

MAGNAQORE · CASE STUDY

NTK MIR

From chaos to system: a 400-person logistics group rebuilt
into an operating model that runs on AI

Logistics & transport · Russia & Kazakhstan · May – August 2026

Starting efficiency 3 / 10 CEO self-assessment	Loss identified roughly \$21M 28% of approximately \$75M revenue	AI opportunities mapped 37 Across 8 departments	Delivered in 90 days 2 Countries, one operating model
--	--	---	---

What this document is
A field record of the Enterprise Evolution System — MagnaQore's implementation methodology for AI Organization Transformation — applied end to end inside a real business, audited, built and running in production. Every claim below traces to an audit transcript, a dated delivery log entry, or a system that is live at the time of writing.

The thesis under test: operations before automation. Technology strategy and implementation sit at steps 9 and 10 of MagnaQore's twelve — deliberately. The Group is where that sequence was proved: eight departments audited, an operating model redesigned, governance and legal risk resolved, and only then a single line of AI shipped.

ABOUT THIS CASE STUDY

This case study describes a real transformation delivered by MagnaQore. Some details have been simplified or anonymised to protect client confidentiality. It is designed to illustrate our approach and its impact, not to disclose sensitive internal information. For deeper, NDA-level discussion, please contact us directly.

Contents

SECTION	WHAT IT COVERS
1 Executive Summary	The before, the after, and what actually caused the change
2 The Company at Zero	Four independent diagnoses of the same chaos
3 The Method	Where the Group sits on the twelve-step Operating Model
4 Phase One — Enterprise Assessment	Operational, organisational, governance and data audits
5 Phase Two — Building the Capability	Eight delivery waves, May to August 2026
6 Build vs. Buy in Practice	What the Group built, bought, and deliberately refused
7 Outcomes	What the company holds today, against the eleven deliverables
8 The Commercial Proof	A client the Group could not previously serve
9 What Comes Next	The AI Department — the one step still open
A Appendix — Delivery Record	The engagement month by month, and confidentiality notice

A note on evidence

Findings are drawn from eight department audit transcripts, an anonymous 82-response employee survey, a compliance review against four statutes, and an engagement log of roughly 430 dated entries spanning four months. All monetary values are stated in US dollars and rounded. Client identity, exact commercial terms and individual performance data have been withheld or aggregated. Where something is unfinished or unproven, this document says so.

1. Executive Summary

In May 2026, the Group was a approximately \$75M transport and logistics company with 400 employees and nearly 30 warehouses spanning the south-west to eastern Siberia — and no system. Its founder rated operational efficiency at three out of ten. Shipments

were tracked in a Google spreadsheet. Three accounting offices could not see each other's books. A hundred vehicles were maintained from handwritten notebooks. Two security officers policed twenty-seven regions with a pen and a fuel-card statement.

Ninety days later the company runs on a single registry holding nineteen thousand migrated shipments, eleven governed roles, live integrations with 1C, the customer portal and a customer's own CRM, an automated control that audits every quotation against a tariff engine, a management analytics application, and a voice AI agent that prices freight during a phone call. The same operating model runs in two countries.

1.1 What actually caused the change

Not the AI. The AI arrived in week eight.

The engagement opened with four audits running in parallel — operational, organisational, legal and technical — before a single system was designed. That sequencing is the entire case. Three findings determined everything built afterwards:

- The operational audit priced the chaos at roughly \$21M a year — 28% of revenue — and identified 37 AI opportunities, of which 19 required no new data at all, because the company was already typing that data in by hand.
- The organisational audit — an anonymous survey of 82 employees across three departments — found a company with an unusually strong human core and severe systemic cracks. Trust inside teams scored 4.5 of 5. Cross-department handover was a systemic failure: all three departments named each other as the primary source of operational breakdowns.
- The legal audit killed the original architecture. Five components were outright statutory blockers, with combined exposure of up to \$833K on a first inspection. The voice platform originally specified would have sent Russian citizens' voice biometrics to servers in the United States. It was replaced before a single call was placed.

THE FINDING THAT SHAPED THE BUILD

The company was not short of demand, trucks, or people. It was short of a single source of truth. Data lived simultaneously in Google Sheets, three disconnected 1C instances, personal Excel files, messenger chats and paper notebooks — so every process crossing a department boundary was completed by a human retyping something. Automating that would have scaled the chaos. It had to be replaced first.

1.2 The shape of the transformation

DIMENSION	MAY 2026	AUGUST 2026
System of record	Google spreadsheet, ~25 monthly tabs, no permissions, no history	Governed registry, nearly 20,000 shipments migrated, full revision history
Access control	Everyone sees everything, or nothing	11 roles, 11 policies, field-level permissions, combined-role model
Cross-system flow	Manual re-entry into 1C, three times per order	Live two-way 1C exchange — 42 orders created by 1C in a single day
Pricing control	Discounts invisible until someone read an invoice	Every quotation scored against the tariff engine automatically on save
Management visibility	Office manager assembling a summary by hand each morning	Analytics application: branch plans, rankings, month-end forecasting
Customer contact	200 cold calls a day for the entire sales team	Voice agent quoting live freight prices, 23-test regression suite
Geography	Russia only; Kazakhstan on a separate spreadsheet	Two synchronised registries, mirrored international shipments
Legal posture	Architecture carrying 5 statutory blockers, up to \$833K exposure	Localised stack — voice, models and parsing inside the perimeter

1.3 Where the programme stands

Eleven of the twelve steps of the Enterprise Evolution System are delivered or in active progress. The twelfth — a permanent internal AI function — was recommended and specified by MagnaQore, and the founder elected instead for MagnaQore to deliver the whole programme. Ownership of that function is now being transferred in stages.

STEPS OF THE SYSTEM ENGAGED

11 of 12 Assessment through continuous evolution Delivered by MagnaQore AI function By the founder's choice; handover staged
Systems in production 2 countries Russia and Kazakhstan

The Group holds the systems. The function that runs them is still ours, by agreement rather than by oversight — and it is being handed across in stages as confidence in the redesigned operating model grows and AI literacy spreads through the teams. Section 9 sets out that transition.

3. The Method

The Enterprise Evolution System — MagnaQore's implementation methodology, applied to every client engagement — runs twelve steps. Technology strategy is step nine; implementation is step ten. Every firm in the market starts at implementation. This engagement started eight steps earlier, and the table below records exactly where the Group landed on each.

	OPERATING MODEL STEP	HOW IT WAS DELIVERED AT THE GROUP	STATUS	SEE
01	Enterprise Assessment	Eight department audits plus an external logistics audit by a former operator at a global logistics carrier; 37 opportunities and roughly \$21M of annual loss quantified against an audited average order value.	Delivered	4.1
02	Operating Model Redesign	The registry schema is the redesigned process: nine tables, a grouped shipment card, conditional fields, a governed status lifecycle. Not a database — a codified way of working.	Delivered	5.2
03	Governance	Seven roles grew into eleven roles and eleven policies with field-level permissions; separately, a full compliance review against Russian data-protection law conducted before any build.	Delivered	4.3, 5.3
04	Leadership Transformation	The director moved from a hand-assembled morning summary to branch-level plan versus actual with month-end forecasting. New management rules emerged directly from data he could finally see.	Partial	5.6
05	Internal AI Function	Recommended and specified. The founder elected for MagnaQore to deliver the full programme rather than stand up an internal unit alongside it; ownership is being transferred in stages as literacy matures.	By agreement	9
06	Department Transformation	Sales, Accounting, Logistics, Warehouse, Legal, Security and Fleet each received specific systems. Coverage is uneven and is mapped honestly in Section 7.	Partial	5.4
07	AI Literacy	Two practical training programmes built and delivered — accounting and prompt-writing — each session ending in a working artefact rather than a certificate.	Delivered	5.7
08	Knowledge Management	A thirteen-document agent knowledge base, revision history on every record, API documentation handed to three external development teams, and operational runbooks.	Delivered	5.4, 5.7
09	Technology Strategy	Explicit build-versus-buy: the ERP was bought, the registry and AI layer were built, the pricing engine was forked and frozen, and model independence was engineered in from the start.	Delivered	6
10	Implementation	Roughly 430 logged delivery entries across two countries, separated test and production contours, staged deployment with one-command rollback.	Delivered	5
11	Measurement	Analytics application, automated discount control, a 23-test agent regression suite, and a drift detector running against the live pricing API.	Delivered	5.6
12	Continuous Evolution	Regression on every change, drift detection, separated environments, and a private repository with promote-and-rollback discipline.	Delivered	5.8

3.1 Two directions, one system

The Enterprise Evolution System runs in two directions. Direction One is the engagement — the twelve steps above, which redesign how the business actually runs. Direction Two is the permanent internal AI function that keeps Direction One alive after the consultants leave. The two can run in parallel or in sequence, and that is a client decision rather than a methodological one.

At the Group they ran in sequence. Direction One was delivered in full; Direction Two is a staged handover now under way. Section 9 explains why the founder chose that order and what the transfer requires.

	THE SYSTEM AS DESIGNED	HOW IT RAN AT THE GROUP
Phase One	Enterprise assessment across every diagnostic track before any technology decision	Four audit tracks executed; 8 of 11 departments covered
Phase Two	Operating model, governance, department transformation, literacy, technology, implementation, measurement	Delivered across two countries in ninety days
Direction Two	Internal AI function established alongside delivery, by placement or by training the client's own people	Recommended; client elected full delivery first, with staged handover
Evidence held	Programme design and deliverable specification	Audit transcripts, survey data, dated delivery log, live systems

The sequencing choice is not a defect in the engagement. A company rating its own operations at three out of ten has no stable platform on which to build an internal AI unit, and no spare capacity to staff one while its operating model is being replaced underneath it. Building the platform first, then transferring ownership onto it, is the slower route and in this case the sounder one.

2. The Company at Zero

Four independent methods were used to describe the same company in May and June 2026: interviews with every department head, an anonymous survey of the workforce, a legal review of the proposed technology, and a technical inventory of where data actually lived. They produced the same picture from four directions — which is the strongest evidence a diagnosis can have.

Annual revenue approximately \$75M ~\$6.4M / month	Employees ~400 ~44 in sales	Warehouses 28 the south-west to eastern Siberia	Average order value ~\$490 Wide range across service types
Shipments per month roughly 13,000 Revenue ÷ average order	Operational efficiency 3 / 10 Founder's own rating	Annual avoidable loss ~roughly \$21M 28% of revenue	Stated goal x3 \$229M revenue

2.1 What the founder said

The founder built the company over eight or nine years, starting as a driver and acquiring businesses until he held a group. He signs nothing personally — everything is delegated. He also has no profit-and-loss statement, no route profitability, and no forecast.

“Gut instinct. That’s all. Nothing else. I’ve been running on that for 8 years. At the level of sensation.”

— The founder, on how strategic decisions are made

“We’re losing money out of thin air when it could be fixed. This is the main reason I can’t stay in a normal working rhythm.”

“The biggest potential is in the sales department, which is not being fully utilised because we can’t let clients track their cargo. That directly affects sales.”

2.2 The five systemic problems

Department-level losses were symptoms. Five root causes generated almost all of them, and every one crossed a departmental boundary — which is why point solutions inside single teams had repeatedly failed at this company before.

#	ROOT CAUSE	ANNUAL SCALE	WHY IT PERSISTED
1	No unified system. Google Sheets, three 1C instances, personal Excel files, messengers and paper, with no single source of truth.	~\$6.9M	1C had been deployed several times and never completed. No platform owned the process end to end.
2	Sales managers were administrators. 50–65% of the working day went to status checks, triple data entry, manual quoting and notifications.	~\$4.6M	No automation of the chain order → registry → 1C → notification → close. Every link was a person retyping.
3	Critical knowledge lived in individual heads. Fleet, sales and warehouse each depended on one irreplaceable person.	~\$2.8M	No formalised processes, regulations or knowledge bases. The fleet manager had not taken a holiday in years.
4	Theft and fraud with no automated control. Fuel reconciled by hand, fake invoices, a branch running “a business within the business”.	~\$2.4M	Two security officers covering twenty-seven regions with no automated reconciliation of any kind.
5	Clients left because of operational chaos. Not price — accumulated friction across a handful of shipments.	~\$2.1M	No tracking, no auto-notification, and errors that compounded until the relationship became too much work for the customer.
The founding maxim, applied “Operations first. AI second.” Automating this company as found would have scaled the chaos — faster triple entry, faster unreconciled fuel, faster untrackable cargo. The first decision of the engagement was not to automate anything yet.			

4. Phase One – Enterprise AI Transformation Assessment

“Before implementing technology, we first understand the business.” In the Enterprise Evolution System the assessment is specified as thirteen parallel diagnostic tracks. At the Group it ran as four, each using a different instrument and each designed to catch a different kind of failure.

TRACK	INSTRUMENT	WHAT IT WAS DESIGNED TO CATCH	HEADLINE FINDING
4.1 Operational	8 department-head interviews plus an external audit	Where time and money leak, and which processes are AI-ready	roughly \$21M / yr, 37 opportunities
4.2 Organisational	Anonymous 82-response survey, six blocks	Whether the people could absorb change at all	Strong core, broken seams
4.3 Governance	Compliance review against four statutes	Whether the intended build was even legal	5 blockers, up to \$833K
4.4 Data & technology	Source inventory, schema design, reverse engineering	Whether the data could support automation	~25 tabs, over 60,000 name variants

The order matters. Had the technical track run alone — as it does in most engagements — the company would have received a well-built system that was illegal to operate, staffed by people nobody had asked whether they could use it.

4.1 The operational audit

Eight of eleven departments were interviewed in depth: the CEO, Sales, Accounting, Legal, Fleet Maintenance, Security, Warehouse Logistics, and Logistics — the last through an external auditor with ten years at a global logistics carrier. Each interview recorded the physical actions people perform, the tool they use, the time it consumes, and the loss it carries.

Department	The finding, in one line	Loss / yr	Opp.	Highest-return quick win	ROI
Sales	Over 40 managers spending 50–65% of the day on administration. No CRM. The same data entered three times.	\$7.5M	9	AI cargo-status agent	\$1.6M
Accounting	Three isolated offices; ~1,000 PDFs a week typed in by hand; a payroll module bought and never switched on.	\$2.7M	7	Payroll module + OCR intake	\$1.0M
Security	Fuel theft reconciled in a notebook against GPS. Fake invoices. A branch head running his own haulage on company work.	\$2.4M	6	AI fuel-theft analytics	\$556K
Logistics	No cut-off times, no record of empty runs, no route profitability. Decisions taken on instinct.	\$2.0M	3	Load dashboard + empty-run alert	\$764K
Warehouse	every warehouse site on manual entry. Measurement disputes settled by filming the pallet. Inventory levels unavailable from 1C.	\$1.7M	5	OCR of acceptance acts	\$583K
Fleet	~100 vehicles maintained from handwritten notebooks organised by brand. One person, no successor.	\$1.1M	5	Digital maintenance log	\$306K
Legal	One lawyer. Contract routing through messenger chats. Litigation deadlines tracked on a hand-drawn grid.	\$764K	6	Counterparty API verification	~\$490K
CEO	No P&L, no forecast, no budget. Daily summary assembled by hand by the office manager.	\$653K	2	Daily AI dashboard	\$167K
TOTAL · 8 of 11 departments	Forecast across all eleven departments: ~roughly \$26M	~roughly \$21M	37	19 deliverable inside 30 days	—

The arithmetic that made it real

Loss was not estimated in the abstract. It was converted into shipments at the audited average order value of ~\$490, which turned an abstraction into something the founder could act on.

LOSS DRIVER	CALCULATION	SHIPMENTS LOST / YR	VALUE
Administration instead of selling, at the audited 50%	Sales headcount × share of day lost × monthly shipments per head × 12	~78,000	~\$38M
The same, on a conservative 30% assumption	Identical method, lower-bound input	~47,000	~\$23M
Slow quoting — the client does not wait	5% of deals lost to response speed	~8,000	~\$4M
Pricing errors leaking as unintended discount	2% of revenue	—	~\$1.5M
The number that moved the decision On conservative assumptions, revenue of approximately \$75M represented roughly half of what the existing sales team could produce with the friction removed. No hiring required. That single sentence is what turned an IT budget conversation into a transformation mandate.			

4.2 The organisational audit

Between 18 and 27 May 2026, before any system was designed, 82 employees answered an anonymous six-block questionnaire built on a spiral-dynamics diagnostic frame: Sales (44), Logistics (26) and Accounting (12). The blocks covered task and role clarity, processes and initiative, KPIs and recognition, culture and meaning, management style and autonomy, and strategy and growth.

This is the track most AI engagements skip entirely, and it is the one that determines whether anything built will actually be used.

What the survey found

DIMENSION	READING	WHAT IT MEANT FOR THE BUILD
Meaning and engagement	High	Over 85% in every department understood how their work affects company results — a rare asset in a distributed team, and the reason change was possible at all.
Trust within teams	Strong · 4.5	Friendly atmosphere and mutual support were the most common answer to “what must not be lost”. Systems had to reinforce this, not replace it with surveillance.
Communication and tooling	Critical	Chat chaos, an unstable 1C and absent cargo tracking were named in all three surveys independently. This validated the registry as priority one.
Cross-department handover	Systemic failure	All three departments named the other two as the source of operational breakdowns. Not personal conflict — an absent standard for passing data between links in the chain.
Recognition and training	Weak	Wins acknowledged only as numbers; training unsystematic. In accounting, recognition scored 2.3 of 5 — the lowest reading anywhere in the survey.
KPI clarity and fairness	Contradictory	Sales understood their KPIs but considered them partly unfair after a recent change. Logistics scored 2.7; roughly 40% saw no clear measures at all.

The three signals that changed the design

— “Managers should stay managers, not a call centre.” Raised independently by four or more people. It corroborated the sales audit

exactly — and it meant the cargo-status agent was not a nice-to-have but the single most requested change in the company.

- Regional information starvation. Several regional staff reported learning about operational changes by accident or after the fact. One noted that the word “department” did not apply to them at all — in the regions, one person is the whole department. This shaped the notification architecture and the decision to give branches their own visibility rather than a head-office dashboard.
- Change fatigue. Multiple respondents described initiatives arriving raw, half-finished, or announced after the fact. One proposed piloting changes with a focus group rather than releasing to everyone at once. That proposal was effectively adopted: every subsequent release went to a test environment first, and the analytics application shipped to a test contour before production.

In their own words

“We need a correct calculator, not one that gives different results for the same parameters at different times.”

— Sales — on tooling that erodes client trust

“In the regions there is information starvation. We find out that something has changed at best in a chance conversation, and more often after the fact, when the absence of information has already become the problem.”

— Logistics

“Value our work.”

— Accounting — the shortest and most precise request in the entire survey

WHY THIS TRACK EARNS ITS PLACE

The survey independently reproduced the operational audit’s three biggest findings — no tracking, unstable 1C, information chaos — from the opposite end of the organisation. When the people doing the work and the people running it describe the same failure without conferring, the diagnosis stops being an opinion. It also surfaced something the department interviews could not: the company had the cultural capacity to absorb change, and would spend it quickly if changes kept arriving half-built.

4.3 The governance and legal audit

The original lead-generation architecture was technically sound and legally unworkable. A compliance review assessed thirteen components of the proposed system against Russian personal-data law, biometrics law, advertising law and the 2025 anti-fraud act. Five were outright blockers.

Blockers 5 Launch impossible as designed	Conditional 1 Depends on configuration	Acceptable 7 No action required	Ceiling exposure \$833K One-time, first inspection
	Component	The problem	Ceiling risk
×	Voice AI platform · US	Decision-makers' voices — legally biometric data — processed on servers in the United States. Primary collection outside Russia plus cross-border transfer of biometrics.	\$278K + shutdown
×	LLM aggregator · UK/US	Call transcripts and lead personal data routed to foreign models through a broker outside the country.	\$208K on breach
×	Scraping service · EU	Raw personal data of Russian citizens landing first on foreign infrastructure before returning to the system.	\$83K + injunction
×	Cold calling itself	Calling with a commercial offer without prior subscriber consent breaches advertising and anti-spam law regardless of the technology stack.	\$7K per episode
×	Call recording	Voice recordings are biometric data. Storage without separate written consent is prohibited.	\$278K
✓	Self-hosted stack	domestic hosting, a self-hosted relational database, a self-hosted data platform, a self-hosted workflow orchestrator, a containerised proxy and domestic SIP telephony — all inside the perimeter and compliant.	—

What was done about it

Three scenarios were priced: full localisation, a hybrid with explicit opt-in, or launch as designed and operate until the first complaint. The third was documented and rejected. Full localisation was chosen, and the consequences ran through the entire build:

- The US voice platform was replaced with an agent platform reached through a Russian telephony bridge, with the SIP layer terminating on infrastructure the client controls.
- The LLM aggregator was replaced, and model routing was later moved again — the design goal being that no single provider could hold the system hostage.
- Foreign scraping was excluded in favour of parsers running on Russian infrastructure, with lead sourcing rebuilt around official registries and a domestic company-data provider.
- Consent, do-not-call handling and caller marking became product requirements rather than legal afterthoughts. When a prospect refuses, the number enters a permanent blacklist enforced on every subsequent call attempt.

THIS IS WHAT GOVERNANCE-FIRST ACTUALLY LOOKS LIKE

The review was produced before the voice system was built — which is why replacing the platform cost a design decision rather than a rebuild and a regulatory file. Most vendors treat governance as a box to tick after the build. Here it functioned as a gate: a component with best-in-class capability was removed from the architecture because it could not lawfully process a Russian citizen's voice.

4.4 The data and technology audit

The fourth track answered a blunt question: is this company's data capable of supporting automation at all? The inventory found a working registry spread across roughly twenty-five monthly tabs in a single spreadsheet, plus contractor forms, exports and documents in no consistent shape.

WHAT WAS ASSESSED	CONDITION FOUND
The operational registry	A Google spreadsheet with a tab per month, no permissions model, no change history, no validation. Duplicate counterparty spellings accumulated freely.
Reference data	Counterparties, cities and drivers existed only as free text inside shipment rows — over 60,000 name mentions that normalised down to over 9,500 distinct counterparties, 200+ cities and 700+ drivers.
Data quality	Dates with typos including impossible years, costs stored as truncated text, weight and volume fused into one column, and formula rows producing 104 phantom records.
The pricing engine	Two separate calculators existed, built by different contractors. Neither was owned by the client team. The working version had to be reverse-engineered from source to establish exactly how a price is formed.
Telephony	The incumbent provider's session border controller would not pass signalling from cloud AI platforms at all — proven by testing two platforms directly before designing a bridge around it.
A parallel ERP programme	The client was simultaneously evaluating a full industry ERP replacement on 1C, with its own express survey running to 29 May. Any AI layer had to integrate with a system that did not exist yet.

The two decisions this track produced

- Structured and unstructured sources were separated, for legal reasons as much as technical ones. Structured tables could be reformatted on Russian infrastructure with no AI involved and therefore no cross-border transfer. Unstructured documents needed a language model to normalise, which was scoped as a one-time process with the automation deleted afterwards, leaving no permanent link to any external service.
- The registry had to precede the ERP, not wait for it. The ERP programme would take quarters. The spreadsheet was losing money every week. The registry was therefore built as the operational system of record while the integration contract with 1C was designed and documented in parallel — which is why, when the ERP work matured, the exchange already existed.

That second decision is the least glamorous and most consequential judgement in the engagement. It is also the one most likely to be got wrong by a firm that starts at implementation: the obvious move is to wait for the ERP, and waiting would have cost the client another two quarters of the roughly \$21M.

5. Phase Two — Building the Capability

Phase Two ran as eight overlapping waves between May and August 2026. They are presented in delivery order rather than by Operating Model step, because the sequence itself is the argument: infrastructure, then the system of record, then governance, then integration, then measurement — and only then AI.

WAVE	WHAT WAS BUILT	PERIOD	OPERATING MODEL STEP	WHY IT HAD TO COME HERE IN THE ORDER
0	Feasibility and infrastructure	May	Technology strategy (09)	Two platforms were tested against the incumbent telephony and failed. The bridge architecture was proven before anything depended on it.
1	The registry — a single source of truth	June	Operating model redesign (02)	Nothing else could be built on a spreadsheet. nearly 20,000 shipments and three reference books had to move first.
2	Governance and the access model	June	Governance (03)	A shared system without field-level permissions would have been less safe than the spreadsheet, not more.
3	The integration layer	June – August	Department transformation (06)	Triple data entry was the single largest loss driver. It ends only when the systems talk to each other.
4	Control and measurement	July	Measurement (11)	Once every shipment was in one place with a price on it, the pricing itself became auditable for the first time.
5	Leadership visibility — the Hub	July	Leadership transformation (04)	The founder could not lead differently until he could see differently.
6	The AI layer	July – August	Implementation (10)	The voice agent needed a registry to read from, a calculator to call, and a legal basis to operate. All three existed by then.
7	Replication to Kazakhstan	July	Continuous evolution (12)	The real test of an operating model is whether it transplants. It did, in weeks rather than months.

Total logged delivery: roughly 430 dated entries. What follows describes each wave and what it changed.

5.1 Wave 0 – Feasibility and infrastructure

Before committing to an architecture, the team established what was physically possible. Two cloud voice platforms were connected directly to the incumbent Russian telephony provider and both failed: the provider's session border controller would not pass signalling from cloud AI platforms. Rather than treat this as a blocker, it defined the design.

- A Russian VPS with a static IP and no carrier-grade NAT was provisioned as a SIP bridge, running an open-source PBX natively rather than in a container — because the RTP port range makes containerisation the wrong tool.
- Registration against the incumbent provider was confirmed and held; dial plans and codec negotiation were configured to avoid transcoding.
- Production container stacks were later stood up on separate servers in Russia and Kazakhstan — database, orchestration, reverse proxy with automatic TLS — with the AI workloads deliberately separated from the operational instances so that experimental work could never destabilise the registry.
- The freight calculator was reverse-engineered from source and documented step by step, then rebuilt as a service with eight tools, validated at 100% parity across over 2,000 historical calculations.

A DETAIL WORTH NOTICING

The team proved the telephony would not work before designing around it, and documented the negative result. Most implementation failures trace to exactly the opposite habit: assuming a component will work because the vendor says so, then discovering otherwise after the architecture depends on it.

5.2 Wave 1 – The registry: one place where the work lives

This is step 02 of the Operating Model, and it is the reason everything after it was possible. The registry is not a database that replaced a spreadsheet; it is the company's process, written down in a form a machine can enforce.

SHIPMENTS MIGRATED

nearly 20,000 From ~25 monthly tabs Counterparties normalised over 9,500 From over 60,000 raw mentions Recognition accuracy 99.6% Counterparties; cities 99% Reference books built 3 Companies, cities, drivers

What the schema encodes

The shipment record is organised into groups that mirror how work actually moves through the company — service fields, finance and manager, participants, pickup, cargo, delivery. That grouping is a process decision disguised as an interface decision: it makes the handover points between sales, logistics, warehouse and accounting visible on a single screen.

- Conditional fields. Actual weight and volume only appear when someone marks the measurement as differing from the booking — which is how the warehouse measurement dispute became a data point instead of an argument.
- Nine status values for cargo and a separate payment lifecycle, with delivery and handover separated, because the audit found those two events were being conflated and hiding failures.
- Transport type as a first-class object, with per-type switches for notifications, client SMS and cross-border synchronisation — so international shipments could follow different rules without forking the system.
- Sequential reference numbering generated by the registry, so a shipment has one identity across every system that touches it.
- Full revision history on every record, which converted the recurring question “who changed this?” from an accusation into a lookup.

The migration itself

Migration was treated as an engineering problem with a quality bar, not a copy-paste. Python normalisation scripts deduplicated reference data with checksum validation on tax identifiers. Date typos were corrected against the source tab month, with the original values backed up. Rows generated by spreadsheet formulas rather than by real business — around a hundred of them — were identified and removed. Where a client name mapped ambiguously to several counterparty records, the name was preserved in a comment rather than guessed at, and flagged for client confirmation.

A MIGRATION PRINCIPLE WORTH STATING

Nothing was silently discarded. Free-text payment notes that did not fit the new status model were normalised and retained verbatim in a comment. Ambiguity was surfaced for a human, not resolved by a machine. This is what makes a migration trustworthy to the people who have to live in the result.

5.3 Wave 2 – Governance: who may see and change what

A shared system without a permissions model is more dangerous than a spreadsheet, because it looks authoritative. Governance was therefore built as part of the registry rather than added afterwards, and it grew in response to how the company actually works.

ROLE	POPULATION	WHAT THE POLICY ACTUALLY ENFORCES
Administrator	~3	Full structural control, user management, restoration of deleted records.
Sales manager	~30+	Sees the whole registry — a deliberate team requirement — but can edit only commercial fields. Logistics, warehouse and accounting fields are read-only. Cannot delete.
Logistician	~16	Owns driver assignment, pickup and delivery status, actual measurements. Attached to cities through a many-to-many relation, so routing follows geography rather than a static list.
Operator	~20+	Terminal-side fields, marking, measurement confirmation.
Accounting	~3	Payment status, invoice and waybill fields; explicitly gated away from operational fields.
View only	On request	Leadership and external auditors — including the external logistics auditor, who needed evidence without write access.
Five combined roles	Regional	Manager-Logistician, Manager-Operator, Logistician-Operator, Universal and Universal Plus. Created because the survey revealed the truth that head office had missed: in the regions one person is the whole department.

That last row is the survey feeding directly into the architecture. A regional employee had written that the questionnaire's use of the word "department" made no sense to them because they work as a single unit. Rather than force regional staff into a head-office role model, the access system was restructured: shared entitlements were extracted into a common policy and five combined roles were composed on top of it, with twenty-two temporary administrator accounts — a security risk created by the workaround — retired in the same pass.

- Eleven roles and eleven policies at final count, with thirty-one permissions on the hired-transport module alone.
- A support-request collection inside the registry itself, so employees raise access and data issues in the system rather than in a chat — with status tracking and automatic notification back to the author. Requests were resolved and logged as a numbered stream.
- Role-based saved views rather than personal spreadsheets, including a "my shipments" bookmark that made the personal colour-coded Excel file obsolete.
- A validated exception: managers may set a payment status to "deferred" and only into an empty field. A narrow, deliberate hole in the wall, because the audit showed the alternative was managers routing around the system entirely.

GOVERNANCE AS AN OPERATING DECISION

None of this is technically difficult. What makes it transformation rather than configuration is that each rule encodes a decision about how the company should work — and several of those decisions were reversed from the original design because the survey and the audits said the original was wrong.

5.4 Wave 3 — The integration layer: ending triple entry

The second-largest loss driver in the audit was the same data being typed into three systems. Closing it required the registry to become the hub of a four-way exchange, each contract designed, documented and handed to a different external development team.

INTEGRATION	DIRECTION	WHAT IT DOES	STATUS	EVIDENCE
1C (ERP)	Two-way	Create, edit and search shipments on one endpoint with three matching keys; full package outbound on creation and a delta on update; anti-echo sensor so the system does not react to its own writes.	Live	dozens of shipments created by the ERP in a single day
Customer portal	Two-way	Order intake from the portal form; edit permitted only while the shipment is still in the opening status, through a whitelist of safe fields; delta updates pushed back.	Live	Full cycle tested end to end
Customer CRM	Inbound + files	Order intake from a customer's own CRM, automatic duplicate document-shipment for border paperwork, label generation, manager assignment and notification.	Live, tested	Accepted in production, labels returned
Russia ↔ Kazakhstan	Two-way mirror	International shipments mirror between the two national registries, matching counterparties and cities on natural keys so nothing duplicates, with loop protection and archiving on delete.	Live	Verified in both directions

What made these integrations hold

- Contracts before code. Each integration began as a written specification handed to the counterpart developer — 1C, the portal team, and the client's own integrators — with field schemas, response codes and status dictionaries. Three external teams were coordinated without any of them having access to the registry internals.
- Idempotency and deduplication everywhere. Duplicate submissions return a conflict rather than creating a second shipment; counterparties resolve by tax identifier before anything is created.
- Defensive validation. Unknown values are rejected with a clear error rather than silently stored, and a reference-data endpoint was published so counterpart systems can fetch valid values instead of guessing them.
- Honest fault attribution. When the portal began sending malformed recipient names — cargo weight arriving in the surname field — the fault was reproduced twice, documented and returned to the portal team rather than patched over. The registry's job is to be right, not to be convenient.

The integration layer is where the operating-model redesign becomes financially visible. A manager who once entered an order into a spreadsheet, then into 1C, then notified a client by hand, now enters it once.

5.5 Wave 4 — Control: making pricing auditable

Once every shipment carried structured parameters and a price, a question became answerable for the first time: were managers discounting responsibly? The audit had estimated 2% of revenue leaking as unintended discount. This wave measured it.

Stage one — the historical audit

An offline pipeline parsed over 5,000 invoices exported from 1C, extracted route, weight, volume and charged amount, and re-priced each one through the frozen calculation engine. The comparison of charged price against list price yields the effective discount.

RESULT	SHARE	NOTE
Matched list price exactly	~15%	Priced to tariff
Discount or surcharge present	~70%	Previously invisible; now measurable
Could not be verified	~15%	No freight line, or priced off-tariff
Total processed	5,000+ invoices	85% verifiable against the tariff engine

Stage two — the live control

The audit was then converted from a one-off exercise into a permanent control. Four read-only fields were added to the shipment record and a sensor fires on every create and update: the system re-prices the shipment through the calculator and writes the reference price, the discount percentage, a verification status and an explanatory comment.

- A dedicated view for leadership filters straight to mismatches, sorted by discount size — so a large discount surfaces the same day rather than at a quarterly review.
- In the first live batch, roughly a third of shipments carried a variance from list price — in both directions, including cases where the customer had been overcharged. The distribution, and its outliers, became visible to management the same day.
- When rigid text parsing produced false positives on managers' free-form tariff notes, the parser was replaced with a language model under a strict schema — with every extracted number verified back against the source text, so a value the model invents marks the record unverifiable rather than fabricating a discount.
- The director introduced a formatting rule for managers off the back of what he saw, and declined a historical backfill — old records correct themselves as they are edited.

WHY THIS IS THE MOST UNDER-RATED SYSTEM IN THE ENGAGEMENT

It has no AI branding and it changed managerial behaviour faster than anything else built. It also demonstrates the anti-hallucination discipline the method requires: the language model is allowed to interpret messy human text, but is not allowed to be the source of a number. Every figure it returns is checked against the original before it can affect anything.

5.6 Wave 5 — Leadership visibility: the Hub

Step 04 of the Operating Model is leadership transformation, and it is usually the vaguest. Here it had a concrete form: the founder who described his decision-making as “at the level of sensation” was given instruments.

A separate web application was launched alongside the registry in late July, deliberately built as a layer rather than a replacement — the governing rule being all data in the registry, all visibility in the Hub. Employees sign in with their registry credentials; blocking someone in the registry closes the Hub to them immediately.

MODULE	WHAT LEADERSHIP CAN NOW SEE
General analytics	Headline figures, city and manager rankings by shipment count and turnover, with month-by-month drill-down on any manager or city.
My analytics	Every employee sees their own shipments, turnover, average order value, twelve-month trend and rank — visibility pushed down, not just up.
Sales plans	Plan versus actual by branch and by manager, with remaining amount, required daily run-rate, month-end forecast and worst-performer sorting.
Tasks	Assignment with deadlines and client linkage, weekly reminder mail, one-click completion.
Pipeline	Client cards, deal stages on a board, and three signal types — overdue steps, deals with no next step, and clients with no contact for five days.
Summary	A single screen assembled from all the above, holding no data of its own.

The detail that shows the method working

When branch-level plans were built, the calculation deliberately exposed what did not fit. In the first month, a substantial share of turnover — comparable in size to the attributable total — could not be assigned to any branch, because employees were not linked to a city and shipments carried no transport type. That gap was displayed as its own block rather than quietly excluded.

The picture is visible rather than hidden.

— the design principle applied to the unattributed turnover block

The same instinct governs the module documentation prepared for the client, which separates metrics that populate themselves from the registry from those requiring manual entry — so that an empty CRM screen reads as “nobody has filled this in” rather than “the system is broken”. Leadership transformation depends on the leader trusting the instrument, and trust is built by showing the holes.

- Rows are driven by who is assigned to a branch, not by which shipments exist — otherwise a failing branch would vanish from its own report, and the table would be empty on the first of the month.
- A copy-forward button for last month's plans, with guards that skip deleted or blocked accounts and unstaffed branches, then reports what it skipped.
- Anomalous single shipments above a threshold are excluded from turnover so a typo cannot distort a branch's month.

5.7 Wave 6 – The AI layer

Eight weeks into the engagement, with a governed registry, a documented pricing engine, a compliant telephony path and a legal basis, the AI layer was finally built. It has two components.

The pricing service

The freight calculator was rebuilt as a machine-callable service with eight tools covering intercity freight, address services, insurance, distance, coverage lookups and zone checks. Two decisions define it:

- It is a fork, not a second engine. The service is a frozen snapshot of the production web calculator, deliberately decoupled from the live data source with validated refresh and snapshot fallback. This prevents the failure the sales survey complained about most: the same parameters returning different prices at different times.
- Drift is detected, not discovered. An automated detector compares the service against the live web calculator on a fixed battery of routes. When it later reported a \$4 gap, the cause was found in hours — the web form silently applies two defaults on Kazakh

routes — and the fix was verified to the cent against the live site.

The voice agent

A Russian-language voice agent was built to handle first contact: introducing itself, establishing shipping needs, answering service questions from a thirteen-document knowledge base, calculating freight cost live during the conversation, and recording outcomes into the CRM.

AGENT TOOLS

5 Including live pricing Knowledge base 13 docs Client materials only Regression tests 23 Run after every change Automation workflows 80+ Across both countries

The engineering discipline that made it trustworthy

- Prices may never be spoken unless the calculator produced them. The agent had occasionally performed a calculation vocally; the prohibition was moved into the final block of the prompt and covered by tests. Verified in live dialogue when the agent quoted figures that matched the engine exactly.
- Gaps are closed with knowledge, not prompt inflation. When the agent invented a document-delivery service that does not exist, the response was to add a knowledge-base document and rewrite the test criterion to something checkable — not to bolt another paragraph onto the prompt.
- The test suite is the product. Nineteen tests grew to twenty-three, covering tool invocation, knowledge-base grounding, gatekeeper handling, hard refusals and edge cases such as a client correcting a city mid-call. Two persistent failures were documented as open defects rather than silently accepted.
- Failure analysis over apology. When the client reported the agent confusing cities, underpricing and pausing too long, all three were measured: the city confusion was speech recognition rather than reasoning and was fixed by loading fifty city names into the recogniser; the underpricing was real and on our side; the pauses were measured across 54 turns with a median response of 2.9 seconds, and the two genuine problems were identified precisely.
- Model independence in practice. When the underlying language model accumulated defects — latency, echoing prompt examples, emitting a tool call as text — it was swapped for another in a single configuration change, with the suite re-run to confirm parity. No rebuild.

THE POINT OF ALL THIS

A voice agent is trivial to demonstrate and very hard to operate. What separates the two is exactly the layer the market routinely skips: a regression suite, a drift detector, a knowledge base with an owner, a prohibition enforced by tests rather than hope, and an architecture where the model is a replaceable part.

5.8 Wave 7 — Replication and continuous evolution

The final test of an operating model is whether it transplants. In July the entire model was replicated to the Kazakhstan branch, which had no comparable system of its own — no ERP, no CRM, nothing but spreadsheets.

- The Kazakh registry was stood up as a mirror of the Russian one — same schema, same roles, same views — then adapted for local specifics: tax identifier formats, company-form suffixes, phone normalisation, Kazakh characters in validation, and border-document fields.

- International shipments were unified into the existing table through a new transport type rather than a parallel structure, keeping one system rather than two.
- Historical Kazakh shipments were migrated with their reference data, and mirrors were created automatically in the Russian registry and linked to the originals.
- The voice agent contour was rebuilt for the local telephony provider and language, with 56 workflows adapted and activated in dependency order.

What makes it continuous rather than finished

PRACTICE	HOW IT WORKS IN PRODUCTION
Separated environments	Test and production contours for both the registry and the Hub. Updates are verified on the test copy first — a direct answer to the survey finding that raw changes were destroying trust in change itself.
One-command promote and rollback	Private repository, staged deployment, and a rollback path that does not touch neighbouring services.
Regression on every change	The agent suite runs after each modification. The pricing service holds 100% parity against its historical corpus and is checked against the live web calculator on every drift run.
Scaling rules written down	Growing lists moved to server-side pagination before volume became a problem, with the rule recorded as a checklist every new screen must pass. Verified by seeding hundreds of test records, then removing them.
Performance treated as a defect	A call-recording player that took 71 seconds and 50MB per open was rebuilt to roughly 2 seconds and 15KB by storing recordings properly rather than proxying them.
Internationalisation	the full interface key set translated so the system can be shown to an English-speaking audience, with an idempotent apply-and-rollback tool and backed-up metadata.

5.9 Throughout – AI literacy, and what the teams did with it

Step 07 of the Enterprise Evolution System ran alongside every other wave rather than after them. Two practical programmes were built specifically for the Group, on a single rule: every session ends in a working artefact, not a certificate.

PROGRAMME	STRUCTURE	WHAT EACH PARTICIPANT LEAVES WITH
AI in accounting	Five modules: the AI already built into 1C (document recognition, assistant); chat AI as a second accountant; end-to-end process automation using the Group's own case; financial literacy and commercial-secrecy risk; an optional advanced module on building an agent for a specific task.	A checklist of documents that no longer need manual entry; a personal library of ten working prompts; a map of their three most repetitive manual tasks with an automation idea for each.
Prompt writing, practical 2	Technique on live examples rather than theory: the prompt formula (task, audience, goal, context, structure); expert-role prompting; question-and-answer patterning; decomposition of complex tasks; chain of thought for calculations; audience framing and explicit boundaries.	Five or six tested prompts built against their own real tasks, plus a personal prompt collection and named sources to keep it current.

The safety rail, taught first

Both programmes open the risk conversation before the capability one. Accounting documents contain commercial secrets — contract terms, prices, supply volumes, client information — and the organisation is legally obliged to protect them. The ethical frame is given at the start: treat the model as the first stage of the work and always verify the result yourself. What may go into a prompt, and what may not, is taught as a practice rather than a policy document.

What the teams produced without being asked

The measure of literacy is not attendance. It is what people build afterwards, unsupervised, on problems nobody assigned them. Within days of the sessions, participants reported the following — in their own words, lightly translated.

WHO	WHAT THEY BUILT OR SOLVED	WHY IT MATTERS
Sales manager	Wrote a lead-generation prompt with real operational constraints — regional metal suppliers, maximum 3–5 tonnes per piece and 6 metres in length, a delivery section on the site — and got back companies with phone numbers that the paid research services had never surfaced, each classified as fitting or not fitting with the reason given.	Opportunity #8 in the matrix, prototyped by a manager
Sales manager	Rebuilt the receivables file the company had failed to reconcile for a long time, allocating every debt to a named manager — and then started building a CRM with a dashboard, receivables, clients and manager performance, having worked out how to rent the server himself.	Opportunities #4 and #16, attempted independently
Regulations lead	Realised she had been briefing the model incorrectly when asking it to revise the staff instructions she was uploading, and restructured her approach.	Knowledge management, self-corrected
Training department	Applied the prompt formula mid-session to a presentation she had been re-editing repeatedly with poor results, and got a usable output immediately.	Technique transferred in real time

In their own words

- “I gave it the task – find suppliers who sell metal products ... the goods must not exceed 3–5 tonnes per piece and 6 metres in length, and the site must have a delivery section. It found a lot of companies that I could not find in the paid services, and with phone numbers. It even explained which ones did not fit my criteria and why.”*

— Sales manager, lead-generation exercise
- “We had a broken file it was impossible to work from. Now the receivables are informative and it is genuinely good to look at – we suffered from this badly, accounting was pushing us and the data was neither current nor clear.”*

— Sales manager, on rebuilding the receivables register
- “While you were explaining, ideas came to me about how to structure the prompt properly. I did it immediately, and out came exactly what I needed – the presentation came out perfectly, right in front of me.”*

— Training department, during the session itself
- “I realised I had been giving the model completely the wrong instruction when asking it to correct the procedure I had uploaded. I am going to practise all weekend.”*

— Regulations lead

WHAT THIS EVIDENCE ACTUALLY SHOWS

Two of these people independently prototyped opportunities that sit in the formal 37-item matrix — AI lead generation and a manager-facing CRM with receivables. They were not told to. They were given a technique and a safety frame, and they went and solved the problem that annoyed them most. That is the difference between training and literacy, and it is the strongest available signal that the transfer described in Section 9 will hold.

6. Build vs. Buy in Practice

The Enterprise Evolution System treats build-versus-buy as one of the first strategic decisions of a transformation: build where internal intellectual property and proprietary operational knowledge are at stake, buy where trusted enterprise platforms already exist, and be explicit about ownership either way. The Group is a clean worked example, including one decision to buy that most AI vendors would have fought.

DECISION	CHOICE	REASONING
Core ERP	BUY	An established industry ERP with 140+ deployments was selected by the client for sales, warehouse and transport logistics. MagnaQore did not compete with it — it designed the integration contract instead. A vendor selling seat licences would have argued the other way.
Operational registry	BUILD	The registry encodes how this company works, including the handover points that were failing. That is proprietary operational knowledge, and it had to be available in weeks rather than quarters.
Pricing engine	FORK	Owned by a third-party contractor and therefore outside the client's control. Forked, frozen, and instrumented with drift detection — the client now owns a stable pricing surface without owning the original code.
Voice platform	BUY	Speech synthesis and recognition are commodities. What matters is the layer around them: prompt, tools, knowledge base, tests. Those were built.
Language models	RENT	Routed through an aggregator specifically so no single provider is load-bearing. Proved out when the primary model was swapped in one configuration change.
Telephony	BUY + BRIDGE	Local carriers in both countries, with a self-hosted PBX bridge the client controls sitting between the carrier and any cloud platform.
Lead sourcing	BUILD	Foreign scraping services were excluded on legal grounds. Parsers run on Russian infrastructure against official registries and a domestic data provider.
Analytics	BUILD	Attempted first as a platform extension; rejected because the stock component could not aggregate across relations. Built as a separate application that reads the registry and stores nothing of its own.

The test a build-versus-buy decision must pass
Every choice above leaves the client able to change its mind. The models are swappable, the voice platform is replaceable behind its own interface, the pricing service runs frozen and independent, and the analytics application holds no data. Nothing in the architecture requires the client to keep paying anyone — including MagnaQore — to keep working.

7. Outcomes

The Enterprise Evolution System specifies eleven deliverables a company should hold at the end of an AI Organization Transformation. Measuring the Group against that list is the fairest available scorecard, because the list is fixed before any engagement begins and does not bend to what a particular project happened to produce.

PROGRAMME DELIVERABLE	WHAT THE GROUP HOLDS TODAY	STATUS	GAP
1. Enterprise-wide AI roadmap	A 37-item opportunity matrix with ROI, complexity, timeline and priority rating, split into 19 quick wins, 13 mid-term and 5 long-horizon initiatives.	Held	Covers 8 of 11 departments
2. AI operating model	The registry schema and role model are the operating model, in force and enforced by software across two countries.	Held	—
3. AI governance framework	Eleven roles and policies with field-level permissions, plus a compliance review against four statutes with a scenario decision on record.	Held	No standing review cadence
4. Internal AI function	Recommended and specified by MagnaQore. The founder elected for us to deliver the full programme rather than build an internal unit alongside it; the function is being transferred in stages as the redesigned model beds in.	By agreement	Handover staged
5. Department transformation plans	Per-department findings, quick wins and opportunity sets for eight departments, several already delivered as systems.	Held	3 departments unaudited
6. AI opportunity matrix	Complete, with each item traced to a quoted operational problem from the audit that produced it.	Held	—
7. AI implementation roadmap	A sequenced top-ten with data-readiness criteria, plus roughly 430 logged delivery entries showing what was actually executed against it.	Held	—
8. Executive KPI dashboard	Live analytics application: branch and manager plans, rankings, forecasting, pipeline signals, per-employee visibility.	Held	Receivables module pending 1C work
9. Change management framework	Test-before-production discipline on every release, in-system support-request channel, client-facing release documentation.	Partial	Practice, not codified framework
10. Employee capability plan	Two practical training programmes delivered; the survey provides a readiness baseline by department.	Partial	No three-tier Champions structure
11. ROI measurement framework	Automated discount control, per-manager and per-branch analytics, agent regression suite, drift detector.	Held	No consolidated ROI ledger
Deliverables held 8 of 11 Fully in the client's hands	Partially held 2 Practice exists, structure does not	Retained by agreement 1 The AI function — handover staged	

7.1 The operational record

Selected figures from the delivery log, all verifiable against dated entries.

WHAT	FIGURE	NOTE
Shipments migrated from spreadsheet to governed registry	nearly 20,000	Counterparty recognition 99.6%, cities 99%
Counterparty records normalised from raw name mentions	over 9,500	From over 60,000 mentions, with checksum validation
Hired-transport and direct-vehicle trips migrated	800+	hired and direct fleet combined; sums reconciled to the cent
Carrier records consolidated	540+	Duplicate spellings merged
Invoices re-priced in the historical discount audit	over 5,000	85% verifiable; ~70% showed a price variance
Roles and policies in force	11 + 11	Including five combined regional roles
Automation workflows in production	80+	56 in the Kazakhstan contour alone
Agent regression tests	23	Run after every change; failures documented, not hidden
Pricing engine parity on historical corpus	100%	every one of more than 2,000 calculations
Shipments created by 1C in a single day	dozens	Measured on a single production day
Call-player load time	71s → 2s	50MB response reduced to 15KB
Call launch latency	43–70s → 2.5s	Moved to queued batch dispatch

A note on what these figures are not: they are delivery evidence, not audited financial results. The roughly \$21M loss figure is an assessment estimate, and the realised recovery against it has not yet been independently measured. Section 9 includes establishing that measurement as a first-order priority.

8. The Commercial Proof

Everything above is internal. The clearest external test of a transformation is whether the company can now win business it could not previously serve. In late June 2026, that test arrived.

8.1 The client the Group could not serve

A consumer furniture manufacturer selling direct to households approached the Group as a carrier. Their requirements were ordinary for a modern logistics buyer and impossible for the Group as it stood in May:

- A price in real time, inside their own CRM, so their sales staff could compare the Group against three competing carriers while the customer waited.
- Order creation by API, not by email or a phone call.
- Printable A6 labels returned automatically — one per package, numbered “piece N of total” — because their warehouse applies them before the driver arrives. The client stated plainly that labels mattered more than the paperwork.
- A separate document consignment for Kazakh routes, so a driver never reaches the border without accompanying paperwork.

In May, the Group had no API, no owned pricing service, no machine-readable registry and no label generation. Three meetings had already passed and the client was still calculating shipments by hand.

8.2 What the transformation made possible

CLIENT REQUIREMENT	WHICH CAPABILITY ANSWERED IT	BUILT DURING
Live price in their CRM	The forked, frozen pricing service exposed as a callable API — the same engine the voice agent uses	Wave 0 and Wave 6
Order creation by API	The registry intake contract, with duplicate protection, counterparty resolution and manager assignment	Wave 3
A6 labels, one per piece	Label generation attached to the shipment record and returned as files	Wave 3
Kazakh document consignment	Automatic duplicate document shipment sharing the parent reference, mirrored to the Kazakh registry	Wave 3 and Wave 7
Status visibility later	The same two-way exchange pattern already running with 1C and the customer portal	Wave 3

8.3 Where it stands

The integration was built and tested in production: intake, the automatic document duplicate, label generation, manager assignment and Kazakh mirroring all confirmed working, with test data cleaned afterwards. Registry API documentation was delivered to the client on 8 July and the calculator API documentation on 13 July. On 14 July the client confirmed receipt and passed both to their own integrators.

STATED PRECISELY

As of the last recorded contact, the customer's integrators were awaiting internal budget approval before starting work, with no outstanding questions raised. The constraint that would have lost this account has been removed and the capability is live and tested — the client's own integration had not yet begun. This case study does not claim a closed contract, because the record does not show one.

8.4 Why this matters more than the integration itself

That manufacturer is one account. What it demonstrates is a change in the class of customer the Group can pursue. The sales audit had recorded the commercial director saying there were large clients the company simply never approached, because it would get one attempt and was waiting until everything was sorted out.

“There are big clients we simply never approach. We get one shot. We’re waiting until everything is sorted out.”

— Commercial Director, May 2026

Three months later the waiting condition is largely satisfied. A prospective client can now be handed API documentation, a live pricing endpoint and a tested intake path in the same week they ask. That is not a feature — it is a different company.

9. What Comes Next — Transferring Ownership

The Group has eleven of the twelve steps of the Enterprise Evolution System delivered or in active progress, and eight of eleven programme deliverables fully in its own hands. The twelfth step — a permanent internal AI function — is the one MagnaQore still performs. That is a decision, not an omission, and this section explains both the decision and how it is being unwound.

HOW THIS WAS DECIDED

MagnaQore recommended standing up an internal AI function in parallel with delivery. The founder chose otherwise, and asked us to deliver the whole programme — on the reasonable ground that a company rating its own operations at three out of ten had no stable platform on which to build such a unit, and no spare capacity to staff one while its operating model was being replaced underneath it. The order was inverted deliberately: build the platform first, then transfer ownership onto it.

The consequence is a real dependency and is worth naming plainly. Every system described in this document is currently operated, extended and debugged by MagnaQore. If the engagement stopped today, the Group would keep working software, complete documentation and trained users — and would not yet have an internal function able to evaluate the next AI capability, govern it, or judge whether a vendor is worth paying.

That transfer is now the primary objective of the engagement's next stage. It is being done gradually rather than by handover date, because the readiness condition is not technical. It is confidence: the teams have to trust the redesigned way of working, and enough people have to be fluent enough with AI tools to own the systems rather than merely use them.

9.1 Why the transfer cannot be a single hire

The predictable shortcut is to hire an AI specialist and hand them the keys. The Enterprise Evolution System is explicit about why that fails: no single person can understand the business, redesign the processes, build the systems, manage delivery, train the teams and produce results. When that person burns out or leaves, the entire capability leaves with them — and at the Group the audit already documented exactly this pattern in three departments, where fleet, sales and warehouse each depended on one irreplaceable individual.

“This knowledge was built up over years, since childhood. How do you transfer it to another person? Only through direct experience.”

— Fleet maintenance, on twenty years of undocumented expertise

Hiring one AI person would recreate that risk in a new department, and this time at the centre of the operating model. The transfer therefore has to be structural.

9.2 What the internal function requires

SEAT	MANDATE AT THE GROUP SPECIFICALLY
AI Transformation Director	Owns the remaining three department audits, the opportunity backlog and the ROI ledger. Runs the audit-before-advice discipline internally, so the next initiative is assessed rather than merely requested. Builds the Champions network.
AI Project Manager	Owns delivery cadence, the charter for each initiative, and the legality check before anything touching personal data goes live — a check this engagement has proved the Group needs and cannot improvise.
AI Architect — Systems	Owns the registry, the integration contracts with 1C, the portal and client CRMs, and the containerised stack. This is the seat carrying the largest knowledge-transfer burden, and therefore the first to be shadowed.
AI Architect — Revenue	Owns the voice agent, the pricing service, the test suite and the drift detector — the customer-facing systems where a hallucination carries commercial risk.

9.3 The Champions layer, and why the Group is unusually ready for it

The function is filled either by placing certified people or by identifying and training the client's own — one senior champion per department, chosen for technical curiosity and peer influence rather than by a department head's preference, then trained through a three-tier programme where every tier ends in a working automation rather than a certificate.

The Group has an unusual advantage here, and it comes from the survey. The organisational audit found engagement and meaning scoring above 85% in every department, and trust within teams at 4.5 of 5. It also found something more specific: several of the highest-value ideas in the whole engagement came from employees, unprompted.

- The lawyer designed the contract workflow himself, and separately proposed the counterparty verification API — two of the highest-return legal opportunities in the matrix.
- The warehouse head endorsed OCR intake on the spot with a specific operational reason, and had already invented his own workaround for missing inventory data using a group chat.
- A sales manager proposed piloting changes with a focus group instead of releasing to everyone — a change-management practice the engagement then adopted.
- The fleet manager had wanted parts barcoding for years and had a clear account of why previous attempts failed.

The AI literacy sessions then produced the harder kind of evidence: not good ideas in an interview, but work delivered afterwards without instruction. Two participants independently prototyped opportunities that sit in the formal matrix — one built a working lead-generation prompt with real freight constraints and recovered prospects the paid research tools had missed; another rebuilt the long-broken receivables register by manager and started assembling a manager-facing CRM, having taught himself how to rent the server. Section 5.9 sets out what each of them produced.

Those two people are the Champions for Sales. They have already demonstrated the three things the role requires — technical curiosity, credibility with peers, and the instinct to fix the process that annoys them rather than complain about it. Naming them is a formality that costs nothing and starts the transfer immediately.

THE CHAMPIONS ARE ALREADY THERE

These are not passive users. They diagnosed their own processes correctly in a single interview, and then — given a technique and a safety frame — went and built things nobody had asked them for. The Champions network at the Group does not need to be recruited; it needs to be named, trained and given authority. That is the cheapest remaining step and the one with the longest payback, and it is where the transfer begins.

9.4 The transfer sequence

Each stage below moves a defined piece of ownership across, and each is gated on readiness rather than on a date.

STAGE	WHAT MOVES ACROSS, AND WHY NOW	HORIZON
1 Name the Champions	One per department, drawn from the people who already proposed solutions during the audit. Costs nothing, starts the transfer immediately, and gives each department an internal first point of contact instead of us.	Now
2 Establish the ROI ledger	The roughly \$21M is an assessment estimate. Until realised recovery is measured against it inside the company, the transformation cannot defend its own budget without us in the room.	0–30 days
3 Deepen AI literacy past the pilot	Training so far has reached accounting and prompt-writing. The three-tier Champions programme extends it to logistics, warehouse and sales — and literacy is the stated readiness condition for everything after this line.	30–90 days
4 Shadow the Systems architect seat	The largest concentration of transferable knowledge sits here. An internal owner works alongside ours on the registry and integrations before taking them.	30–90 days
5 Complete the assessment	Customer Service, the 1C function and HR remain unaudited, carrying an estimated further \$4.9M — and running that audit internally is the best possible training exercise for the incoming Director.	30–90 days
6 Stand up the full function	Four seats, with the Director owning the backlog and the legality gate institutionalised rather than improvised. MagnaQore moves from delivery to advisory.	90+ days

9.5 The conclusion this case study supports

The Group proves the sequence works. A company rating itself three out of ten, running a approximately \$75M business on a spreadsheet, was rebuilt in ninety days into an operating model with a governed system of record, live integrations across four external systems, automated pricing control, management visibility and a compliant AI layer — in two countries.

It also shows what the second direction of the system is for. Systems can be delivered to a company in a quarter. The capability to keep evolving them has to be grown into it, and that runs at the speed of confidence rather than the speed of software. The Group chose to build the platform before the people who will own it — and is now, deliberately and in stages, closing that distance.

THE MEASURE THAT MATTERS

Our success is measured not by how much work remains for us, but by how capable the client becomes without us. The Group is a complete Direction One and a Direction Two now under way by design. This document is, deliberately, as much a statement of the transfer ahead as a record of the work behind it.

Appendix – Delivery Record

The engagement log holds roughly 430 dated entries. Summarised here at month level to show the sequence rather than the volume, and to keep operational specifics out of a public document. Every claim in this case study traces to entries of this kind.

PERIOD	FOCUS	WHAT WAS DELIVERED
Month 1	Assessment and feasibility	Anonymous employee survey run across three departments. Department-head interviews completed. Two cloud voice platforms tested directly against the incumbent national carrier and both rejected — the carrier would not pass cloud signalling — so a self-hosted bridge architecture was adopted and documented. Compliance review completed against four statutes. The client's parallel ERP evaluation concluded, fixing the integration scope.
Month 2	System of record and governance	Production stacks deployed in both countries. Pricing engine reverse-engineered, rebuilt as a service and validated at full parity across more than two thousand historical calculations. The registry went live: shipment collection, reference books, automatic sequential numbering. Nearly twenty thousand shipments migrated with reference data deduplicated. Role-based access established, then extended with combined roles for regional staff after the survey showed head-office role models did not fit the regions. Event notifications and an SLA monitor activated. An in-system support-request channel replaced ad-hoc chat requests.
Month 3	Integration and control	Written integration specifications issued to three external development teams. Two-way ERP exchange completed and validated across all methods. Customer-portal intake and status return delivered. Customer-CRM intake tested in production, with automatic document-consignment duplication, label generation and manager assignment. Cross-border mirroring of international shipments activated with loop protection. The voice agent was created, grounded in a client-materials knowledge base, and connected to the pricing service under a strict rule that prices may only come from the calculator. Historical pricing audit completed across five thousand-plus invoices.
Month 4	Measurement, visibility and replication	The pricing audit converted from a one-off exercise into a live control firing on every shipment change, with a management view filtered to variances. Hired-transport and direct-fleet registers migrated and reconciled. The analytics application launched across test and production contours, followed by branch-level sales plans with month-end forecasting, task management and a pipeline with exception signals. Interface fully localised to English. The agent test suite hardened and extended; a drift detector added against the live pricing engine. The complete operating model replicated to the second country, including a locally adapted telephony contour.
Ongoing	Continuous evolution	Regression suite run after every change, with unresolved failures documented as open defects rather than accepted. Drift detection against the live pricing engine. Separated test and production environments with staged promotion and one-command rollback. Performance treated as a defect class — the call-recording player was rebuilt from roughly seventy seconds and fifty megabytes per open to about two seconds and fifteen kilobytes. Scaling rules recorded as a checklist every new screen must pass.

About this case study

CONFIDENTIALITY NOTICE

This case study is based on a real client engagement. Certain company names, figures and operational details have been anonymised or aggregated to respect confidentiality obligations. It is shared for informational purposes only, to illustrate the type of AI Organization Transformation work MagnaQore delivers. It does not constitute a public offer, a guarantee of results, or a full representation of any client's internal systems, strategy or financial performance. For deeper, NDA-level discussion, please contact us directly.

Monetary values are stated in US dollars and rounded. Percentages, ratios and sequencing are reported as measured. Client and employee identities, exact commercial terms, pricing depth, individual performance data and system-level operational detail have been withheld or generalised. Loss and return figures are conservative assessment estimates prepared during the engagement and are not audited financial results.

MagnaQore · AI Organization Transformation · Enterprise Evolution System · Case study compiled 2026.

About this case study. This case study describes a real MagnaQore engagement. Certain company names, figures, technical details and operational information may be simplified, anonymised or aggregated to respect confidentiality obligations. It does not constitute a guarantee of results, a public offer, or a complete representation of a client's internal systems, strategy or financial performance. Additional detail may be shared for qualified due-diligence purposes under a mutual non-disclosure agreement.