nalar.ai

BUSINESS PROPOSAL · CONFIDENTIAL

FleetCare AI

Agentic AI fleet maintenance & inventory platform — built for coal shipping, tug & barge operations.

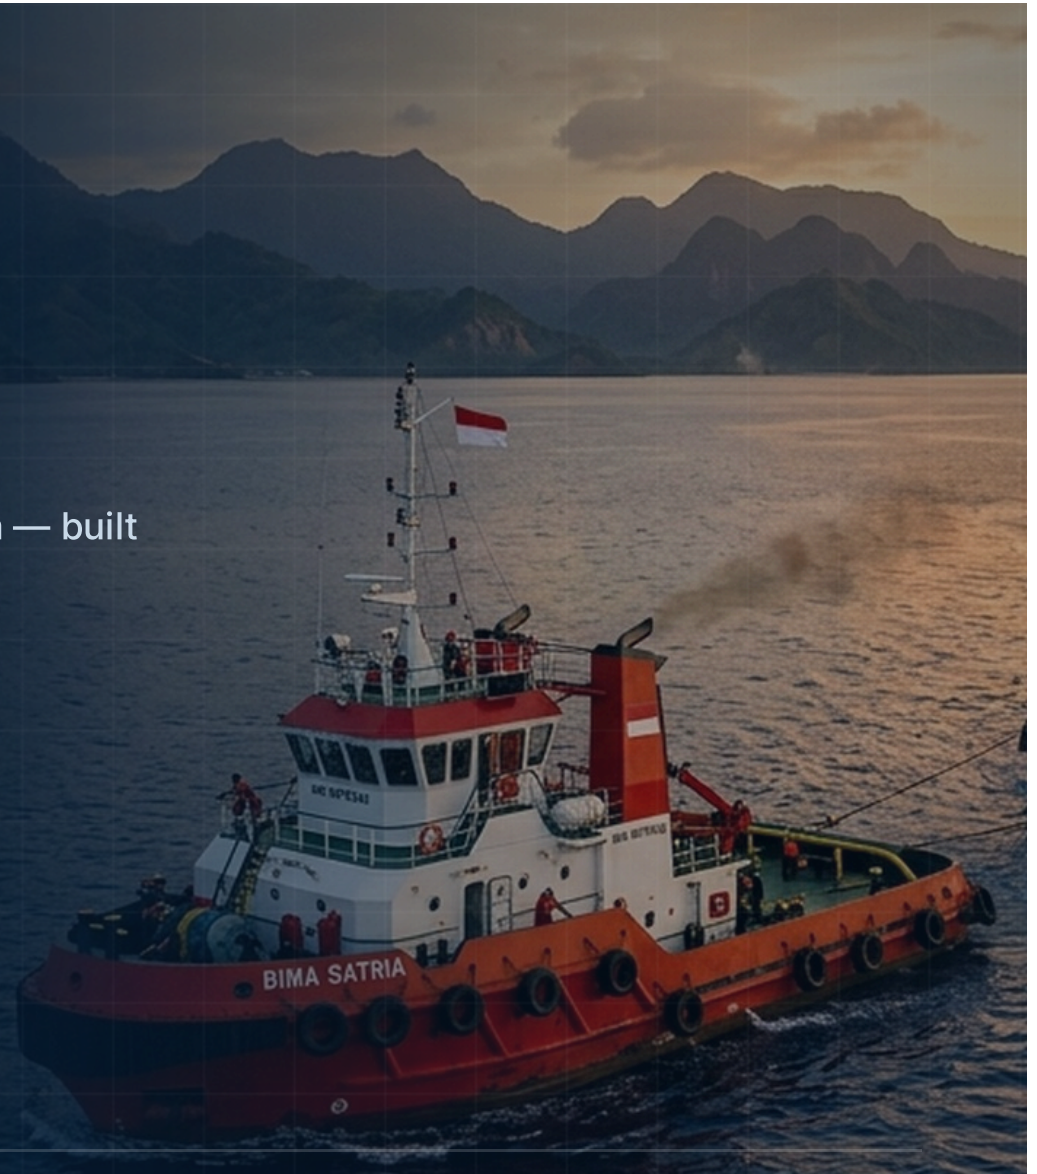

- AI does the heavy lifting — plans, reads, triages, predicts
- Designed for 27 vessels · tugs & barges
- Four phases — Maintenance Plan → Overhaul
- ISM & class compliant — less typing, more visibility

Prepared by Nalar.ai
Agentic AI Solutions for Indonesian Industry

10 AUGUST 2026 · DRAFT v0.8

Contents

01 Executive Summary

FleetCare AI is a **lightweight, AI-first fleet maintenance platform** for Indonesian coal shipping — tug & barge operations. A lean web app plus an **agentic AI (Artificial Intelligence) layer** that plans, reads, triages, and predicts — so crews photograph and move on, and AI does the admin.

25–40%	20–38%	2–4 wks	27
unplanned downtime reduction (industry case studies) [1]	maintenance cost savings [2]	advance failure warning [3]	vessels — tugs & barges

The problem

- Incumbent maritime ERPs (Enterprise Resource Planning systems — DNV, MariApps smartPAL, SERTICA) are **heavy, expensive, data-entry-centric** — USD (US dollars) 100k–500k+ implementations (Nalar.ai market estimate) [4].
- Fixed running-hour maintenance **misses real wear** from variable loads, fuel quality, and weather routing.
- Job libraries are typed by hand; readings and oil reports are logged on paper or spreadsheets; stock lives on clipboards.

The solution — four phases, one platform

- **P1 Maintenance Plan** — AI generates the job library and schedules around voyages.
- **P2 Maintenance Monitoring** — photo OCR (Optical Character Recognition) for hours & oil reports, anomaly alerts, defect triage.
- **P3 Inventory** — photo stock-taking, smart reorder, job-to-spares matching.
- **P4 Overhaul Plan & Monitoring** — AI-drafted scope, progress & budget tracking.

POSITIONING

Not “another shipping ERP”. A lean, AI-first platform that **complements the ISM Code (International Safety Management) & class compliance** — at a fraction of incumbent cost, with AI included.

NOTES & SOURCES — EXECUTIVE SUMMARY

Abbreviations — all introduced at first use in this document: AI (Artificial Intelligence) · ERP (Enterprise Resource Planning) · PMS (Planned Maintenance System) · ISM (International Safety Management Code) · OCR (Optical Character Recognition) · LLM (Large Language Model) · UAT (User Acceptance Testing) · SLA (Service Level Agreement) · ROI (Return on Investment) · BYO (Bring Your Own AI keys) · SaaS (Software as a Service) · MVP (Minimum Viable Product) · PPN (Indonesian value-added tax, 11%) · UU PDP (Indonesia's Personal Data Protection Law) · OEM (Original Equipment Manufacturer) · API (Application Programming Interface) · PWA (Progressive Web App) · CI/CD (Continuous Integration / Delivery) · AWS (Amazon Web Services) · DR (Disaster Recovery) · HA (High Availability) · IDR (Indonesian Rupiah) · USD (US Dollar).

1. **Downtime reduction (25–40%):** McKinsey & Company (2020), "Prediction at scale" — mckinsey.com (30–50%); PRIME Marine (2026), Orion Ship Management case — primemarine.com (40%).
2. **Maintenance cost (20–38%):** McKinsey (2020) — 18–25%; ATS (2025) — up to 40% vs reactive — advancedtech.com; PRIME Marine (2026) — 38% corrective reduction.
3. **Advance warning (2–4 wks):** iFactory (2026), maritime predictive maintenance — ifactoryapp.com; OXMaint (2025) — 2–6 weeks — oxmaint.com.
4. **Incumbent implementation USD 100k–500k+:** **Nalar.ai market estimate** — vendors (DNV, MariApps, SERTICA) do not publish pricing; see Sources in Section 05.

02 Scope at a Glance

Four modules, delivered as one platform over four phases — each building on the last. All scopes cover the full fleet of **27 vessels**.

Maintenance Plan — P1

AI builds the job library from OEM (Original Equipment Manufacturer) manuals and schedules work around voyages.

`AUTO JOB CARDS` `SMART SCHEDULING` `ISM/CLASS CHECK`

Maintenance Monitoring — P2

Photo OCR for running hours & oil reports; AI flags degradation and triages defects.

`PHOTO OCR` `OIL INTELLIGENCE` `DEFECT TRIAGE`

Inventory — P3

Photo stock-taking with AI label reading; shortages flagged weeks before a job is due.

`PHOTO STOCK-TAKING` `SMART REORDER` `JOB→SPARES`

Overhaul Plan & Monitoring — P4

AI-drafted overhaul scope from class surveys; progress & budget tracked automatically.

`SCOPE GENERATION` `PROGRESS EXTRACTION` `BUDGET VS ACTUAL`

Phasing

P1 — Maintenance Plan
Foundation & compliance driver — 8–10 wks

→

P2 — Monitoring
Highest operational pain — 10–12 wks

→

P3 — Inventory
Depends on P1 job library — 8–10 wks

→

P4 — Overhaul
After core platform is proven — later

Phases can be contracted separately or as one program. A low-entry pilot (2–3 vessels) is available to prove value before scaling — see Commercial.

03 Solution Overview

Web-based and mobile-friendly — the crew's phone is the capture tool. Photos of gauges, labels, and reports feed an AI layer that **extracts, interprets, and acts**, with every action logged for class and ISM readiness.

Platform traits

- **Web + mobile-friendly** — crews use phones they already carry; superintendents get a shore-side dashboard.
- **Photos as first-class data** — OCR and vision extract readings, labels, and quantities automatically.
- **Offline-tolerant capture** — works on board with poor connectivity; syncs when back in range.
- **Audit trail built-in** — every reading, photo, and change logged for class/ISM readiness.
- **AI is agentic, not dashboards** — generates, classifies, predicts, drafts; humans review and approve.

What the AI does

Automated plan generation
OEM manuals, PDFs, and Excel job lists become a structured job library — tasks, intervals, spares, tools, safety notes — with human review.

Effortless capture
A photo of the hour meter, an oil report, or a spare-parts shelf is digitized automatically — no typing, no re-keying, no faxes.

Data → decisions
Trends monitored, anomalies flagged, predictive alerts 2–4 weeks before failure — with recommended jobs auto-generated.

Intelligent triage & matching
Defect photos + text classified by system and severity; every scheduled job matched against spares availability before it's due.

What this means for your crew: less paperwork, fewer spreadsheets, fewer surprises. The crew captures with a photo — the AI handles the record-keeping — and the superintendent sees the fleet picture in real time.

PROTOTYPE — FLEET VIEW

Fleet dashboard — fleet KPIs, plan compliance, due/overdue, AI alert board, 8-week workload.

03 App Deep Dive — P1 · Maintenance Plan

The compliance backbone: a complete, audit-ready planned maintenance system (PMS) — **without the manual typing**. Foundation for everything else.

TODAY

- Engineers build job libraries from OEM manuals by hand
- Job cards typed manually; intervals calculated by spreadsheet
- Schedules break when voyages change

WITH FLEETCARE AI

- AI reads manuals / PDFs / Excel → structured job library for review
- Smart weekly work lists per vessel — hours, due dates, voyages, crew, spares
- Voyage change → schedule re-proposed automatically

OUTCOME

- Audit-ready ISM/class coverage with gap flags
- Job-library build time roughly halved
- Overdue & due-soon visibility fleet-wide

PROTOTYPE — PROOF IN PRACTICE

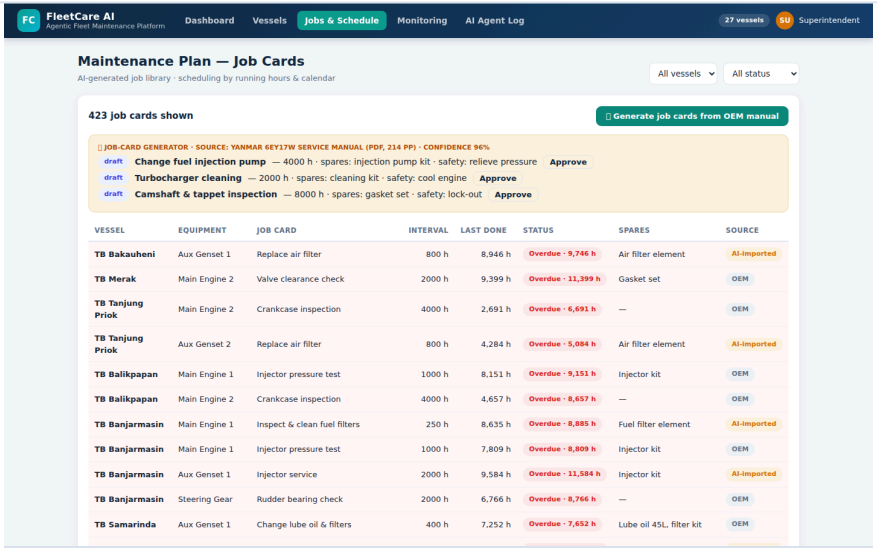

AI job-card generation — agent drafts job cards from OEM manuals; superintendent reviews & approves.

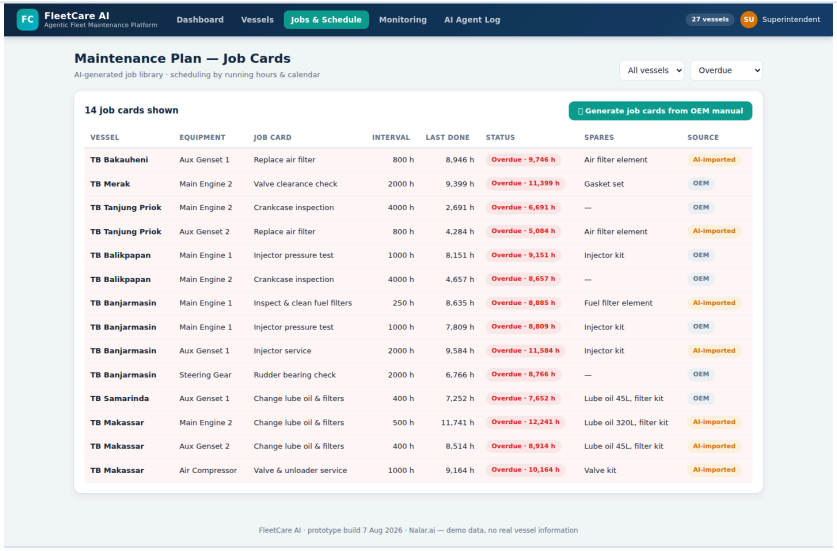

Smart scheduling — weekly work lists per vessel, with overdue/due-soon flags and voyage-aware replanning.

03 App Deep Dive — P2 · Maintenance Monitoring

The highest operational pain — and the **most visible AI value**: readings and reports captured in seconds, degradation caught before damage.

TODAY

- Running hours read from gauges, logged by hand
- Oil lab reports typed, filed, never trended
- Defects reported by phone/fax, triaged manually

WITH FLEETCARE AI

- Photo of hour meter → OCR auto-fills reading
- Oil reports digitized — TBN (Total Base Number), viscosity, water, debris trended per engine
- Defect photo + text → AI classifies, suggests fix, auto-creates work order

OUTCOME

- Predictive alerts 2–4 weeks before failure
- Degradation flagged with recommended action & lead time
- Anomaly detection on readings (backwards/missing entries)

PROTOTYPE — PROOF IN PRACTICE

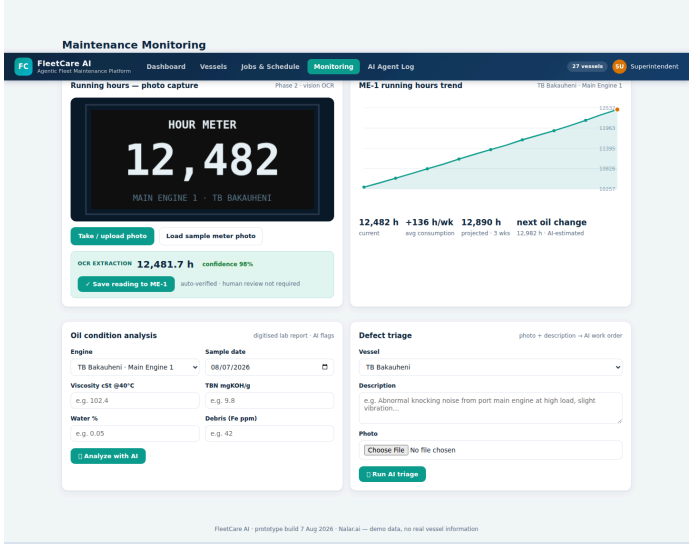

Effortless capture — photo of the hour meter OCR'd with confidence scoring; readings auto-fill trends.

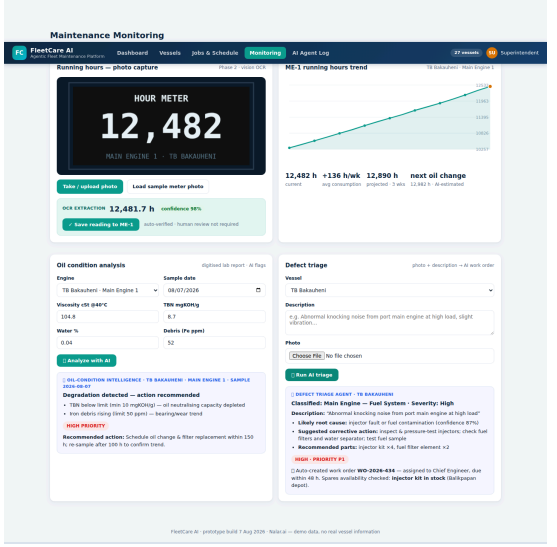

AI monitoring flows — oil-condition analysis with AI flags, running-hour trends, defect triage → auto work order.

03 App Deep Dive — P3 · Inventory

Close the gap between **stock on the shelf** and **jobs on the plan** — shortages solved weeks before a job is due.

TODAY

- Physical stock counts on clipboards
- Photos scattered in phone galleries
- Shortages discovered only when a job is due

WITH FLEETCARE AI

- Photo of shelf/bin → AI reads labels, counts, updates stock with photo evidence
- Consumption forecast vs min/max & lead time → reorder suggestions
- Every scheduled job checked against stock, weeks ahead

OUTCOME

- Shortages flagged before they stop a job
- Receiving verified by photo vs PO (Purchase Order)
- Full stock audit trail

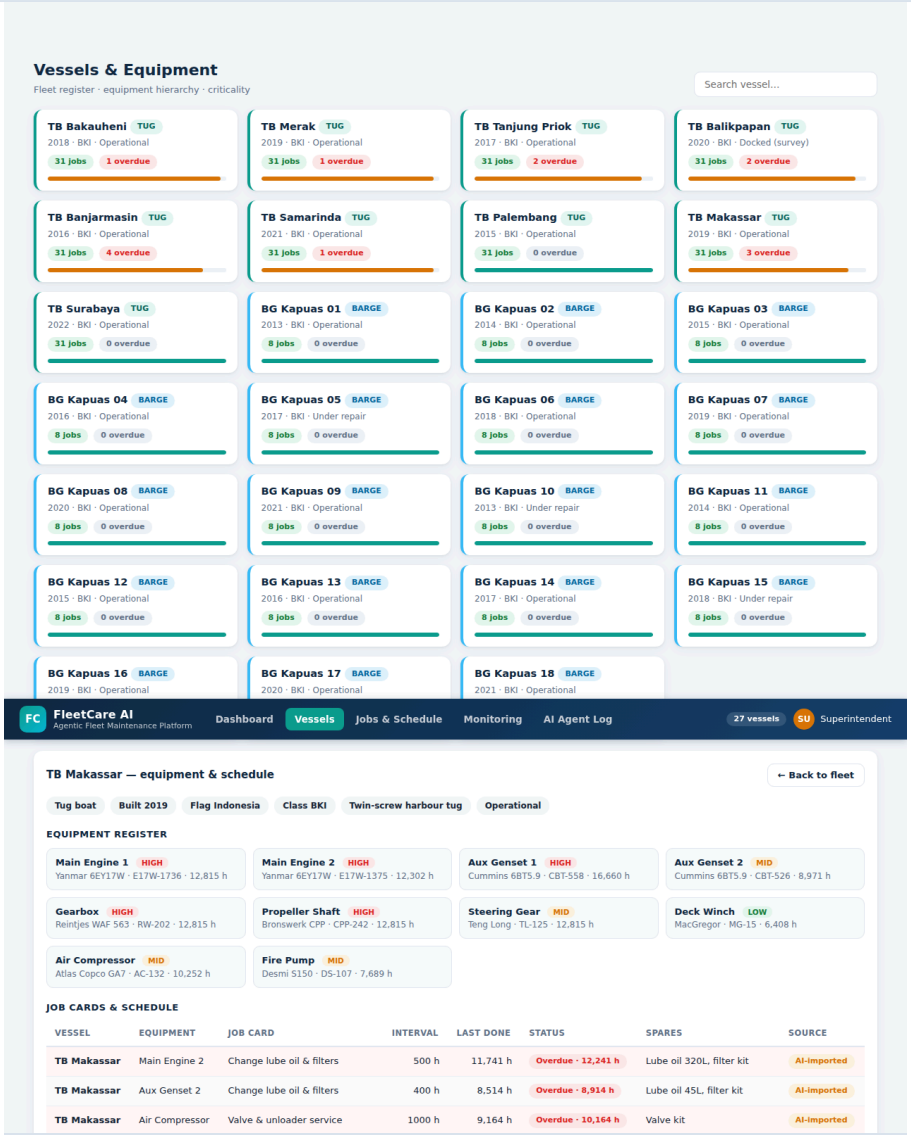

Vessel context — each vessel's assets, readings & stock in one place; shortages visible against the plan.

Mobile-first capture — the same effortless photo flow powers stock-taking on the crew's phone; AI reads labels & counts with photo evidence.

03 App Deep Dive — P4 · Overhaul Plan & Monitoring

Dry-dock and overhaul is expensive and easy to overrun — **AI keeps scope, progress, and budget under control**. Delivered after the core platform is proven.

TODAY

- Overhaul scope built from class surveys & experience
- Progress tracked manually, cost overruns common

WITH FLEETCARE AI

- AI drafts scope from class recommendations, surveys, history, OEM intervals
- Daily reports/photos → progress % per work package, critical-path flags
- Budget vs actual reconciled; overruns flagged early

OUTCOME

- Scope ready faster, anchored to evidence
- Slippage visible before it costs
- Fleet-wide lessons-learned knowledge base

PROTOTYPE — PROOF IN PRACTICE

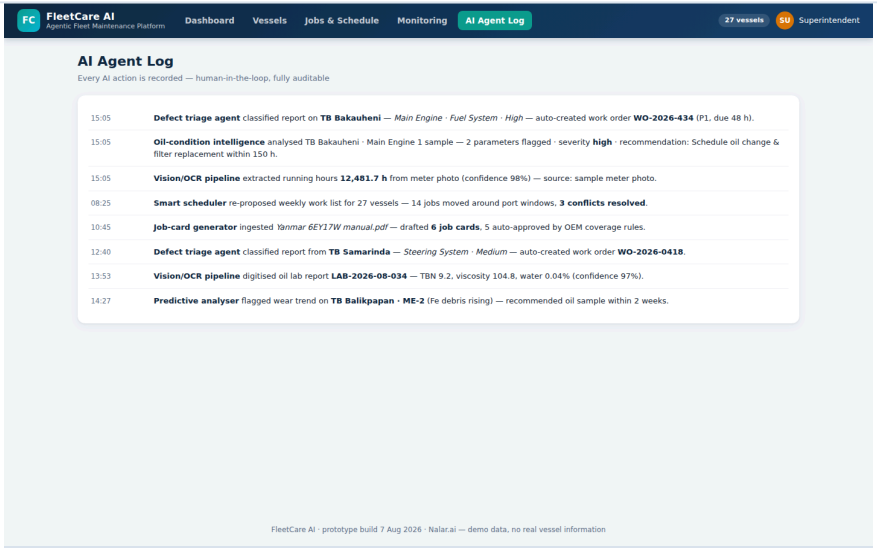

Transparent agent log — every AI action (drafting scope, extracting progress, flagging overruns) is recorded, human-in-the-loop, audit-ready.

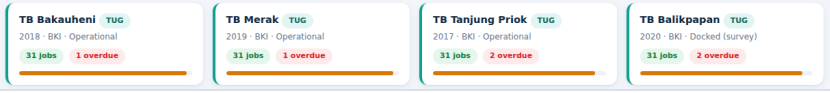

Fleet overview — all vessels & phases tracked from one place; overhaul progress visible per vessel.

04 Solution Architecture

A **lean, AI-first web platform** — not another heavyweight ERP. Four principles: **agentic AI does the heavy lifting** · **mobile-first, offline-tolerant** · **phased and pragmatic** · **lean by default** (cost is a first-class constraint — target well under IDR 10–15 juta/month all-in at 27 vessels).

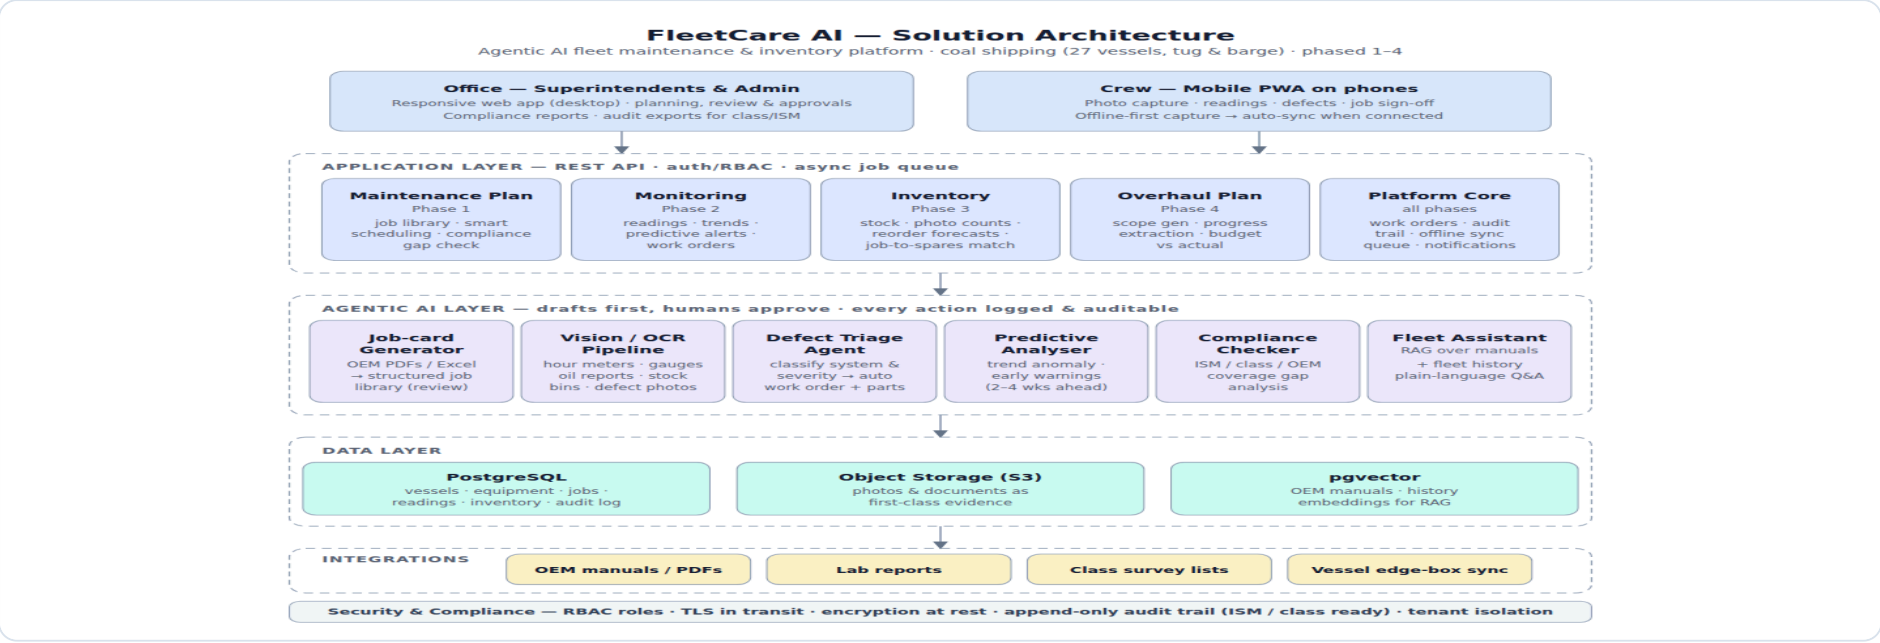

Five-layer architecture: client · application · agentic AI · data · integrations.

Key decisions

Decision	Choice	Why
Platform shape	Responsive PWA (Progressive Web App), not native apps	One codebase for office + crew phones; installable with offline support
AI integration	Agentic layer via API (Application Programming Interface) + async job queue	OCR/document jobs run in background; app stays responsive on vessel-grade devices
Human-in-the-loop	AI drafts → superintendent approves	Trust + ISM/class record-keeping (who did what, when, why)

Decision	Choice	Why
Multimodal extraction	Vision-LLM (Large Language Model) for structured extraction, classical OCR as fallback	Hour meters, oil reports, bin labels are messy real-world documents
Offline	Local queue + auto-sync with conflict rules	Vessels have poor/satellite-only connectivity; capture never blocked
Data model	Relational core + typed JSON for AI payloads	Flexible for messy OCR, strict enough for audit & reporting
Deployment	Containerised (Docker), CI/CD (Continuous Integration / Continuous Delivery)	Repeatable, cheap, easy to scale

04 Architecture — AI Models, Routing & Hosting

Model routing — cheapest adequate model per task

Text / agents (scheduling, triage, forecasting, RAG)	DeepSeek V4 Flash	DEFAULT
Clean captures — meters, gauge digits, stock labels	Classical OCR (free, on-box)	DEFAULT
Messy docs — oil reports, defect photos	Gemini Flash-Lite class	DEFAULT
Hard docs / low confidence	Gemini 2.5 → 3.5 Flash	ESCALATE
Offline edge box (Phase 2+)	Qwen2.5-VL 7B local	FALLBACK

LEAN AI spend @27 vessels ≈ IDR 25–60k/month (text agents ≈ 15–30k · classical OCR ≈ 0–5k · vision-LLM ≈ 10–25k · one-time onboarding ≈ 5–10k). Even a 5× usage spike stays under IDR 300k/month.

Hosting — Lean by default

Tier	Spec	Monthly cost (IDR, ≈)	Providers	When to use
Lean RECOMMENDED	2 vCPU · 4 GB — Postgres+pgvector on-box, no GPU (Graphics Processing Unit)/K8s (Kubernetes)/Redis	75–190k	Hetzner CX22 (Singapore) · IDCloudHost · DO	**Default for the full 27-vessel fleet**
Standard (upgrade)	4 vCPU · 8 GB + object storage	220–400k	Hetzner CX32/CX42 · IDCloudHost · DO	Headroom when usage proves it
Production / HA (High Availability)	8 vCPU · 16 GB + managed DB (Database) + CDN (Content Delivery Network)/WAF (Web Application Firewall)	500k–1.5M	AWS (Amazon Web Services) Jakarta · Hetzner + managed PG · IDCloudHost · Hetzner dedicated AX41 (cost-optimised)	Phases 3–4, SLA (Service Level Agreement), strict residency, DR (Disaster Recovery)

MONTHLY PLATFORM COST — LEAN / STANDARD / PRODUCTION (27 VESSELS)

Tier	What's included	Singapore cost-tier	Indonesian region
Lean RECOMMENDED	1 box · Postgres+pgvector on-box · object storage · classical-OCR-first AI	≈ 150–350k	≈ 350–750k
Standard	4 vCPU/8 GB · separate backups · heavier batch & office load	≈ 275–710k	≈ 560k–1.66M
Production / HA (High Availability)	8 vCPU/16 GB + managed DB (Database) · CDN (Content Delivery Network)/WAF (Web Application Firewall) · DR (Disaster Recovery)	≈ 500k–1.5M	≈ 1.0–2.5M

PLATFORM OPEX — THE HEADLINE NUMBER

Lean tier, Indonesian region (proposal default): ≈ IDR 350–750k / month all-in (Singapore Lean: ≈ 150–350k/mo). **Phase 1 alone on Lean runs ≈ IDR 200–400k/month** — the cheapest credible ISM-compliant maintenance platform in the market, with AI included. Benchmark: incumbent implementations run USD 100k–500k+ (≈ IDR 1.6–8 miliar) before licensing.

05 Commercial — Fees at a Glance

One platform. Four phases.
Two-part pricing: one-time build + monthly subscription to run.

Recommended entry:
Pilot → Standard (Option B)
Year-one all-in ≈ **IDR 620M**

HOW PRICING WORKS — ONE-TIME + SUBSCRIPTION

DEVELOPMENT & IMPLEMENTATION — ONE-TIME

Pays for **building it**: AI job library, OCR & integrations, UAT, go-live, 90-day warranty — per phase.

SUBSCRIPTION — MONTHLY (SAAS)

Pays for **running it**: hosting (Indonesian region), support, AI usage (Option A), updates — per active vessel, tiered by scope. Fixed = recommended; hybrid optional.

A — Fully Managed B — BYO AI keys **(recommended)** C — Client-Hosted (lowest recurring)

Implementation fees identical across options; subscription differs.

Fee component	Pilot 2–3 vessels	Phase 1 Maintenance Plan	Phase 1+2 RECOMMENDED Plan & Monitor	+ Phase 3 Enterprise	Full Program P1–4
Scope	Proof of value on real vessels — AI job library, running-hour OCR, oil-report digitization, defect triage 4–6 wks	AI job-library generation, asset hierarchy (27 vessels), smart scheduling, ISM/class gap check 8–10 wks	Phase 1 + photo OCR of hours, oil intelligence + trend alerts, defect triage → work orders, predictive early-warning 8–12 wks/phase	+ Photo stock-taking, smart reorder, job-to-spares matching, receiving verification +8–10 wks	All four phases as a **flat bundled package**; Phase 4 (Overhaul) delivered later within the program in phases
Development & Implementation one-time	**IDR 50–110M** 100% credited vs P1+2 (90 days)	**IDR 140–200M**	**IDR 340–500M**	**IDR 480–690M**	**IDR 750M flat** bundled package — all four phases, one agreement
Training	Included (basic)	Included — 2 sessions + e-learning Premium add-on IDR 32M	Included — 2 sessions/phase Premium P1+P2 IDR 55M	Included Premium 55M P1+P2 + 32M P3	Included Premium full program IDR 100M
Subscription per vessel / month	Platform usage included during pilot	**A 700k · B 500k · C 250k**	**A 1.0M · B 750k · C 350k**	**A 1.4M · B 1.0M · C 450k**	Per tier of active scope (see Phase 1+2 / +Phase 3)
Infrastructure	Included — no separate charge	Included in subscription (A/B) · client-provided (C) · Indonesian region · Hetzner dedicated (AX41 €42–49/mo ≈ IDR 0.75–0.9M/mo) if residency waived	Included in subscription (A/B) · client-provided (C) · Indonesian region · Hetzner dedicated (AX41 €42–49/mo ≈ IDR 0.75–0.9M/mo) if residency waived	Included in subscription (A/B) · client-provided (C) · Indonesian region · Hetzner dedicated (AX41 €42–49/mo ≈ IDR 0.75–0.9M/mo) if residency waived	Included in subscription (A/B) · client-provided (C) · Indonesian region · Hetzner dedicated (AX41 €42–49/mo ≈ IDR 0.75–0.9M/mo) if residency waived
Fleet of 27 per month	—	**A 18.9M · B 13.5M · C 6.75M**	**A 27.0M · B 20.25M · C 9.45M**	**A 37.8M · B 27.0M · C 12.15M**	Per tier of active scope

COMMERCIAL MECHANICS — APPLIES TO ALL OPTIONS

Options A/B/C = delivery & ownership: A Fully Managed · B BYO AI keys (recommended) · C Client-Hosted (lowest recurring). Implementation fees identical across options; subscription differs. · **Package discount:** 10% off implementation when 2+ phases are contracted in a single agreement at list (Standard, Enterprise, custom combos); **Full Program = flat bundle at IDR 750M** — the package price, no additional discount stacks · **Annual prepay:** 12 months for the price of 10 · **Hybrid option (Standard scope):** base IDR 450k/vessel/mo + IDR 750 per AI-processed item + IDR 2,500 per AI job card · **Fixed plans include:** up to 8,000 AI items/mo fleet-wide (overage IDR 750/item) · **Payment:** 30/40/30 per phase (40% at UAT incl. live agentic-AI acceptance); pilot 50/50 · **Prices exclude VAT/PPN (11%)** · valid 90 days · FX illustration ≈ IDR 16,000/USD · **Included in any subscription:** Indonesian-region hosting (AWS Jakarta); Hetzner dedicated variant (AX41 €42.30–49/mo ≈ IDR 740–860k/mo) if residency waived; 99.5% uptime SLA (A/B), backups/security, AI model updates, support, 100 users, data export; extra vessels at 85% of per-vessel rate. Full detail in 5.1–5.9; sources in 5.10.

05 Commercial Proposal

Lean, phased, AI-first pricing — **three delivery options**, **fixed or hybrid subscription**, and a **low-entry pilot** to prove value before scaling. All prices in **IDR** for 27 vessels.

VALUE ANCHOR

Incumbent ERPs cost **USD 100k–500k+ in implementation alone** (≈ IDR 1.6–8+ billion) (Nalar.ai market estimate — vendor pricing not published) [3]. Our **entire 4-phase program (IDR 750M flat) is less than half the incumbent's USD 100k starting price** and roughly **one-sixth of its typical cost** (≈ IDR 4.8B) — AI included, no per-module license fees.

Why the low price is structurally possible (not discounting): template-based delivery — tug & barge equipment is highly standardized, so job libraries are built once and parameterized across the fleet; a lean continuous delivery team ships in phases; and the AI layer runs on cost-efficient hosted models with automatic routing (no GPU hardware). The low entry point, pilot path, and flexible ownership options are deliberate design.

At a glance

Item	Value
Pilot (2–3 vessels) — **credited against scale-up**	IDR 50-110M
Phase 1 + 2 — Plan & Monitor (recommended first delivery)	IDR 340-500M
Full program (Phases 1–4) — flat bundled price	IDR 750M
Recurring (per vessel / month)	IDR 250k - 1.4M
Fleet recurring (27 vessels / month)	IDR 6.75M - 37.8M
Payment milestones (per phase) — “until agentic AI is built”	30% / 40% / 30%

Item	Value
Implementation duration	`8–12 wks/phase · pilot 4–6 wks`
Price validity	`90 days`

vs. incumbent maritime ERPs

Dimension	Incumbent ERP (DNV, smartPAL, SERTICA-class)	FleetCare AI
Implementation cost	USD 100k–500k+ (≈ IDR 1.6–8B+)	**IDR 750M flat all four phases**; IDR 340M–500M for P1+2; pilot from IDR 50M
Timeline	12–24 months typical	8–12 wks/phase; P1+2 in 4–6 months; pilot 4–6 wks
Licensing	Per-module + per-vessel annual fees	Single subscription; AI included, no module fees
Data entry	Manual, form-centric	AI-assisted: photo OCR, auto job cards, voice/photos for defects
Crew workload	High — structured forms	Low — “photograph and move on”
Deployment	Often on-prem, heavy integration	Fully managed cloud, BYO (Bring Your Own) keys, or client-hosted; offline-tolerant

Benchmarks indicative, based on vendor product documentation (DNV ShipManager, MariApps smartPAL, SERTICA) and Nalar.ai market estimates (vendors do not publish pricing) [3]; conversion ≈ IDR 16,000/USD. Full sources in 5.10.

5.1 DELIVERY & OWNERSHIP OPTIONS

Dimension	Option A — Fully Managed	Option B — BYO AI Keys RECOMMENDED	Option C — Client-Hosted
Platform hosting	Nalar.ai cloud — Indonesian region (AWS Jakarta) · cost-optimised Hetzner dedicated (AX41 €42.30–49/mo ≈ IDR 740–860k/mo) if residency waived	Nalar.ai cloud — Indonesian region (AWS Jakarta) · cost-optimised Hetzner dedicated (AX41 €42.30–49/mo ≈ IDR 740–860k/mo) if residency waived	Client's own VPS (Virtual Private Server)/cloud
AI/LLM API keys	Nalar.ai's keys (AI included)	Client's keys (OpenAI/Gemini/DeepSeek; billed by provider)	Client's keys
Client effort	Minimal	Medium (manage AI accounts & quotas)	High (operate infrastructure)
Data control	Nalar.ai cloud	Nalar.ai cloud	Client infrastructure (strongest residency)
Recurring (Standard, 27 vessels/mo)	`IDR 27.0M`	`IDR 20.25M`	`IDR 9.45M`

Recommended: Option B (BYO AI keys) — best value/control balance; the client keeps AI-cost transparency and can use cost-efficient models like DeepSeek. Option A for zero AI-vendor management; Option C if data residency / infrastructure control is a hard requirement.

05 Commercial — Options, Tiers & Pricing

5.2 PRICING TIERS (SCOPE)

Tier	Modules in scope	Implementation (one-time)
Basic	Phase 1 — Maintenance Plan	IDR 140–200M
Standard RECOMMENDED	Phase 1 + 2 — Plan & Monitor	IDR 340–500M
Enterprise	Phase 1 + 2 + 3 — Plan, Monitor & Inventory	IDR 480–690M
Full Program	Phases 1–4 (adds Overhaul, delivered later) — **flat bundled price**	IDR 750M

Package discount: 10% off implementation fees when two or more phases are contracted under a single agreement at list (Standard, Enterprise, custom combos). **Full Program = flat bundle at IDR 750M** — the package price, no additional discount stacks.

RECURRING SUBSCRIPTION BY OPTION & TIER (PER VESSEL / MONTH)

Tier	Option A — Fully Managed	Option B — BYO Keys	Option C — Client-Hosted
Basic	IDR 700k	IDR 500k	IDR 250k
Standard	IDR 1.0M	IDR 750k	IDR 350k
Enterprise	IDR 1.4M	IDR 1.0M	IDR 450k

FLEET TOTALS (27 VESSELS, MONTHLY / ANNUAL)

Tier	Option A	Option B	Option C
Basic	18.9M / 226.8M	13.5M / 162M	6.75M / 81M
Standard	27M / 324M	20.25M / 243M	9.45M / 113.4M
Enterprise	37.8M / 453.6M	27M / 324M	12.15M / 145.8M

YEAR-ONE ALL-IN (IMPLEMENTATION MIDPOINT + ANNUAL SUBSCRIPTION, 10% DISCOUNT)

- **Pilot path (recommended entry):** IDR 50–110M pilot (credited) → Standard (Option B) ≈ **IDR 620M year one**
- **Standard (Option B) direct:** ≈ IDR 620M · **Standard (Option A):** ≈ IDR 700M · **Full Program (Option B):** ≈ IDR 920M year-one (750M list − 10% = 675M + first-year sub 243M)[†]

[†] Review flag: if the 10% package discount does not apply to the Full bundle (750M as final package price), Full B year-one ≈ IDR 993M (750M + first-year sub) — for confirmation at review.

5.3 IMPLEMENTATION FEES BY PHASE — "UNTIL THE AGENTIC AI IS BUILT"

Completion milestone = the AI layer live and accepted (job-card generation, OCR, prediction quality verified) — not merely UI deployment.

Phase	Scope summary	Duration	Fee range	Midpoint
P1 — Maintenance Plan	AI job-library from OEM manuals/Excel, asset hierarchy (27 vessels), smart scheduling, ISM/class gap check	8–10 wks	IDR 140–200M	IDR 170M
P2 — Maintenance Monitoring	Photo OCR of hours, oil-analysis digitization with trend alerts, defect triage, predictive early-warning	10–12 wks	IDR 200–300M	IDR 250M
P3 — Inventory	Photo stock-taking, smart reorder forecasting, job-to-spares matching, receiving verification	8–10 wks	IDR 140–190M	IDR 165M
P4 — Overhaul (later)	AI-drafted overhaul scope, progress extraction, budget vs actual	8–10 wks	IDR 140–190M	IDR 165M
Full program	All four phases as a **flat bundled package** (P1–P4, one agreement) — the package price	in phases	IDR 750M	flat

PILOT — PROOF OF VALUE (2–3 VESSELS)

Scope: subset of P1 + 2 — asset hierarchy, AI job library for pilot vessel types, running-hour OCR, oil-report digitization, defect triage. · **Duration:** 4–6 weeks · **Fee: IDR 50–110M** (platform usage included).
Scale-up credit: 100% of the pilot fee credited against full-fleet P1+2 implementation if the client proceeds within 90 days.

Each build phase includes:

- Project management & dedicated delivery team; AI job-library build-out (27 vessels)
- Data migration support (up to 5 person-days/phase); UAT (User Acceptance Testing) **including AI acceptance criteria**
- Training: 2 sessions per phase (virtual or Jakarta) + e-learning guides
- Go-live support + 30-day stabilization · documentation & handover
- **90-day defect warranty** after each phase go-live

05 Commercial — SaaS (Software as a Service), ROI & Terms

5.4 RECURRING SUBSCRIPTION — FIXED VS HYBRID

Fixed (flat monthly): platform + support + AI included (Option A) within a fair-use allowance of up to **8,000 AI-processed items/month fleet-wide**; overage IDR 750/item. Options B/C: flat platform fee; AI billed directly by the client's provider (B) or self-hosted (C).

Hybrid (lower base + usage): cost tied to actual use.

Hybrid component	Price
Base platform fee (per active vessel / month) — platform, hosting, support, 100 users	IDR 450k (Standard scope)
AI per processed item (photo OCR, oil report, defect triage, stock count)	IDR 750 / item
AI job-card generation (build-time)	IDR 2,500 / card
Illustrative fleet month (27 vessels, ~10,000 items)	**≈ IDR 19.65M**

Hybrid example sits below fixed Option B Standard (IDR 20.25M/month).

Annual prepay (all options): 12 months for the price of 10. Hybrid: prepay base annually, usage monthly.

Included in any subscription:

- **Cloud hosting (Indonesian-region infrastructure — AWS Jakarta ap-southeast-3, UU PDP (Undang-Undang Pelindungan Data Pribadi — Indonesia's Personal Data Protection Law)-compliant; multi-tenant)** — cost-optimised variant (only if the client waives in-country residency; EU/US DCs — no Asia): Hetzner dedicated **AX41 €42.30–49/mo ≈ IDR 740–860k/mo — recommended quote line** (Ryzen 5 3600 · 64 GB · 2×512 GB NVMe · 1 Gbit unmetered — ~4× production RAM headroom vs min-spec 8 vCPU/16 GB/160 GB+, no GPU); AX42 €57.30/mo ≈ IDR 1.0M/mo (ECC); AX102 €122.30/mo; server auction ~€30–39/mo ≈ IDR 525–680k. ID→EU latency ~150–200 ms — fine for batch/back-office pipelines (vessels offline-first); CDN mitigates office UX.
- SLA (Service Level Agreement) 99.5% monthly uptime (A/B) · daily backups, security patches (A/B)
- **AI model updates** — continuous OCR, job-card, and prediction improvements
- Support: business-hours helpdesk (ID + EN), next-business-day response
- Up to **100 users** included; +IDR 25k/user/mo · fleet growth beyond 27 at **85% of per-vessel rate**

5.5 TRAINING & ENABLEMENT

Package	Content	Price
Basic (included)	2 sessions per phase (virtual/Jakarta), e-learning, quick-start guides	`Included`
Premium Enablement	5 on-vessel visits, train-the-trainer, superintendent workshop, 30-day hypercare, SOP (Standard Operating Procedure) playbooks	`IDR 32M/phase · 55M P1+P2 · 100M full`
A-la-carte	Additional on-vessel visits	`IDR 6M / visit`

5.6 PAYMENT MILESTONES

Build phases: 30% signing · 40% UAT acceptance **including live agentic AI** · 30% AI go-live + 30-day stabilization. **Pilot:** 50/50. **Subscription:** monthly in advance or annual prepay (10 months' price). IDR; payable within 14 days.

5.7 ASSUMPTIONS & EXCLUSIONS (summary)

1. **Fleet:** 27 vessels, tug & barge (assumed ≈ 10–12 tugs / 15–17 barges; mix affects job-library volume, not price).
2. **Users:** up to 100 named users included.
3. **Hosting (A/B):** Indonesian-region cloud — AWS Jakarta (ap-southeast-3), UU PDP-compliant, ~5–15 ms (IDCloudHost Jakarta alt); Singapore only if client waives in-country residency.
4. **AI usage:** fixed plans include 8,000 AI items/month; overage IDR 750/item. Hybrid billed per item. B/C AI costs borne by client.
5. **Data:** OEM manuals, job lists, oil reports, stock data, surveys provided by client (machine-readable or scanned).
6. **Connectivity:** vessels use existing onboard connectivity; platform is offline-tolerant; Nalar.ai supplies no hardware.
7. **BYO keys (B/C):** client manages provider accounts, billing, quotas; provider outages managed by client (we support on request).
8. **Training:** 2 sessions/phase virtual or Jakarta; travel outside Jakarta at cost.
9. **Currency:** IDR (Indonesian Rupiah); USD benchmarks ≈ IDR 16,000/USD for illustration.
10. **Taxes:** exclude VAT/PPN (Pajak Pertambahan Nilai — Indonesian value-added tax, 11%) and withholding, payable by client where applicable.
11. **Scope:** covers documented scope; customization quoted as change requests.
12. **Warranty:** 90-day defect warranty per phase; support liability capped at phase fee.

5.8 ILLUSTRATIVE ROI (RETURN ON INVESTMENT) & PAYBACK (indicative)

AI-assisted fleet maintenance typically delivers **25–40% unplanned-downtime reduction** and **20–38% maintenance-cost reduction**, with **2–4 weeks' failure warning** [1]. Illustrative annual maintenance & downtime spend of **IDR 8–12 billion** (Nalar.ai estimate, validated at kickoff):

Annual baseline	Savings @ 15%	Savings @ 20%	Savings @ 25%
IDR 8B	IDR 1.2B	IDR 1.6B	IDR 2.0B
IDR 10B	IDR 1.5B	IDR 2.0B	IDR 2.5B
IDR 12B	IDR 1.8B	IDR 2.4B	IDR 3.0B

Payback (Standard, Option B, year-one ≈ IDR 620M): approximately **3–8 months**; net-positive **IDR 0.9–2.8B/year** from year two (savings minus Option B subscription ≈ IDR 243M/yr).

RISK REDUCTION

A single unplanned main-engine overhaul on a tug typically costs **IDR 1–3 billion** (repairs, off-hire, replacement charter) (Nalar.ai estimate — marine-diesel overhaul benchmarks) [2]. Even one avoided major failure per year is a significant return across 27 vessels.

5.9 PRICE VALIDITY

Prices valid for **90 days** from the proposal date. Nalar.ai may adjust pricing if fleet size, scope, or FX (foreign exchange)/cloud costs change materially.

SOURCES & BASIS — COMMERCIAL SECTION

1. **Downtime / cost / warning ranges (25–40%, 20–38%, 2–4 wks):** synthesis of peer-reviewed and industry research — Kalafatelis et al. (2025), "Towards Predictive Maintenance in the Maritime Industry", *JMSE* 13(3):425 — doi.org/10.3390/jmse13030425; "AI-Driven Predictive Maintenance in Modern Maritime Transport" (2024), *Applied Sciences* 14(20):9439 — doi.org/10.3390/app14209439; McKinsey & Company (2021), "A smarter way to digitize maintenance and reliability" — mckinsey.com (15–30% cost reduction; ~20% downtime reduction).
2. **Main-engine overhaul cost (IDR 1–3B):** Nalar.ai estimate built from published marine-diesel overhaul benchmarks — Neptunus Power (2025), "Marine Diesel Engine Overhaul Cost" — neptunus-power.com; Con-tractor Equipment (2026), marine diesel rebuilds — con-tractor.com — plus off-hire & replacement-charter economics.
3. **Incumbent ERP implementation USD 100k–500k+; 12–24-month rollouts; per-vessel license USD 5k–20k/yr:** Nalar.ai market estimates (indicative, to be validated in a competitive quote) — vendors do not publish pricing. Vendor product documentation for module scope: DNV ShipManager (7,000+ vessels / 300 customers) — dnv.com; MariApps smartPAL (600 vessels at Bernhard Schulte) — mariapps.com; SERTICA (SVITZER case, 400+ vessels) — sertica.com.
4. **Hybrid illustrative fleet month (≈ IDR 19.65M):** Nalar.ai model — base IDR 450k/vessel × 27 + usage at IDR 750/item for ~10,000 items; to be validated with client data at kickoff.

06 About Nalar.ai

Nalar.ai builds agentic AI solutions for Indonesian industry. Our name comes from the Indonesian word **nalar** — "reason" — and our conviction that AI should **reason with your data**, not just display it.

What we do

- **Agentic AI platforms** — AI that generates, classifies, predicts, and drafts, with humans in the loop.
- **Practical AI for operations** — OCR, computer vision, and LLM agents applied to real industrial workflows.
- **Prototype-first delivery** — working demos on real data before full build-out.
- **Indonesian market focus** — built for local realities: connectivity constraints, remote sites, lean teams.

Why Nalar.ai for FleetCare AI

Deep AI engineering

Agentic AI is our core competency — we design the AI layer to do real work: generate job cards, read photos, triage defects, forecast failures.

Operational realism

We design for vessels with poor connectivity and crews with limited time. Capture must be effortless — a photo, not a form.

Audit-ready by design

Class and ISM readiness are built in: photo evidence, full audit trails, compliance checks as standard features.

Transparent, phased partnership

You see working demos at every phase gate. Contract what delivers value — expand when you're ready.

OUR PROMISE

We measure success by what your crew no longer has to type, what your superintendents can now see in real time, and what your maintenance budget does over the next 12 months.

07 Next Steps

1. **Review & alignment discussion**
   We walk through the proposal, demo the Phase 1 prototype, and align on priorities and success metrics.

2. **Confirm Phase 1 scope & pilot vessels**
   We agree the pilot vessel set, data sources (OEM manuals, job lists), and the first working demo target.

3. **Agreement & kickoff**
   Commercial terms confirmed in IDR, phase plan locked, discovery begins.

4. **Phase 1 delivery begins**
   Working demo within weeks — not months.

Nalar.ai — Agentic AI Solutions for Indonesian Industry

A Appendix — Working Prototype

A **working, self-contained prototype** demonstrates the platform end-to-end — built with mock data for 27 vessels. Capture in seconds; AI does the admin. **Screenshots appear inline in each app's deep-dive** (Sections 03) so proof sits where you read about that app.

> **LIVE DEMO**
>
> A live, clickable demo can be arranged for your team — walk through job-card generation, photo OCR, defect triage, and the fleet dashboard on real vessel scenarios.

All prototype data is mock/illustrative. The OCR shown is simulated for the demo — production uses vision-LLM extraction as described in Section 04. Screenshot mapping: P1 job cards & scheduling · P2 photo OCR & AI flows · P3 mobile capture & vessel context · P4 agent log & fleet overview · fleet dashboard in Section 03.