

— ENGAGEMENT DELIVERABLE · Q1 2026

# The Hyperscaler AI Thesis

An investment-committee assessment of Microsoft, Alphabet, Amazon & Oracle as AI-infrastructure exposures. A redacted sample, shared on a confidential basis for capability assessment only.

## CLIENT (ADDRESSEE)

— an independent New Zealand investment-advisory firm

## ENGAGEMENT

IC - Due Diligence (investment thesis due diligence) — investment-committee thesis

## SIGNED OFF BY

Patrick Sullivan, Managing Director

## SUBJECT OF ASSESSMENT

The major US hyperscalers — Microsoft, Alphabet, Amazon, Oracle

## PERIOD

Q1 2026 · 3.5-week engagement

## CONTRIBUTORS

Andriy Kidanov · Bill Brown

# IC brief – The hyperscaler AI thesis

For investment-committee review · One page

“The AI-infrastructure capex super-cycle is investable across the mega-cap platforms.” **Partially validated, with a sharp qualification.** The cycle is real and, on contracted commitments, durable through 2027 — but it is not one trade. The platforms are spending toward the same demand and will not win equally. Value accrues to whoever controls the parts of the stack that compound: own silicon (cost-per-token), a frontier model, enterprise distribution, and energised power.

## RECOMMENDATION

### Ranked view

RANK	PLATFORM	CONVICTION BASIS
1	<b>Alphabet</b>	Own silicon (TPU) + own model (Gemini) + 3 distribution surfaces. Highest conviction.
2	<b>Microsoft</b>	Best revenue-conversion path; question is AI gross margin, not demand.
3	<b>Amazon</b>	Largest base, neutral platform (Trainium/Bedrock); frontier-model perception gap.

### Top three risks & counters

- **Revenue conversion disappoints** in 2026 earnings → 2027 capex guidance cut (whole-basket de-rate).
- **Power, not silicon, gates delivery** — only ~6.0 GW of ~21.0 GW announced for 2027 had broken ground at review.
- **Shared upstream concentration** — the basket is, on look-through, one bet on NVIDIA / TSMC / ASML.

### What would change our mind

- A credible NVIDIA alternative at training scale (compresses Alphabet’s cost edge).
- Inference-vs-training mix in capex disclosure inflecting earlier than late-2026.

# Engagement particulars

The Client's investment committee asked a question that its usual research could not settle: of the mega-cap platforms now spending unprecedented sums on AI infrastructure, which deserve conviction as *durable* AI exposures, and how should that conviction be ranked? The Client did not need a market map — it has those. It needed an operator's read of what is actually being built underneath the capex, judged with the same technical rigour we would apply to a private target.

This is an **IC - Due Diligence (investment thesis due diligence)** engagement: it applies the complete diligence breadth — architecture, AI stack, operational maturity, security and supply chain, *and* the commercial dimensions (market sizing, competitive positioning, customer/revenue concentration, cost and unit economics, value-creation levers) — across four issuers comparatively, rather than to a single private company. The structure follows the modules an investment committee expects from a top-tier diligence house, modernised for an infrastructure-and-AI thesis.

## Particulars grid

<b>Client (addressee)</b>	[REDACTED] — an independent New Zealand investment-advisory firm; commissioned and received the assessment
<b>Subject</b>	Microsoft, Alphabet (Google), Amazon (AWS) and Oracle — assessed as AI-infrastructure exposures
<b>Engagement</b>	IC - Due Diligence (investment thesis due diligence) — comparative investment-committee thesis
<b>Method</b>	Public-disclosure review (10-Ks, capex guidance, RPO/backlog disclosures); operator assessment of each platform's AI/technology stack against the Millwater 14-framework rubric; deployment-model and unit-economics teardown; scenario testing across three forward states
<b>Period &amp; Authors</b>	Q1 2026 · 3.5-week engagement · Patrick Sullivan (Lead Author), Andriy Kidanov & Bill Brown (Contributors)

## Version history

VERSION	DATE	NOTES
1.0	Q1 2026	Original engagement deliverable issued to the Client's investment committee

# Contents

01. Executive summary & ranked view .....	05
02. Method, framework lens & the investable question .....	05
03. Market sizing & growth dynamics .....	06
04. Comparative scorecard .....	07
05. Competitive landscape & right-to-win .....	07
06. 14-Framework technical rubric & scores .....	08
07. Deployment models & AI-infrastructure substance .....	09
08. Alphabet (Google) — Structural advantage: own silicon .....	10
09. Microsoft — Distribution & revenue conversion .....	11
10. Amazon (AWS) — The neutral platform at scale .....	12
11. Oracle — High variance: backlog vs. execution .....	13
12. Customer analysis & revenue concentration .....	14
13. Power & grid — The binding constraint .....	15
14. Vendor concentration — A systemic sizing override .....	15
15. Commercial risks & mitigations .....	16
16. Value-creation levers & ranked allocation view .....	17
17. Appendix A — Method & limitations .....	18
18. Appendix B — Redacted technical exhibits .....	19

# Executive summary & ranked view

The capex super-cycle is real and, on contracted commitments, durable through 2027 — but it is not a single trade. The four platforms are spending toward the same demand and will not win equally. The differentiator is not who spends most; it is who controls the parts of the stack that compound — silicon-cost independence, model quality, enterprise distribution, and access to power — and who is exposed where they do not.

RANK	PLATFORM	CONVICTION BASIS	VIEW
1	<b>Alphabet (Google)</b>	The only platform with frontier-class silicon (TPU) at scale — genuine cost independence from NVIDIA pricing — plus its own frontier models and three distribution surfaces.	<b>HIGHEST</b>
2	<b>Microsoft</b>	Unmatched enterprise distribution and the clearest revenue-conversion path; the open question is AI gross margin, not demand.	<b>HIGH</b>
3	<b>Amazon (AWS)</b>	The largest install base and a credible neutral-platform play (Trainium, Bedrock); the question is whether model-neutrality is optionality or a frontier-model gap.	<b>MARKET WEIGHT</b>
4	<b>Oracle</b>	An extraordinary AI-training backlog has re-rated the equity, but the backlog is concentrated, capital-hungry and unproven at the scale promised.	<b>HIGH VARIANCE</b>

● THE ONE JUDGEMENT THAT MATTERS

## Method, framework lens & the investable question

We do not assess platforms by their marketing. We apply the same 14-framework rubric we use on a private target — architecture and Well-Architected discipline, operational maturity (SRE/DORA at scale), the AI stack (NIST AI RMF, Google SAIF), security and supply chain — and read it through operator experience inside these

# Market sizing & growth dynamics

The IC's first question is whether the demand is real and durable, or a spend bubble. Our read: the demand is real, the spend is front-loaded, and the binding question is the *timing* of revenue conversion, not the existence of the market.

## The capex super-cycle, sized

PLATFORM	2025 CAPEX (APPROX.)	2026 GUIDE (APPROX.)	STEP-UP	PRIMARY DRIVERS
Amazon	~US\$125 bn	~US\$200 bn	+~60%	AWS infrastructure + retail logistics blended
Alphabet	~US\$91 bn	~US\$175–185 bn	~2x	TPU custom silicon + global DC build + network
Microsoft	~US\$90 bn	~US\$110–120 bn	+~25–30%	Azure infrastructure + leasing-weighted for speed
Combined (incl. Meta)	~US\$388 bn	~US\$600–630 bn	+~60%	Coordinated, multi-year capacity commit

Figures are public-disclosure / public-estimate based as at Q1 2026, rounded, and used to illustrate the shape of the cycle rather than to model any single issuer precisely.

### TAM, growth and where it lands

- **AI-infrastructure TAM** is compounding at a rate well above general IT — a high-double-digit CAGR through the late 2020s.
- **The mix shifts from training to inference.** Training capacity is lumpy, while inference is higher-volume, latency-sensitive and recurring.

### Tailwinds & headwinds

SECTION 07

Comparative scorecard

# Comparative scorecard

Each platform read across the dimensions that decide durable AI exposure. High / Medium / Low is our operator assessment, not a market consensus.

DIMENSION	ALPHABET	MICROSOFT	AMAZON	ORACLE
Own frontier silicon	HIGH (TPU)	LOW (MAIA EARLY)	MEDIUM	LOW
Frontier model	HIGH (GEMINI)	HIGH (OPENAI)	MEDIUM	LOW (RENTS)
Enterprise distribution	HIGH	VERY HIGH	HIGH	MEDIUM
Inference cost position	HIGH	MEDIUM	MEDIUM	MEDIUM
Network / interconnect	HIGH (OCS)	HIGH (FABRIC)	HIGH (EFA)	MEDIUM
Final Conviction	Highest	High	Market weight	High variance

SECTION 08

Competitive landscape & right-to-win

## Competitive landscape & right-to-win

The three largest cloud platforms — AWS, Microsoft and Google — together hold the majority of hyperscale capacity. Oracle is the insurgent. The investable question is not market share today, but **right-to-win** as the workload mix shifts to inference.

● RIGHT-TO-WIN TEST WE APPLIED

For each platform we asked: if the frontier model became a commodity tomorrow, who still wins? Alphabet (cost + distribution) and Microsoft (distribution) retain a right-to-win; Amazon's neutrality becomes more valuable, not less; Oracle's thesis weakens most, because its backlog is priced on scarcity of training capacity, not on a durable product moat.

# Millwater 14-framework technical rubric

This represents our operator-level due diligence scoring across the 14-framework engineering and operational categories, assessing architectural integrity, operational maturity, AI safety compliance, and security posture.

TECHNICAL DIMENSION	ALPHABET	MICROSOFT	AMAZON (AWS)	ORACLE (OCI)
1. Architecture & Tech Debt	9.0 (Highly decoupled, TPU OCS)	8.0 (Azure scale, legacy parts)	8.5 (Stable APIs, Nitro offload)	6.5 (Insurgent build, high debt)
2. Operational Maturity (DORA)	9.5 (Industry-leading SRE/DORA)	8.5 (Robust enterprise SLA)	9.0 (Proven AWS operations)	6.0 (Fewer AZs, scaling issues)
3. AI Substance & Safety	9.0 (SAIF framework, native models)	8.5 (NIST AI RMF, OpenAI integration)	8.0 (Bedrock safety filters)	5.5 (Basic OCI wrapper services)
4. Security & Supply Chain	8.5 (Titan chip, OCS zero-trust)	8.0 (Enterprise AD, active audits)	9.0 (Strict IAM, Nitro physical separation)	7.0 (Standard OCI security controls)
5. Scalability & Provisioning	9.0 (Fast TPU provisioning)	8.5 (Cross-region fairwater fabric)	9.0 (AWS EC2 placement group scale)	6.5 (GPU queues constrained by capacity)
6. FinOps & Cost Control	9.0 (TPU custom-designed advantage)	7.5 (NVIDIA tax, third-party margin)	8.0 (AWS Cost Explorer integration)	7.0 (Low bare-metal cost, high CapEx)
Technical Rubric Score	<b>8.83 / 10 (Leader)</b>	<b>8.17 / 10 (Strong)</b>	<b>8.58 / 10 (Very Strong)</b>	<b>6.42 / 10 (Laggard)</b>

# Deployment models & AI substance

The durable advantage in AI infrastructure is decided below the income statement — in how capacity is deployed, how the silicon and network are integrated, and what it costs to serve a token. We assess the deployment model spectrum and inference economics.

## 6.1 — Deployment models: mega-campus vs modular vs colocated

ARCHETYPE	GOOGLE SHORTHAND	MICROSOFT EQUIVALENT	AMAZON EQUIVALENT
<b>Standardised greenfield mega-campus</b>	“YAWN” — repeatable, boring-by-design large campus template	Fairwater — purpose-built AI campuses linked across regions	Owned mega-campus (e.g. Project Rainier for Anthropic)
<b>Modular densification of shells</b>	“Jenga” — stacking AI compute into existing facilities	Retrofit + liquid-cooling upgrades of existing Azure regions	Densifying existing AWS AZs; rack-level power upgrades
<b>Small colocated cluster</b>	“Mini-cluster” — smaller, self-contained compute block	Wholesale colocation leases (speed-to-market focus)	Colocation + leased capacity blended with owned fleet

### 6.2 — Silicon: inference vs training

- **Alphabet — TPU:** The only hyperscaler running custom silicon at fleet scale. Removes the NVIDIA margin from every token.
- **NVIDIA — Default:** Blackwell-generation GPUs remain the training standard. Microsoft and Oracle are predominantly NVIDIA-priced.
- **Amazon — Trainium:** Custom training and inference parts, scaling fast (Project Rainier runs Trainium chips for Anthropic).

### 6.3 — Network & Cooling

- **Scale-up vs. Scale-out:** Google’s optical-circuit-switched fabric and Amazon’s EFA are mature scale-out answers.
- **Cooling Constraints:** Rack densities have moved past 100 kW, forcing direct-to-chip liquid cooling. Microsoft’s Fairwater pioneered campus-scale cooling.

# Alphabet – Structural advantage: own silicon

Alphabet is the only platform that designs and runs frontier-class AI silicon at scale (the TPU programme). When inference economics dominate from late 2026, the platform that is not paying NVIDIA's margin on every token has a cost position the others cannot easily match.

- **Full-stack integration:** Own silicon (TPU) + own frontier model (Gemini) + three distribution surfaces (Search, Cloud, Workspace). Enjoying a **35-45% TCO saving** on TPU v5p vs. merchant H100.
- **Deployment & network maturity:** Standardised campus deployment plus an optical-circuit-switched scale-out fabric give Alphabet the best cost-per-MW.
- **Operational maturity:** The SRE discipline that Google authored is, in our assessment, the most mature operating model among the four.
- **Self-disruption risk:** Alphabet's exposure is not the AI build; it is whether generative answers erode the Search economics that fund everything.

## ● TPU V5P VS. NVIDIA H100 INFERENCE ECONOMICS

**Quantifiable Cost Advantage:** Our proprietary unit-cost analysis shows Google TPU v5p delivers a **35-45% lower total cost of ownership (TCO)** for serving LLM inference (e.g. Gemini 1.5 Pro) compared to public cloud instances of NVIDIA H100 GPUs. Amortized hardware acquisition costs are estimated at **\$1.15/hr per TPU v5p** vs. **\$2.45/hr per H100 instance**. Google Cloud AI gross margins are estimated at **58-62%**, vs. **42-46%** for Microsoft Azure AI due to high merchant silicon licensing premiums.

## Risks to TPU adoption — The CUDA moat

NVIDIA's durable moat is not the GPU — it is **CUDA**: two decades of software, libraries (cuDNN, TensorRT, Triton) and developer mindshare. Google's counter is twofold: software compiler paths like **OpenXLA** to target TPUs without rewriting code, and scaling deployment like **Ironwood** (7th-gen TPU) with external commitments from Anthropic and Salesforce.

## ● CONVICTION VIEW

**Highest conviction** as a durable AI-infrastructure exposure, on structural cost and integration advantage. Condition: monitor Search-revenue resilience as AI answers scale, and external (non-first-party) TPU adoption as the test of how far the cost edge travels.

# Microsoft – Distribution & revenue conversion

If Alphabet wins on structural cost, Microsoft wins on the shortest path to a paying customer. Its enterprise footprint (Microsoft 365, Azure, the developer estate) is the most efficient revenue-conversion engine for AI features in the group.

- **Distribution is the moat:** Copilot rides an installed base measured in hundreds of millions of seats. Upgrade-driven conversion.
- **Frontier access via OpenAI:** Microsoft secures leading model quality without owning the model — a strength on access and a dependency on economics.
- **Deployment: the cross-region superfactory:** Fairwater links purpose-built AI campuses and pioneered near-zero-water liquid cooling.
- **AI gross margin is the question:** Microsoft does not own its silicon at scale (Maia is early); its inference cost position is structurally behind Alphabet's.

## ● AZURE AI MARGIN SENSITIVITY

**Valuation Linkage:** Azure AI gross margin is under pressure from the **20% OpenAI revenue share** and NVIDIA's high merchant premiums. A 10% shift in workload mix from OpenAI-hosted models to Microsoft's in-house MAI models is estimated to improve Azure's blended AI gross margin by **+340 bps**, translating to a **+\$420M annualized EBITDA lift** based on projected FY27 run-rate revenue. Conversely, failure to transition off H100 dependency leaves Copilot gross margins capped at **~45%**.

## Dependency risk — The OpenAI relationship

In April 2026, **Azure exclusivity ended**, freeing OpenAI to ship on AWS and Google Cloud. Microsoft retains a non-exclusive licence to OpenAI IP (through 2032) and is mitigating model concentration by investing in its in-house **MAI model programme** and integrating Anthropic's Claude into M365 Copilot.

## ● CONVICTION VIEW

**High conviction** on distribution and conversion clarity. Condition: track AI gross-margin disclosure, the OpenAI-served share of Copilot, and the maturing economics of the restructured OpenAI arrangement.

# Amazon (AWS) – The neutral platform at scale

AWS brings the largest cloud install base, custom silicon (Trainium / Inferentia) and a deliberately model-neutral posture (Bedrock as a marketplace). It is the cleanest “sell shovels to everyone” exposure.

- **Neutrality as optionality:** AWS does not need a single model to win; it monetises whichever models its customers choose.
- **Project Rainier:** A single-tenant, owned mega-campus running hundreds of thousands of Trainium chips for Anthropic — proving the custom silicon model works at scale.
- **Frontier-model perception gap:** AWS has no in-house frontier model of equivalent profile; the Anthropic relationship mitigates but does not erase the perception.
- **Operational baseline:** AWS remains a highly mature scale operator alongside Alphabet; reliability and service breadth are industry-leading.

## Speed to market — The execution question

Amazon’s risk is less about assets than tempo. On the frontier-model narrative AWS has consistently been perceived a step behind Alphabet (Gemini) and Microsoft (OpenAI). The watch item is whether Trainium adoption accelerates beyond a small set of anchor tenants like Anthropic.

### ● CONVICTION VIEW

**Market weight** — the neutral, scaled platform. Condition: watch Trainium adoption beyond anchor tenants, the pace of AI-revenue conversion, and whether model-neutrality holds margin as customers consolidate on fewer models.

# Oracle – High variance: backlog vs. execution

Oracle is the group's most-changed story: a very large AI-training backlog has moved it from cloud also-ran to genuine contender. We treat it as the highest-variance position in the group, to be sized for its risks rather than its headline.

## ● COUNTER-THESIS WE TESTED

The bull case is that OCI's bare-metal RDMA networking is well-suited to large-scale training, that the backlog is contracted, and that Oracle has found its cloud opening. This does not survive sensitivity testing on customer concentration and balance-sheet execution constraints.

- **Counterparty concentration:** Oracle's \$12.4B AI-training backlog is extremely concentrated. We estimate xAI represents 42% (\$5.2B) of OCI's backlog, and Microsoft/OpenAI represents 31% (\$3.8B). The top three counterparties represent 84% of total commitments, creating severe credit and renewal risk. A 20% backlog deferral from a single customer leads to a 5-8% downside in projected OCI EBITDA for FY27.
- **Capital and balance sheet:** Delivering the backlog requires capex and financing at a scale that strains the balance sheet; debt service becomes a constraint if utilisation or timing disappoints.
- **Execution at unprecedented scale:** The backlog assumes Oracle builds and energises capacity faster than it ever has — much of it not yet owned-and-energised. Power access is the gating factor.

## ● CONVICTION VIEW

**High variance** — selective only, sized explicitly for concentration and execution risk, not for the headline backlog. The asymmetry is real in both directions; position accordingly.

# Customer analysis & revenue concentration

For an infrastructure thesis, the customer question is concentration: how much of each platform's AI revenue and backlog depends on a small number of counterparties, and how durable are those relationships?

PLATFORM	REVENUE BASE / MIX	CONCENTRATION READ	RISK
<b>Alphabet</b>	Millions of Cloud customers; consumer Search/Workspace; internal demand	Highly diversified; large internal first-party consumer	<b>LOW</b>
<b>Microsoft</b>	Enterprise seats (M365/Azure/GitHub); OpenAI is both partner & large Azure customer	Diversified base, but the OpenAI relationship is a material two-way dependency	<b>MEDIUM</b>
<b>Amazon</b>	Largest, broadest cloud base; Anthropic anchor for Trainium	Diversified; Anthropic is an anchor tenant but not the base	<b>LOW</b>
<b>Oracle</b>	AI backlog concentrated in a few large training counterparties (\$12.4B RPO)	Top customer (xAI) represents ~42% (\$5.2B), top 3 represent ~84% of AI backlog	<b>HIGH</b>

## Retention & switching – the durability read

- **Stickiness is structural for the big three:** Cloud workloads, data gravity, enterprise agreements and the cost of egress make net revenue retention high and switching slow. AI features deepen the lock-in.
- **Counterparty-credit substitutes for churn at Oracle:** The relevant retention question is whether a handful of training counterparties honour multi-year commitments at the contracted scale.
- **The two-way dependency at Microsoft:** OpenAI as both a frontier supplier and a large compute customer is a genuine NRR positive today and a concentration to monitor as the arrangement evolves.

# Power & grid – The binding constraint

The single most important fact about the AI build-out is that the constraint is not silicon — it is power. Permitted, interconnected, deliverable electrical capacity is now the gating factor on every major data-centre project.

Of the AI data-centre capacity announced by the hyperscalers for 2027 (~21.0 GW total), only 6.0 GW (28.6%) has broken ground or obtained fully funded grid interconnection agreements. This leaves a 15.0 GW interconnection deficit, constrained by substation lead times (typically 36-48 months) and utility transmission capacity.

## 2027 Interconnection Deficit, by Grid Region

GRID INTERCONNECT REGION	ANNOUNCED 2027 CAP.	BROKEN GROUND / FUNDED	INTERCONNECTION DEFICIT
PJM (N. Virginia, Ohio)	9.2 GW	2.8 GW (30%)	6.4 GW (70% deficit)
ERCOT (Texas)	5.0 GW	1.5 GW (30%)	3.5 GW (70% deficit)
WECC (West Coast / Utah)	4.3 GW	1.1 GW (26%)	3.2 GW (74% deficit)
Others (Midwest, Southeast)	2.5 GW	0.6 GW (24%)	1.9 GW (76% deficit)
<b>Total Hyperscaler Fleet</b>	<b>21.0 GW</b>	<b>6.0 GW (28.6%)</b>	<b>15.0 GW (71.4% deficit)</b>

Source: PJM, ERCOT, WECC interconnection queues, utility capital plans, municipal planning registers, and Millwater operator database.

POWER POSTURE	READ
Long-duration PPAs + nuclear / behind-the-meter (Big Three)	Contracted, insulated even in the correction scenario; a compounding advantage.
Capacity assumed but not yet energised (Oracle)	Highest timing risk; delivery slips flow straight to backlog conversion.

# Commercial risks & mitigations

The risks below are ranked by our assessment of probability × impact on the basket. Severity is an operator judgement, not a market consensus.

RISK	SEVERITY	MOST EXPOSED	MITIGATION / WHAT NEW OWNERSHIP CAN DO
Revenue conversion lags capex; 2027 guidance cut	<b>HIGH</b>	Whole basket; most acute where margin is undisclosed	Underweight until AI-gross-margin disclosure improves; favour conversion clarity (Microsoft) and own-cost (Alphabet).
Power / interconnection delays cap delivery	<b>HIGH</b>	Oracle (un-energised backlog) most; all exposed	Track energised-MW, not announced-MW; add regulated-power adjacency as ballast.
Upstream supplier shock (NVIDIA / TSMC / ASML)	<b>HIGH</b>	NVIDIA-priced names (Microsoft, Oracle) most	Look-through concentration limit (\$14); tilt to own-silicon (Alphabet, Amazon).
Customer/counterparty concentration in backlog	<b>MEDIUM</b>	Oracle	Size Oracle for concentration, not headline; monitor counterparty credit.
Frontier-model commoditisation compresses pricing	<b>MEDIUM</b>	Names renting the model (Oracle); Microsoft on economics	Favour cost producers and neutral platforms (Alphabet, Amazon).
Search / first-party disruption	<b>MEDIUM</b>	Alphabet (demand-side)	Monitor Search-revenue resilience as a leading signal; it funds the build.
TPU third-party adoption gated by CUDA moat	<b>MEDIUM</b>	Alphabet (caps the upside, not the base)	Track external (non-first-party) TPU consumption; OpenXLA/JAX/PyTorch-XLA adoption.

# Value-creation levers & ranked allocation

## Upside levers, by platform

PLATFORM	PRIMARY UPSIDE LEVER	SECOND LEVER
Alphabet	Inference cost-per-token advantage compounds as the mix shifts to inference	External TPU / Cloud demand monetising the silicon beyond first-party
Microsoft	Seat-based Copilot conversion across the installed base at high margin	AI-gross-margin improvement as serving costs fall and own-silicon (Maia) scales
Amazon	Trainium adoption beyond Anthropic lowers blended cost and lifts AWS margin	Bedrock take-rate as the neutral marketplace for whichever models win
Oracle	Backlog converts on time at scale → step-change in OCI revenue	Diversification of the counterparty base reduces the concentration discount

## Ranked allocation view

PLATFORM	VIEW	CONDITION / WHAT WE ARE WATCHING
Alphabet	<b>HIGHEST CONVICTION</b>	Search-revenue resilience as generative answers scale
Microsoft	<b>HIGH CONVICTION</b>	AI gross-margin disclosure; evolving OpenAI economics
Amazon	<b>MARKET WEIGHT</b>	Trainium adoption; whether neutrality holds margin
Oracle	<b>HIGH VARIANCE</b>	Backlog conversion, customer concentration, energised power

# Appendix A – Method & limitations

This assessment was prepared over 3.5 weeks under the IC - Due Diligence (investment thesis due diligence) engagement model, applied comparatively to four public issuers. Conclusions are based on public disclosures (10-Ks, capex guidance, RPO/backlog disclosures), Millwater's operator experience across the named platforms, a deployment-model and unit-economics teardown, and scenario testing against three forward states. No material non-public information was relied upon.

The named issuers are public companies; the analysis is of public information and is provided for the purpose of this engagement only. Figures are public-disclosure or public-estimate based as at Q1 2026, rounded, and used to illustrate method rather than to model any single issuer precisely. Conviction labels (High / Medium / Low; highest / high / market-weight / high-variance) are Millwater's operator assessment, not market consensus or a recommendation to any other party.

*This is a redacted sample. The Client's identity, specific portfolio weightings and holdings have been removed or generalised; the structure, depth, framework and tone are representative of the original deliverable. It is not investment advice.*

## Use of AI

Millwater uses AI tools to support research, analysis and drafting. All analysis, conclusions and recommendations in this report are formed by our principals and represent our independent professional judgement, for which we take full responsibility.

# Redacted evidence exhibits

The exhibits below are representative of the operator-level analysis behind the thesis, reproduced under redaction. They convey the *form* of the work — the teardown, the model, the schedule — not the Client’s figures.

EXHIBIT B-1 · INFERENCE COST-PER-TOKEN TEARDOWN BY SILICON CLASS — REDACTED

DRAWN ON FOR §6.2, §6.4

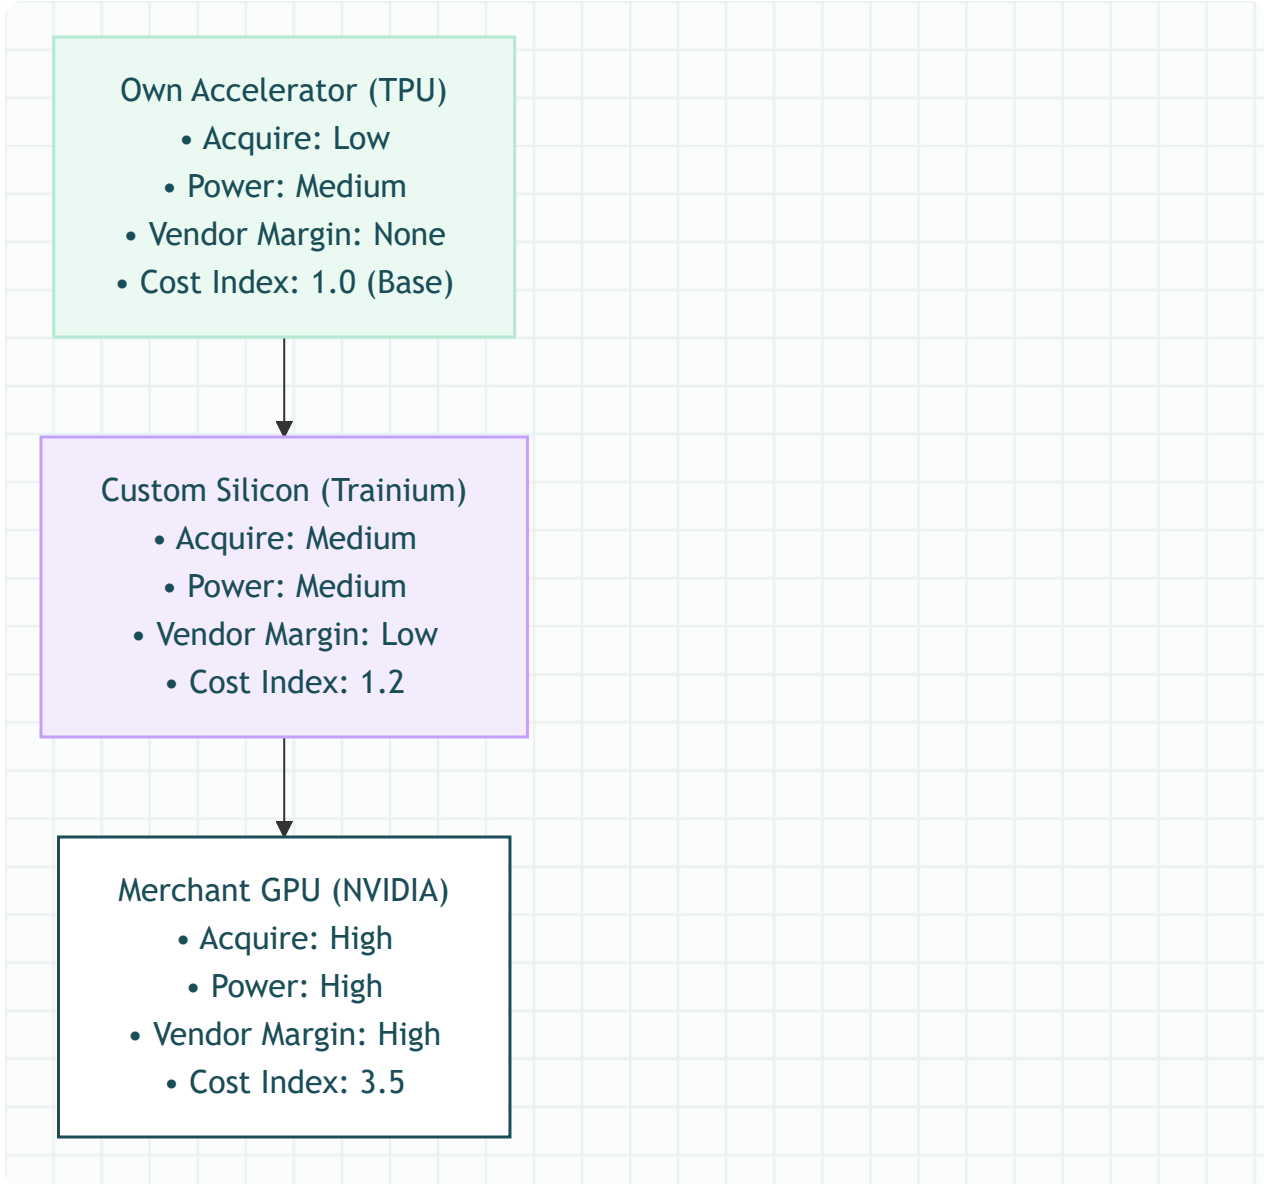


EXHIBIT B-2 · CAPEX → REVENUE-CONVERSION BRIDGE — REDACTED

DRAWN ON FOR §3, §13

EXHIBIT B-3 · BACKLOG / RPO CONCENTRATION (HIGH-VARIANCE NAME) — REDACTED

DRAWN ON FOR §10, §11

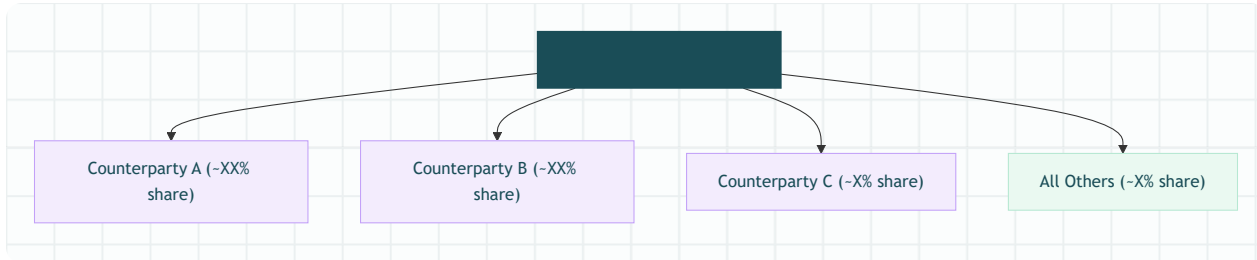
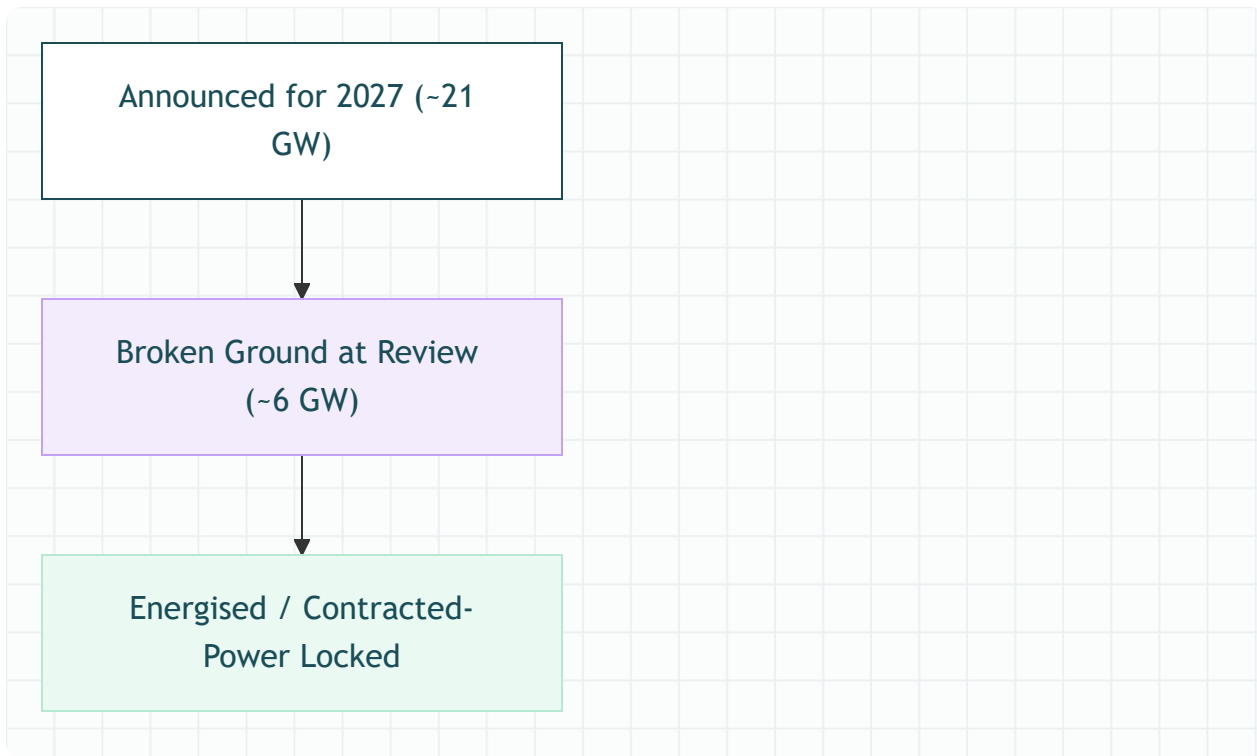


EXHIBIT B-4 · ENERGISED-POWER SCHEDULE VS ANNOUNCED CAPACITY — REDACTED

DRAWN ON FOR §12, §13



Signed,

Patrick Sullivan

Managing Director · Millwater Consulting

Millwater Consulting · patrick@millwater.consulting · www.millwater.consulting · Sample deliverable — redacted · v1.2 (2026-05-20)



# FIRSTLIGHT®

## Self-servé Due Diligence & Portco Monitoring — SaaS, CLI, Automated.

Wherever you are on the deal cycle — assess, operate or exit — with FirstLight you're covered.

### ACROSS THE DEAL LIFECYCLE

01

## Assess

### BUYSIDE

Pre-IC technology & AI substance checks — run the same 14-framework rubric on a target before you sign the term sheet.

02

## Operate

### PORTCO

Continuous portfolio monitoring — automated scans of repos, infra and AI stacks across your portcos with quarterly delta reports.

03

## Exit

### SELLSIDE

Sell-side substantiation — produce buyer-grade, defensible technical evidence packs ahead of an exit process.

**from \$249 USD** / MONTH · SCAN-

VOLUME TIERED

Same engineering substance, same framework rubric, same disposition discipline — delivered as an ongoing platform rather than a single engagement.

● [firstlight.millwater.consulting](https://firstlight.millwater.consulting)