Local AI compute for Thailand

Findings, a staged path and a first data center

Prepared for the investor by **B4L (buildfor.life)**, consultant
October 2026

2 MW

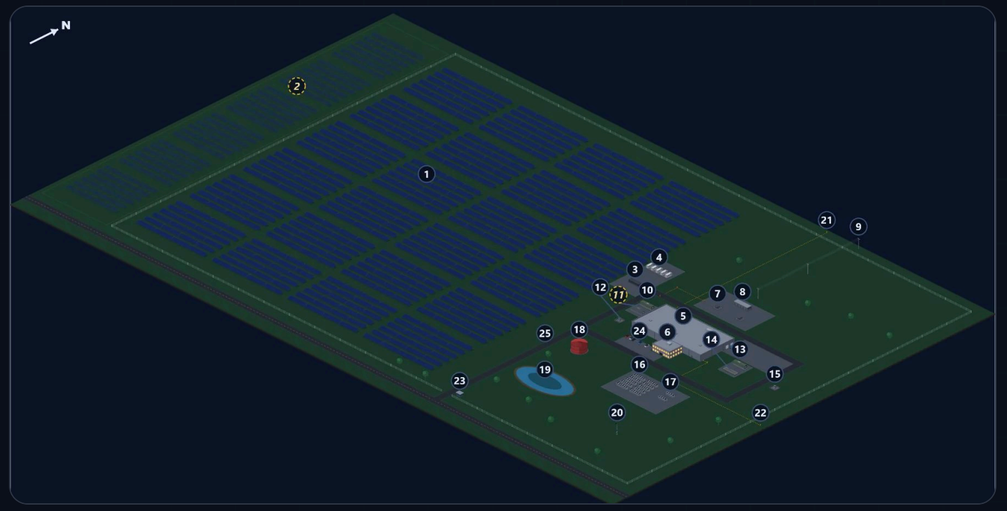
of IT in the first design, the smallest data center the Board of Investment (BOI) promotes

0.86 to 1.27bn

THB of facility capex for phase 1 before any GPU (USD 26 to 38 million, our estimate)

H2 2028

first IT load at the earliest, with a pilot at Chiang Mai University's Information Technology Service Center (CMU ITSC) before then (our estimate)



The first data center as designed, on a 150-rai plot at Sikhio in Nakhon Ratchasima, seen from the southeast. The render is a concept, not a construction drawing; slide 9 names its parts.

Sources: summary; first data center (In short, schedule); model/first_dc_results.md (capex by system). Render: docs/img/first-dc-site.svg. Money in THB, with USD at 33 THB per USD, the study's planning rate.

WHAT THE STUDY COVERS

The study covers ThaiLLM, the Thai models, where they run, three ways to get compute and a first data-center design

SCOPE

- It describes **ThaiLLM**: who runs it, which models it releases, and what it trains and serves them on.
- It surveys **the Thai models** and the strong general models, with their sizes, their licenses and the hardware they need.
- It maps **where Thai AI runs today**: the national supercomputer, the Thai GPU clouds, the hyperscaler regions and BOI's rules.
- It compares **three ways to get compute**: renting as the baseline, the investor's servers at CMU ITSC (option A), and a new data center (option B).
- It designs **the building itself**: the site, the plant, what it can sell, the open software it runs and who supplies it in Thailand.

METHOD

1,281 sources

in one numbered list. Every outside figure in the notes carries its citation.

7 cost models

written in Python on those sources: renting or buying, site energy, hosting, storage, connectivity, server packages and the first data center.

One design

with renders, a 3D model, a business calculator and seven simulated demos.

ROLES

B4L (buildfor.life) is the consultant that prepared the study. **The investor** would build, own and operate everything that follows from it and makes every investment decision.

Outside figures are cited in the study's notes. Our own estimates are marked "est." or "our estimate," and judgments are labeled as ours.

Money is in THB, with USD in parentheses at 33 THB per USD, the study's planning rate.

FOURTEEN NOTES ThaiLLM · Thai models · Thai compute · Hardware · Option A at CMU ITSC · Option B, a new data center · Site selection · Hosting and agent servers · VPS and dedicated servers · Open stack · Facility systems · Connectivity · The first data center · Sourcing

Sources: brief; summary; sources (1,225 numbered entries); the cost models in model/ (rent_vs_buy, site_energy, hosting_prices, open_infra_costs, connectivity_costs, server_packages, first_dc); the calculator and demos on thaiai.lol.

ThaiLLM releases two Qwen3-based models, 8B and 30B, trained on the LANTA supercomputer and served by NT and SIAM.AI CLOUD

WHO The Big Data Institute (BDI) leads ThaiLLM with NECTEC, VISTEC, AIEAT and AIAT, and the DE Fund pays for it. Phase 1 cost about 80 million THB (USD 2.4 million).

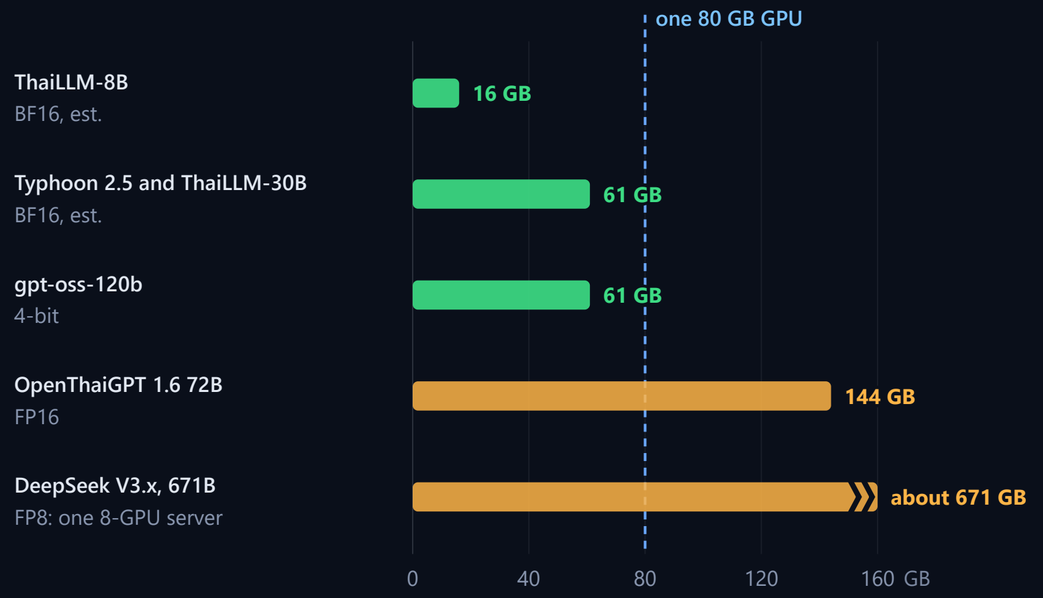
MODELS ThaiLLM-8B and ThaiLLM-30B are Alibaba Qwen3 models trained further on about 63 billion tokens and released under Apache 2.0.

TRAINING They were trained on LANTA, the national supercomputer, which has 704 NVIDIA A100 GPUs in service since November 2022. We estimate the final runs at about 9,000 to 23,000 A100-hours.

SERVING The free Playground and API run on compute from NT and SIAM.AI CLOUD. Their outputs may not be used commercially without BDI's written consent, so a business would host the weights itself.

THE OTHERS Every Thai model of 2025 and 2026 continues a foreign open model, and the teams trained on one 8-GPU server or less. SCB 10X trained every Typhoon 2 size, up to 70B, on one node of eight H100s.

MEMORY FOR THE WEIGHTS ALONE, IN GB



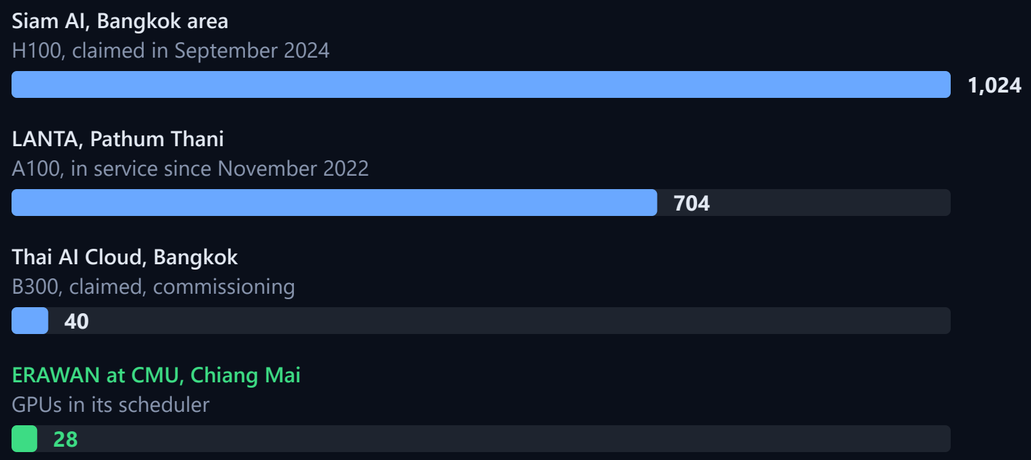
> The best Thai models, such as Typhoon 2.5, each run on one 80 GB GPU. Only frontier models such as DeepSeek V3 need a whole 8-GPU server.

Sources: summary; ThaiLLM (In short); Thai models (the landscape, how they were trained, hardware to serve them); calculator assumptions (DeepSeek V3.x weights at FP8). Memory marked est. is our estimate of the weights at two bytes per parameter.

THE GAP IN THE MARKET

No Thai model maker sells a paid, production API hosted in Thailand, and Chiang Mai has no GPU cloud

GPUS BY PLACE, AS THEIR OPERATORS STATE THEM



> Every commercial GPU offer we found is in Bangkok or the Eastern Economic Corridor (EEC). In Chiang Mai, ERAWAN at CMU is the only sizeable GPU system.

THAI MODEL APIS

Typhoon's free beta allows 5 requests a second and 200 a minute, and its AWS listing still reads "Coming soon." SEA-LION's free API allows 10 requests a minute, and ThaiLLM's API is free but not for commercial use.

HYPERSCALERS

AWS has run a Thai region since January 8, 2025, and Google Cloud one in Bangkok since January 21, 2026. Neither lists a GPU instance there, and Microsoft's region is not live.

THAI GPU CLOUDS

Siam AI lists the H100 at USD 2.69 and the H200 at USD 2.99 per GPU-hour, Float16 the H100 at USD 4.32, and Thai AI Cloud the B300 from USD 8.90. All of them sit in the Bangkok area.

CHIANG MAI

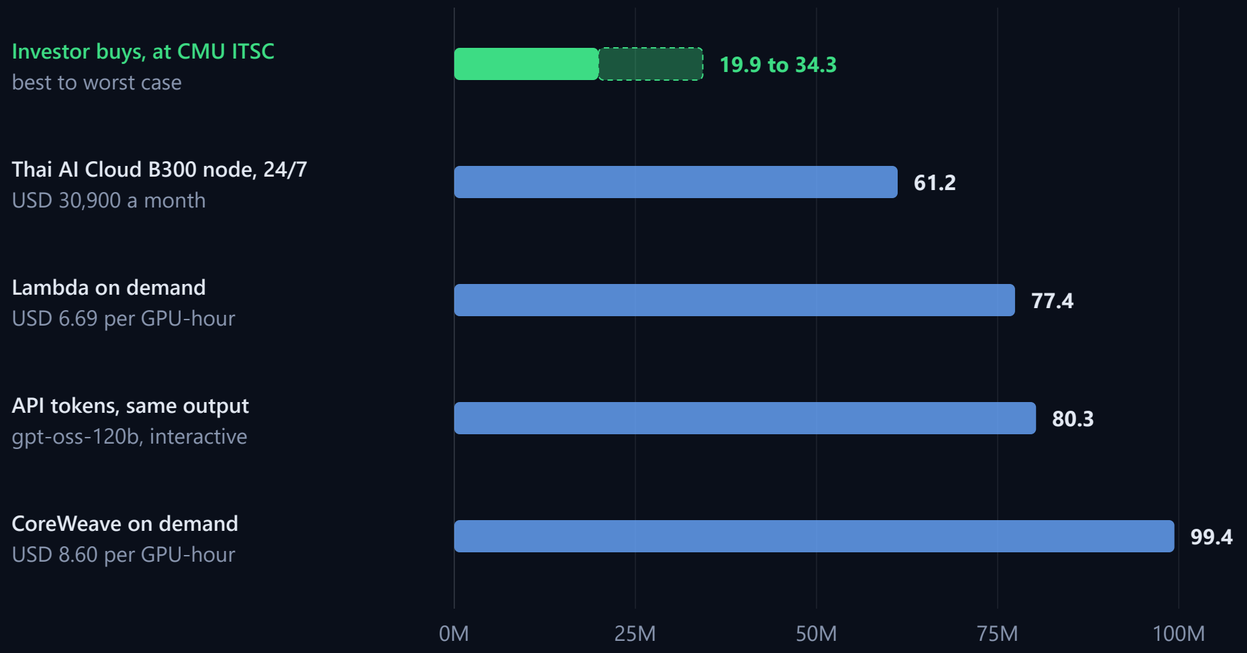
NT's data center there lists 25 racks and is under renovation. ERAWAN, CMU's cluster, has 28 GPUs and charges outside users 200 THB (USD 6.06) per GPU-hour.

Sources: summary; Thai models (hosted APIs today); Thai compute (Thai GPU clouds, hyperscalers, Chiang Mai today). GPU counts are the operators' own claims; prices are list prices read on October 27, 2026.

RENTING IS THE BASELINE TO BEAT

An owned B200 server costs less than renting only above 23% to 43% utilization over five years

FIVE YEARS OF ONE 8X B200 SERVER AT FULL USE, IN MILLION THB



> At full use, owning is the cheapest bar. Utilization decides the case, because rent and tokens are paid only when used, while a server costs much the same whether it works or not.

USD 0.03 / 0.15

per million input and output tokens for gpt-oss-120b on OpenRouter, the hardest baseline.

23% to 43%

utilization above which a B200 server owned at ITSC beats renting at Lambda's USD 6.69 per GPU-hour over five years (our model).

7 to 13 billion

tokens a day, or 23% to 41% interactive load, before the same server beats API tokens. RTX PRO 6000 and H200 servers need over 80% at best (our model).

Rents move fast. H200 rents rose 143% and B200 rents 111% from January 1, 2026, according to a news report. Owning wins on data control, on Thai models that nobody sells as a service, and at high, steady load.

Sources: summary; hardware (In short, buy versus rent); model/results.md (S2, 8x B200, five years, inputs of October 27, 2026). Owning includes the server, support, the ITSC rack and power; it leaves out staff, duty and financing.

THE OPTIONS SIDE BY SIDE

Renting costs nothing up front, ITSC has room for one or two servers, and a new data center must reach 2 MW to earn BOI's incentive

	RENTING (BASELINE)	OPTION A: SERVERS AT CMU ITSC	OPTION B: A NEW DATA CENTER
UP-FRONT COST	none	one server: 5.4 to 6.5M THB (USD 164,000 to 197,000) with 8x RTX PRO 6000, or 11.9 to 18.6M THB (USD 360,000 to 565,000) with 8x B200	320 to 510M THB (USD 9.7 to 15.4M) per MW of IT for the facility (our estimate), plus the servers
RUNNING COST	B200 at USD 6.69 per GPU-hour on demand (Lambda); gpt-oss-120b tokens from USD 0.03 / 0.15 per million	about 38,000 to 84,000 THB (USD 1,150 to 2,530) a month for one 8x B200 server, power included (our estimate)	electricity of about 40M THB a year per MW of IT at PUE 1.2 (our estimate), plus staff and maintenance
TIME TO SERVICE	days	set by the CMU agreement; the RTX PRO 6000 server had a lead time of 3 to 5 weeks	about 18 to 24 months from the start of design (our estimate)
CAPACITY	as much as is paid for	one or two servers (our reading)	sized by design; 2 MW is about 130 servers of 15 kW
STANDARD	the provider's	an ISO/IEC 27001 room; no tier rating found	certified Tier III at PUE 1.15 to 1.2, our proposed target
DATA CONTROL	the provider's terms; ThaiLLM Playground outputs not for commercial use	the investor's hardware in Chiang Mai	the investor's hardware and building
MAIN RISK	price rises (H200 rents +143% in 2026, a news report) and terms	the power cap and access rules set by ITSC	utilization, and a scale below BOI's 2 MW

Sources: summary (the options side by side); hardware; option A (In short); option B (cost and time; BOI and scale); model/results.md. ITSC's room runs on UPS of 120 and 125 kVA and one 320 kVA generator, which is why option A stays small.

OUR RECOMMENDATION

We recommend four stages: preparation, a pilot server at ITSC, a second server on evidence, and a site study in parallel



Prepare

MONTHS 0 TO 3

Define the workloads and customers, test the models on free APIs and rented GPUs, have a Thai lawyer read the Typhoon terms, and open talks with ITSC, BDI and NT, all without buying hardware.

Stage 1: pilot

MONTHS 3 TO 12

One 8x RTX PRO 6000 server at CMU ITSC, for 5.4 to 6.5 million THB, serves Thai models with the data kept in Chiang Mai and measures real power, load and demand.

Stage 2: on evidence

FROM MONTH 12

An 8x B200 or B300 server follows only past the API break-even of 7 to 13 billion tokens a day, with a committed partner such as CMU, or for a frontier-size model in Thailand.

Option B study

IN PARALLEL, FROM MONTH 6

A site study for 1 to 2 MW on the Korat and Chaiyaphum shortlist aims at certified Tier III and a PUE of 1.15 to 1.2; the building follows when anchor tenants commit enough load to approach 2 MW.

WHAT WOULD BRING OPTION B FORWARD

A committed anchor tenant, or a public partner such as NT, BDI or CMU that takes a large share of the capacity.

WHAT WOULD WEAKEN THE CASE FOR OWNING

A fall in API prices, or a Thai production API for the Thai models.

Sources: summary (our recommendation: a staged path); hardware (buy versus rent); site selection (ranked shortlist). The path and its triggers are our judgment; months count from the investor's go-ahead.

WHERE TO BUILD

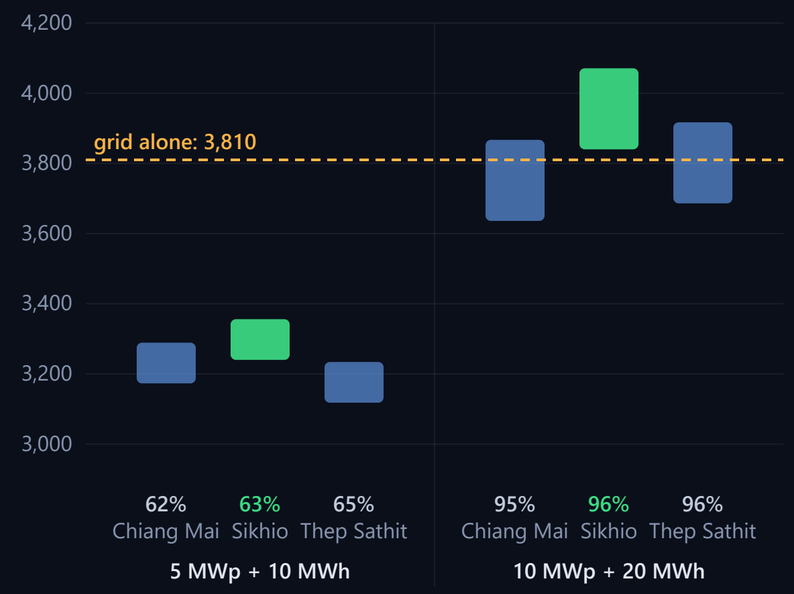
The Korat and Chaiyaphum wind belt is a better place for option B than Chiang Mai

	CHIANG MAI (MAE HIA)	C1 SIKHIO, KORAT	C4 THEP SATHIT, CHAIYAPHUM
Seismic zone	Zone 3	none	none
Wind capacity factor (model)	1%	21%	21%
Solar yield, kWh per kWp a year	1,418	1,452	1,499
Farmland listed, THB per rai	not found	about 300,000	about 129,000
Straight line to Bangkok	581 km	170 km	209 km
Covered hour by hour, 5 MWp + 10 MWh	62%	63%	65%
Covered hour by hour, 10 MWp + 20 MWh	95%	96%	96%

Our shortlist: (1) Sikhio and the Lam Takhong corridor, Korat; (2) Thep Sathit and Sap Yai, Chaiyaphum; (3) Muak Lek, Saraburi; (4) Khon Kaen; (5) the 304 Industrial Park, Prachinburi. Outside the EEC a qualifying data center gets 8 years of BOI tax exemption, and Chaiyaphum adds the extra benefits of a lowest-income province.

> **Two risks remain:** water, since the Lam Takhong reservoir held 16% in May 2025, and the pause of 166 data-center projects on September 4, 2026.

ENERGY COST PER MWH OF LOAD, THB, PER MW OF IT



Bars: low to high solar cost. %: share covered hour by hour.

Sources: summary; site selection (the candidates compared, how much on-site power can cover, ranked shortlist); model/site_results.md (1 MW of IT at PUE 1.2, PVGIS and NASA POWER hourly data 2019 to 2023). Cover and costs are our model's estimates; surplus solar counts as lost.

The first design is a 2 MW building on 150 rai at Sikhio, built in two phases



Site plan of the 150-rai plot, seen from the southeast; north points to the upper right. Concept render, not a construction drawing.

SITE C1, Sikhio and the Lam Takhong corridor in Nakhon Ratchasima: 150 rai with an option on 100 more, about 45 million THB at 300,000 THB per rai. Thep Sathit in Chaiyaphum is the fallback.

SCALE Designed and wired for 2.0 MW of IT, BOI's floor, at a PUE near 1.2 against BOI's limit of 1.3.

HALLS Hall A holds 1.4 MW of liquid-cooled GPUs in pods A1 and A2, hall B 0.4 MW of hosting and agent servers in 60 racks, and the core room 0.2 MW of storage and network.

PHASE 1 Fits out about 1 MW: pod A1 at 0.7 MW and 30 racks in hall B. Phase 2 follows when pod A1 is about 70% committed.

BUILDING One story of 80 by 40 m, with a two-story office and operations block on its south face and the staff parking beside it. A generator yard and a fuel farm stand at each end.

1.27	1.20	1.19
PUE at the opening load	phase 1 at full use	full build, 2 MW (our model)

Sources: first data center (In short, design basis, site and plot, building and halls, schedule); model/first_dc_results.md (PUE by phase). Render: docs/img/first-dc-site.svg. PUE is total facility power over IT power.

POWER AND COOLING

Two independent power paths (2N) reach every rack, and dry coolers cool the GPUs with warm water in a closed loop

GRID A 22 kV connection of about 3.5 MVA from the Provincial Electricity Authority (PEA) feeds two 3,150 kVA transformers, each able to carry the whole site, and lithium-ion UPS sits on both the A and the B path (2N).

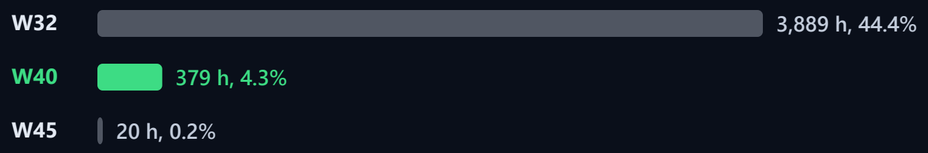
GENERATORS Three continuous-rated sets of about 1,600 kVA in a 2+1 arrangement, two of them in phase 1, meet BOI's call for the whole load plus a spare. They stand in two yards, one per power path: G1 at the west end of the building feeds board A, G2 at the east end feeds board B, and G3 joins G1 in phase 2. The A/B bus tie runs them in parallel.

FUEL Bulk diesel stays within the 15,000 L notification limit, as 7,500 L in each yard's bunded farm, plus 2,500 L built into each set. Phase 1 holds 20,000 L, about 64 hours at full load; the full build lasts about 36 hours at 2 MW.

SOLAR 10 MWp with a 20 MWh battery covers about 96% of phase 1 hour by hour. It is bought in tranches: at the opening load the full field would lose 69% of its output, while 2.5 MWp with 5 MWh would cover 73%.

COOLING Warm water at 40 °C (W40) runs straight to the chips and out to dry coolers, with small trim chillers for the hottest hours, and hall B is cooled by air on chilled water. No outside air enters; gas-phase filters and corrosion monitors guard against haze.

HOURS A YEAR WHEN DRY COOLERS NEED HELP, AT SIKHIO





Building cutaway, concept render: halls A, core and B from west to east, with the electrical rooms along the north wall.

Sources: first data center (power; cooling and air); facility systems; model/first_dc_results.md (hours above the dry-cooler limit, solar and battery coverage, generators and fuel). Render: docs/img/first-dc-building.svg. Hours, cover and fuel times are our model's estimates.

NETWORK AND SOFTWARE

Two separate fiber routes reach Bangkok, and the software is open except for NVIDIA's CUDA and firmware

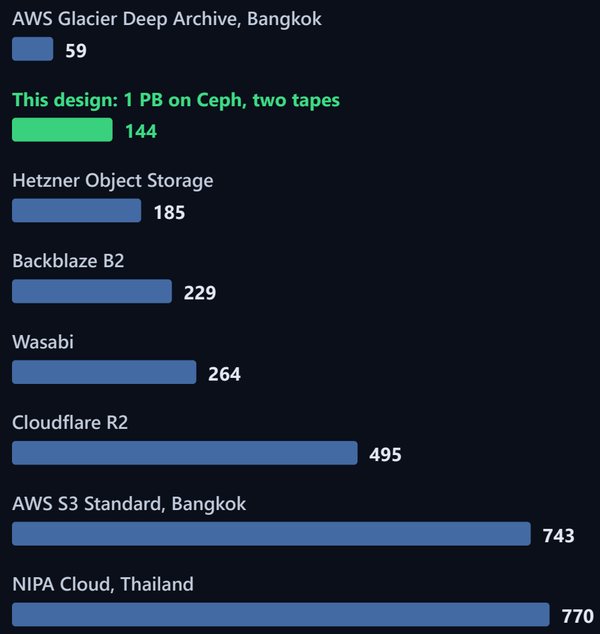
NETWORK

- Two 100G backhauls on separate rights of way, one a wavelength on the fiber of the Electricity Generating Authority of Thailand (EGAT) and one along the railway, end at a one-rack point of presence at **Telehouse Bangkok**, where six Thai exchanges meet.
- The investor holds **its own autonomous system (AS) number** from APNIC: a /23 of IPv4 with a /32 of IPv6 costs AUD 500 once and AUD 2,256 a year.
- Two transit carriers bring at least 10 Gbps each to the site, which with the backhauls meets BOI's four systems and 60 Gbps.
- We estimate the budget at **0.54 to 1.63 million THB a month** at opening, plus 1.8 to 3.9 million once, rising to 1.6 to 3.5 million a month at 2 MW.

SOFTWARE

- **Open software** runs nearly everything: NixOS and clan.lol on the hosts, Incus for tenant machines, SONiC on white-box switches, FRRouting at the edge, Ceph with LTO-9 tape, vLLM and SGLang behind LiteLLM, Kanidm, Keycloak and OpenBao, NetBox, and Prometheus with Grafana.
- **Closed software** remains only where nothing open works: CUDA and NVIDIA's user-space driver, the firmware of GPUs, drives and tape, and safety-certified controllers.
- NVIDIA AI Enterprise lists at USD 4,500 per GPU a year, about 28.5 million THB a year for the 192 GPUs of the reference fit-out, which the open stack avoids.

OBJECT STORAGE, THB PER TB-MONTH



Sources: first data center (network and connectivity; the software stack); connectivity (In short); open stack (a sized example: 1 PB usable plus two tape copies; storage benchmarks read October 2026). The 144 THB covers hardware and power only; staff, space and transit come on top.

WHAT THE BUILDING CAN SELL

The building can sell colocation, servers, private AI and GPU time, and BOI requires the hosting in any case

PRODUCT	INDICATIVE PRICE A MONTH (OUR ESTIMATE)	THE MARKET IN THAILAND AND THE REGION
HOSTING		
Colocation by rack, 3 to 5 kW	**at market rates**	True IDC sold a full rack with 3 kW at 38,800 THB a month (a 2025 promotion)
Shared-vCPU VPS	**from 490 THB (USD 15)**	NIPA Cloud charges 2,060 THB for 2 vCPU and 8 GB
Dedicated-vCPU VPS	**from 1,650 THB (USD 50); 8 vCPU with 32 GB for 5,300 THB**	the same 8 vCPU cost 6,599 THB at Hetzner Singapore and 8,240 THB at NIPA
Dedicated servers	**from 9,900 THB (USD 300)**	the Thai offers found are old Xeons and desktop Ryzens
AGENT SERVERS AND GPUS		
Agent server, Starter, up to 25 users	**13,000 THB (USD 394)**	25 ChatGPT Business seats at USD 20 cost 16,500 THB a month
Agent server, Business, up to 100 users	**54,000 THB (USD 1,636)**	100 such seats cost 66,000 THB a month
A whole 8x RTX PRO 6000 server	**467,000 THB (USD 14,152)**	Thai AI Cloud's managed tier starts at 1.12 million THB with four B300s
GPU time and tokens	**at Thai cloud prices**	Siam AI H100 at USD 2.69, Thai AI Cloud B300 from USD 8.90 per GPU-hour

about 71%

of one RTX PRO 6000 agent-server pod must be sold to break even (high case).

40% to 68%

of capacity sold before the VPS lines break even, depending on what the hardware costs.

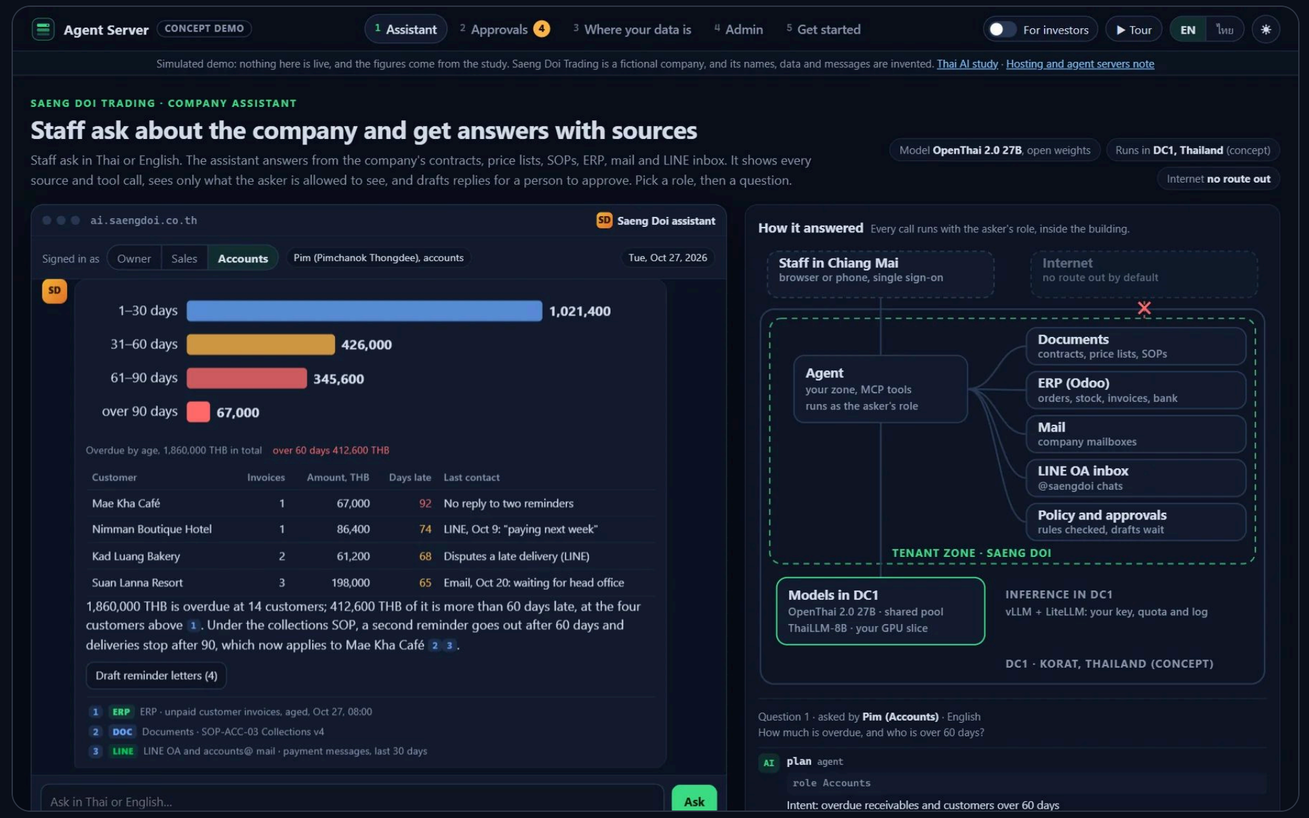
56% to 94%

for the dedicated servers, priced against the regional market.

> Memory contract prices rose 93% to 98% in the first quarter of 2026, and memory is now about half the cost of a host, so the server margins are thin. A BOI-promoted data center must offer colocation, managed services, backup and data hosting.

Sources: summary; hosting and agent servers (what the investor could sell; packages and indicative prices); VPS and dedicated servers (In short, packages); model/hosting_results.md; model/server_packages_results.md. Agent-server and VPS prices are 1.5 times the high-case cost, and dedicated servers are priced against the market; none is a quote.

An agent server runs a company's data, its AI agents and the model in a separate zone inside the data center



From the simulated agent-server demo: a clerk at a fictional Chiang Mai firm asks which customers are overdue; the answer cites its sources, and the letters wait for approval. Nothing is live, and the data are invented.

One zone per company

Each customer gets its own virtual machines on its own network segment, with no route to the internet unless it allows one.

Agents work through tools

Agents reach mail, files and the accounting system through tools inside the zone, using the Model Context Protocol (MCP), and the models only through a gateway that holds the customer's key, quota and log.

People approve

Staff sign in with the company's own single sign-on, the agent sees only what the asker may see, and its drafts wait for a person.

The honest limit

A model reads only what it can decrypt, so data in use is visible to the host unless the GPU runs in confidential-computing mode.

Why it can sell

AWS has run a Thai region since January 2025, so residency alone protects nothing. What no hyperscaler offers in Thailand is GPUs next to the data, Thai models and people who wire agents into a company's systems (our judgment).

Sources: hosting and agent servers (why data and AI belong in the same building; the agent server: how it would work; licenses). Screenshot: the agent-server demo on thaiai.lol, simulated.

WHAT IT COSTS

Phase 1 costs 0.86 to 1.27 billion THB before any GPU, and solar and battery are about a third of it

phase 1, low to high, million THB all our estimates



Phase 1, facility	**860 to 1,269 M THB**	USD 26 to 38M
Full build, 2 MW	**1,415 to 2,029 M THB**	USD 43 to 61M
Example IT fit-out for phase 1	**266 to 482 M THB**	USD 8 to 15M

RUNNING COSTS A YEAR, MILLION THB

CASE	OPEX	PUE
Opening load, 0.39 MW of IT	42 to 88	1.27
Phase 1 at full use, 1 MW of IT	44 to 91	1.20
Full build, 2 MW of IT	74 to 146	1.19

Staff and connectivity are about 60% of phase-1 opex at full use, while grid power after solar is under a tenth.

THE EARLY YEARS ARE THE WEAK POINT

At the opening load the halls are half empty. PUE rises to 1.27, each of the two generators would run at 19%, and the full 10 MWp would lose 69% of its output, so energy would cost 8,993 to 10,963 THB per MWh against 3,810 from the grid. Buying the solar in tranches and filling pod A1 first answer most of this.

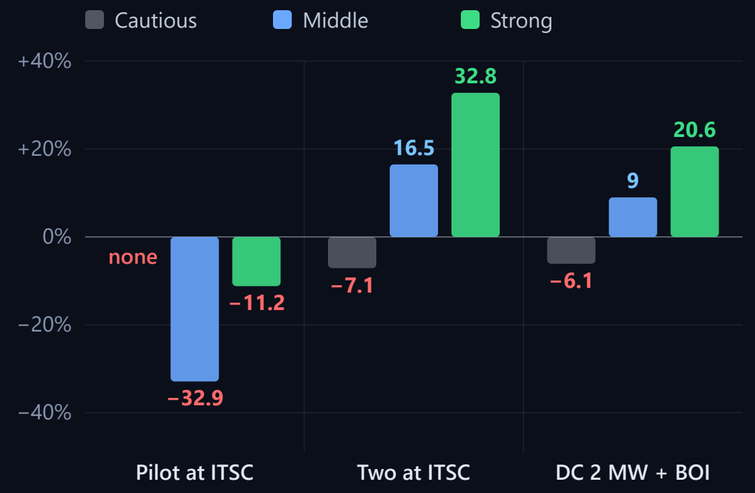
Sources: first data center (cost; risks; what this means for the investor); model/first_dc_results.md (capex by system, opex, energy check, sensitivities). The building and its mechanical and electrical plant use option B's 321 to 508 million THB per MW of IT; the IT itself is not in these figures.

WHAT IT COULD EARN

The 2 MW case returns 9% only at middle demand, about 664 billion tokens a day, and the ITSC pilot never pays back

	STAGE 1: PILOT AT ITSC	STAGE 2: TWO SERVERS AT ITSC	OPTION B: DC 2 MW WITH BOI
FLEET	one 8x RTX PRO 6000	one 8x RTX PRO 6000 and one 8x B200	12 GB300 NVL72 racks, BOI tax exemption
UPFRONT	8.3M THB USD 253k	29.7M THB USD 901k	3.76bn THB USD 114M
PAYBACK, CAUTIOUS / MIDDLE / STRONG DEMAND	beyond year 10 in all three	beyond year 10 / 3.6 yr / 2.6 yr	beyond year 10 / 7.8 yr / 3.5 yr
INTERNAL RATE OF RETURN (IRR), 10 YEARS, AFTER TAX	none / −32.9% / −11.2%	−7.1% / 16.5% / 32.8%	−6.1% / 9% / 20.6%
TOKENS A DAY THE API MUST SELL, MIDDLE DEMAND	about 246 M	about 6.4 bn	about 664 bn
UTILIZATION AT WHICH NET PRESENT VALUE IS ZERO	above 100%	44%	56%

IRR BY DEMAND CASE, 10 YEARS AFTER TAX



No Thai market data backs token volumes like these. The 2 MW preset prices the facility at option B's rate per MW with placeholder solar; it does not follow the first design's phases.

The next slide shows what VPS and dedicated servers add, on their own in a Bangkok rack and inside the 2 MW case.

Sources: the calculator on thaiai.lol, presets "Pilot at ITSC," "Two at ITSC" and "DC 2 MW + BOI," read from the page on October 28, 2026 at its default settings (midpoint prices, 10 years, PEA tariff) for each demand case. Demand is the decisive unknown; every utilization is an estimate.

NORMAL HOSTING

VPS and dedicated servers in hall B would add 64.7 million THB of revenue a year and leave the 2 MW case's IRR at 9%

	START: SERVERS IN A RENTED BANGKOK RACK	GROWTH: SERVERS IN HALL B OF THE 2 MW CASE
FLEET	8 VPS hosts and 22 dedicated servers, no GPUs, possible from 2027	40 VPS hosts and 110 dedicated servers beside the GPUs
UPFRONT	30.3M THB USD 917k, the customer portal included	+145M THB USD 4.4M, added to the case's 3.76bn THB
SERVER REVENUE A YEAR AT 50 / 70 / 90% SOLD	9.2M / 12.9M / 16.6M THB	46.2M / 64.7M / 83.2M THB
SERVER EBITDA A YEAR AT THE SAME SHARES	2.7M / 6.3M / 9.8M THB	24.3M / 42.1M / 59.9M THB
IRR, 10 YEARS AFTER TAX, CAUTIOUS / MIDDLE / STRONG	−22% / −1.4% / 10.9%	−6.2% / 9% / 20.6% for the whole case, against −6.1% / 9% / 20.6% without the servers
NET PRESENT VALUE AT 10%, SAME CASES	−35.6M / −15.5M / +1.3M THB	−107M / −8M / +89M THB, the change to the case
SHARE SOLD AT WHICH THE SERVERS ADD ZERO NPV	88%	about 72%

about 6%

of the 2 MW case's revenue would come from the servers at middle demand, for under 4% of its upfront cost.

8 times

the revenue of colocation per kW of peak draw at 70% sold: 108,396 against 12,933 THB a month.

3,773

IPv4 addresses needed at 70% sold in hall B, against the 450 that APNIC leaves for customers, so IPv6-only plans become the default and the rest are leased.

> Memory sets the margin: on the server note's own basis the hall B fleet breaks even at 50% sold with memory at index prices and at 75% at retail. The rescue and reinstall engine, 0.96 to 2.7 million THB of engineering (est.), is not yet in the 1.5 million THB portal budget.

Sources: the calculator on thaiai.lol, presets "Servers in a Bangkok rack," "DC 2 MW + BOI" and "DC 2 MW + servers," read on October 28, 2026 at its default settings for each demand case, which sets the share of server capacity sold at 50, 70 and 90% (est.); the changes to the 2 MW case are the difference between the last two presets. VPS and dedicated servers (A ramp in hall B; IPv4 and IPv6; What this means for the investor); rescue system (cost).

BUILDING IT IN THAILAND

Most of the plant is made in Thailand; GPU servers, generators, batteries and large UPS come through Thai dealers

MADE IN THAILAND

Seven Thai makers of 22 and 33 kV cable are on PEA's list. Ekarat builds transformers of up to 30,000 kVA, Nidec coolant distribution units in Ayutthaya, CIG BluSolutions a dry cooler rated for 40 °C water, Trane chillers in Samut Prakan, and Cheval racks in Chiang Mai.

IMPORTED THROUGH THAI DEALERS

GPU servers, battery containers, generator sets, the large UPS, the air handlers for hall B and most fire suppression. For these, the local partner's service matters more than the brand.

SERVICE NEAR THE SITE

We found no service center for generators, UPS or switchgear in Nakhon Ratchasima or Chaiyaphum. Contracts should fix response times, and the site should hold its own critical spares.

LONG LEAD TIMES

Coolant distribution units take 30 to 50 weeks and large chillers 40 to 70 (US quotes, May 2026); hard disks have been on backorder for up to two years. The electrical and cooling plant is ordered in the second quarter of 2027.

SCHEDULE TO THE FIRST IT LOAD (OUR ESTIMATE)



Sources: sourcing (In short; IT, network and construction); first data center (schedule; risks). Thai lead times were not published. The schedule is our estimate; the critical path runs through the grid connection and the long-lead electrical plant.

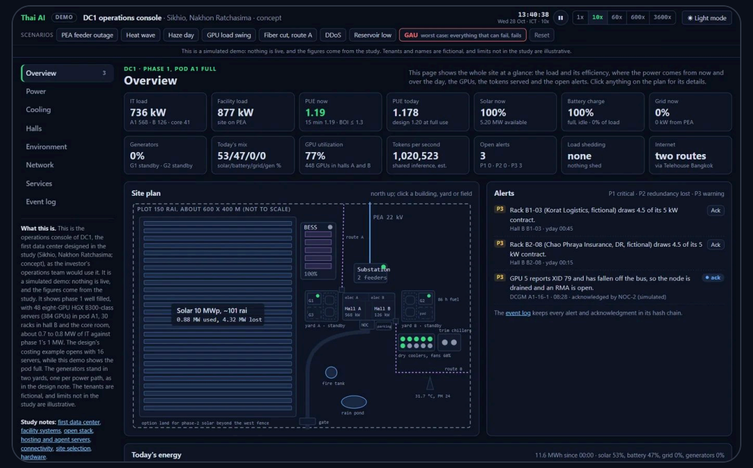
RISKS AND HOW THE PLAN ANSWERS THEM

The largest risks are the new data-center rules, grid capacity at the site and demand that nobody has measured yet

RISK	EVIDENCE	HOW THE PLAN ANSWERS IT
Approval pause and new rules	166 projects paused on September 4, 2026; new minimum rules and seven tiers by demand, with applications from early 2027	apply in the first window; design to the stricter rules now; a small first tier
Grid connection	data-center requests of about 35,000 MW against 8,800 MW in the draft power plan; a 4.5M THB guarantee per reserved MW proposed	3.5 MVA is small; solar and battery cut grid energy; a PEA inquiry in the first quarter of 2027
Fuel permits	the Department of Energy Business (DOEB) reported in September 2026 to have suspended new fuel-tank permits for data centers	bulk fuel kept under the 15,000 L notification limit in two farms; farms laid out for a later license; a refueling contract
Water	the Lam Takhong reservoir held 16% in May 2025	dry coolers and closed loops, no evaporation, a rainwater pond; water rights checked before purchase
Low early load	at the opening load, PUE 1.27, generators at 19% and 69% of the solar output lost (our model)	fill pod A1 first; load-bank tests; solar and battery bought in tranches
GPU and memory prices	H200 rents up 143% and B200 111% in 2026 (news report); memory contract prices up 93% to 98% in the first quarter of 2026; B300 lead times of 30 weeks	IT bought as demand arrives; the pilot first
Demand and competition	no paid Thai model API exists to measure demand against; AWS in Thailand since January 2025; Thai AI Cloud sells private AI from USD 34,000 a month	stage the money; hall B sells hosting and agent servers if GPU demand lags
Staff	thin talent outside Bangkok; BOI's rule on Thai staff in senior posts	hire and train at the ITSC pilot; Suranaree University nearby; support contracts for the open stack

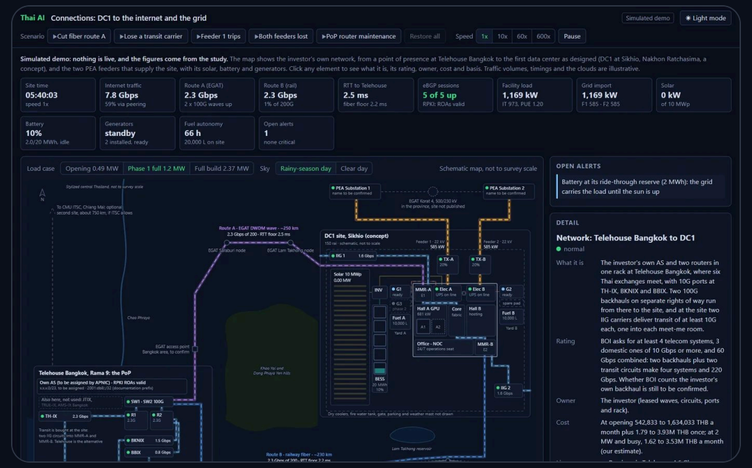
Sources: first data center (risks and how the design answers them); summary; site selection (grid, water and hazards); VPS and dedicated servers (In short); hosting and agent servers (In short); Thai models (hosted APIs today).

Three simulated demos show how the design would run: the operations console, the network and an assistant for the staff



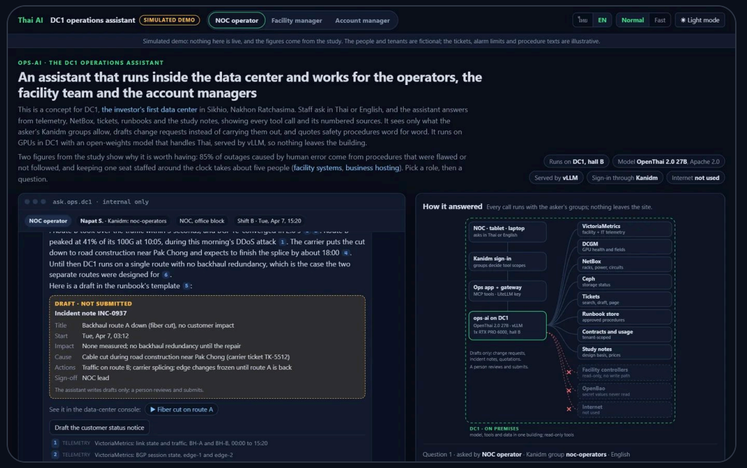
Operations console

Phase 1 with pod A1 full: power, cooling, halls, outside feeds, the network and a log that cannot be rewritten, with failure scenarios and a drill in which everything fails at once.



Connections

The point of presence at Telehouse Bangkok, both fiber routes, the transit circuits and the two PEA feeders, with five failure scenarios to play.



Operations assistant

Answers operators, the facility team and account managers in Thai or English from telemetry and runbooks, and drafts changes for a person to approve.

TRY THEM The demo hub on thaiai.lol also holds the agent server, the hosting business's website, customer portal and staff backend, the calculator and the 3D model, side by side as tabs.

Sources: the demos on thaiai.lol (console, connections, operations assistant), all simulated: nothing is live, tenants and names are fictional, and the figures come from the study; first data center; connectivity; facility systems.

WHAT THIS MEANS FOR THE INVESTOR

The next decisions buy no hardware: the preparation phase, talks with ITSC, BDI and NT, the 2 MW question and site checks

DECISIONS NOW

1. **Approve the preparation, months 0 to 3:** define the workloads and customers, test the shortlisted models on free APIs and rented GPUs, and have a Thai lawyer read the Typhoon terms.
2. **Authorize talks** with CMU ITSC about colocation, with BDI about a Chiang Mai role in ThaiLLM, and with NT, which serves ThaiLLM and runs a Chiang Mai site.
3. **Decide whether option B is planned for 2 MW** and BOI's incentive or gives it up; option A avoids the question.
4. **Fund the site checks at C1** in the first quarter of 2027, so that applications can go in when Thailand reopens them.

FIRST STEPS, WHICH B4L CAN PREPARE FOR THE INVESTOR

- Ask ITSC about power per rack, spare backup capacity, 24/7 access, transit and the terms of the agreement.
- Ask BDI and NT who hosts ThaiLLM and where, and whether they want an inference partner in Chiang Mai.
- Get quotes for the first server from a Thai distributor and by import, with duty, VAT and on-site warranty.
- Ask DOEB about the fuel-tank suspension and PEA about a second feeder, and get quotes for continuous-rated generators and SONiC switches.
- Make land and grid inquiries in Nakhon Ratchasima and Chaiyaphum, and look into water rights.

OPEN QUESTIONS IN THE DESIGN

- PEA must say how much spare 22 and 33 kV capacity reaches the C1 parcel, whether a second feeder is possible, and whether the proposed guarantee of 4.5 million THB per MW applies at 3.5 MVA.
- Nobody yet knows which of the seven new tiers a 2 MW site falls in, or whether a phased build qualifies.
- DOEB must say whether the fuel-tank suspension is formal and whether it counts the two fuel farms together against the 15,000 L limit.
- Water for a closed-loop site, the title, the zoning and the 100 m setback for ground solar all need checking at the parcel.

> The study, its notes, the calculator, the demos and all 1,281 sources are at thaiai.lol.

Sources: summary (our recommendation; next steps); first data center (what this means for the investor; open questions); Thai compute (what this means for the investor). The order of the decisions is our judgment.