GROUP CHIEF AI & TECHNOLOGY TRANSFORMATION OFFICER · KEPPEL LTD.

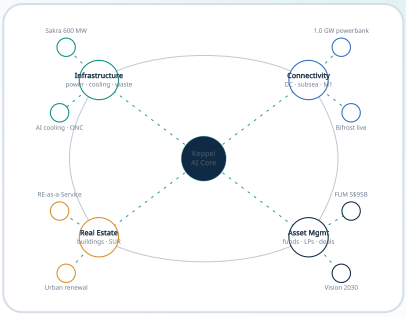

From isolated pilots to embedded intelligence.

A three-horizon plan to make artificial intelligence an embedded capability across Keppel's asset management and operating platform, written as I would write it in my first week in the role: where we stand, what I believe, what we build, where the value is, and how we will be measured.

By **LI Jiang** Working draft · Day 1 perspective

September 2026

3 horizons, 100 days to 36 months

6 lighthouse initiatives in year one

4+1 business divisions plus the Group office

1 shared asset ontology under everything

This is a personal working draft prepared for the role, not a Keppel document. Keppel figures are taken from its published FY 2025 results and public announcements; targets marked "indicative" are to be calibrated in the first 100 days.

01 / STARTING POINT

Where Keppel stands on Day 1

Keppel is no longer a conglomerate; it is a global asset manager and operator with three operating divisions and a fast-growing funds platform. The AI mandate has to be read against that shape: the value pools are in assets, decisions and the infrastructure Keppel sells to the AI economy itself.

FY 2025 NET PROFIT

S$1.1 bn

Up 39% year on year; recurring income S$941 million, up 21%. ROE 18.7%.

FUNDS UNDER MANAGEMENT

S$95 bn

Target S$100 billion by end-2026 and S$200 billion by 2030; asset management net profit S$189 million.

DATA CENTRE POWERBANK

1.0 GW+

Asia Pacific, including 720 MW secured for an AI data centre in Melbourne; KDC SGP 9 AI-ready campus breaks ground mid-2026.

ENERGY TRANSITION

600 MW

Sakra Cogen, Singapore's first hydrogen-compatible CCGT, in operation; renewables 4.7 GW toward 7 GW by 2030.

CONNECTIVITY

Bifrost

Trans-Pacific cable in commercial operation since December 2025; about S$200 million O&M fees per fibre pair over 25 years.

AI ALREADY IN THE ASSETS

20 %

Energy savings from the AI-powered cooling solution launched in 2025; the Operations Nerve Centre at Changi runs ML-driven monitoring of cooling and energy assets.

INFRASTRUCTURE DIVISION

S$703 M

Record recurring income; decarbonisation solutions EBITDA S$130 million, up 32%; S$7.1 billion of long-term contracts.

EFFICIENCY PROGRAMME

S$120 M

Annual cost-savings target by end-2026 (S$98 million achieved 2023–2025); M1 on a 90-day AI-supported efficiency plan.

READ 01

AI is present, but as point solutions

Cooling optimisation, the ONC, M1's digital platform and pockets of analytics exist. What is missing is a shared data and semantic foundation, a common engineering practice and a path that moves a good pilot into every relevant asset. The job is to turn islands into a system.

READ 02

The value is in assets and decisions

Keppel's economics are recurring income from operating assets and fee income from managing capital. AI that improves dispatch, cooling, maintenance, capacity and deal decisions moves those lines directly. Generic productivity tools matter, but they are not where the S$ is.

READ 03

Keppel sells the AI economy its infrastructure

Over 1 GW of data centre powerbank, hydrogen-ready generation, subsea capacity and AI-driven cooling mean AI is also a product line. Every capability we build for our own assets should be designed to be packaged for tenants, partners and LPs.

02 / PRINCIPLES

Five operating beliefs

These come from building an industry large model and a 100-agent platform at a listed group, from founding a hardware company, and from a year of shipping agentic systems for manufacturers. They decide what we fund and what we stop.

01

A business capability, not an IT project

Every initiative has a business owner with a P&L, a value hypothesis in S$ and a named decision it improves. The AI office builds platforms and squads; the divisions own outcomes.

Glodon: AI revenue booked through business units, disclosed separately

02

Ontology before agents

A Keppel Asset Ontology is the executable contract between data teams and decision makers: what a plant, a chiller, a rack, a lease, a fund and a commitment mean, with constraints the machine checks before any recommendation.

String2AI method: 7+1 semantic specification, provenance binding

03

Deterministic core, generative shell

Optimisers, physics and control produce the plan; language models understand, explain and orchestrate. Anything that changes a commitment, a dispatch or a bid stops in front of a person. Autonomy is graded L0–L3 and earned.

Agentic APS and ultrasound-robot work: safety outside the model

04

Production is the metric

No proof of concept without a path to production, a data owner and an evaluation set. "No eval, no deploy." Value is tracked in the division's P&L, not in a slide.

Glodon: RMB 17M to 100M+ direct AI revenue in one year

05

Governance as an accelerator

Responsible-AI foundations are built once, as reusable controls, evaluation pipelines and model-risk tiers, so that each new use case inherits compliance instead of renegotiating it.

Glodon: national LLM registration and MLPS Level 3 opened public-sector clients

03 / TARGET ARCHITECTURE

The Keppel AI Stack

Five layers, one ontology, and governance cutting across all of them. Each division plugs its assets into the same platform; each product, internal or external, is assembled from the same engines and agents.

Products & channels

L5 · WHERE VALUE IS BOOKED

Copilots for Keppel teams; Result-as-a-Service for assets in Keppel-managed funds; external offerings packaged with the operating businesses.

- `Investment Copilot`
- `Asset Intelligence RaaS`
- `AI Cooling-as-a-Service`
- `AI-ready DC services`
- `M1 enterprise AI`

THE AI + DATA DOUBLE LOOP

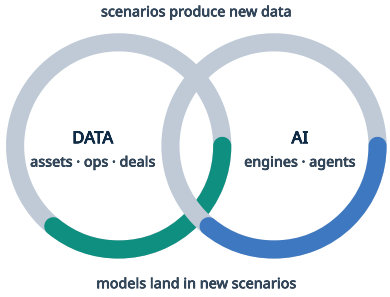

Group platform at the base, division models in the middle, scenario applications at the edge. Data flows back into the ontology and models; models land in the next asset. Systematic AI is the moat; point solutions are not.

Domain agents

L4 · PLAN · TOOLS · EVIDENCE

Agents that decompose tasks, call engines and tools, cite evidence and stop at human checkpoints. Autonomy graded L0 (advise) to L3 (act within guardrails).

Energy Operations

DC Operations | Asset Health

Investment Analyst

Sustainability & ESG

Customer & Network (M1)

Intelligence engines

L3 · DETERMINISTIC CORE

Forecasting and optimisation for dispatch, cooling and capacity; computer vision for inspection and site safety; digital twins; document and knowledge models for deals, contracts and regulation.

Forecast & optimise

Computer vision | Digital twin

Document AI | Anomaly & health

Data & ontology

L2 · THE EXECUTABLE CONTRACT

Keppel Asset Ontology across power, cooling, data centres, buildings, funds and contracts; a unified data platform over SCADA, BMS, IoT, ERP and fund administration; lineage, quality gates and decision-readiness checks.

Asset Ontology | IoT / OT data

Fund & deal data

Readiness gate

Foundations

L1 · MODELS · COMPUTE · SECURITY

Multi-model by design: frontier APIs, open-weight models and Singapore-hosted options chosen per task; ML infrastructure on Keppel's own data centres; MLOps, evaluation harnesses, secrets and OT/IT segregation.

Multi-model routing

Keppel DC compute

MLOps & evals | Cyber & OT

DESIGN CHOICES I WILL DEFEND

- **Run our own compute where it pays.** Keppel builds AI-ready data centres; our ML platform should live in them and become a reference customer.
- **One ontology, many models.** Model vendors will change every year; the semantic layer and evaluation sets are the durable assets.
- **Evidence fields are mandatory.** Every agent recommendation carries its sources, assumptions and confidence, or it is not shown.
- **OT stays deterministic.** Plants, chillers and switchgear take set-points from validated optimisers and controllers, never from a language model directly.

Responsible AI & model risk

CROSS-CUTTING · BUILT ONCE

Policy, model-risk tiers, human oversight, data protection, IP, transparency and resilience, implemented as reusable controls and evaluation pipelines rather than per-project reviews.

04 / VALUE MAP

Where AI creates measurable value

Twenty-two opportunity areas mapped by value at stake and readiness to reach production, based on public information and the first-100-day assessment that will replace these estimates with Keppel's own numbers.

Infrastructure

largest earnings share

Dispatch, bidding and fuel optimisation across Sakra, Merlimau and BESS	1–3% margin · indicative
Predictive maintenance for turbines, boilers and chillers	15–30% maint. cost
AI cooling scaled across all district cooling and EaaS sites, then sold as a service	up to 20% energy
Waste-to-energy combustion and availability optimisation	availability +1–2 pts

Connectivity

AI-ready campus

Data centre PUE, cooling and capacity planning; incident prediction	PUE −0.05 to −0.1
Floating data centre design and operations twin	de-risk first-of-kind
Bifrost capacity sales analytics and network operations	fibre-pair pricing
M1 customer-service and network-operations agents supporting the 90-day plan	cost to serve −20–30%

Real Estate

RE-as-a-Service

Building energy and tenant-experience optimisation	8–15% energy
Sustainable urban renewal: AI audit to retrofit scope and business case	weeks, not months
Leasing and asset-plan analytics	faster decisions

Asset Management

FUM S$95bn → 200bn

Investment copilot: screening, data-room Q&A, IC memo drafting with evidence	DD cycle −30–40%
Portfolio and asset-performance monitoring across funds	same platform as ops
LP reporting, ESG and climate data automation	days to hours

05 / LIGHTHOUSE PORTFOLIO

Six initiatives for year one

Chosen for value, readiness and what they teach the platform. Each has a division owner, a value hypothesis, a production path and a 12-month milestone. The bar shows the share of the initiative that is foundation work versus new capability.

LIGHTHOUSE 01 · CONNECTIVITY

Asset Intelligence for AI-ready data centres

Owner: Keppel Data Centres · with Keppel DC REIT assets

Cooling and PUE optimisation, predictive maintenance and capacity planning across the Genting Lane campus first, then the portfolio. Designed so tenants running AI workloads see it as a service.

VALUE	PUE −0.05 to −0.1; unplanned downtime −30%; capacity released for AI racks
12 MONTHS	Two campuses in production; RaaS offer defined for REIT assets

foundation — new capability

LIGHTHOUSE 02 · INFRASTRUCTURE

Energy Operations Agent

Owner: Keppel Infrastructure · integrated power

A deterministic dispatch and bidding optimiser for Sakra, Merlimau and storage, wrapped by an agent that explains, runs what-ifs and prepares decisions for traders and plant managers. Plant health models feed the same loop.

VALUE	Spark-spread capture +1–3%; maintenance cost −15%
12 MONTHS	Live in the control room at L1 (advise), L2 (propose) for storage

foundation — new capability

LIGHTHOUSE 03 · ASSET MANAGEMENT

Investment Intelligence Copilot

Owner: Keppel Asset Management · deal and portfolio teams

Screening, data-room question answering, investment-committee memo drafting with mandatory evidence, and portfolio monitoring that reads the same asset data the operators use.

VALUE	Due-diligence cycle −30–40%; more deals screened per analyst
12 MONTHS	Used in every IC paper of two flagship funds; LP reporting automated

foundation — new capability

LIGHTHOUSE 04 · GROUP

Keppel Asset Ontology & data platform

Owner: Group AI & Technology Office · with division data leads

The semantic layer for power, cooling, data centres, buildings, funds and contracts, with source-system mappings, constraint axioms and a decision-readiness gate. The contract every lighthouse is built on.

VALUE	Second use case at a fraction of the first; data disputes settled once
12 MONTHS	Ontology v1 covering three divisions; readiness gate in front of every agent

foundation — new capability

LIGHTHOUSE 05 · INFRASTRUCTURE → EXTERNAL

AI Cooling and EaaS productisation

Owner: Keppel Infrastructure · decarbonisation solutions

Package the Operations Nerve Centre and the award-winning AI cooling into a repeatable Cooling-as-a-Service product, following the Chennai model, with an AI performance guarantee and a partner delivery kit.

VALUE	New recurring revenue; higher margin on decarbonisation EBITDA (S$130M, +32%)
12 MONTHS	Productised offer; three new sites under contract outside Singapore

foundation — new capability

LIGHTHOUSE 06 · GROUP

Responsible AI & model-risk foundation

Owner: Group AI & Technology Office · with Risk, Cyber, Data, Legal

Policy, model-risk tiers, evaluation pipelines, human-oversight rules and an AI inventory, aligned with Singapore's Model AI Governance Framework, AI Verify and ISO/IEC 42001. Enables the other five to move fast.

VALUE	Approval time per use case measured in days; zero unmanaged models
12 MONTHS	Framework adopted by the Board; all production agents tiered and evaluated

foundation — new capability

06 / ROADMAP

Three horizons, 36 months

Listen and assess, build foundations while shipping lighthouses, scale across assets and funds, then embed. Quarters from the start date; milestones in red.

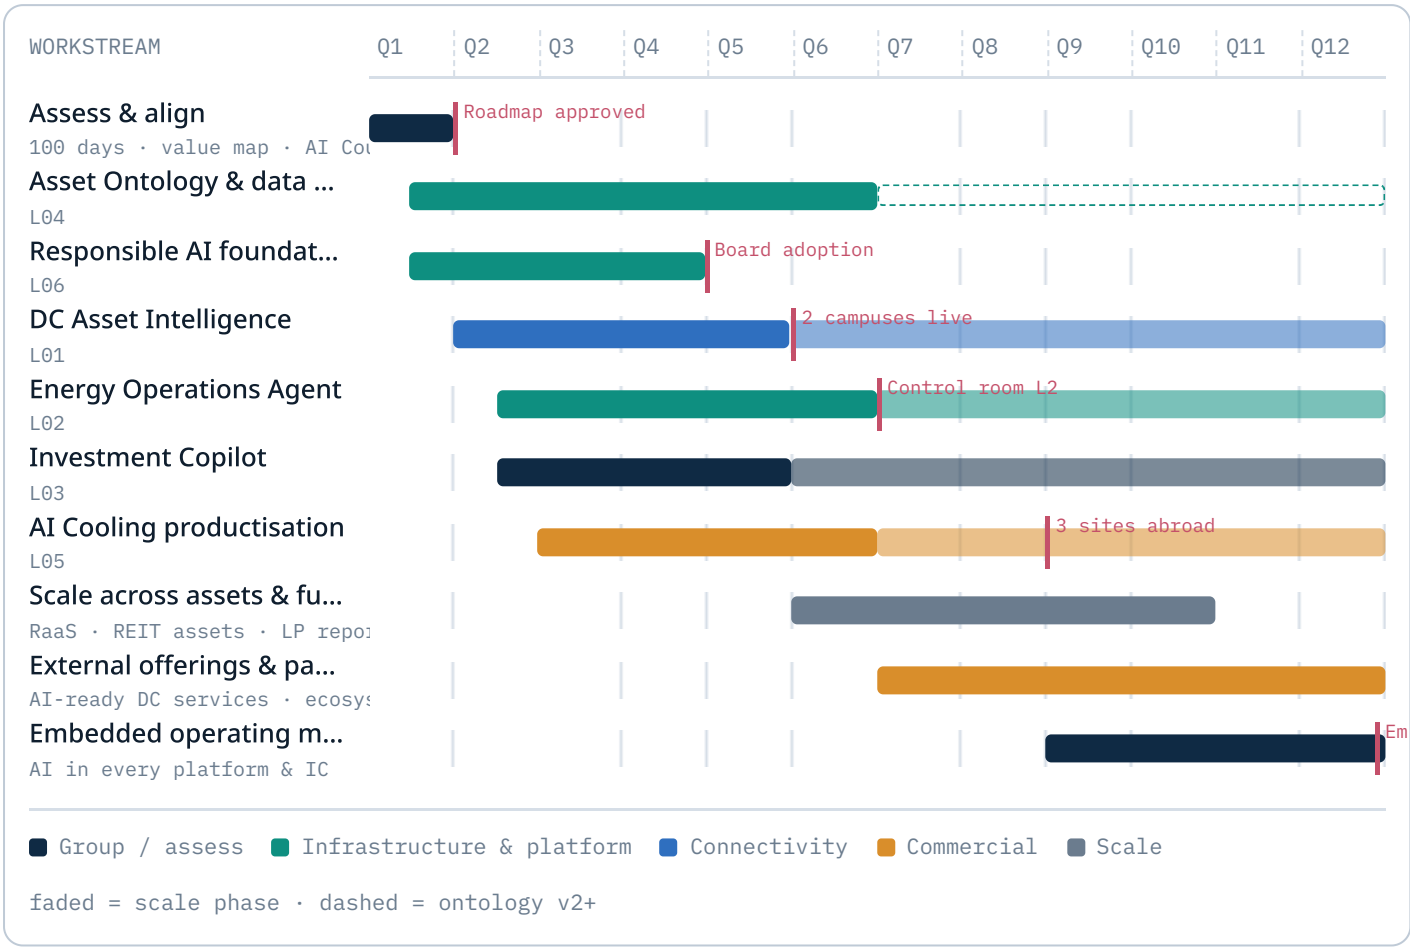

HORIZON 0 · DAYS 1–100

Listen, assess, decide

- Asset and technology landscape across all divisions; data audit
- Value map with Keppel's own numbers; six lighthouses funded
- AI Council chartered; hiring plan; two quick wins shipped

HORIZON 1 · MONTHS 3–12

Foundations and lighthouses

- Ontology v1, data platform, MLOps, model-risk tiers
- All six lighthouses in production at first sites
- First S$ value tracked in division P&Ls

HORIZON 2 · MONTHS 12–24

Scale and productise

- Roll-out across Keppel-managed assets and REIT portfolios
- RaaS for LPs; external Cooling-as-a-Service and DC services
- Partner ecosystem and academic programmes running

HORIZON 3 · MONTHS 24–36

Embed

- AI in every operating platform and every investment committee
- AI-native offerings contributing to recurring income
- Group-level value and risk reported like any other line

07 / OPERATING MODEL

Hub, spokes and a stage gate

A small Group AI & Technology Office that builds the platform and lends forward-deployed engineers; division AI squads that own use cases; business leaders who own value. Ideas move through one funnel with clear exits.

1 · Idea	Any division; logged against the value map with a named decision it improves	`2 weeks`
2 · Value case	Business owner, S$ hypothesis, data owner, ontology coverage check	`gate: AI Council`
3 · Proof	Time-boxed build on the platform with an evaluation set; no bespoke stack	`≤ 8 weeks`
4 · Pilot	Production-grade at one site, tiered for model risk, value tracked weekly	`≤ 6 months`
5 · Scale	Roll-out kit, RaaS packaging, cost and value in the division P&L	`run-rate`

TALENT AND ECOSYSTEM

- **Forward-deployed engineers** paired with plant managers, DC engineers and deal teams: the people who turn tacit expertise into ontology and agents.
- **Hire for the platform, borrow for the frontier:** ML, ontology and data engineering in-house; frontier research through NUS, NTU, A*STAR and AI Singapore, with joint Ph.D. and industry-attachment tracks.
- **Partners with skin in the game:** hyperscalers already working with Keppel on data centres and cables, model providers, and start-ups evaluated through the same stage gate.
- **Keppel as customer zero:** our own AI workloads run in Keppel data centres and become the reference for AI-ready DC services.

08 / RESPONSIBLE AI

Fast because it is governed

Built with Risk, Cybersecurity, Data and Legal from day one, aligned with Singapore's Model AI Governance Framework (including the 2024 generative-AI edition), AI Verify, the PDPA, ISO/IEC 42001 and the NIST AI RMF, and with the operational realities of plants, networks and funds.

CONTROL 01

Model-risk tiers and human oversight

Every model and agent is tiered by the decision it touches. Higher tiers require validation, monitoring and a named human approver; the tier decides the autonomy level an agent may reach.

CONTROL 02

Data protection, IP and provenance

Data classification and residency rules for LP, tenant and customer data; provenance recorded for training and retrieval sources; IP terms settled in every vendor and partner agreement before data moves.

CONTROL 03

Cyber, OT resilience and evaluation

OT/IT segregation for plants and data centres; adversarial and safety evaluation before deployment; drift and incident monitoring in production; fallback to deterministic control on any failure.

TIER 1	Decisions affecting safety, physical assets, regulatory filings, capital commitments or bids	Human approval mandatory · autonomy capped at L2 (propose) · independent validation · full audit trail
TIER 2	Operational recommendations, customer-facing content, financial analysis used in decisions	Human-in-the-loop by sampling · evaluation set and drift monitoring · L3 only inside tested guardrails
TIER 3	Internal productivity, drafting, search and summarisation	Self-service on approved platforms · usage policy · periodic review

09 / SCORECARD

How the role should be measured

Indicative 36-month targets, to be calibrated with the CEO and CFO in the first 100 days. The measure is value delivered and transformation embedded, not models trained.

AI-ATTRIBUTABLE VALUE

S$100 M+ run-rate

EBITDA and cost impact tracked in division P&Ls, audited like any other line. Indicative.

PRODUCTION USE CASES

30 +

Agents and models in production across all divisions, each with an owner, a tier and an evaluation set.

CONCEPT TO PRODUCTION

< 6 months

Median time from approved value case to first production site, down from the years typical of enterprise AI.

ASSETS ON THE PLATFORM

70 %

Share of Keppel-operated capacity (power, cooling, data centres, buildings) covered by the ontology and asset-intelligence services.

EXTERNAL AI REVENUE

2 offerings

Productised AI services (Cooling-as-a-Service, AI-ready DC services) contributing to recurring income.

MODEL RISK

0 unmanaged

Every production model in the inventory, tiered and monitored; zero Tier 1 decisions without a human approver.

TALENT

40 + engineers

Platform, ontology, ML and forward-deployed engineers across the Group office and division squads; joint programmes with two universities.

INVESTMENT PROCESS

100 % of IC papers

Investment-committee papers for flagship funds prepared with the copilot, with evidence and assumptions machine-checked.

10 / FIRST 100 DAYS

What I will do first

Listening before building, but not listening instead of building: two quick wins ship inside the first quarter.

DAYS 1–20

Listen

- Site visits: Sakra, ONC, Genting Lane campus, M1, two funds' deal teams
- 30 conversations with division heads, CIOs, risk and investment leaders
- Inventory of existing models, pilots, vendors and data

DAYS 20–45

Assess

- Data and architecture audit; OT/IT and cyber posture
- Value map rebuilt with Keppel numbers; lighthouse short-list
- Responsible-AI gap analysis against the Singapore framework

DAYS 45–70

Decide

- Roadmap, budget and hiring plan to the CEO and Board
- AI Council chartered; division AI leads named
- Platform choices: compute, models, ontology tooling

DAYS 70–100

Ship

- Quick win 1: AI cooling extended to two more sites with value tracking
- Quick win 2: investment-copilot pilot on one live deal
- Ontology v0 for one asset class; first FDEs embedded

DAY 100

Report

- Board update: baseline, targets, first value, risks
- Public narrative for LPs, tenants and talent
- Horizon 1 starts with all six lighthouses funded

11 / THE LENS

Why this plan looks the way it does

The mandate asks for deep technology, commercial judgement and the ability to take AI from concept to production at scale. This is the experience behind each part of the plan.

At Glodon, a listed construction-technology group, I was Assistant President and Head of AI: I set the three-year AI roadmap, defined the three-tier architecture that became the group standard, built a 50-person research–engineering–product organisation, took the industry large model through national registration and the platform through Level 3 security certification, and grew direct AI revenue from RMB 17 million to over 100 million in a year across 100+ enterprise customers.

Before that I founded Touchjet, a computer-vision hardware company that shipped products to 30+ countries and won an Edison Gold Award. Since 2026 I have advised CEOs on AI strategy and shipped agentic systems for manufacturers, medical-device and industrial-software companies, built on an ontology-first method

KEPPEL MANDATE	EXPERIENCE
AI strategy & transformation	Group AI roadmap adopted at President's Office level; vendor-to-platform transition
Enterprise technology leadership	Three-tier architecture; industry LLM and multi-agent framework; on-premise deployment

that treats semantics as the executable contract between data and decisions. That is the method proposed here for Keppel.

Executive CV · Selected work · Glodon case

Applied AI & intelligent systems	Agentic scheduling, decision engines, computer-vision products, embodied robotics
Productisation & commercialisation	RMB 100M+ direct AI revenue; hardware shipped to 30+ countries
Innovation & partnerships	National Key R&D Programme PI; university-linked AI companies advised
Responsible AI & governance	National LLM registration; MLPS Level 3; group AI governance system
Leadership & organisation	50-person Center of Excellence; three-country company from zero