% of flags right

Why the selected model? Within 0.06 minutes of the best, 10 times faster to train than the ensemble, one library, 48 ms per order. In production it must first beat the built-in estimate on our own deliveries (backtest and trust gates).

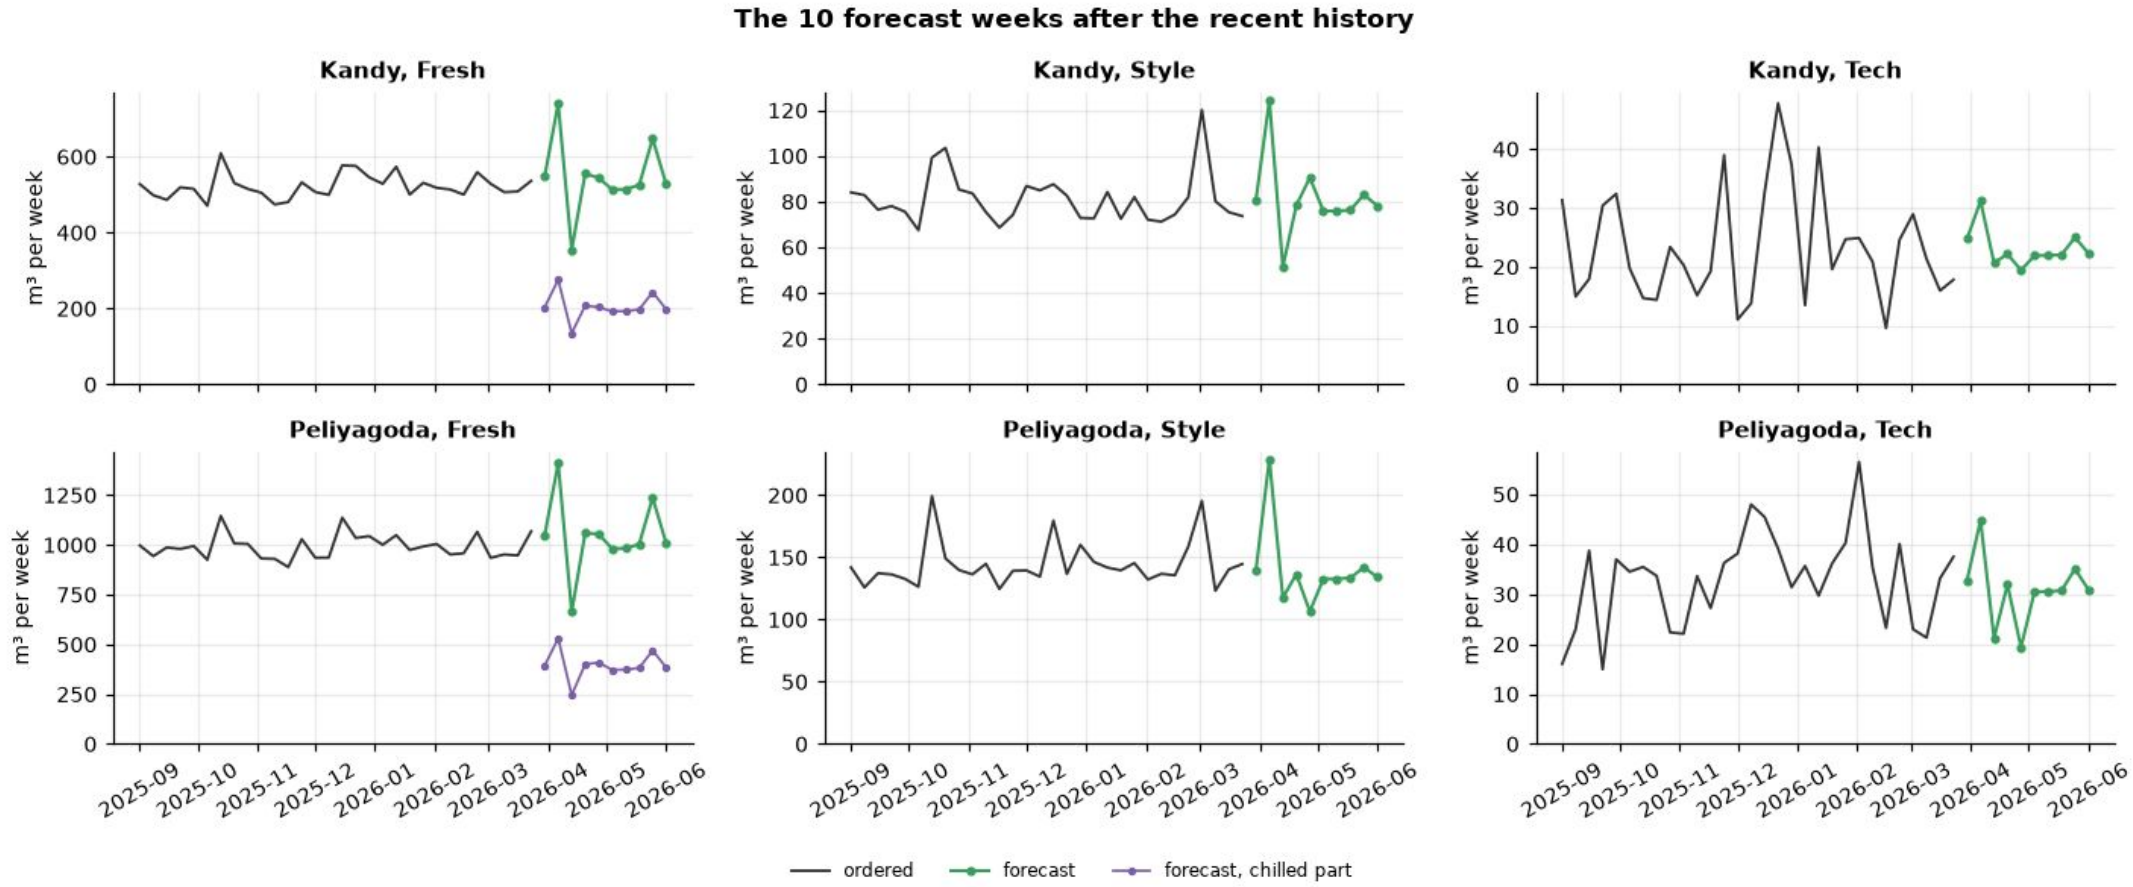
0.976

AUC: a late stop ranks riskier 97.6% of the time

0.957

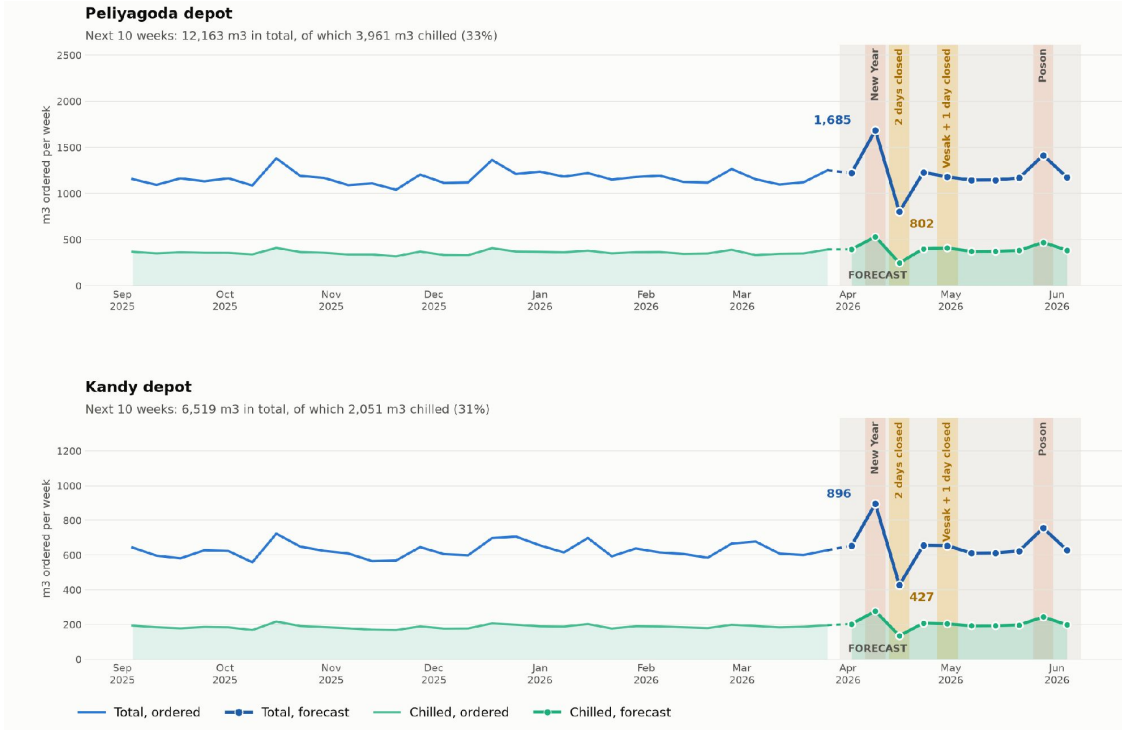
AUC with no road data (production figure)

Task 2A · **Demand,** seen ten weeks ahead

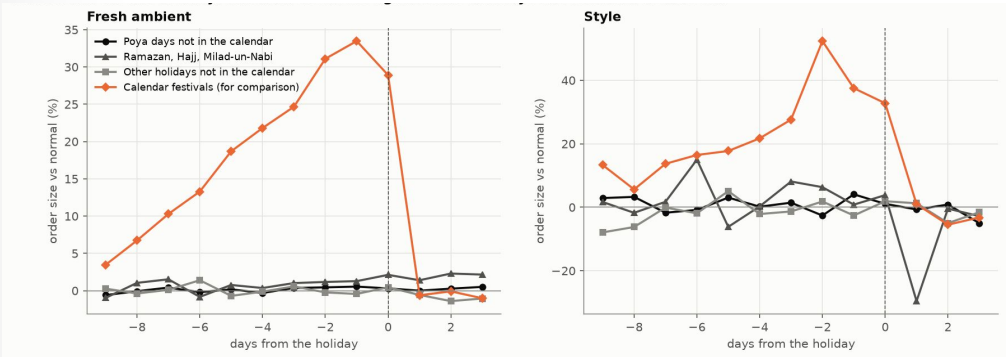


Task 2A · **Demand,** seen ten weeks ahead

Weekly forecast per depot, total and chilled, m^3



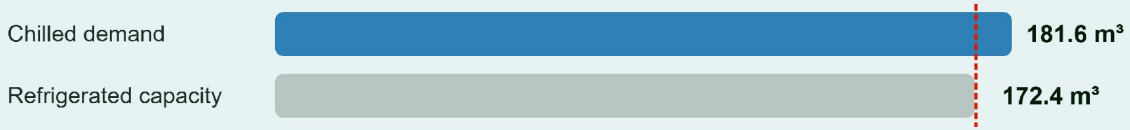
Holidays not in the calendar: no lift



Task 2B · The peak-day plan

Too few trucks? The best plan, explained

THE REAL LIMIT: REFRIGERATED SPACE



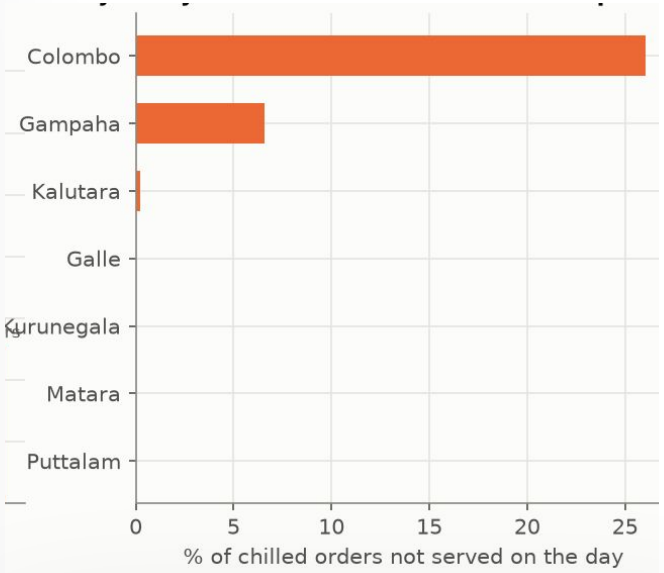
TASK 2B · GOALS SOLVED IN ORDER BY AN EXACT OPTIMISER (CP-SAT)

1. **Yesterday's deferred orders** — All 10 served, so no store is skipped twice
2. **Most orders** — 77 of 85 · 79 proven maximum
3. **Most chilled m³** — 131.0 of 181.6 m³
4. **Fewest trips** — 27 trips
5. **Fewest vehicles** — 17 vehicles

Each goal is solved, then locked in, so a later goal can never trade away an earlier one. The seven booklet rules are hard constraints.

Yesterday's deferrals go first, then as many orders as fit. **Every order left waiting comes with a reason the dispatcher can give the store.**

Where chilled orders wait



4 of 9

reefers on the road; 5 in the workshop

Trust, safety and privacy

Advice, never autopilot; every AI answer checked; every action audited

Models

Predictions are advice: they never change a plan

Trust gates before going live: precision ≥ 60%, recall ≥ 70%, calibrated

Outcomes recorded, daily metrics, drift alerts

Shadow candidate on every plan; worse model rolled back

Competition models off in production until own-data model passes

Language AI

Off by default; a daily call cap

Only allow-listed facts are sent; names masked

Any number not in the facts → the plain sentence stands

Zero retention; prompt version logged

An outside AI (Jev) sorts problem reports, in shadow only

Access and data

Role + depot, store or vehicle scope on every request

Row-level security in the database

Every command and AI call audited

Sessions 12 h, 2 h idle; login throttling

AI assistants (MCP): own permissions, confirm-before-change

Ready for Sri Lanka's PDPA (in force 1 January 2027): data agreements before any cloud AI, hosting in Sri Lanka as an option

Works offline, fails gracefully

Designed on Day 5 as our failure screen; built and working today

Driver, loader and store keep working with no signal

1. Saved on the phone with a command ID; sent exactly once
2. Survives a reload or an expired session
3. ‘Saved on this device’ is shown apart from ‘confirmed’
4. Conflicts are flagged for review, never overwritten

If this fails…	GO keeps going
Mobile signal	Queue on the phone, sync later
Warehouse API	Orders kept as ‘stock not checked’
Model service	Built-in estimate, labelled
Language AI	‘Suggestions: rules only’
Push service	In-app notifications
Vehicle breakdown	Driver reports it; plan revised

Business case

Team estimate from public data, 5 Oct 2026; rows are rounded

LKR million a year	Low	Expected	High
Fuel, better loads	4.5	11.0	20.0
Fresh before 8 AM	2.2	6.0	13.9
Fewer vehicle-days	1.4	2.7	5.4
Damaged-goods write-offs	1.2	2.4	3.6
Festival hires avoided	0.5	1.1	1.8
Fuel-log reconciliation	0.5	1.1	2.7
Planning, calls, paperwork, disputes	1.2	2.3	3.5
Total	**11.5**	**26.5**	**50.8**

Option, 3 years	Cost and gap
Waypoint GO	**≈ LKR 20.3m all-in; full scope**
Routific	LKR 6.6–8.5m; route planning only
Shipday	≈ LKR 19.4m; gaps below
Onfleet Enterprise	LKR 36.7m+; gaps below
Track-POD	LKR 34.8–63.0m; gaps below

Every bought option also needs ≈ LKR 6.2m of integration, and none covers multi-brand allocation with explained deferrals, reefer/van/mall rules, offline Sinhala and Tamil, store receipt and disputes, the Fuel Pass quota, or data kept in Sri Lanka.

Costs: build LKR 8.2–11.8m · run ≈ LKR 19,000 a month · maintain ≈ LKR 3.3m a year · 3-year total ≈ LKR 20.3m · payback 11 / 5 / 3 months (low / expected / high) · 3-year value ≈ 3.9× cost; expected net gain ≈ LKR 59m

Rollout plan and risks

22 weeks with 4 engineers; the critical path is collecting real outcomes, not code

Weeks	Work
1	Hygiene and safety: secrets, backups, restore drill
2–3	Baseline: today's KPIs measured
3–5	Quick wins: early capacity plan
5–9	Foundation fixes
6–13	**Kandy pilot: 4 weeks shadow, 4 weeks live**
10–17	Field features
≈ 19	Model trust gates on real deliveries
17–22	AI features, after go/no-go and data agreements

Risk	Answer
Drivers don't adopt	Depot champion, induction card, paper fallback
Models not trusted	Shadow first, trust gates, reasons on every flag
Data protection (PDPA)	Agreements and masking before any cloud AI
New AI vendor	Behind our own interface, pinned version, rules fallback
Exchange rate on cloud costs	Invoice in LKR with yearly indexation

Pilot success: Fresh before 8 AM +5 points · calls −30% · write-offs −20% · 90%+ of stops in the app · zero cross-store data leaks · zero lost delivery records

Beyond Waypoint: market and revenue

Waypoint is the reference customer; the same platform fits other fleets

Tier	Per vehicle a month	Includes
Core	LKR 4,900	Dispatch, offline driver app, proof of delivery, store receipt, Sinhala and Tamil
Pro	LKR 8,200	Core + constrained allocation with explained deferrals, reefer windows, fuel quota, issues
Enterprise	LKR 13,100	Pro + arrival times, late risk and forecast, capped AI, single sign-on, hosting in Sri Lanka

Year	Customers	Vehicles	Revenue
1	5	125	**LKR 14.4m**
2	14	420	**LKR 42.0m**
3	28	980	**LKR 90.9m**

LKR 525m

a year within reach: ≈ 300 chains, manufacturers and distributors with 10+ vehicles (≈ 6,000 vehicles)

≈ 95%

gross margin on subscription (≈ LKR 330 per vehicle a month); min. 10 vehicles; setup LKR 1.0–4.9m

Next

Prerequisite: multi-tenancy, decided after the pilot proves the value

From design to product, and our honest limits

The Figma file is the specification; every departure is documented

Design → product

Day 5 failure screen (‘proof saved with no signal’) is built and works

All four roles, end to end, on phone, tablet and desktop

Departures from the design listed in the README with reasons

Sinhala and Tamil reviewed by a native speaker

Explanations, what-if and suggestions added after the Datathon

What we state openly

Models were trained on synthetic competition data; production retrains on Waypoint’s own deliveries first

Tech forecast error is ≈ 30% (only 3% of volume)

Without live road data, late-risk AUC drops from 0.976 to 0.957

Value figures are estimates until the pilot measures them

Language AI and the report-sorting AI (Jev) are off or in shadow until agreed; multi-tenancy not built

Engineering quality

Built to be run, not just demonstrated

217
backend test files: unit, architecture, integration

93
frontend unit tests

98
browser test specs across every role

12
model-service tests, incl. exact reproduction

188
versioned database migrations

41
feature plans written before the code

One command (docker compose up) starts the whole stack: database, backend, model service, frontend and seed data

Every pull request: backend tests, type check, unit tests, build and browser suites per role; secret scanning; model-file scanning

Prometheus, Grafana, Loki and Alertmanager, with alerts on model drift and refused AI-assistant calls

AI coding tools were used and are disclosed in the repository (docs/ai-disclosure.md)

Numbers at a glance

Every rounded number, with its exact figure and source

Topic	Say	Exact figure and source
Late deliveries flagged	**7 in 10**	488 of 678 late deliveries, 6 unseen weeks (3 Jan–14 Feb 2026)
Warnings that were right	**8 in 10**	488 of 583 flags (95 false alarms)
False alarms	**74% fewer**	95 vs 368 for the statistical baseline
Unloading-time error	**under 4 minutes**	3.72 min mean absolute error
Demand forecast error	**about 3%**	3.03% / 3.18% / 2.99% WAPE; chilled 2.07% / 1.97% / 2.65%
Same week last year	**8%**	8.26% WAPE on the same selection forecasts
Peak day	**77 of 85, 79 max**	27 trips, 17 vehicles; 8 deferred (6 unavoidable, 2 by choice)
One reefer back (app)	**+9 orders, +48% chilled**	81 vs 72 orders; 161.0 vs 108.9 m^3 (what-if, VEH004)
Delayed orders	**nearly all chilled**	1,911 of 1,956 not dispatched on order day were Peliyagoda chilled
Value and cost	**LKR 26.5m · LKR 19,000 a month**	expected value a year; running cost; payback ≈ 5 months