NURA AI

EXCHANGE NETWORK

NURA data center | AIX network & token

An open AI compute marketplace anchored by modular infrastructure in Abu Dhabi, designed for public access, independent GPU participation and token-coordinated network growth.

Strategic presentation

Al Ain, Abu Dhabi

The investment thesis

The project is deliberately positioned as a public compute economy, managed by de-centralized framework.

Central Holding Group LLC will lead NURA as a private commercial platform. The physical anchor is NURA, the Central Holding Group data-center and managed compute brand: a sponsor-provided Al Ain site with a target 20 MW power allocation, beginning with a 4 MW modular AI campus. The network anchor is AIX, the token and decentralized compute protocol that opens participation to independent GPU hosts.

The business is built to create real usage before token scale: NURA sells AI services and GPU access; data-center GPUs provide dependable capacity; independent hosts add low-cost and long-tail supply; AIX token coordinates staking, rewards, reputation and settlement.

4 MW
Phase-I site power target

20 MW
Expansion power application

40,000 m²
Approximate land area

0.10 AED
Target kWh electricity input

256–512
Initial high-end GPU fleet

nuraix.com
Public marketplace and AI portal

Not a smaller Stargate

National AI campuses validate regional demand, but NURA serves a different market: open, retail-accessible and participant-owned compute.

Government-level sovereign AI projects are built for national capability, strategic model access, regulated enterprise use and geopolitical infrastructure. Their scale is important, but their customers, controls and participation model are not the same as a public compute market.

NURA is designed for the segment these mega-projects do not naturally serve: developers, creators, small companies, AI agents, GPU owners, studios and communities that want direct access to open models and rentable compute without enterprise sales cycles.

Sovereign AI programs

Centralized funding
National and enterprise demand
Strategic model controls
Limited public participation
Infrastructure as state capability

NURA

Open public marketplace
Independent host participation
Open-source model choice
Tokenized incentives
Infrastructure as economic network

OpenAI and G42 public Stargate announcements; management positioning

AI compute is abundant, but access is fragmented

Cloud capacity exists, open-source models exist and idle GPUs exist. What is missing is a trusted market that connects all three.

The current market forces users into separate channels. Enterprise GPU clouds sell managed virtual machines. GPU marketplaces sell raw capacity. Open-model applications sell consumer AI. Decentralized networks sell tokenized infrastructure. NURA combines these use cases into one product path: use AI, rent compute, list a GPU host and settle usage through the AIX economy.

Three disconnected markets today

Market	What it sells	What remains missing
GPU cloud	Managed GPU instances	Difficult for small suppliers to participate
GPU marketplace	Independent host rentals	No strong consumer AI demand engine
AI apps	Open or closed model access	Users cannot earn or provide compute
DePIN compute	Protocol incentives	Often weak user experience and weak SLA

NURA is built as the missing bridge: one account for consumer AI, dedicated GPU rental, independent host listing, staking, credits and settlement.

RunPod pricing pages; Vast.ai hosting documentation; Akash documentation

NURA and AIX must stay distinct

NURA is the data-center and managed compute brand. AIX is the token and decentralized compute network that expands supply beyond NURA-owned infrastructure.

The revised positioning separates the physical brand from the protocol brand. NURA should represent the Al Ain AI campus, the managed GPU cloud, the customer portal at nuraix.com, service quality and commercial operations. AIX should represent the open network: staking, host collateral, routing incentives, usage settlement, node rewards and decentralized governance.

This distinction matters for investors and regulators. NURA can be valued and financed as an infrastructure and cloud-services company. AIX can be designed as a utility network that coordinates independent compute supply without promising ownership of NURA assets.

NURA

Data center brand

- Al Ain modular AI campus
- Managed GPU cloud and API service
- nuraix.com customer portal
- Enterprise contracts, SLA and billing
- Owned H100/H200/B200 capacity

AIX

Token + decentralized AI compute network

- Host staking and collateral
- Independent GPU node onboarding
- Verified usage, settlement and rewards
- Network reputation and slashing
- Governance over marketplace parameters
- Accessible through NURA Portals, Apps or Terminals

NURA provides trusted base capacity. AIX opens participation and expands the compute market beyond one data center.

NURA creates the demand engine

The NURA service layer creates recurring workload that monetizes both the owned data center and the AIX open compute network.

Every chat request, image generation, voice task, video job and AI agent workflow becomes paid demand. The routing layer sends each request to the most suitable endpoint: the company-owned data center for reliability, enterprise partner nodes for high capacity, or community GPUs for low-cost and flexible jobs.

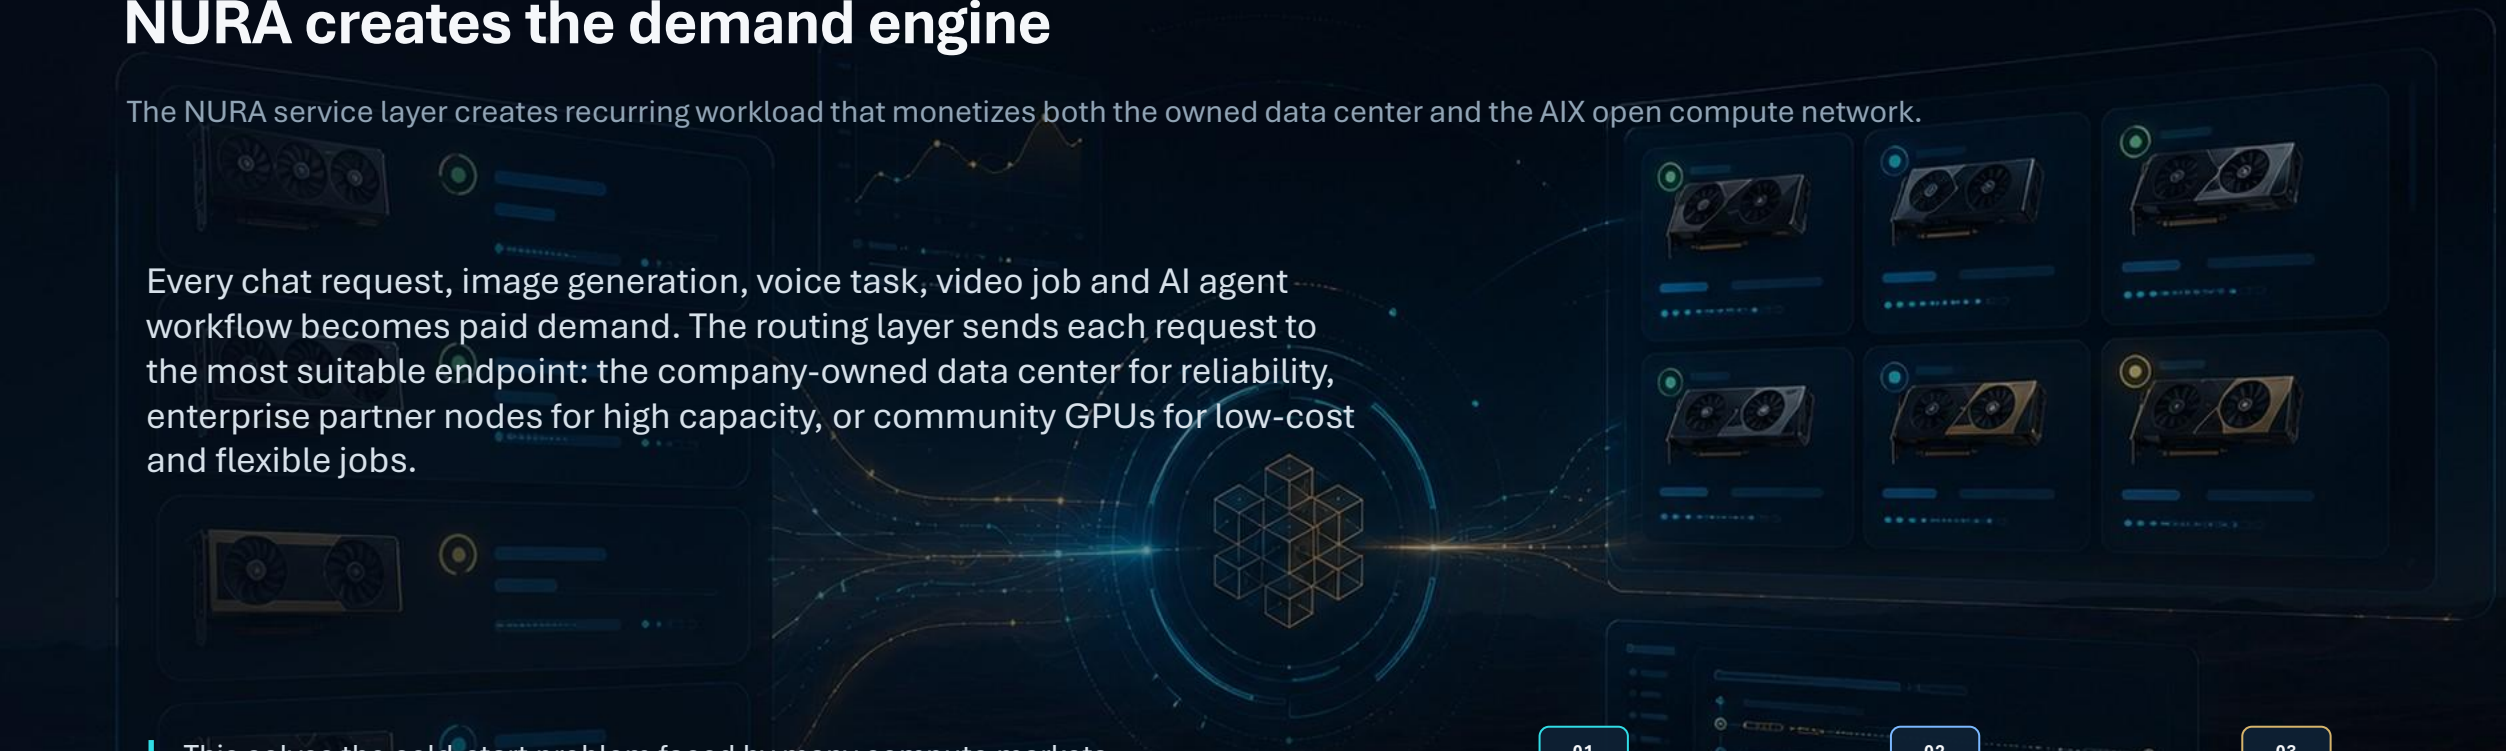

> This solves the cold-start problem faced by many compute markets. Independent hosts do not join only because a token exists; they join because the consumer product creates real jobs, visible income and a reason to keep hardware online.

01

Use AI

One NURA account for models and tools.

02

Route workload

Scheduler chooses the best warm endpoint.

03

Settle usage

Usage is billed and host rewards are calculated.

From owned GPUs to open nodes

AIX should not pretend that all compute is equal. The network must price and route different tiers differently.

The Al Ain data center provides the trusted base load: high-SLA GPUs, controlled cooling, stable networking, and enterprise-grade reliability. Independent supply then expands the curve with lower-cost 4090/5090 hosts, professional workstations, enterprise spare capacity and partner data-center racks.

The scheduler should route by task type, trust tier, GPU memory, model cache, uptime, latency, price, security rating and user preference. A 5090 node is not a training cluster, but it is highly useful for 7B–35B quantized LLMs, image generation, TTS and many batch jobs.

Tier	Supplier type	Best workloads	Trust model
S	NURA data center	Enterprise APIs, private endpoints, H200/B200 workloads	High SLA, full control
A	Enterprise / colocation partners	Longer jobs, batch inference, dedicated rentals	Certified hardware and contracts
B	Prosumer 4090/5090 nodes	Images, TTS, small LLMs, agents, low-cost batch	Stake, rating and sandboxing
C	Experimental spot nodes	Cheap interruptible tasks and test workloads	Low SLA, low price

> **The technical goal is not to merge many home GPUs into one supercomputer. The goal is intelligent task routing across many single-node and multi-node compute endpoints.**

Nura AI Exchange Network Framework

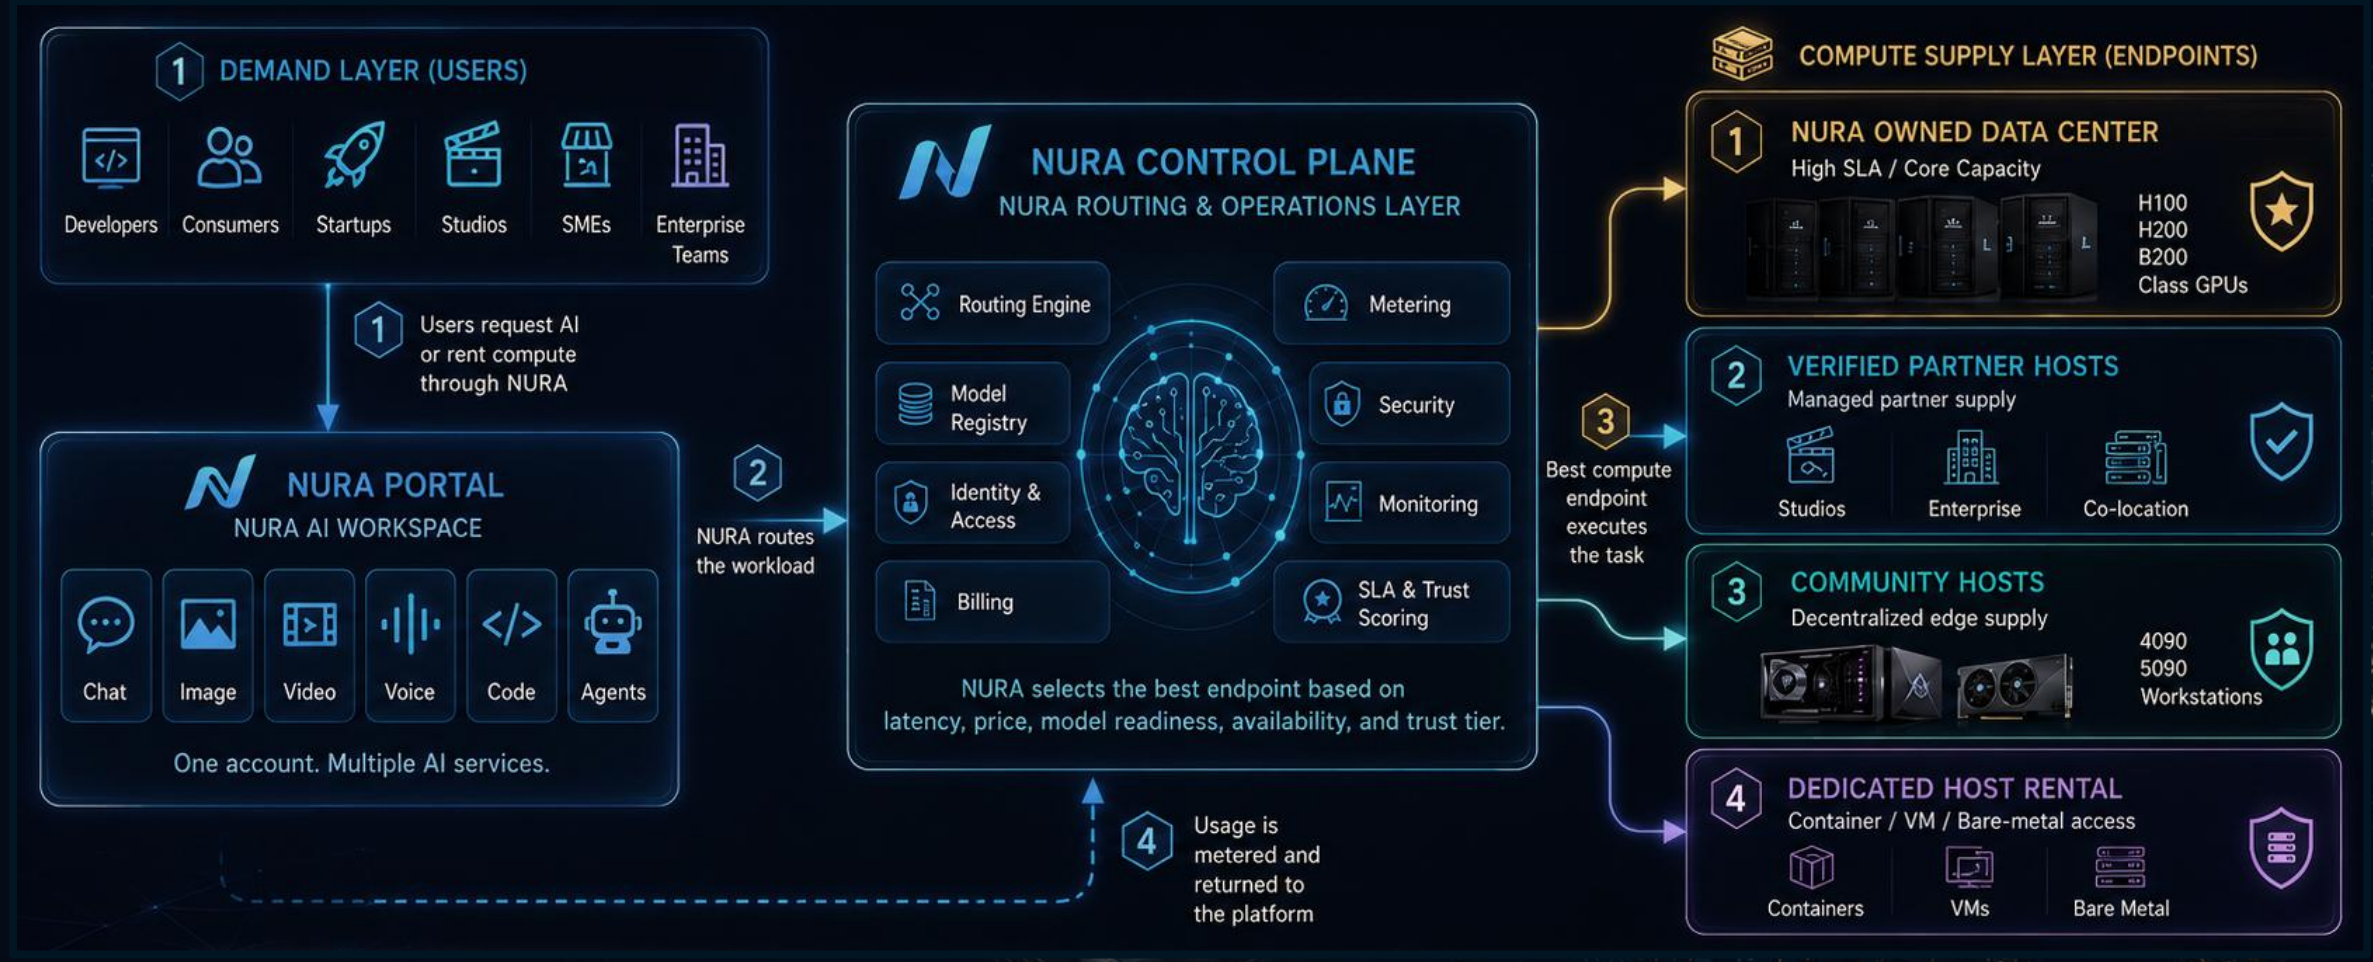

Al Ain modular AI campus

The initial land parcel is described as approximately 200 m × 200 m, or about 40,000 m^2. This is suitable for a phased modular campus: incoming electrical infrastructure, compact AI halls, liquid-cooling plant, security, network rooms, maintenance lanes and future expansion space.

The recommended master plan is to build Phase-I at 4 MW while preserving trenching, transformer bays, cooling loops and fiber paths for the 20 MW end state. The commercial objective is to prove NURA operations first, then scale AIX participation after real demand exists.

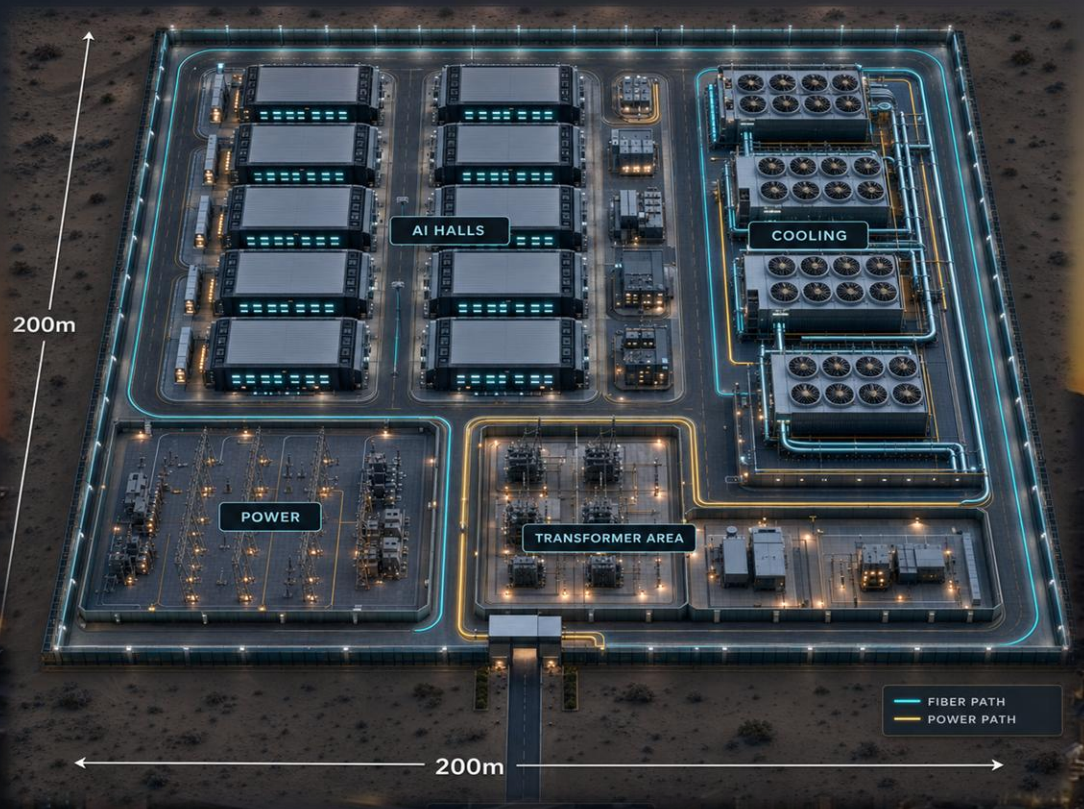

200 × 200 m	4 ha	4 MW	20 MW
Plot geometry	Site area	Phase-I	End state

Provided satellite image and coordinates; management assumptions

Build a profitable 4 MW proof point

The first campus is not a shell-only data center. It must include enough owned GPU capacity to prove the NURA service layer and revenue model.

Phase-I should be divided into two linked workstreams. Infrastructure investors fund power distribution, modular halls, cooling, network rooms, fire systems and security. Compute investors fund the first GPU fleet. The operating company sells NURA managed compute, GPU rental and AIX marketplace access.

A staged commissioning model reduces risk: energize the first 0.5–1 MW IT tranche, deploy 256–512 high-end GPUs, test billing and utilization, then expand rack modules as demand and financing improve.

01	02	03
0.5–1 MW	**4 MW**	**20 MW**
MVP IT tranche	**Phase-I utility capacity**	**Master-plan expansion**
256–512 GPUs, NURA API, first customers	~3.0–3.2 MW IT load at target PUE	Additional halls, partners and GPU financing

> **The 4 MW campus is the credibility layer that makes the token economy believable. It proves that AIX is backed by real power, real machines and real user demand.**

Engineering and financial planning assumptions

Liquid-ready from day one

Desert AI infrastructure must be designed around heat, dust, serviceability and warranty. Retrofitting consumer assumptions later is expensive.

The site will not copy a crypto-mining container layout without modification. AI servers carry higher power density, more expensive networking, stricter uptime requirements and more sensitive airflow needs. The right approach is modular, but with professional AI data-center controls.

Phase-I will support two zones. The air/hybrid zone handles RTX PRO, L40S and lower-density H100/H200 systems through hot-aisle containment, filtration and rear-door or in-row cooling. The liquid-ready zone handles high-density H200/B200 and future GB-class systems using direct-to-chip liquid cooling, CDUs, dry coolers and N+1 pump loops.

Recommended design choices

01 Buy OEM liquid-ready servers where possible; avoid warranty-risk field conversions for production clusters.

02 Use positive-pressure filtered air for the remaining non-liquid components: PSUs, SSDs, switches and service aisles.

03 Prefer dry or hybrid cooling loops; add chiller support only for peak desert conditions or specific SLA zones.

04 Track PUE, coolant temperature, inlet temperature, dust load and failure rates from the first rack.

NVIDIA DGX/HGX power references; ASHRAE AI data-center design principles

Power is the strategic margin

If the target AED 0.10/kWh tariff is secured in writing, electricity becomes a durable advantage against many commercial compute locations.

At 4 MW continuous facility power, annual energy use is 35.04 GWh. At AED 0.10/kWh, the annual electricity bill is about AED 3.5M. At 20 MW, the annual bill is about AED 17.5M. This is meaningful, but small relative to GPU rental revenue if utilization is achieved.

The risk is not the arithmetic; the risk is whether the project can contract the tariff, secure connection capacity, maintain uptime and finance GPU assets fast enough to fill the power envelope.

Case	Facility power	Annual kWh	Power cost @ AED 0.10/kWh	Indicative role
Phase-I	4 MW	35.04 GWh	AED 3.50M	Proof campus and first GPU cloud
Expansion	20 MW	175.20 GWh	AED 17.52M	Regional compute campus
Public commercial reference	20 fils/kWh	—	2× target cost	Shows value of sponsor tariff
Public government reference	29.4 fils/kWh	—	2.94× target cost	Shows margin sensitivity

> **The investment case contains a long-term written power agreement, connection timeline, tariff treatment, and expansion rights before major GPU procurement.**

ADDC 2025 tariff page; management power assumptions

Use the right GPU for each market

AIX should avoid a single-GPU narrative. Different customers require different price, memory, SLA and isolation levels.

The first fleet should mix enterprise accelerators and lower-cost inference GPUs. H100/H200 class systems provide credible enterprise capacity and support high-memory workloads. B200 should be introduced through liquid-ready racks where capital and supply permit. RTX 5090/RTX PRO supply is attractive for marketplace inference and creative workloads, especially when hosted by independent participants.

GPU class	Role in NURA	Strength	Constraint
H100 / H200	Owned data-center cloud	Reliable high-SLA inference, fine-tuning and enterprise API capacity	High capital cost and supply planning
B200 / Blackwell	Premium liquid-cooled zone	Large-model throughput and next-gen enterprise demand	Power density and cooling complexity
RTX PRO / L40S	Managed inference pool	Image, video, rendering and medium model workloads	Lower training suitability
4090 / 5090 community	Open host marketplace	Low-cost small models, image, voice, agents and batch jobs	Lower SLA; requires sandbox and reputation

> **The best margin comes from matching workload to hardware. Not every job needs a B200. Many consumer and creator tasks can run profitably on lower-cost GPUs routed through NURA.**

NVIDIA public specs; RunPod pricing pages; management assumptions

The control plane is the company moat

The data center is important, but the durable software advantage is routing, sandboxing, billing and verified settlement.

Every external host runs AIX Node Agent. The agent verifies hardware, installs Docker-based runtimes, joins an encrypted tunnel, reports uptime, downloads approved model images, starts isolated jobs and uploads usage proofs. Home nodes do not need public IPs; they connect outward to the AIX gateway.

The central router decides where to send work based on GPU type, VRAM, model cache, price, latency, trust tier, customer privacy requirements and current utilization. Marketplace rentals receive temporary credentials through the AIX Gateway, not permanent host passwords.

Security design principle: customers receive temporary, revocable access to isolated environments; hosts receive settlement only after verified uptime and usage.

Renting a host must be safe

A public host market works only if access, payment, data cleanup and dispute handling are standardized by the platform.

The marketplace should look simple to users, but the transaction should be controlled end-to-end. A host listing shows GPU type, VRAM, CPU, RAM, storage, region, price, minimum rental time, online rate, success rate, trust tier and allowed environments. After payment, the platform provisions access automatically.

This avoids the dangerous model of direct password exchange. NURA should use temporary SSH keys, container credentials, Jupyter URLs, VPN tunnels or reverse proxy access, then revoke credentials and wipe the environment at the end of the rental period.

01 Browse

Compare hosts by GPU, price, region and reliability.

02 Escrow

Payment locks until access and runtime are verified.

03 Provision

Agent creates isolated VM/container and short-lived credentials.

04 Use

Customer runs models, Docker workloads or notebooks.

05 Settle

Host receives payout; reputation updates automatically.

> **The experience can resemble an NFT marketplace visually, but the asset is usable time on verified compute, not a speculative picture.**

NURA COMPUTE MARKETPLACE

A marketplace for real GPU machines, not abstract cloud SKUs

NURA lists verified hosts and data-center capacity, while **AIX** coordinates trust, staking and network incentives.

NURA | Marketplace

LIST YOUR GPU

Find compute

GPU, region, VRAM, price, model...

Available now

8x H200 SXM

NURA CORE CAPACITY

NURA • Abu Dhabi

1,128 GB VRAM • 400G fabric

99.95% uptime

$34.80 / hr RENT

4x RTX PRO 6000

VERIFIED PARTNER

Studio DC • Dubai

384 GB VRAM • 100G network

99.91% uptime

$7.60 / hr RENT

1x RTX 5090

COMMUNITY HOST

Independent • EU

32 GB VRAM • 1G network

98.7% uptime

$0.69 / hr RENT

FROM ORDER TO ACCESS

1. **Select**
   Choose machine, duration and environment
2. **Escrow**
   Pay by card, bank transfer, USDC or AIX
3. **Isolate**
   NURA provisions a container, VM or dedicated host
4. **Access**
   User receives an ephemeral SSH key or private URL
5. **Settle**
   Metered usage and automated host payout

NURA is the marketplace experience; **AIX** is the economic layer that secures and scales the network.

NURA | AIX

Freedom with operational discipline

The brand promise should be broad model choice and minimal centralized gatekeeping, while still protecting the network from illegal abuse and infrastructure attacks.

Product mode	User freedom	Platform responsibility
NURA Managed AI	Simple public access to supported text, image, video, voice and agent tools	Billing, model operations, basic abuse controls
Private endpoint	Dedicated model instance, private data space and user-selected open model	Isolation, uptime, access control and logging
Raw compute rental	Customer deploys own containers, notebooks or applications	Infrastructure, metering, revocation and legal minimum controls

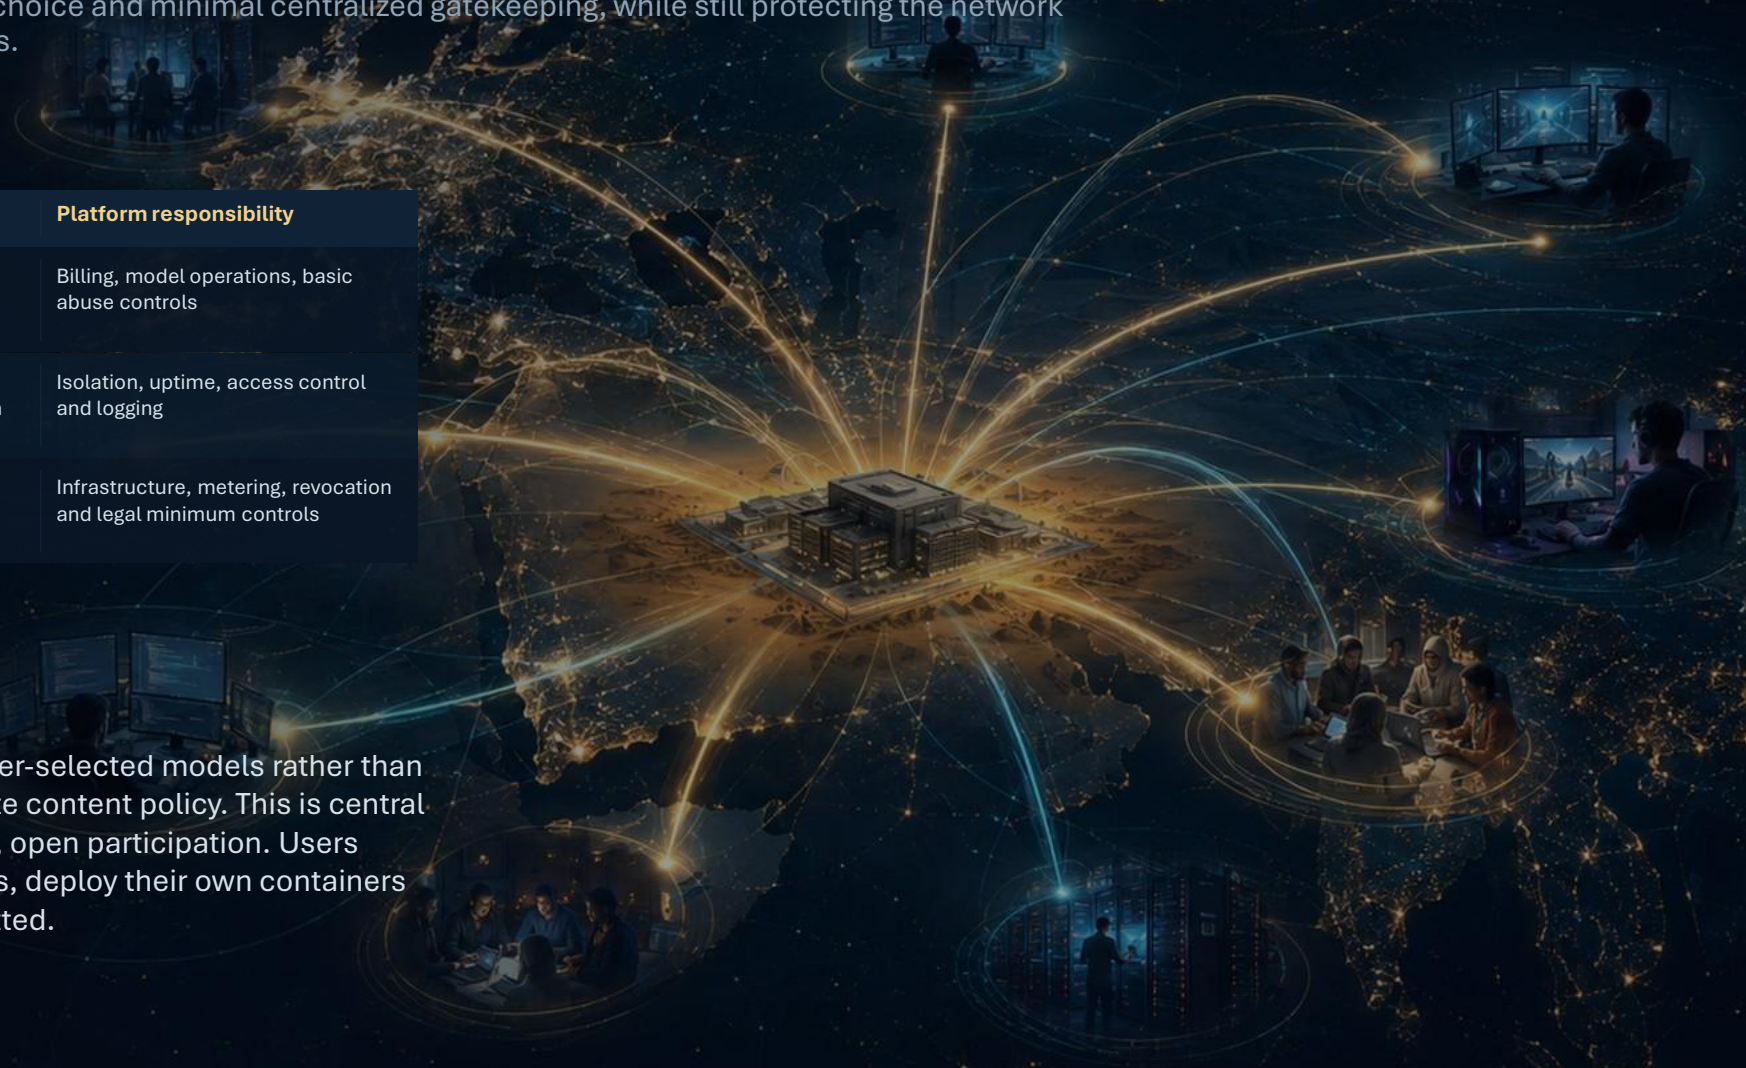

NURA should support open-source and user-selected models rather than forcing all customers through one corporate content policy. This is central to the brand: open compute, open models, open participation. Users should be able to choose private endpoints, deploy their own containers and rent raw compute where legally permitted.

NURA AI PORTAL

The consumer portal converts everyday AI usage into recurring network demand

Users interact with NURA; the platform routes workloads across NURA infrastructure, verified partners and community hosts in the AIX network.

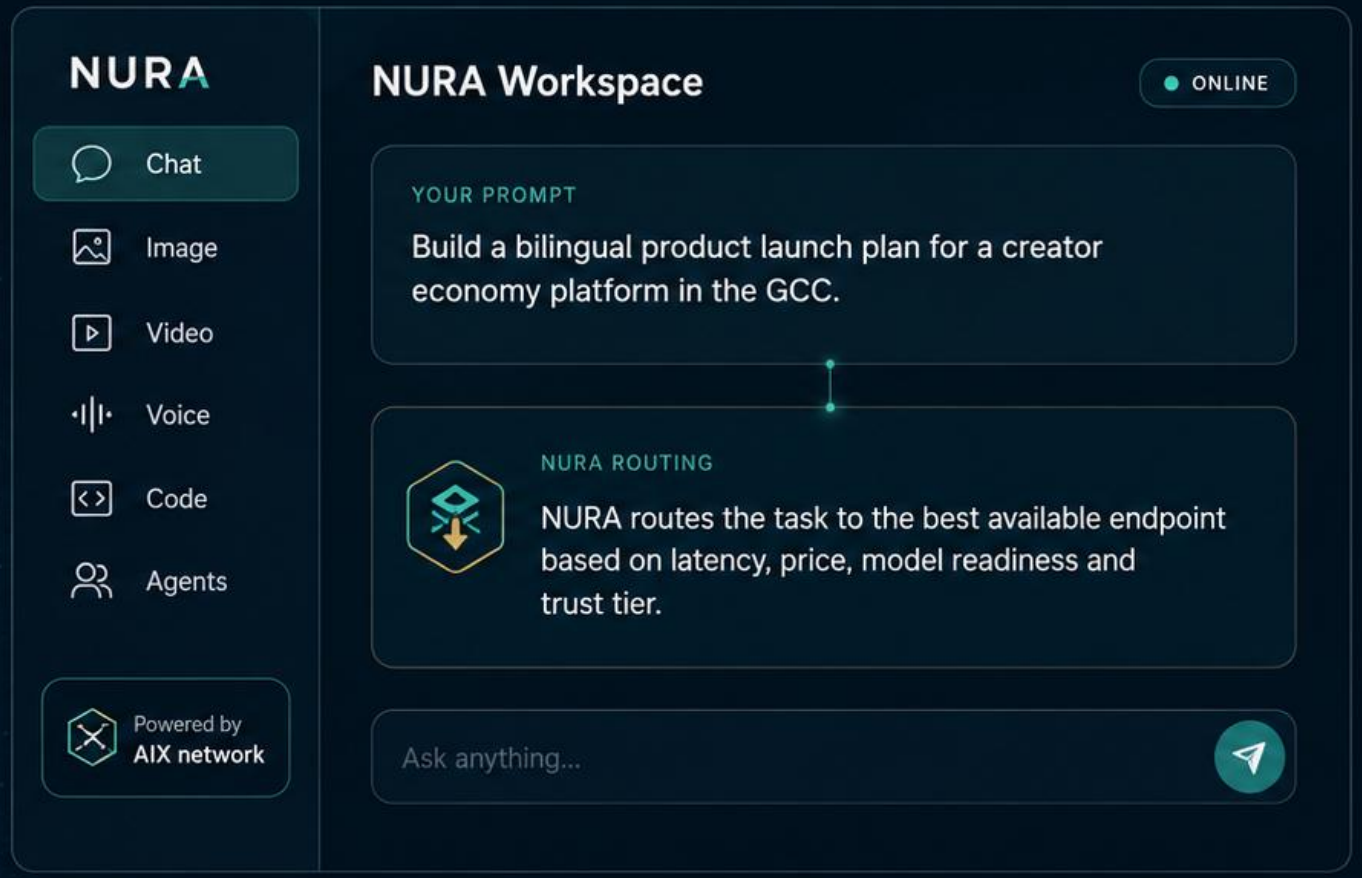

ONE ACCOUNT. MANY AI MODES.

Chat	Billed by tokens
Image	Billed by images
Video	Billed by seconds
Voice	Billed by minutes
Code / Agents	Billed by GPU-time + tokens

1 → 2 → 3

Use AI
User submits a request through NURA Portal.

Route workload
NURA selects the best endpoint across its infrastructure and AIX network.

Settle usage
Usage is metered and settled transparently back to the network.

Every consumer request becomes real demand for NURA capacity and AIX-connected nodes.

AIX coordinates the open network

AIX is the token and decentralized compute network. It is useful inside the system before it becomes an exchange-traded asset.

The useful lesson from Venice VVV is that a token can be connected to AI access when a real product already exists. AIX should follow the same discipline: first launch NURA as a working AI data center and service layer, then use AIX to coordinate users, hosts, developers and rewards.

AIX will not be presented as a claim on land, electricity contracts, GPU ownership or dividends. The token's role is utility: staking, host collateral, routing priority, compute access benefits, reputation weight, governance parameters and verified-work rewards.

Layer	What it does	How value is created
NURA	Operates the data center and managed compute services	Sells AI access, APIs, GPU rentals and enterprise capacity
AIX Network	Onboards external GPU hosts and verifies usage	Expands supply beyond owned infrastructure
AIX Token	Coordinates staking, rewards, trust and governance	Turns useful compute into a measurable economic network
NURA Credits	Closed-loop usage credit inside the platform	Converts token or fiat value into actual AI compute consumption

> **AIX will be marketed as a utility and coordination token for AI compute access, not as a tokenized profit right.**

AIX Economic Structure

> Download full AIX Whitepaper

VVV proves the AI-access token pattern

Venice VVV is not a template to copy blindly, but it shows how a live AI product can turn staking into access, emissions and demand.

Venice launched VVV after its AI product and API were already operating. Its public materials describe a no-presale launch, a 100M genesis supply, a 50% community airdrop, staking-linked inference access, ongoing emissions and later revenue-funded buy-and-burn mechanics. The important point is sequencing: product first, token second.

AIX should adopt the principle, not the exact parameters. The AIX version must fit a compute marketplace with physical data-center capacity, independent GPU hosts, escrow settlement, slashing, host reputation and a UAE legal pathway.

What VVV demonstrates

01	**Product before token**	AI users and API usage existed before token scale.
02	**Staking for access**	Token ownership becomes useful through inference access.
03	**Community distribution**	Users and AI-aligned communities receive meaningful allocation.
04	**Emissions discipline**	Rewards bootstrap supply, then should decline as demand matures.
05	**Buy-and-burn option**	Revenue can support token scarcity if legally approved.

> **For AIX, the equivalent of “staking for inference access” will become staking for NURA Compute Credits, host priority, marketplace trust and verified-work rewards.**

Venice.ai VVV launch and VVV public pages

Proposed genesis allocation

AIX should reserve meaningful supply for users and host operators, while still funding data-center expansion, liquidity and long-term protocol development.

The following structure is a planning proposal, not a final issuance term sheet. It mirrors the VVV idea of community-first distribution, but adapts it to a physical compute marketplace where data-center capacity, GPU financing and host reliability are essential.

Genesis supply assumption: 1,000,000,000 AIX

A 1B supply gives enough unit granularity for consumer pricing, rewards and exchange liquidity while remaining simple for public communication. Final supply can be adjusted by counsel and market makers before launch.

The allocation will favor real users, real hosts and real capacity growth — not short-term insider liquidity.

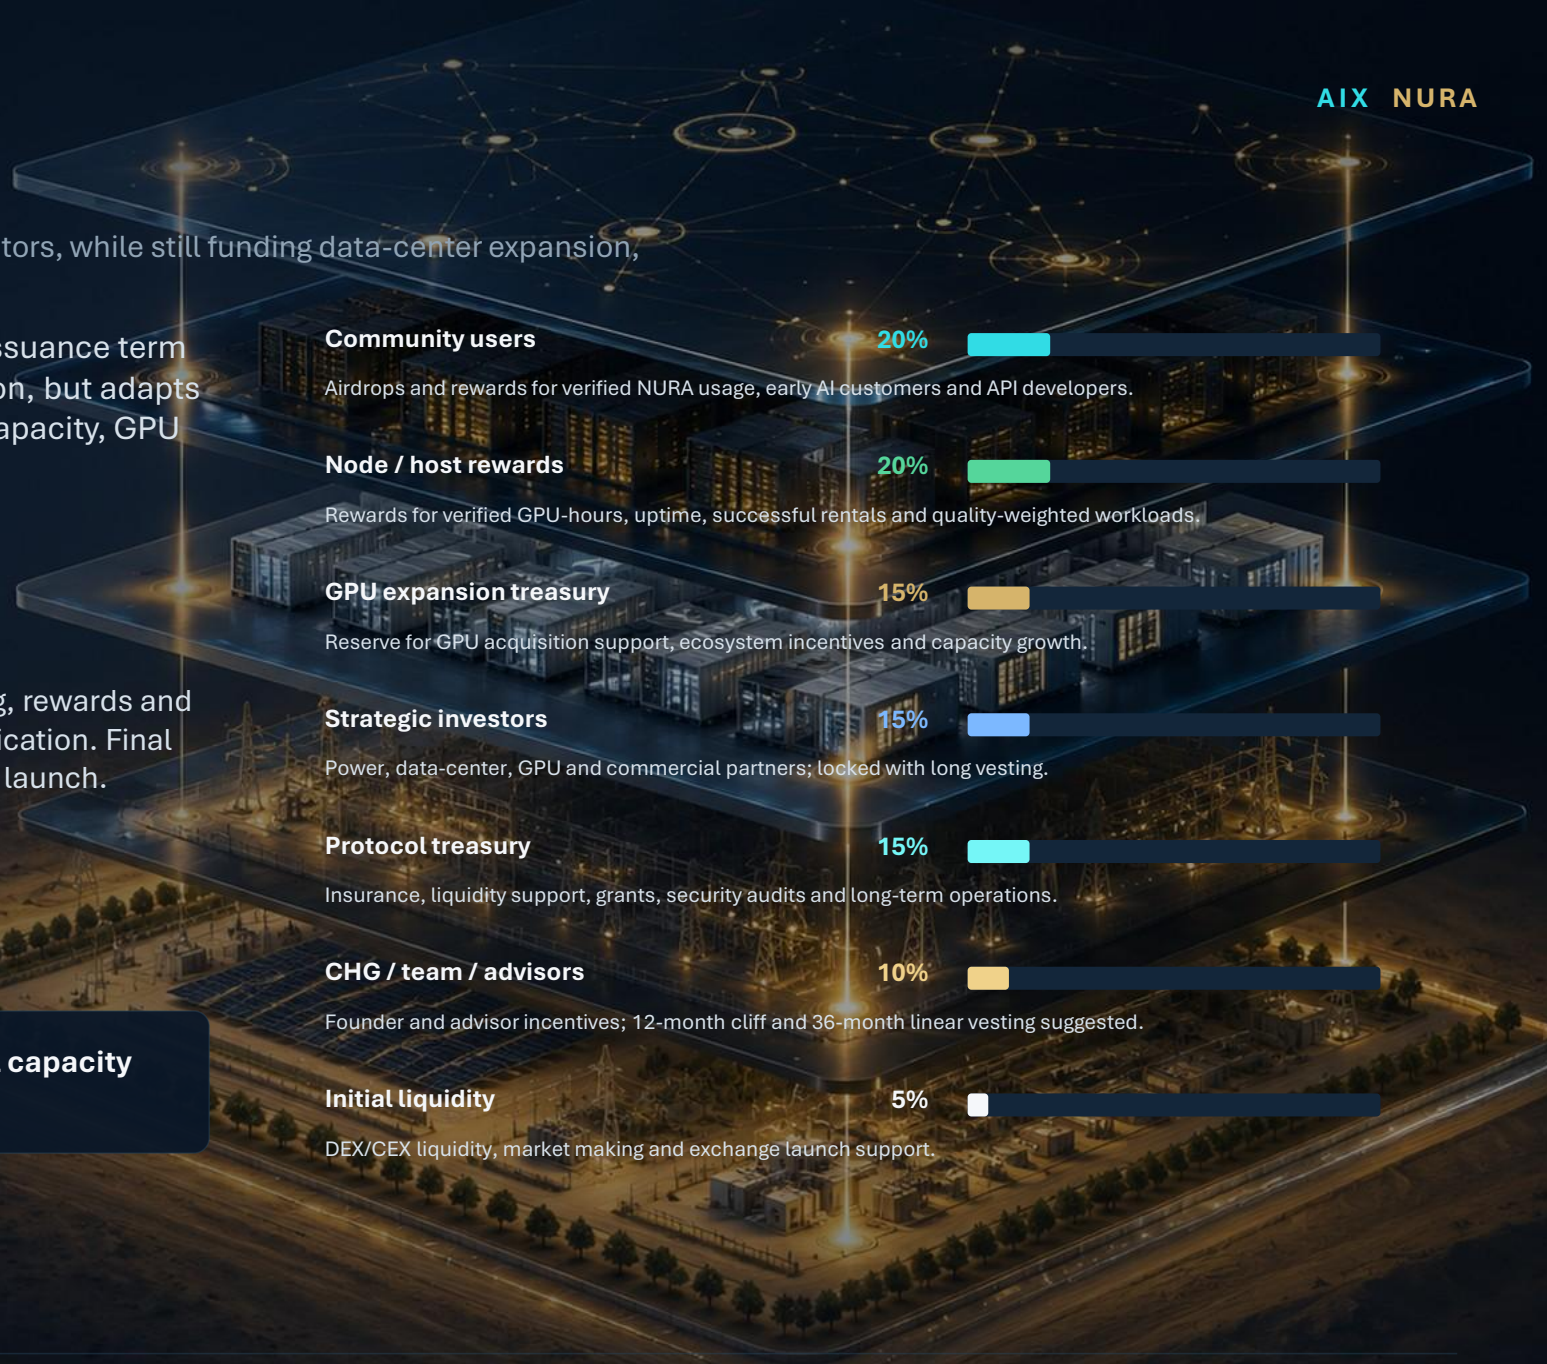

Community users	20%	Airdrops and rewards for verified NURA usage, early AI customers and API developers.
Node / host rewards	20%	Rewards for verified GPU-hours, uptime, successful rentals and quality-weighted workloads.
GPU expansion treasury	15%	Reserve for GPU acquisition support, ecosystem incentives and capacity growth.
Strategic investors	15%	Power, data-center, GPU and commercial partners; locked with long vesting.
Protocol treasury	15%	Insurance, liquidity support, grants, security audits and long-term operations.
CHG / team / advisors	10%	Founder and advisor incentives; 12-month cliff and 36-month linear vesting suggested.
Initial liquidity	5%	DEX/CEX liquidity, market making and exchange launch support.

Emissions tied to useful compute

The token should reward verified productive work: online GPUs, completed jobs, API demand, model services and reliable host behavior.

AIX emissions should not look like passive mining. New rewards should be released only against measurable contribution: verified GPU-hours, successful rental time, model inference volume, uptime, low dispute rate and quality-weighted customer usage. This directly separates AIX from wasteful proof-of-work narratives.

Indicative emissions framework

- Initial network emissions: up to 4% of genesis supply per year
- Annual step-down: reduce emissions as marketplace demand matures
- Dynamic adjustment: rewards favor under-supplied model categories and trusted regions
- No guaranteed yield: rewards depend on real usage and protocol rules

Reward pool	Indicative share	Purpose
Host operators	60%	Pay for verified GPU-hours, uptime, successful rentals and SLA score
User staking / credits	20%	Encourage demand, API usage and long-term platform participation
Developers / model providers	10%	Reward model endpoints, tools, agents and ecosystem integrations
Insurance / security reserve	10%	Cover disputes, slashing appeals, audits and liquidity support

> **Users stake AIX to receive NURA Compute Credits or priority tiers. Hosts stake AIX to list machines, receive high-value jobs and secure customer trust. Bad behavior can be slashed.**

Indicative reward model; subject to legal review

Credits convert token value into usage

The system should be easy for normal users: they can pay with card, bank transfer or stablecoin; AIX works underneath as the coordination asset.

NURA Compute Credits should be closed-loop platform credits. They can represent API usage, image jobs, video seconds, voice minutes, reserved GPU-hours or host rental time. They should not be marketed as an investment instrument or transferable yield product.

AIX can still create a strong economic flywheel: users stake or buy AIX for access benefits; hosts stake AIX to receive work; platform fees support rewards, insurance, liquidity and GPU expansion; a legally approved share of revenue may fund buyback-and-burn or capacity purchase programs.

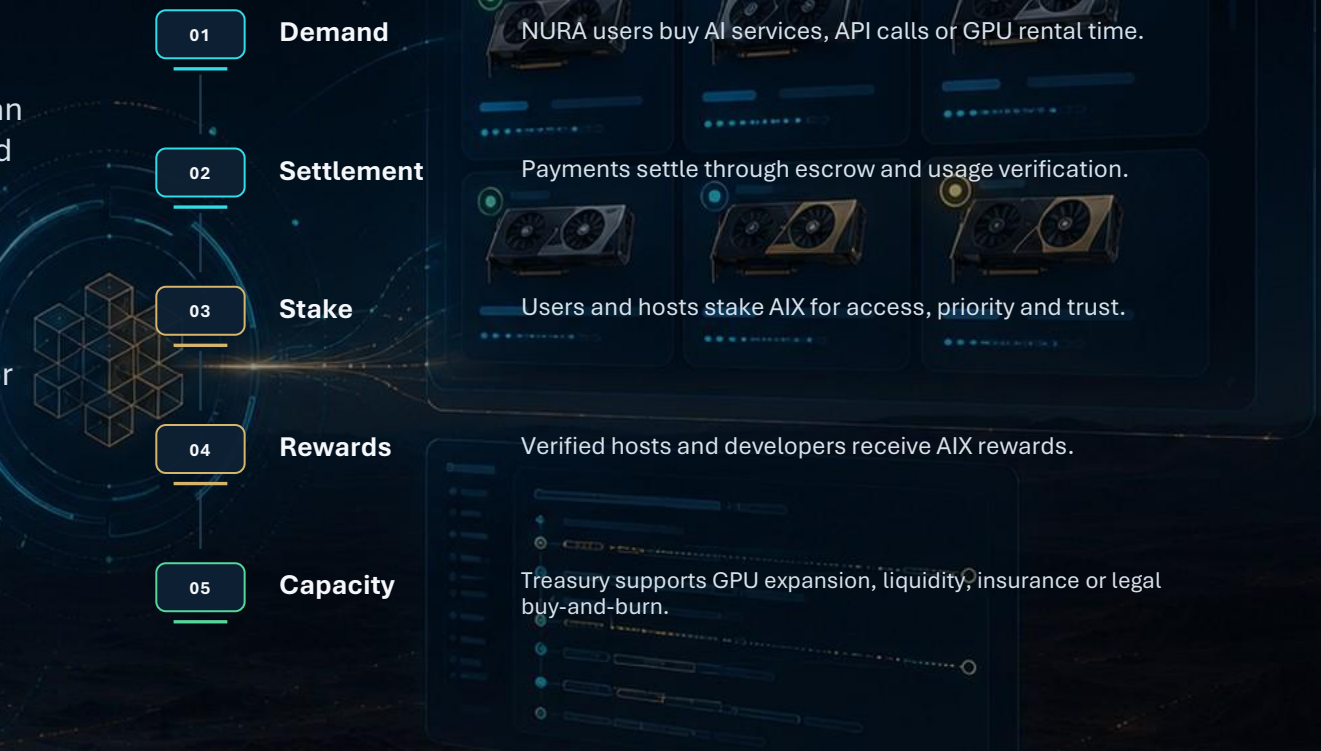

01	**Demand**	NURA users buy AI services, API calls or GPU rental time.
02	**Settlement**	Payments settle through escrow and usage verification.
03	**Stake**	Users and hosts stake AIX for access, priority and trust.
04	**Rewards**	Verified hosts and developers receive AIX rewards.
05	**Capacity**	Treasury supports GPU expansion, liquidity, insurance or legal buy-and-burn.

The message is simple: AIX is useful because it unlocks access, trust and rewards in a real compute market anchored by NURA.

How value moves through the network

The economic model must be understandable to users, hosts and investors before any token launch.

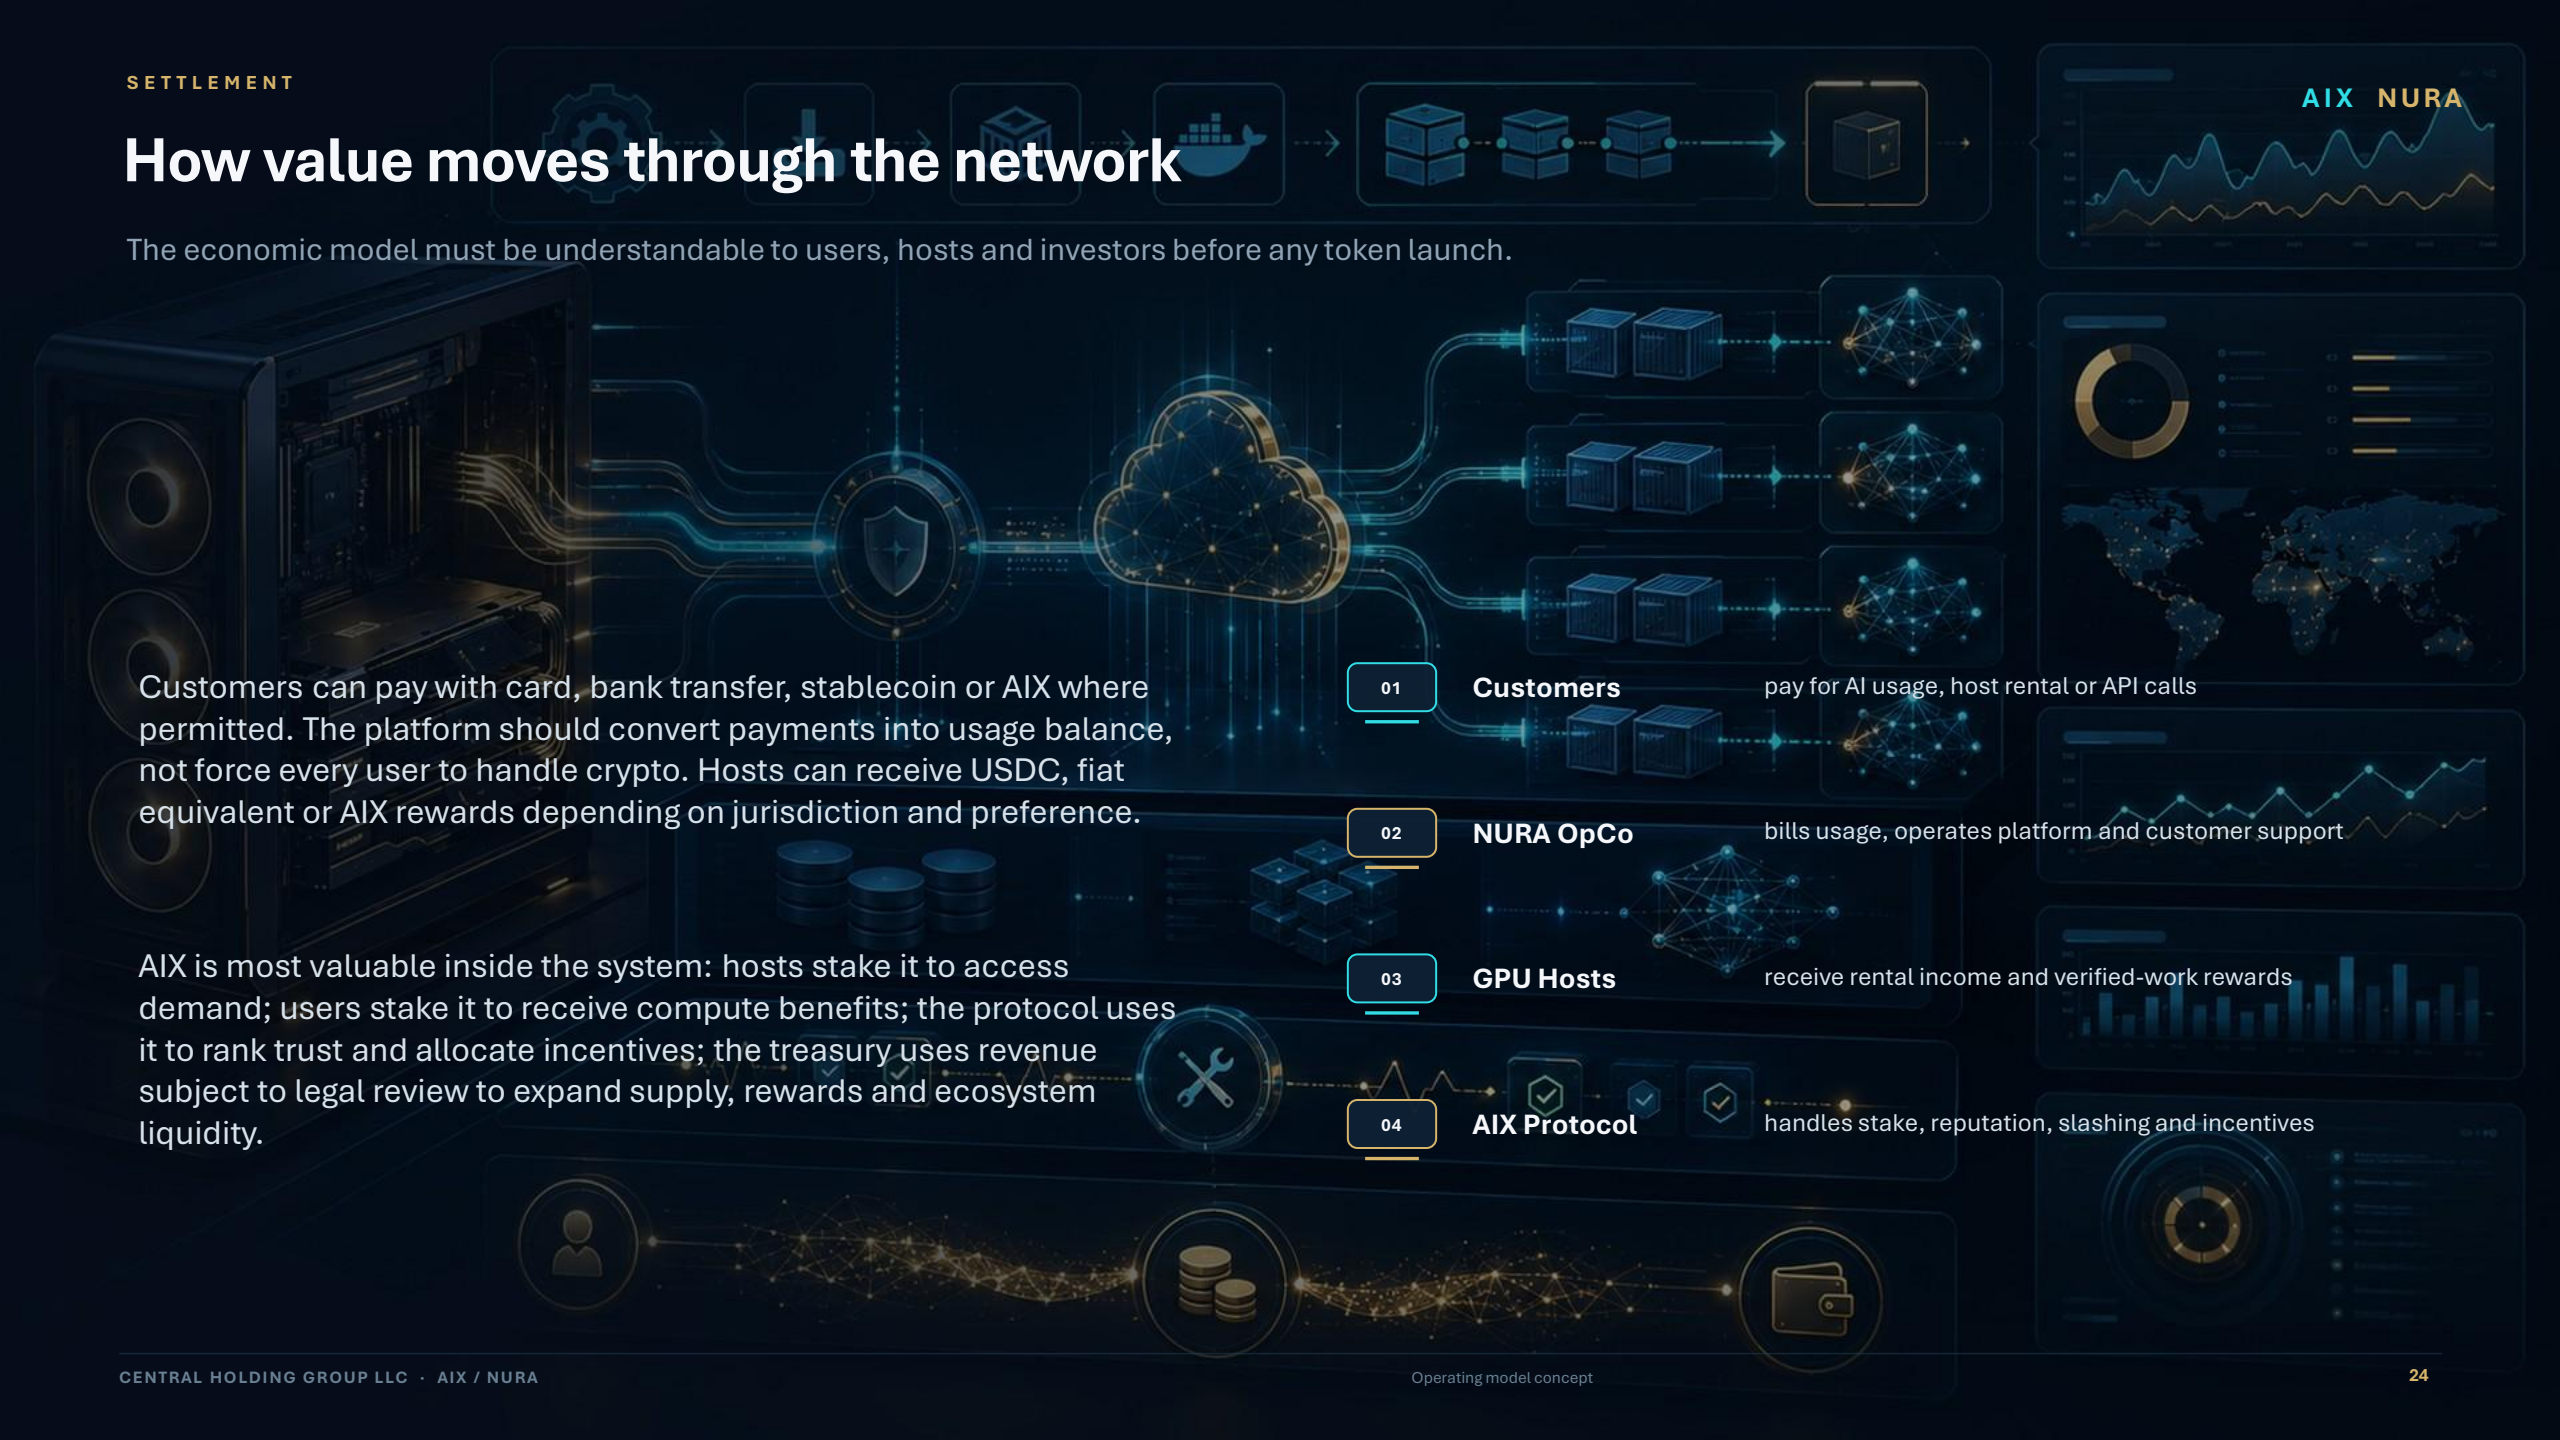

Customers can pay with card, bank transfer, stablecoin or AIX where permitted. The platform should convert payments into usage balance, not force every user to handle crypto. Hosts can receive USDC, fiat equivalent or AIX rewards depending on jurisdiction and preference.

AIX is most valuable inside the system: hosts stake it to access demand; users stake it to receive compute benefits; the protocol uses it to rank trust and allocate incentives; the treasury uses revenue subject to legal review to expand supply, rewards and ecosystem liquidity.

01	**Customers**	pay for AI usage, host rental or API calls
02	**NURA OpCo**	bills usage, operates platform and customer support
03	**GPU Hosts**	receive rental income and verified-work rewards
04	**AIX Protocol**	handles stake, reputation, slashing and incentives

Multiple ways to monetize compute

The project should not depend on only one revenue line. The strongest model combines owned-GPU margin with marketplace commission.

Owned data-center GPUs produce the highest control and enterprise credibility. Marketplace hosts add asset-light supply and expand model coverage. The consumer AI product creates recurring retail usage and helps keep the supply side utilized.

Revenue line	Customer	Pricing basis	Margin logic
NURA AI apps	Consumers, creators, agents	Tokens, images, minutes, video seconds	Usage margin over routed compute cost
GPU cloud	Developers and AI teams	GPU-hour, VM, container or notebook	Owned fleet margin; high utilization required
Dedicated host rental	Power users and enterprises	Reserved hours, days or months	Premium for isolation and guaranteed access
Marketplace commission	Independent host transactions	Take rate on GMV	Asset-light expansion of supply
Managed deployment	Enterprises and studios	Setup fee + monthly service	Higher support margin and stickier customers

> **Only renting empty data-center racks caps the opportunity. Participating in compute operations, software routing and marketplace settlement captures more of the AI value chain.**

10–20%

Potential marketplace take-rate range after escrow, support and settlement costs

60–75%

Target utilization range for mature owned GPU fleet

Revenue depends on GPU scale and utilization

The numbers below are planning scenarios, not guarantees. They show why GPU participation is materially stronger than passive rack leasing.

At the target AED 0.10/kWh power input, facility electricity is not the limiting cost. The decisive variables are GPU purchase price, financing structure, utilization, blended GPU-hour rate, maintenance and sales execution. The model should be updated once exact GPU quotes and signed power terms are available.

Scenario	Owned GPU count	Blended rate	Utilization	Annual compute revenue	Strategic meaning
MVP	256 GPUs	US$3.00/hr	65%	~AED 16.1M	Proves NURA usage and billing
Base Phase-I	512 GPUs	US$3.25/hr	70%	~AED 37.5M	Strong first commercial fleet
Expanded Phase-I	1,024 GPUs	US$4.25/hr	75%	~AED 105.0M	Shows path beyond AED 70M revenue
Power cost check	4 MW site	AED 0.10/kWh	100%	~AED 3.5M cost	Power is small vs compute revenue

> **The model use “annualized compute revenue potential,” not guaranteed profit. Net profit will depend on GPU financing, depreciation, staffing, cooling, bandwidth and utilization.**

Model uses USD/AED 3.6725, 8,760 hours/year and management utilization assumptions

Separate assets, risk and upside

A clean SPV structure lets different investors fund the layer they understand: power, data-center infrastructure, GPUs or protocol growth.

The project should not mix all risks into one company from day one. Land and power rights, electrical infrastructure, modular data-center infrastructure, GPU assets, operating platform and token protocol each have different capital requirements, depreciation profiles and investor preferences.

Central Holding Group LLC can lead the operating and sponsor relationship while using long-term contracts between entities: site lease, power capacity agreement, infrastructure lease, GPU lease, platform service agreement and marketplace settlement rules.

01	**Sponsor contribution**	Land, power pathway and strategic relationships
02	**Power Infra SPV**	Substation, distribution, backup and metering
03	**Data Center SPV**	Modular halls, cooling, security and network rooms
04	**Compute SPV**	H100/H200/B200/RTX fleet and GPU financing
05	**NURA Data Center OpCo AIX Network**	Marketplace, software, billing, token and community

This structure makes the pitch clearer: infrastructure investors fund predictable physical assets; compute investors fund higher-yield GPUs; token investors fund network coordination and marketplace growth.

Management financing structure concept

Token launch after product proof

AIX can be ambitious, but the issuance path must be designed before public marketing, exchange listing or token sale.

The practical path is to launch NURA as a normal AI data center and managed compute service first, then launch AIX after legal structuring, whitepaper preparation, risk disclosure and distributor review. The token should avoid asset-reference, fiat-reference, profit-sharing or guaranteed-yield language unless a regulated path is chosen.

The safest architecture separates platform revenue from token ownership. AIX can support staking, access benefits, governance and host collateral. NURA Credits can represent usage rights inside the platform. Investor documents should state that the token is not a claim on land, power contracts, GPUs or dividends unless legal counsel explicitly approves.

Area	Design principle
Token classification	Aim for utility / access / coordination; avoid profit-share promises
Distribution	Use licensed and legally reviewed distribution channels where required
KYC / AML	Apply tiered KYC for fiat ramps, high-value users and host payouts
Abuse control	Sanctions, malware, fraud, illegal content and infrastructure attacks must be enforceable
Data protection	Private endpoints, isolation, logging controls and deletion workflows

VARA Virtual Asset Issuance Rulebook; ADGM digital asset framework

A staged path to launch

The plan should prove utility, revenue and operations before expanding power draw or token exposure.

01	0–90 days	**Legal, power and design lock**	Confirm site rights, power tariff, connection pathway, cooling strategy, SPV structure, token counsel and first GPU procurement list.
02	3–6 months	**NURA MVP**	Launch website, account system, payments, API billing, model catalog, host listing prototype and first owned GPU cluster.
03	6–12 months	**4 MW Phase-I build**	Commission modular halls, liquid-ready cooling, network core, monitoring, security and initial 256–512 GPU fleet.
04	9–15 months	**Open host beta**	Release AIX Node Agent, onboard controlled 4090/5090 and enterprise hosts, test staking, escrow, slashing and payout.
05	12–18 months	**AIX launch readiness**	Publish whitepaper, complete legal review, finalize token utility, build liquidity plan, activate rewards and scale to more capacity.

The hard risks are manageable if designed early

The project becomes bankable only when power, cooling, software, compliance and demand risks are handled as one system.

Risk	Why it matters	Mitigation
Power contract risk	The AED 0.10/kWh assumption drives margin and investor confidence	Obtain written tariff, capacity, term, escalation and connection schedule
GPU capex risk	High-end GPUs depreciate and can be supply-constrained	Stage purchases; combine owned GPUs, leases and partner supply
Cooling risk	Desert heat and dust can destroy uptime and hardware life	Liquid-ready design, filtration, positive pressure, N+1 pumps and monitoring
Marketplace trust risk	Bad hosts can harm customers and brand reputation	Stake, verification, escrow, sandboxing, temporary credentials and slashing
Regulatory risk	Poor token design can block exchanges, banks and investors	Utility design, legal review, licensed distribution and clear risk disclosure
Demand risk	Compute supply without users creates idle assets	Launch NURA consumer AI first; use marketplace only after demand signal

Central Holding Group leads the platform

The project needs a credible local sponsor that can coordinate land, power, investors, operators, legal advisers and the public product.

Central Holding Group LLC will be presented as the project lead and commercial coordinator. Its role is to organize the platform structure, investor consortium, power infrastructure partners, modular data-center vendors, GPU financing partners, software buildout and token legal path.

The sponsor contribution is not only land access. It is the ability to turn local infrastructure advantages into a market product: discounted electricity, expansion optionality, regional credibility, a UAE brand, and access to capital partners who understand both AI and digital assets.

Leadership responsibilities

01 Secure site, power and expansion rights.

02 Form SPVs and align investor economics.

03 Select EPC, cooling, GPU and network partners.

04 Launch NURA service layer and operating company.

05 Coordinate AIX token counsel, whitepaper and go-to-market.

NURA AI EXCHANGE NETWORK

Open Compute.
Open Models.
Open Participation.

AIX can become the UAE-born open compute economy: NURA as the physical AI campus, a public model service, a marketplace for independent GPU supply and a token layer that rewards useful work rather than wasted energy.

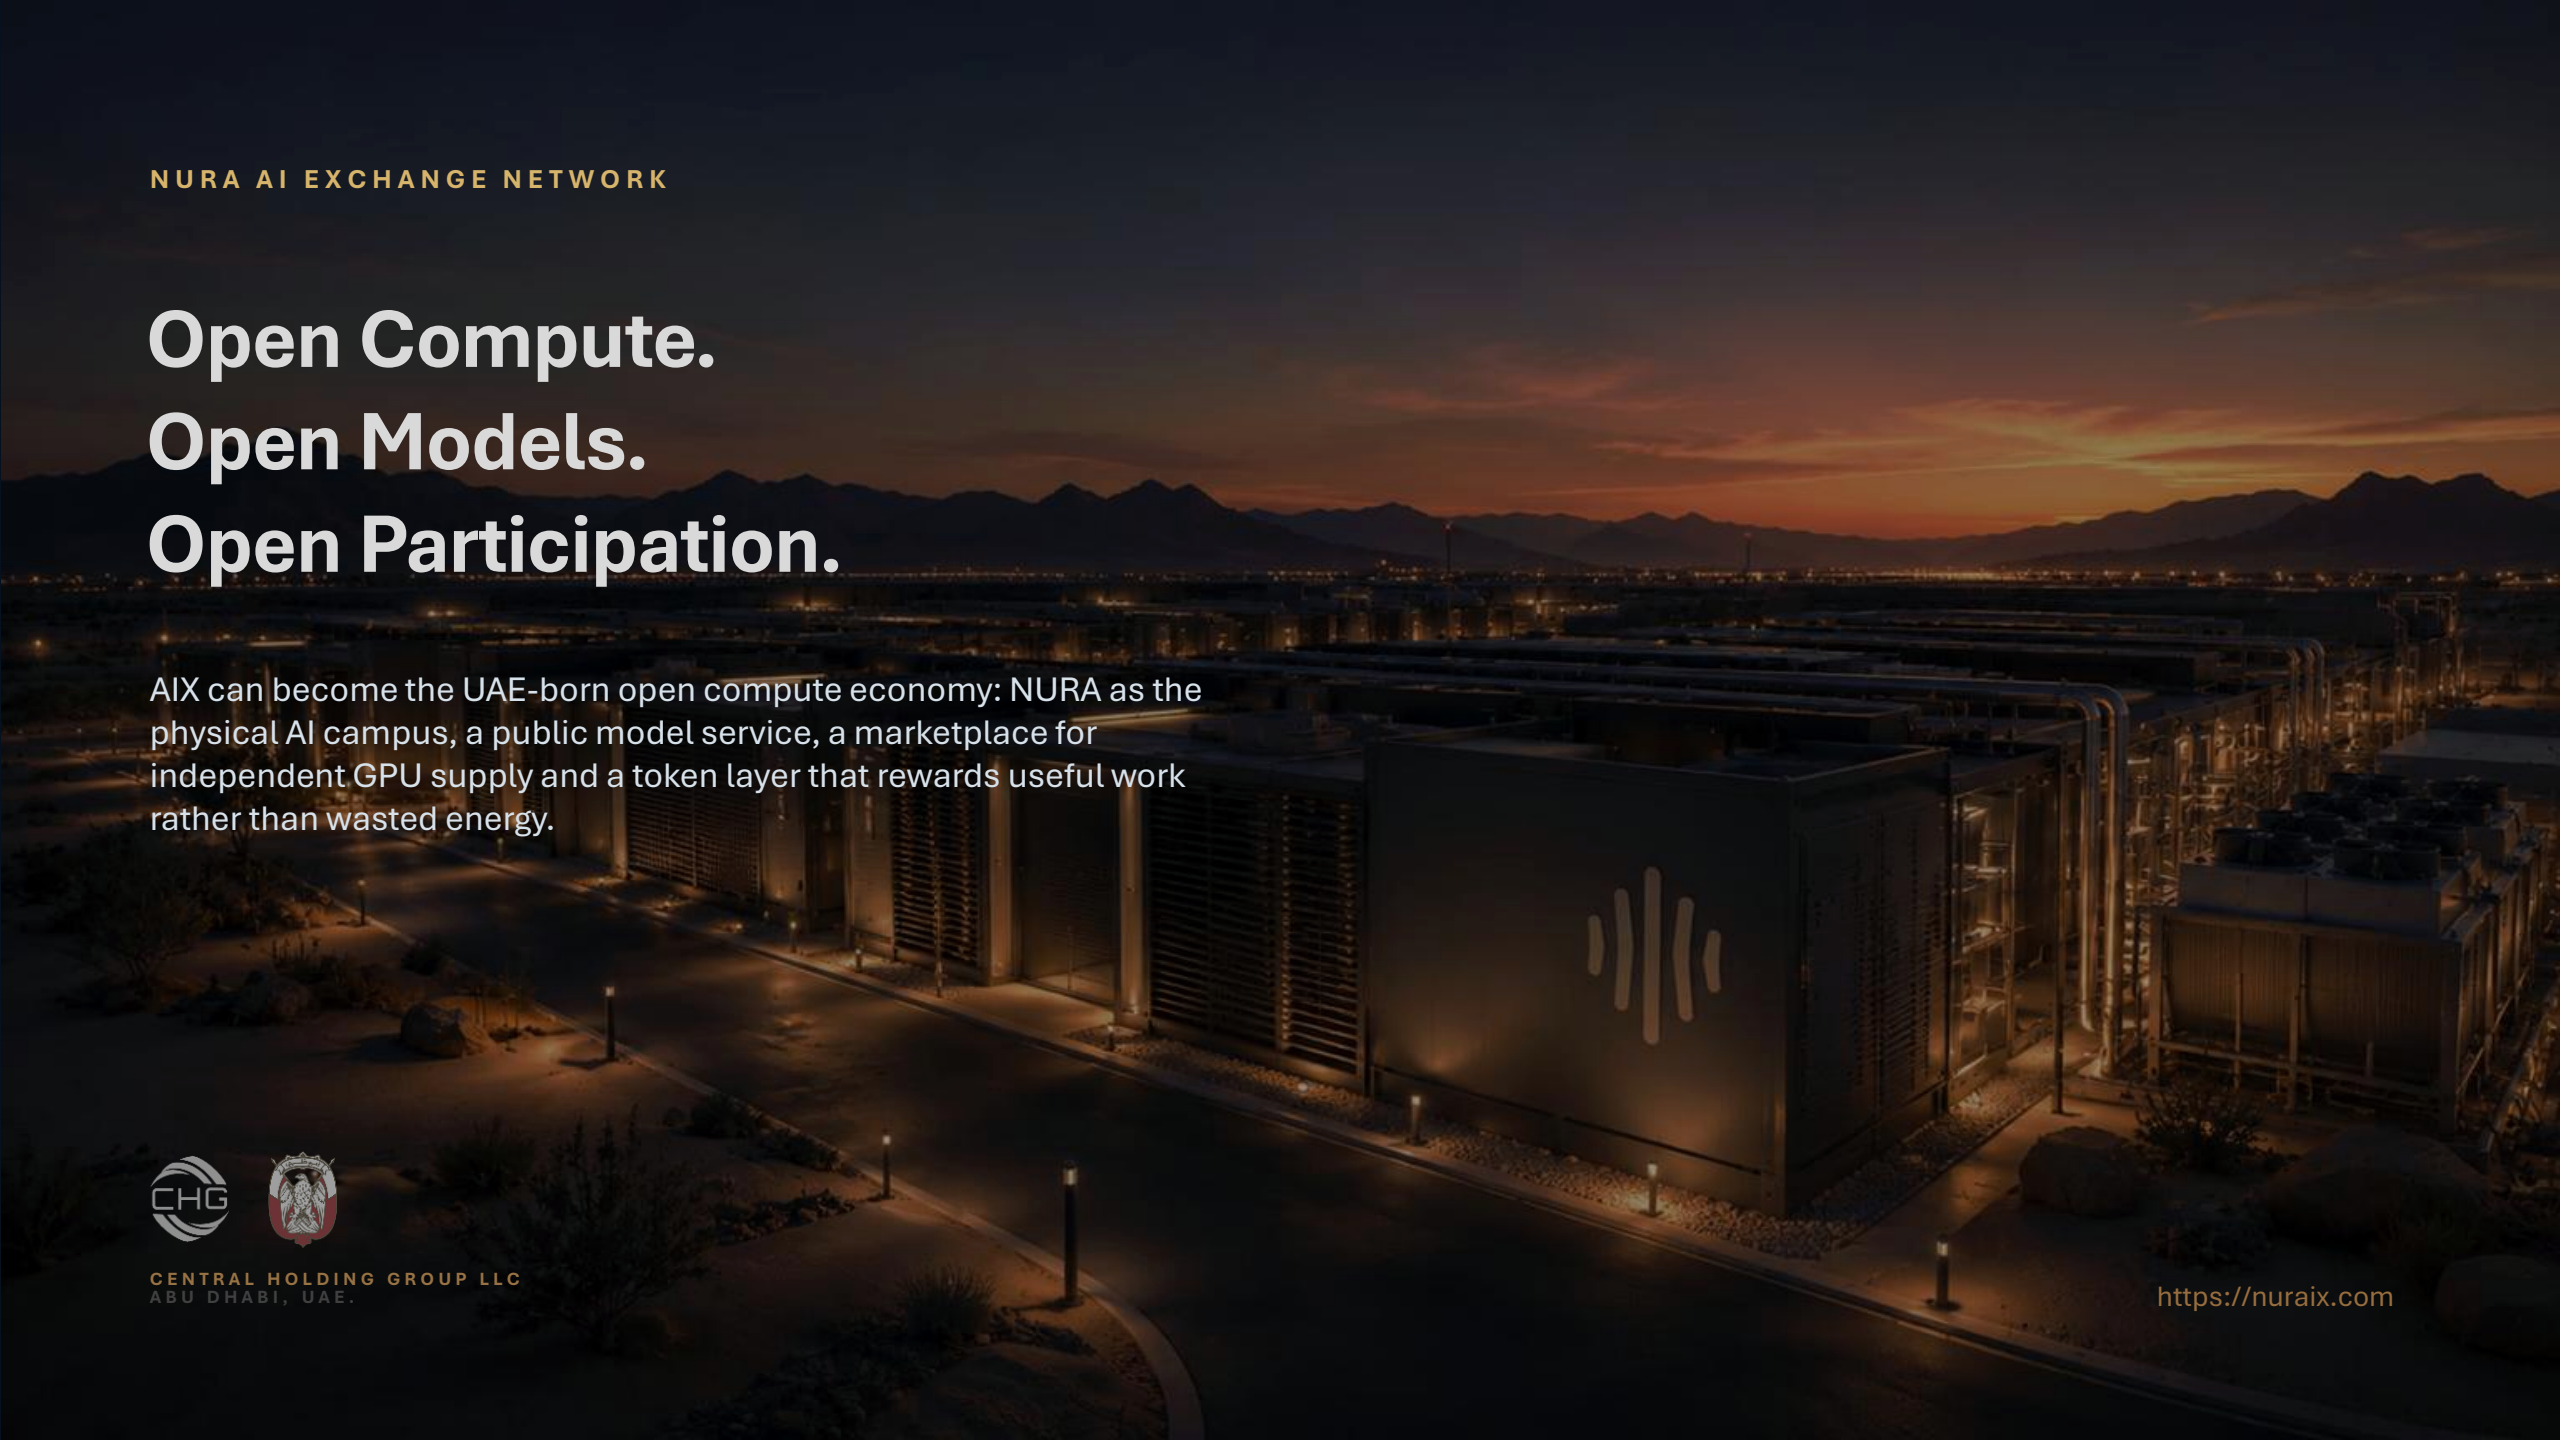

CENTRAL HOLDING GROUP LLC
ABU DHABI, UAE.

https://nuraix.com