DLResearch × SHAGA

From Idle GPUs to Global Infrastructure

The Shaga Model

From Idle GPUs to Global Infrastructure

The Shaga Model

Executive Summary

Key Findings

Edge over incumbents

Shaga delivers compute at the true edge through idle GPUs, achieving latency under 40ms and operating at a cost base twenty-six times lower than centralised providers.

Dual revenue streams

By monetising synchronised gameplay data for AI training, Shaga turns a by-product into a high-margin asset, allowing gaming sessions to be subsidised and creating a compounding growth flywheel.

Efficient and verifiable incentives

Unlike many DePIN projects, Shaga ties emissions to two concrete anchors: geography (via the H3 mapping) and measurable outputs. For infrastructure, rewards are tied to verifiable outputs such as electricity consumption, creating a clear link between real-world contribution and rewards. On the adoption side, incentives are directed to areas with potential and growth can be easily tracked through the activity and geolocation of Globbers and clients.

Surplus potential at modest adoption

Even in a base-case scenario of 5% adoption of the underserved gaming market, gameplay data monetised at $5 per hour generates an annual surplus of $1.4 billion, demonstrating that sustainability and scale are achievable without extreme assumptions.

Geographically validated growth

Our geospatial analysis highlights clear regional dynamics. Interest is strongest in Uganda, Indonesia, and India, while conversion ratios are highest in parts of Latin America, North Africa and a few key countries in Europe and Asia where early adoption is already translating into active usage. At the same time, node strains are emerging in Bangladesh and Nigeria, with similar pressures expected in some regions showing high conversion, such as Latin America, or strong interest, such as Indonesia and Uganda.

The Opportunity

Cloud gaming and artificial intelligence are two of the fastest-growing digital markets of the decade. By 2025, the cloud gaming industry is expected to generate $15.7 billion in revenues, with projections reaching up to $140 billion by 2032. AI training data, a less visible but equally critical market, is forecast to grow from $2.6 billion in 2024 to $12.7 billion by 2032. Both sectors face fundamental bottlenecks: cloud gaming suffers from high costs and limited coverage, while AI development is constrained by a shortage of annotated, high-quality datasets.

Shaga positions itself at the intersection of these two challenges. It transforms idle GPUs in everyday PCs into a decentralised edge network that can both power cloud gaming and generate high-value datasets. In doing so, it simultaneously reduces the cost of gaming infrastructure and creates a new revenue stream from AI data markets.

Technical Validation

Shaga's architecture has already shown strong performance in practice. During beta testing, with more than 400,000 registered users and around 700 active nodes, the network achieved median round-trip latency of under 40 milliseconds in underserved regions, a level that outperforms many centralised providers outside major hubs.

A major strength lies in hardware diversity. The network spans everything from high-end RTX 4090 GPUs to older GTX 1050s and mainstream laptop CPUs. This inclusivity ensures both resilience and efficiency: high-performance devices take on compute-intensive workloads, while lower-power hardware supports lighter processes or off-peak demand. By leveraging the full spectrum of available devices, Shaga maximises utilisation and avoids reliance on a narrow class of machines.

Shaga also incorporates key innovations that enhance scalability. Neural Game Codecs reduce bandwidth consumption by up to 75%, making streams more efficient and affordable, while control bifurcation cuts multiplayer delay by as much as 45 milliseconds by routing inputs directly to game servers. Together, these features strengthen Shaga's ability to scale effectively and position it as a credible alternative to centralised incumbents, for the latest stats, check glob.shaga.xyz.

Economics and Cost Advantage

The economics are equally compelling. Shaga's base cost of operation is tied almost entirely to electricity consumption. At an average tariff of $0.20/kWh, running a 500W gaming rig equates to about $0.10 per hour, or just $0.025 per player-hour when concurrency is factored in.

In contrast, centralised providers operate with far higher overheads. AWS GPU instances are priced at roughly $0.52 per hour, while Xbox Cloud Gaming charges around $0.50 per hour on an effective usage basis. The result is a twenty-six-fold cost advantage for Shaga.

Rather than capturing this spread directly, Shaga monetises gameplay data. Synchronised video and input streams are uniquely valuable for AI training. Broker quotes range from $10 to $30 per annotated hour, with high-end cases reaching $100. Even at a conservative floor of $5 per hour, Shaga comfortably covers costs and generates surplus. One hour of data resale can subsidise multiple hours of free gaming, creating a powerful growth flywheel: more players produce more data, which drives more revenue, which lowers costs and attracts more users.

Adoption Signals and Geospatial Growth

Shaga's ability to direct growth is reinforced by its H3 geospatial framework. Among the 220,000 Globbers analysed, Uganda leads with over 43,000, followed by Indonesia (24,085), the United Kingdom (12,887), China (11,496), India (8,039), and Serbia (7,602). High per-capita traction in smaller nations such as the Maldives and the Faroe Islands highlights the depth of Shaga's penetration beyond absolute numbers.

Conversion data highlights where adoption is strongest. In the United States, 110 active clients are already connected, with Nigeria (92) and India (91) close behind. Ratios of clients to Globbers show particularly strong adoption in Latin and Central America, where Brazil (5.2%), Argentina (8.2%), and Cuba (5.8%) stand out. North Africa also shows meaningful traction, with Morocco (2.0%), Algeria (8.9%), and Egypt (2.9%). In Europe and Asia, countries such as Portugal and Thailand also record strong conversion, further underscoring Shaga's global reach.

At the infrastructure level, the dataset shows 604 active nodes, with concentrations in the United States (70), Singapore (59), Japan (32), and India (31). Pressure points are already visible in high-interest regions such as Bangladesh and Nigeria, where client-to-node ratios suggest early strain. On a more speculative note, Globber-to-node ratios indicate that future bottlenecks could arise in Uganda, Indonesia, and parts of the Balkans if infrastructure growth does not keep pace with interest.

Tokenomics and Incentives

Shaga's tokenomics address one of the key weaknesses of many DePIN projects: inefficient or misaligned incentives. Rewards are tied to two anchors: geography and measurable outputs.

- **Geography:** Through H3 indexing, rewards can be directed to underserved but high-potential regions, ensuring that incentives follow demand signals.
- **Measurable outputs:** Rewards are tied to verifiable work, such as electricity consumed and frames rendered, rather than idle staking. This ensures that only productive contributions are compensated.

This combination of spatial precision and output-based fairness creates a token economy that is both transparent and efficient.

Valuation and Market Potential

The underserved total addressable market for cloud gaming is estimated at 150 million players. Based on average playtime of 40 hours per month, even a 1% conversion at a conservative $1 per hour data price allows Shaga to break even.

At $5 per hour, the base case of 5% adoption yields an annual surplus of $1.4 billion. Applying revenue multiples observed in comparable firms (4.6x for CoreWeave, 30x for Rescale) implies valuations between $6.6 billion and $43 billion. However, these figures must be treated with caution. Adoption will take time, and high multiples reflect early-stage speculation that typically flattens as markets mature.

Risks

Despite strong fundamentals, risks remain. Regulatory uncertainty over gameplay data, potential node shortages in high-demand regions such as Nigeria or Bangladesh, rising energy costs, hardware supply constraints, and volatility in AI data pricing all present challenges.

However, Shaga's design, combining spatially targeted incentives, output-based rewards, and adaptive scaling, offers natural buffers, making these risks significant but ultimately manageable.

Closing Takeaway

Shaga sits at the intersection of two rapidly expanding markets: cloud gaming and artificial intelligence. Its validated architecture, cost advantage, and geospatially anchored incentives position it as a credible and sustainable alternative to centralised providers.

By transforming idle hardware into productive infrastructure and gameplay into valuable AI datasets, Shaga creates a model that is both efficient and sustainable. Risks remain, from regulation to infrastructure scaling, but its design offers natural buffers.

If executed effectively, Shaga has the potential to establish itself as a sustainable DePIN model and a critical layer in decentralised infrastructure, addressing key bottlenecks in both gaming and AI.

Introduction

Cloud gaming and artificial intelligence are two of the fastest-growing digital markets of the decade. Cloud gaming promises to make AAA-quality play accessible anywhere, while AI development is defined by an insatiable demand for training data. Both industries are on steep growth trajectories, yet both face structural bottlenecks. Cloud gaming struggles with high infrastructure costs and uneven service quality, while AI faces scarcity of the high-quality behavioural datasets needed to improve decision-making.

Shaga positions itself at the intersection of these challenges. By transforming idle GPUs into a decentralised mesh, it creates a distributed infrastructure that delivers compute at the edge, reducing latency and costs. At the same time, it generates gameplay datasets enriched with real-time inputs, a by-product uniquely suited for AI training. This dual capability allows Shaga to compete directly with centralised cloud incumbents while unlocking an entirely new revenue stream.

This report analyses the foundations of the Shaga model. It explores the market context, technical design, geospatial distribution, and cost economics of the network, before examining tokenomics, valuation, and long-term risks. The aim is to provide a comprehensive view of how Shaga can grow from an early-stage experiment into a global infrastructure layer for gaming and AI.

Market & Competitive Context

Cloud-gaming fundamentals in 2025

Cloud gaming delivers interactive entertainment through remote servers instead of local hardware. Rendering and computation happen in the cloud, with games streamed in real time to laptops, tablets, or smartphones. This lowers the hardware barrier and expands access to AAA-quality content.

ONLINE GAMING / CLOUD GAMING

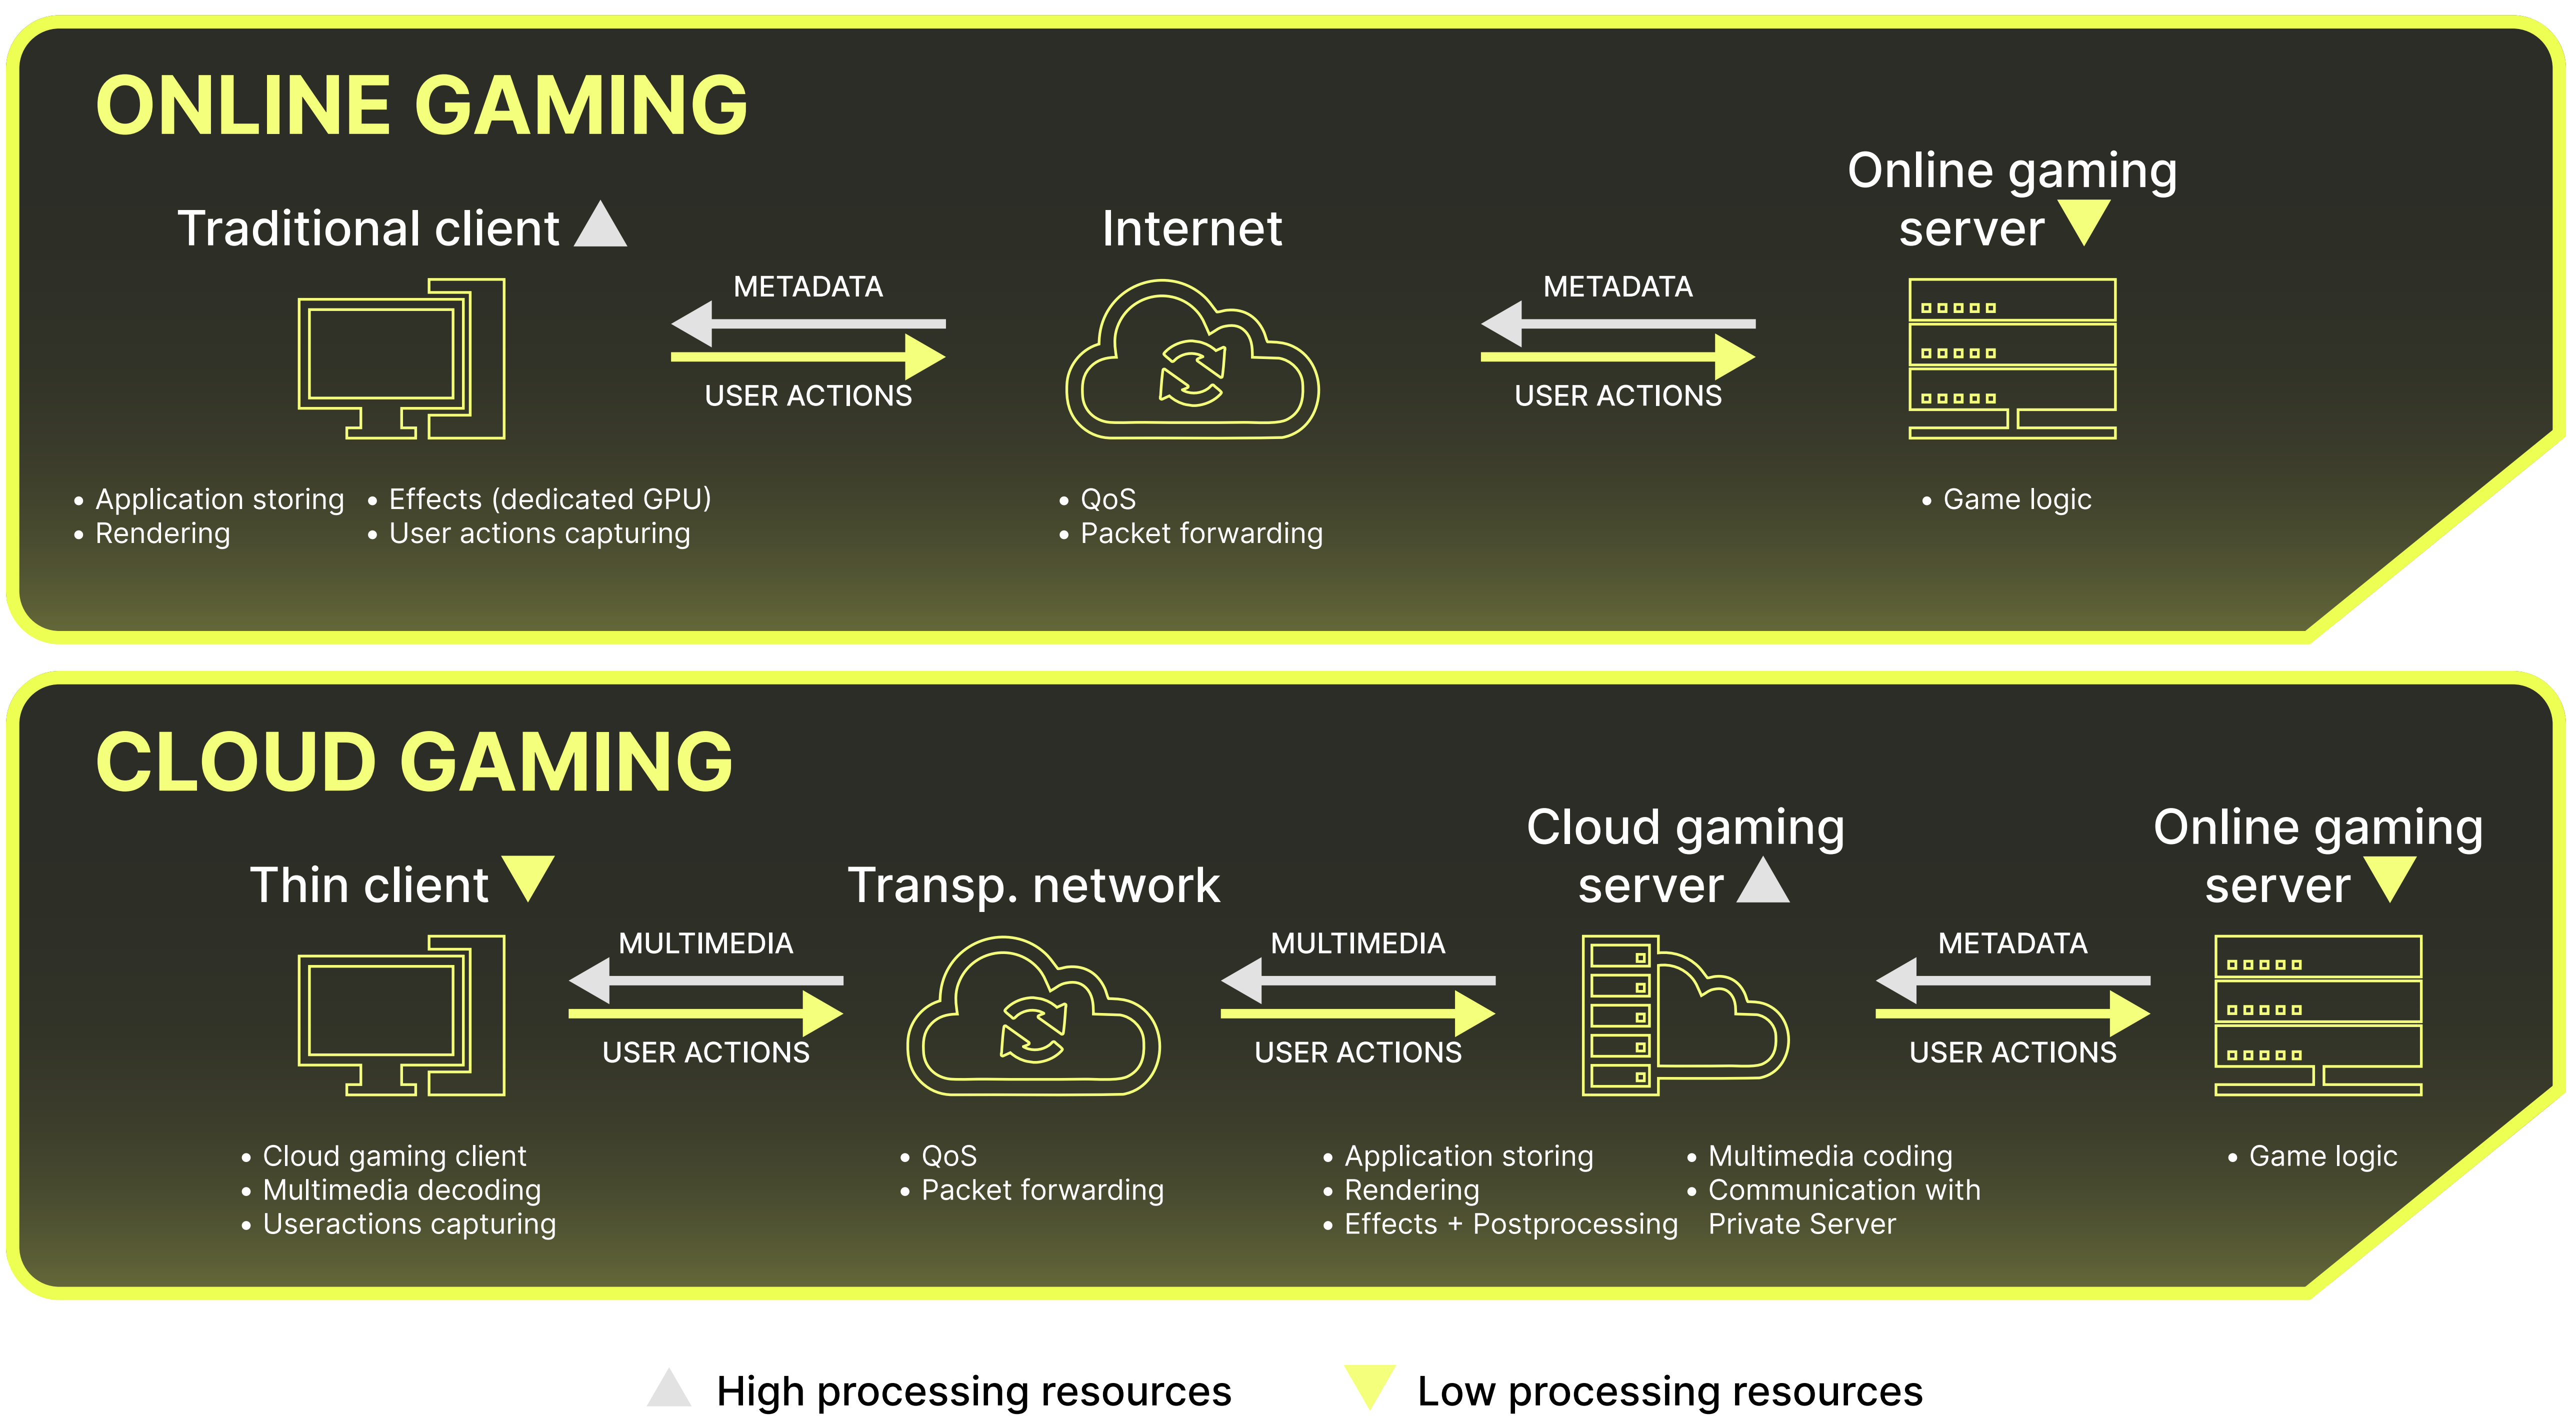

Source: Online-and-Cloud-Gaming-features

The real shift came in the early 2020s: the pandemic accelerated demand for digital entertainment, while advances in broadband, fibre, and 5G materially reduced latency. From that point, cloud gaming evolved from an experiment into a commercial reality. EY estimated the market at $2.4 billion by 2022, marking the first true phase of scale after a decade of incubation, with around 32 million paying users that year.

GLOBAL CLOUD GAMING MARKET FUNNEL — SEGMENTATION OF THE MARKET IN 2022

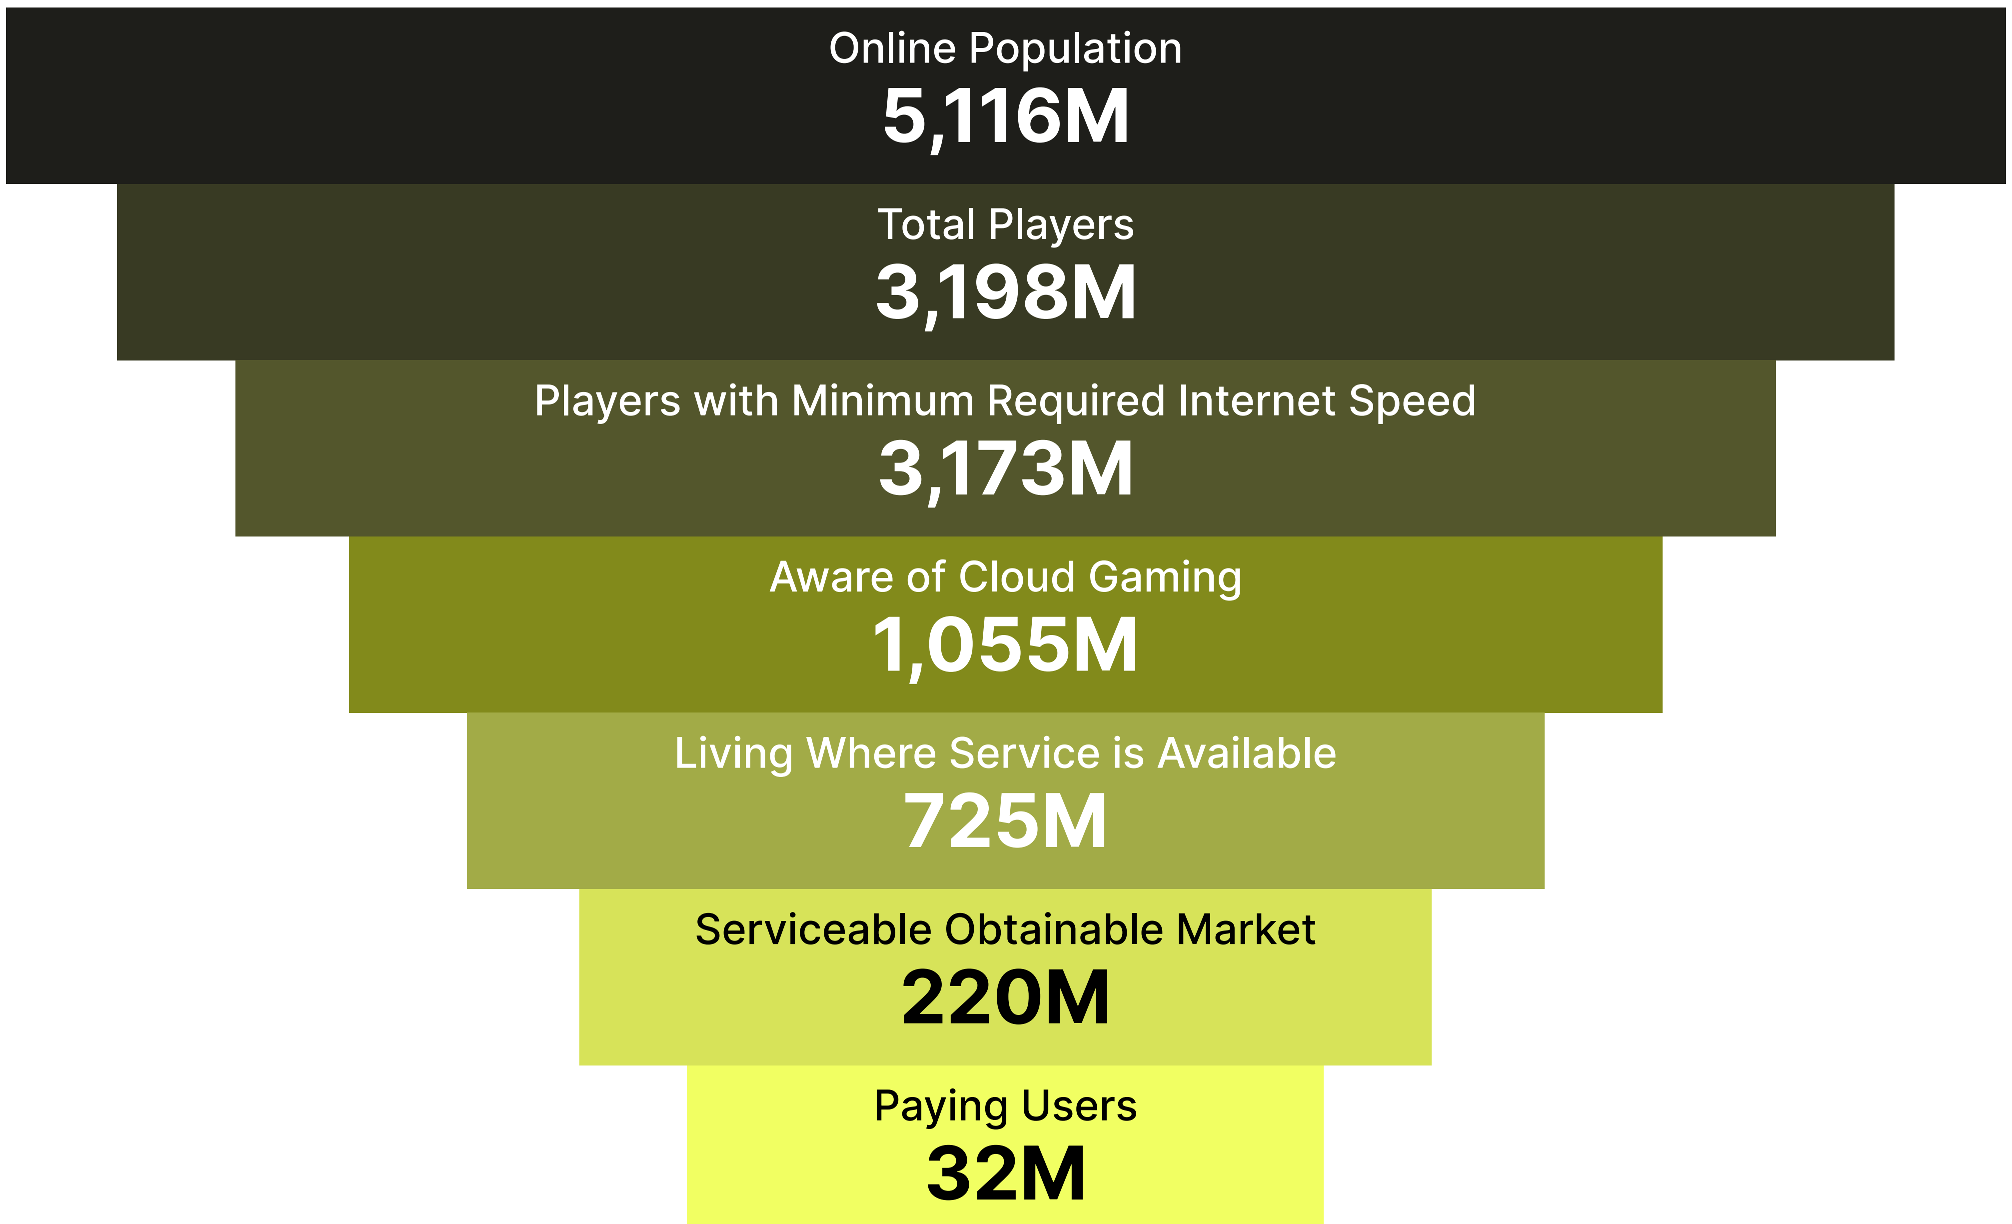

Source: Newzoo

Revenues accelerated to $9.7 billion in 2024 and are forecast to reach $15.7 billion in 2025, with long-term projections of $120-140 billion by 2032. Adoption is rising in parallel: the number of cloud gaming users (free and paid) grew from 182 million in 2022 to 396 million in 2024, with 455 million expected in 2025. Considering the global gaming industry includes more than 3.2 billion players, the growth potential remains substantial.

GROWTH IN CLOUD GAMING USER ADOPTION AND PENETRATION RATE (2017-2027)

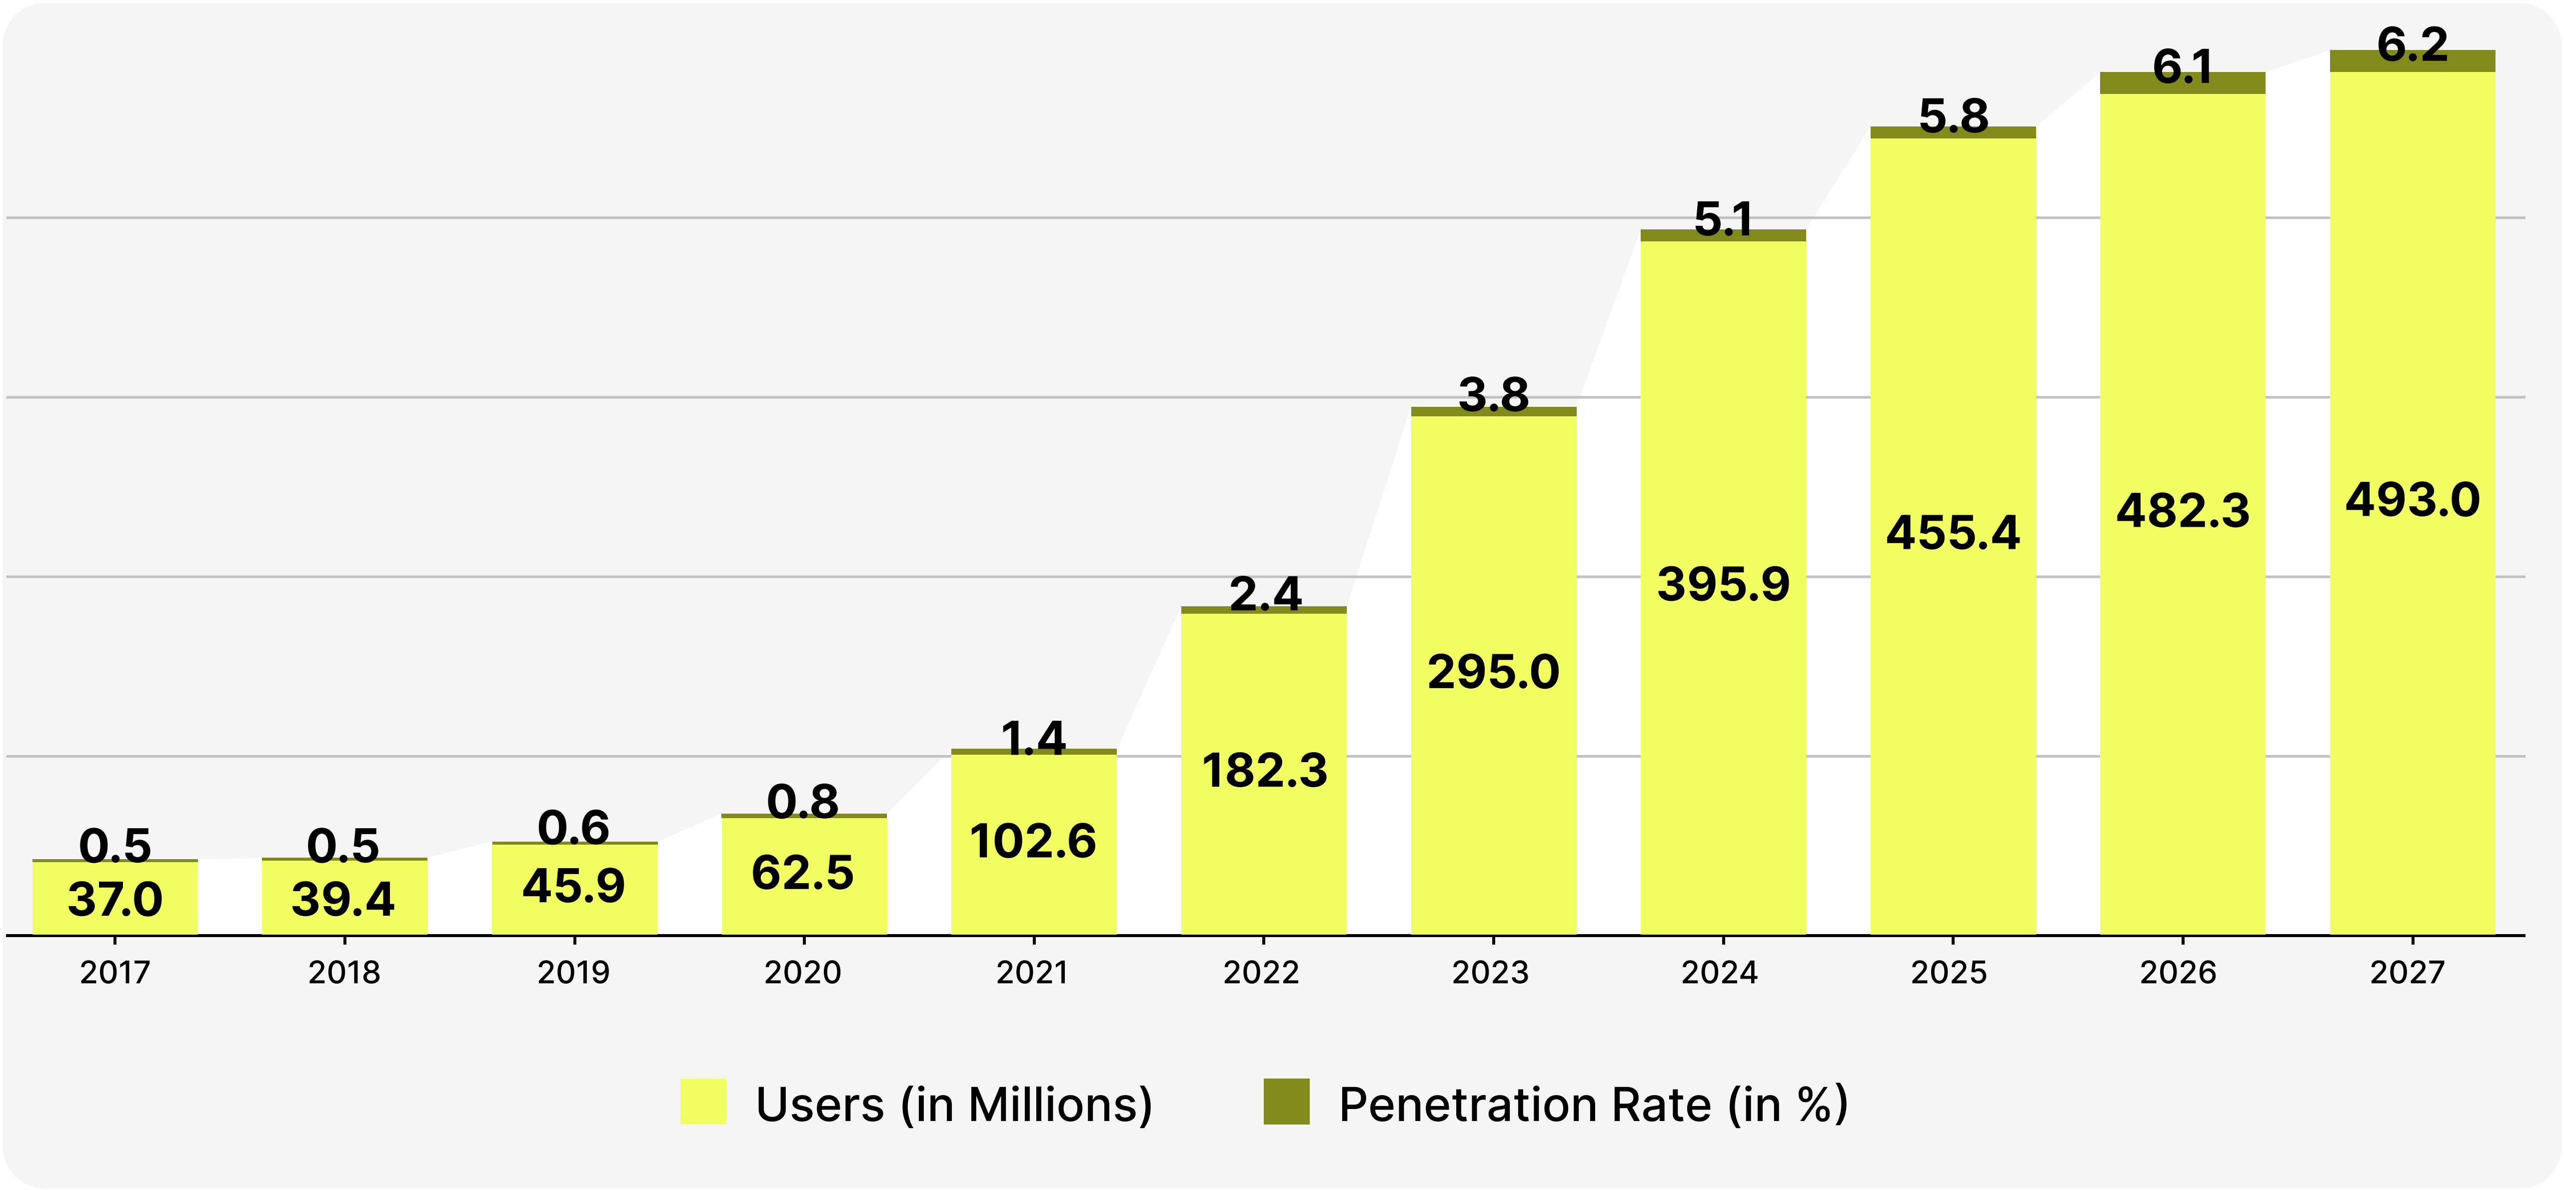

Source: Market.us Scoop

Looking at adoption, the cloud gaming market is largely built around two models. The first is the subscription library, exemplified by Xbox Cloud Gaming (Game Pass) and Amazon Luna. For a flat monthly fee, players gain access to extensive catalogues, often bundled with other digital services. This model drives recurring revenue and reduces entry barriers, but catalogues remain fragmented as publishers negotiate licences independently.

CLOUD GAMING MARKET

Source: Newzoo

The second is bring-your-own-game, led by NVIDIA GeForce NOW. Players stream titles they already own on platforms like Steam or Epic, extending libraries across devices. It appeals to core gamers, but the economics rely on hourly or tiered fees, typically with limits on session length or concurrent use. Both models prove strong demand, yet both are constrained by centralised infrastructure.

These two business models are supported by a broader ecosystem of providers, with the most critical being data centres and edge networks. To serve users, operators rely on a hybrid of large-scale data centres and regional edge servers. Data centres provide scale, while edge facilities reduce latency in urban areas.

This setup improves responsiveness but remains inefficient: each session requires a dedicated GPU, leading providers to overprovision for peak demand and leaving hardware underutilised during off-peak hours. Extending edge capacity to less dense regions is also costly, resulting in uneven service quality.

CLOUD TO DEVICE LAYERS STRUCTURE

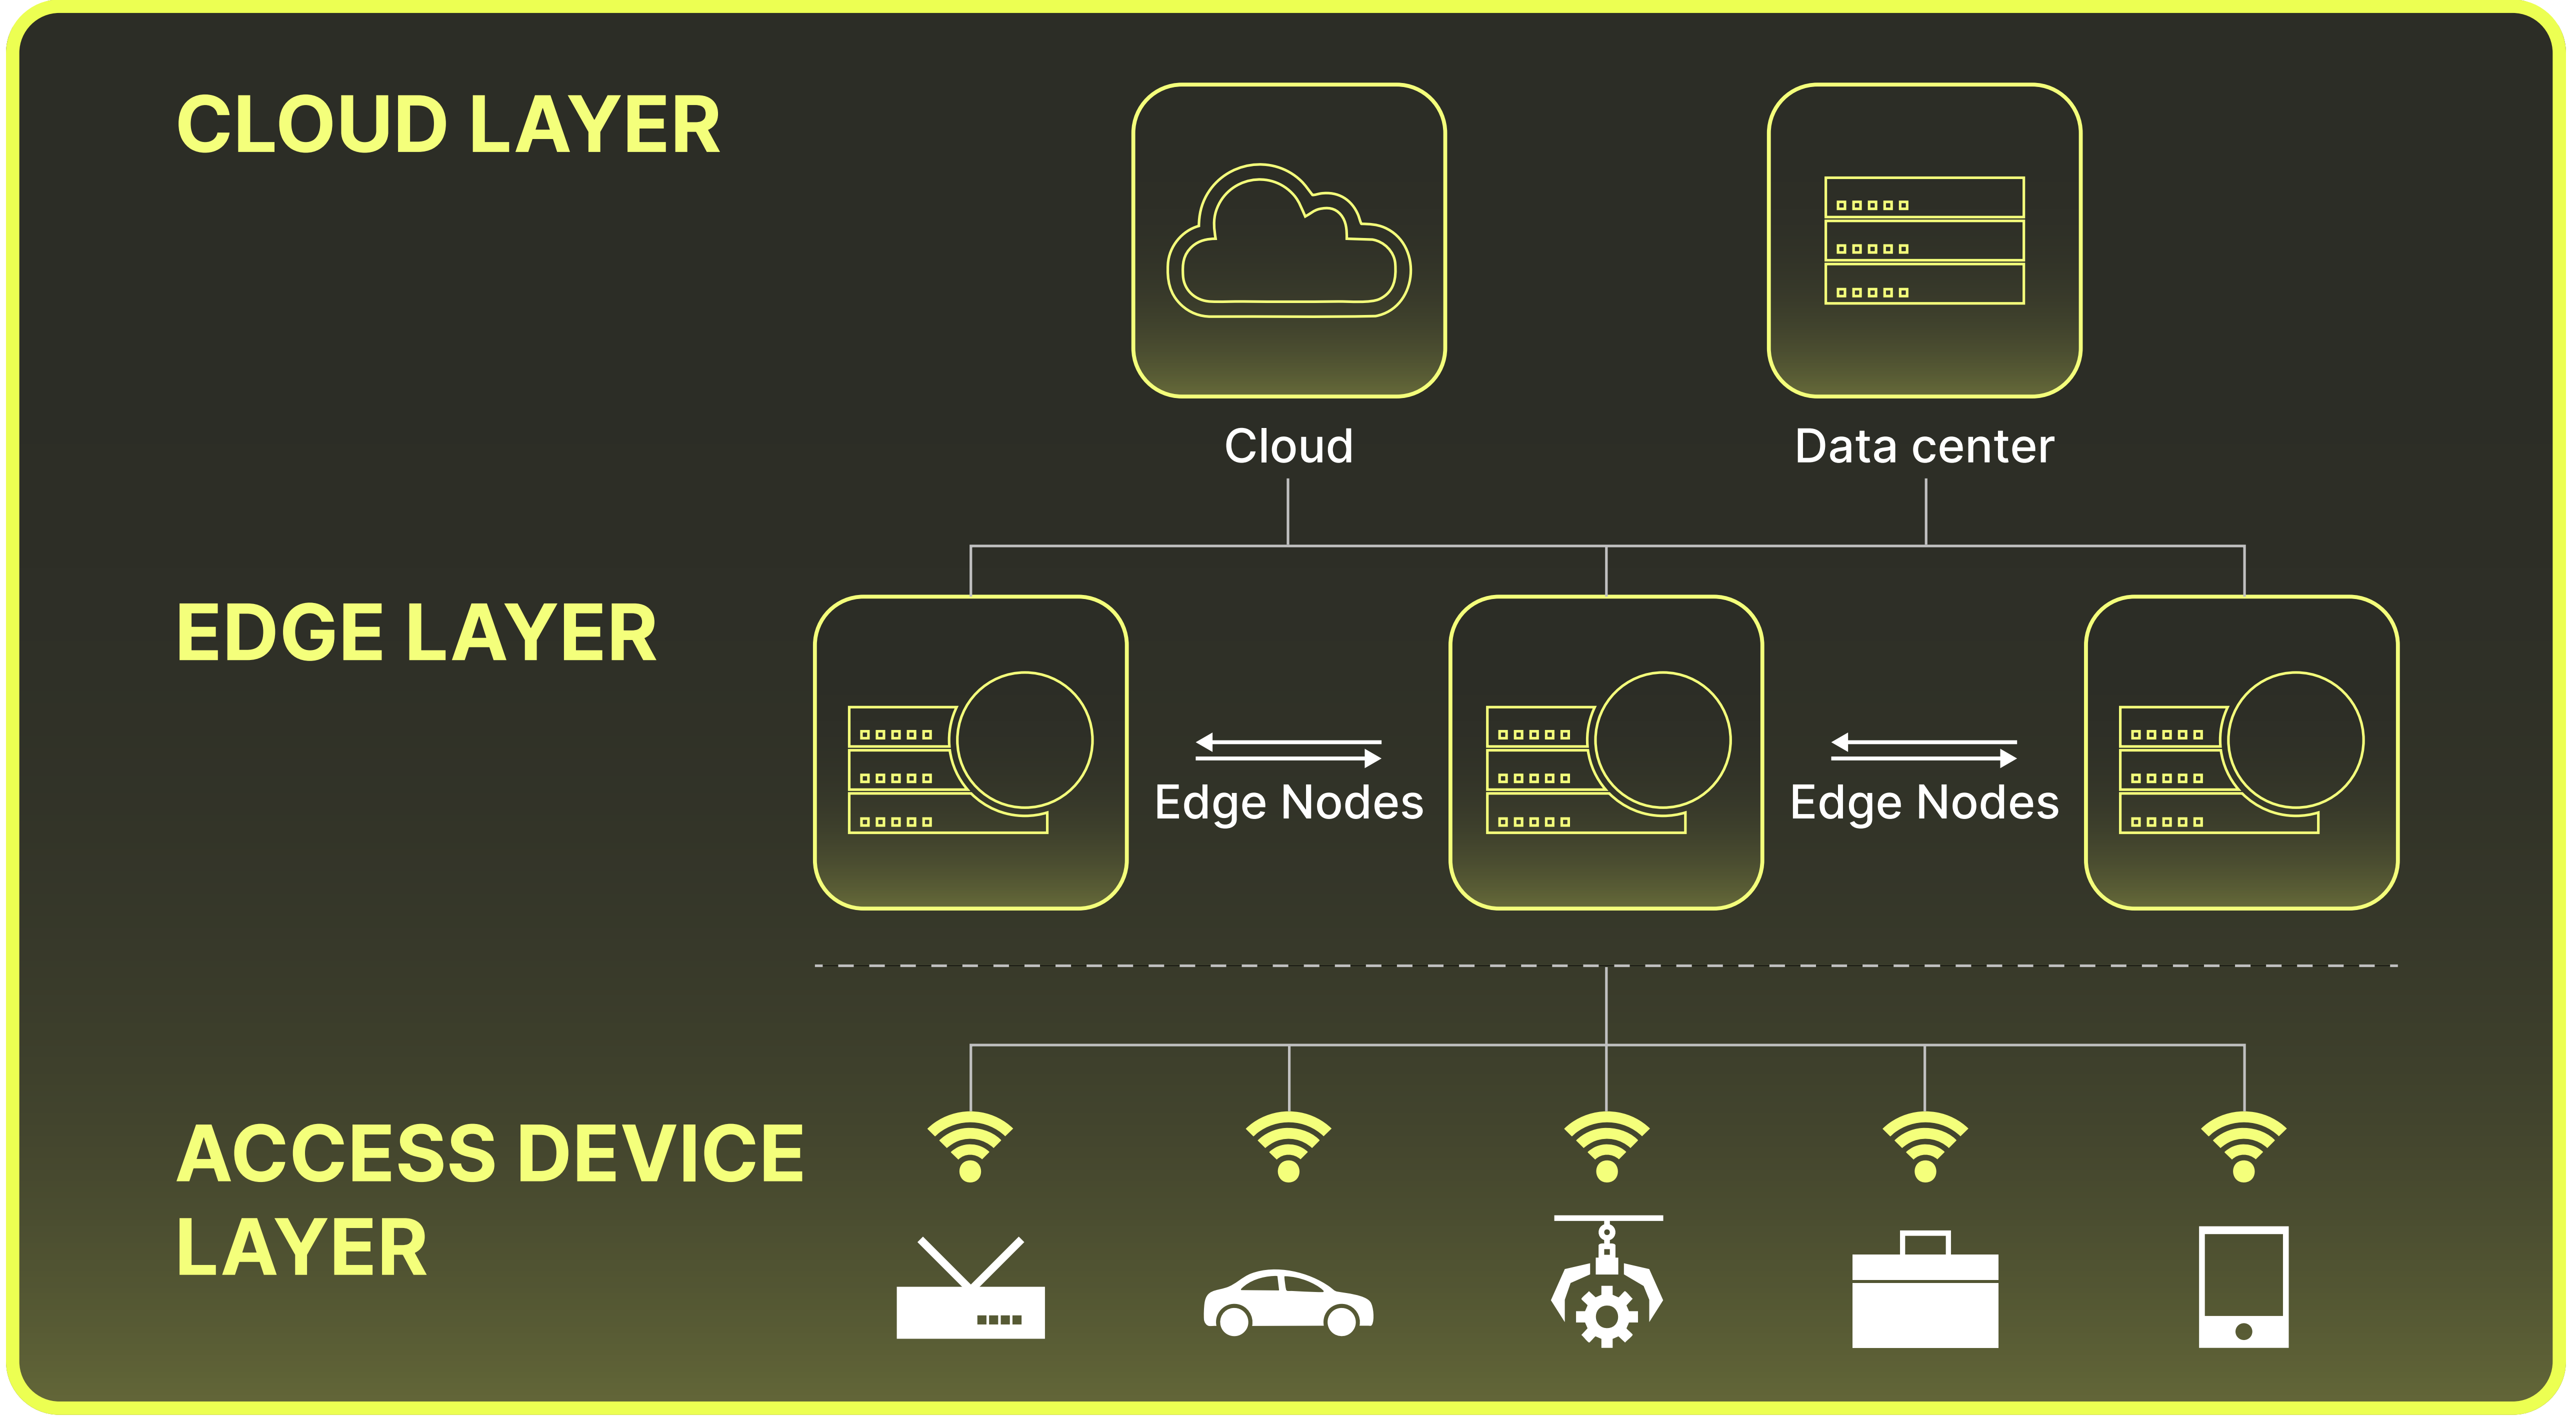

Source: Newzoo

These constraints revealed a gap in the market. Instead of pouring resources into bigger data centres or trying to extend costly edge networks, capacity could come from the idle GPUs already sitting in gamers' PCs.

Shaga built its model around this idea, using a peer-to-peer system where players run lightweight software to share spare GPU power. Other users can then stream games from these nearby nodes, delivering lower costs and better latency than centralised data centres or regional edge facilities.

SHAGA A PEER-TO-PEER SYSTEM

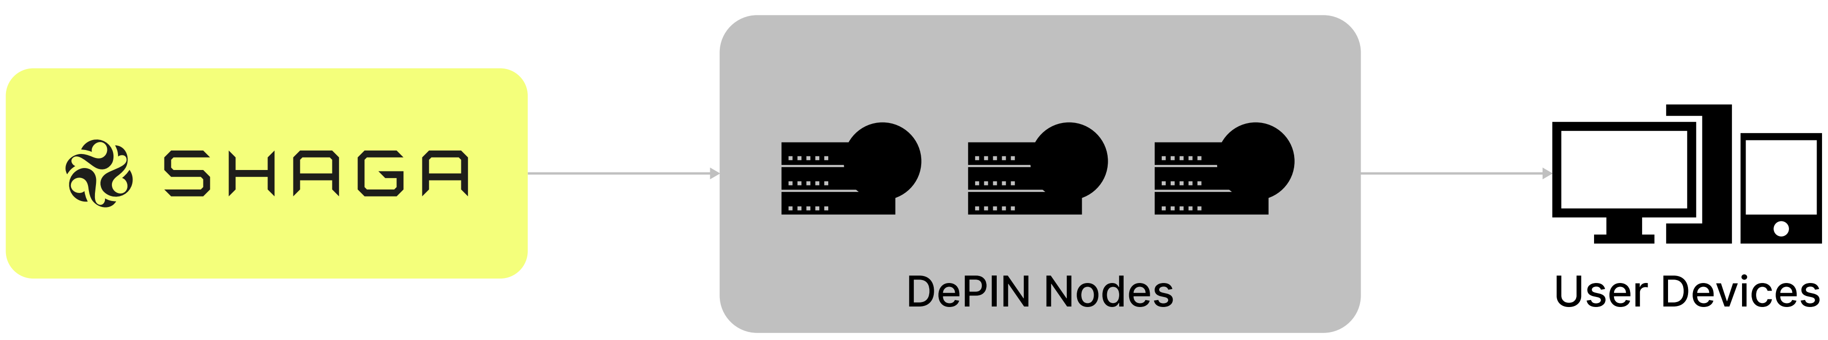

By transforming gaming PCs into micro data centres, Shaga creates a hyper-local network powered by DePIN. Compute is delivered at the true edge, directly within the same cities and neighbourhoods where players live. Because the network leverages existing hardware, it avoids the heavy capital expenditure of building new facilities or partnering with telecom providers, allowing capacity to be offered at a structurally lower cost than traditional designs.

AI-Ready Gameplay Data Market

Alongside its traditional business model, cloud gaming creates an additional layer of value. Beyond entertainment, it generates a rare by-product: video streams synchronised with real-time player inputs. These control-annotated datasets are highly sought after for AI development.

In robotics, they provide sequential decision-making signals, while in generative media, they supply behavioural data that improves realism. This combination makes gameplay streams significantly more useful than raw video alone.

AI TRAINING DATASET MARKET SIZE

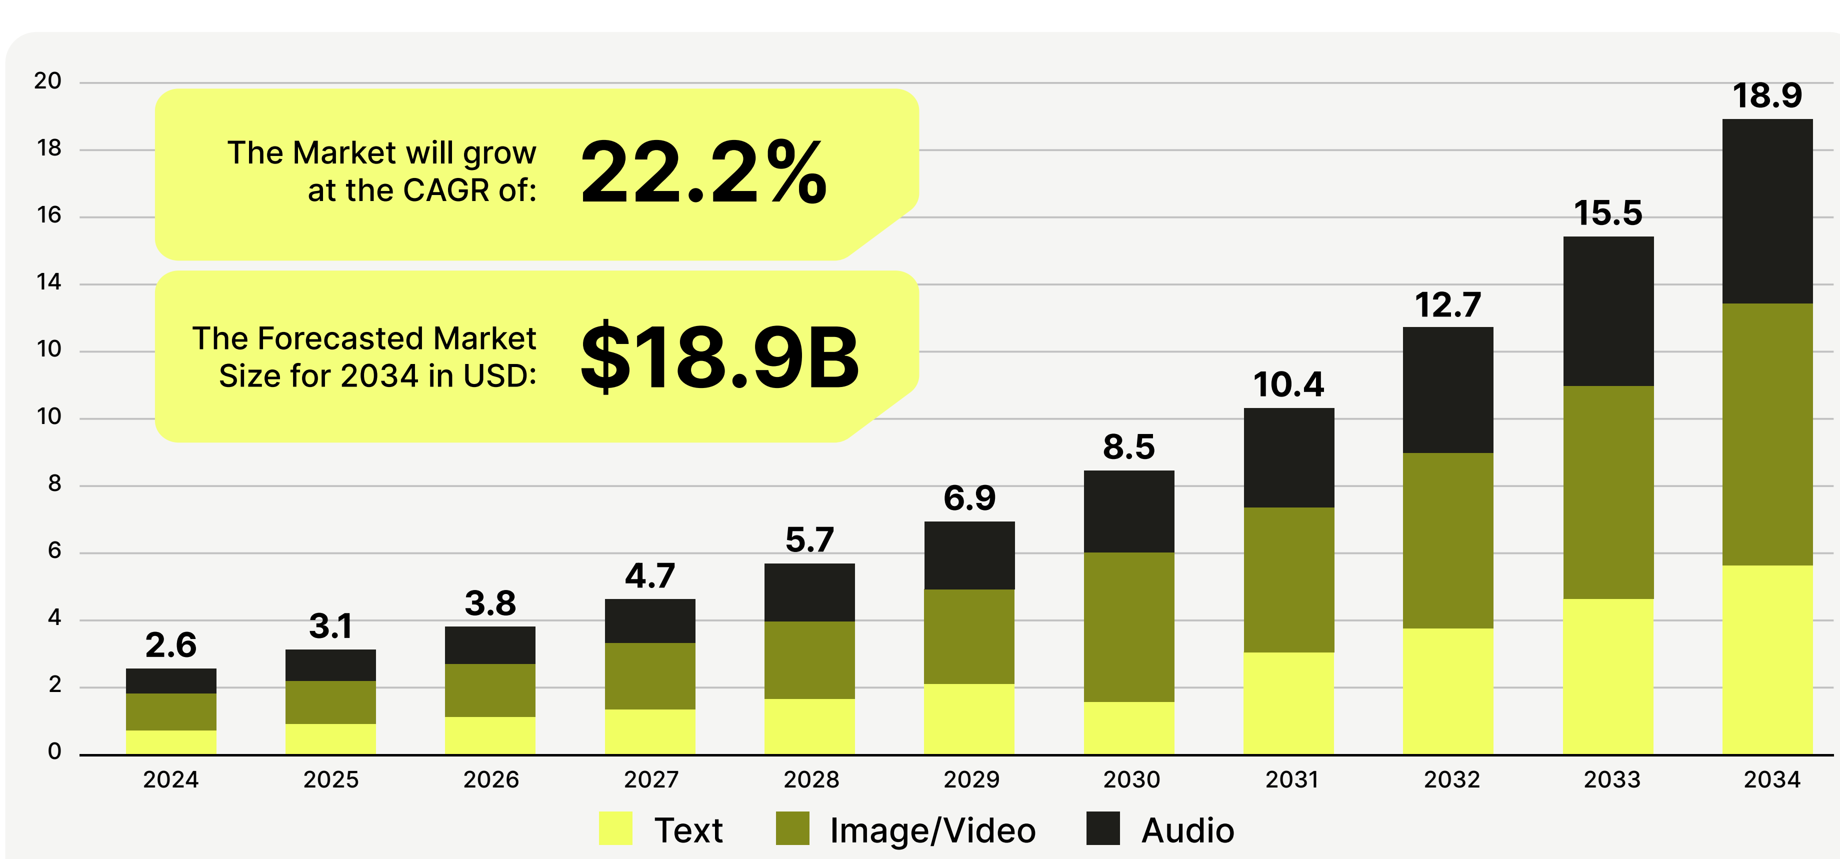

Source: Market.us

The rise of AI has created fierce competition for unique datasets, as companies look for any edge in training their models. This surge in demand has pushed the training data market to around $2.6 billion in 2024, with forecasts suggesting it could reach $12.7 billion by 2032.

Pricing benchmarks from providers such as Defined.ai show rates of $100-$300 per hour for well-annotated video with control metadata, with even higher prices for exclusive or richly tagged streams.

Shaga's architecture is designed to capture and monetise this opportunity. With an opt-in model, players can resell anonymised gameplay data, turning a by-product into a high-margin revenue stream.

This creates a circular economy: players enjoy cheaper sessions, node operators earn from both compute and data resale, and AI developers gain reliable access to premium datasets. Together, these dynamics make Shaga's business model far more sustainable.

Business Model Comparison

To understand where Shaga fits, it is useful to contrast the three main approaches to cloud gaming. Subscription libraries, bring-your-own-game platforms, and decentralised peer-to-peer networks each solve the same problem in different ways.

Lining them up on cost, latency, and monetisation reveals why Shaga's model points toward a potentially superior path.

BUSINESS MODEL COMPARISON

MODEL	EXAMPLES	HOW IT WORKS	STRENGHTS	WEAKNESSES
Subscription Library	Xbox Cloud, PlayStation Plus, Amazon Luna	Players pay a monthly fee to access a library of games, often bundled with other services.	Easy entry for users, predictable revenue for providers.	Catalogues are fragmented, expensive licensing, not all games available.
Bring-Your-Own-Game	NVIDIA GeForce NOW	Players stream games they already own on Steam or Epic, paying hourly or tiered fees.	Let's core gamers use existing libraries across devices.	Usage fees add up, sessions often capped, still costly to run.
Decentralised Peer-to-Peer	Shaga	Players share idle GPU power through lightweight software. Other users stream games directly from these peers.	Lower costs, local latency improvements, and adds revenue from reselling gameplay data.	Relies on community adoption and enough nodes in each region.

When viewed side by side, the comparison highlights a clear gap in the market. Subscription and BYO-game services are weighed down by licensing costs and infrastructure overhead, while smaller startups lack the scale or consistency to compete seriously.

Shaga, if it can scale its peer-to-peer network, offers a model with lower costs and reduced latency that could provide a superior alternative to both established incumbents and emerging challengers.

Technical & Network Architecture

Cloud gaming presents a technical paradox: delivering AAA-quality interactivity requires enormous compute power, yet any additional latency undermines the very experience it tries to provide. Centralised architectures address this by overbuilding data centres and edge facilities, but this leaves capacity underutilised and fails to cover large parts of the globe.

Shaga takes a different approach. Its architecture is designed as a cognitive grid, a decentralised mesh of high-performance gaming PCs that execute compute, storage, and bandwidth functions within the same unit.

This section analyses how the grid is structured, how the streaming pipeline is engineered, and how new mechanisms such as control bifurcation and geospatial incentives position the system to scale.

Cognitive Grid Design

The core of Shaga's architecture is the reconfiguration of idle consumer gaming PCs into micro data centres. These machines already contain the key elements required for cloud gaming: powerful GPUs, multi-core CPUs, high-speed RAM, NVMe storage, and broadband connectivity. Instead of replicating these at enormous capital cost inside centralised data centres, Shaga aggregates them into a peer-to-peer mesh. Each node operates as a vertically integrated package, delivering every component of the game session locally.

TRADITIONAL VS. SHAGA: NETWORK ARCHITECTURE

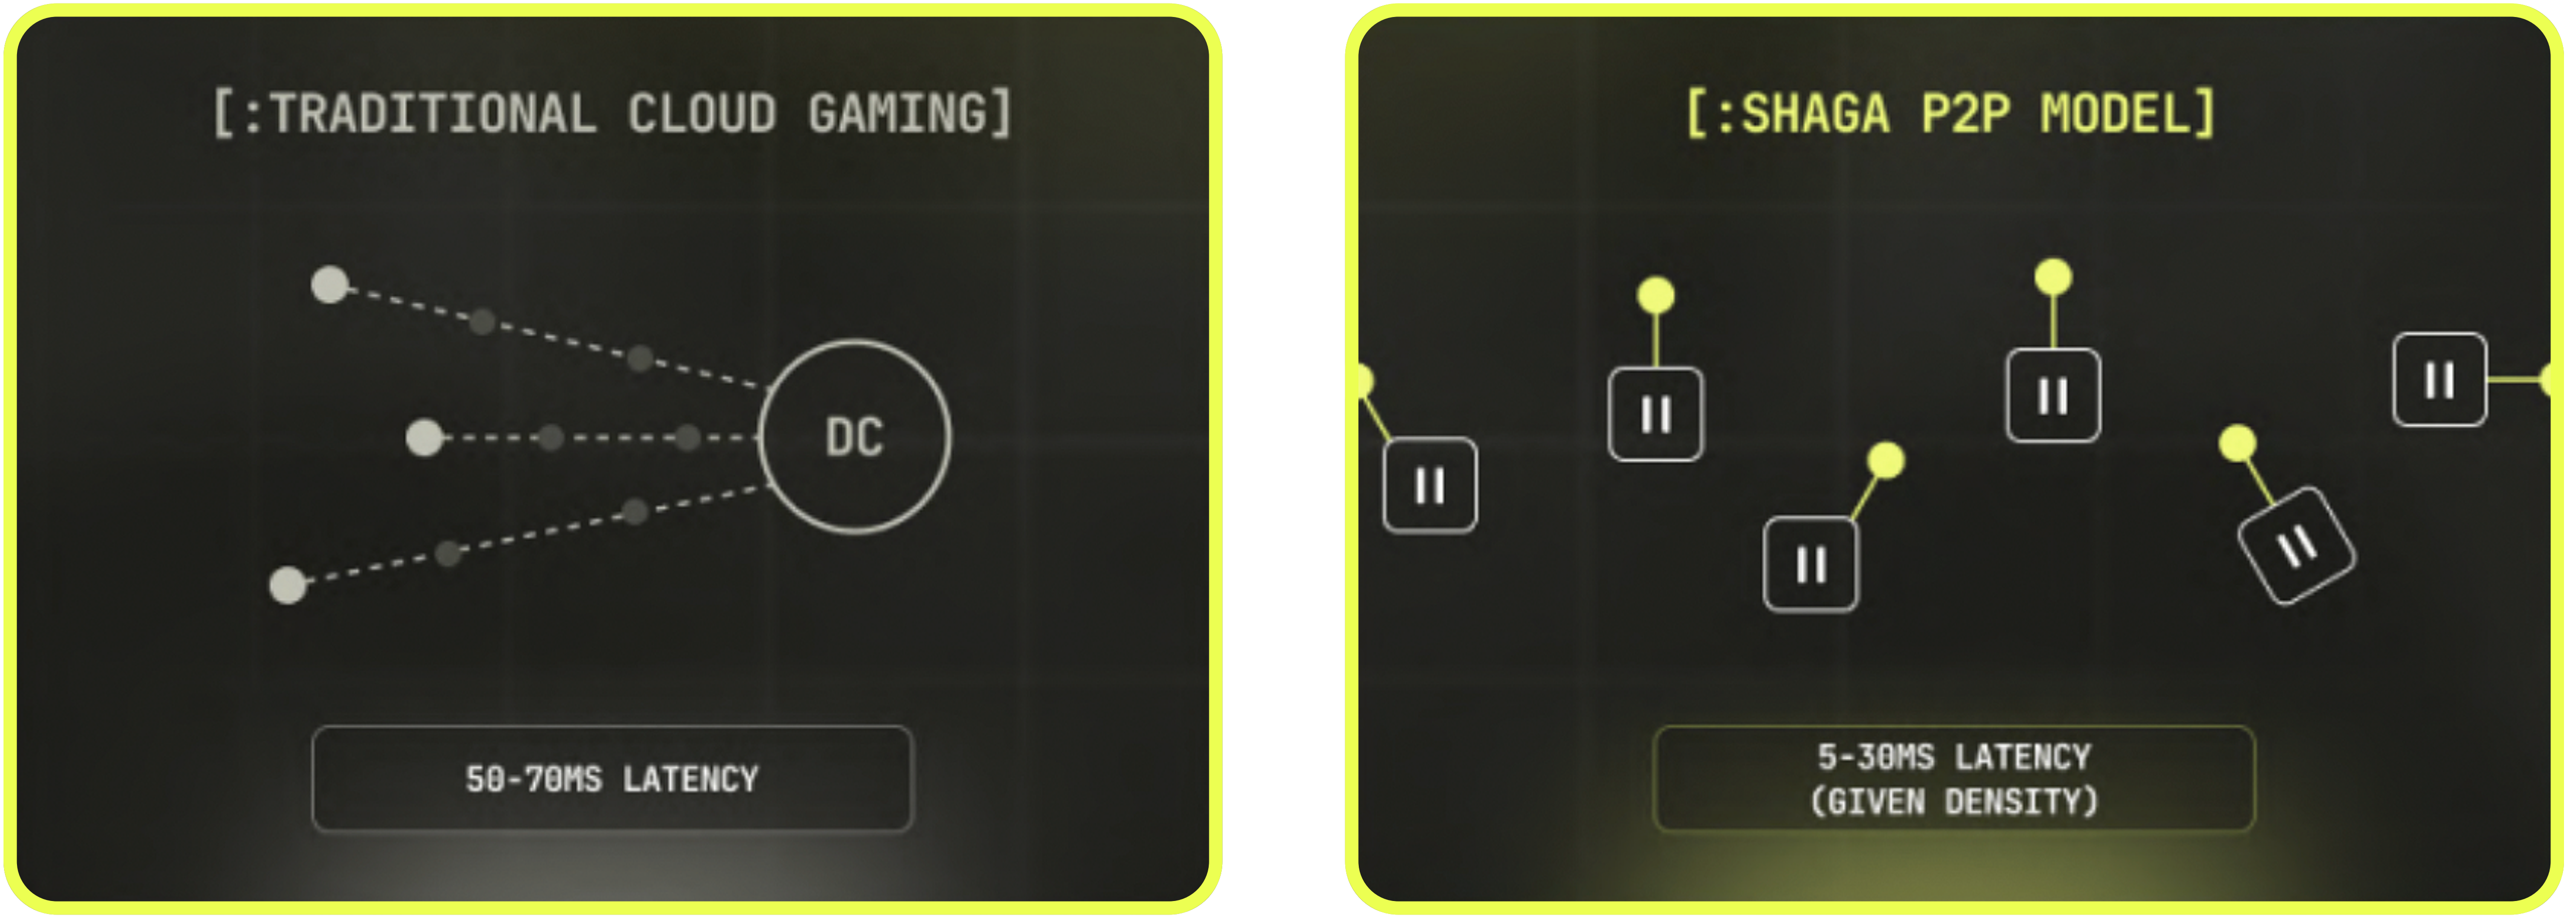

Participation is structured around three roles. Globbers are users who register their hardware with the network. As of mid-2025, more than 404,000 wallets have signed up, reflecting strong latent supply. From this pool, active nodes, currently 767 in operation, deliver real compute to the network. On the demand side, clients access sessions via lightweight applications on iOS, Android, or desktop, lowering the barrier for entry into AAA gaming.

The network uses latency-aware routing to connect these actors. When a client requests a session, the system probes available peers and assigns the connection to the lowest-latency node. Shaga employs the H3 hexagonal indexing framework, developed by Uber for geospatial analytics, to divide the globe into uniform cells (more information in section 4). Each hexagon functions as a micro-market for supply and demand, with the protocol capable of monitoring utilisation and incentivising new nodes where capacity falls short.

SHAGA: DUAL-LAYER ARCHITECTURE

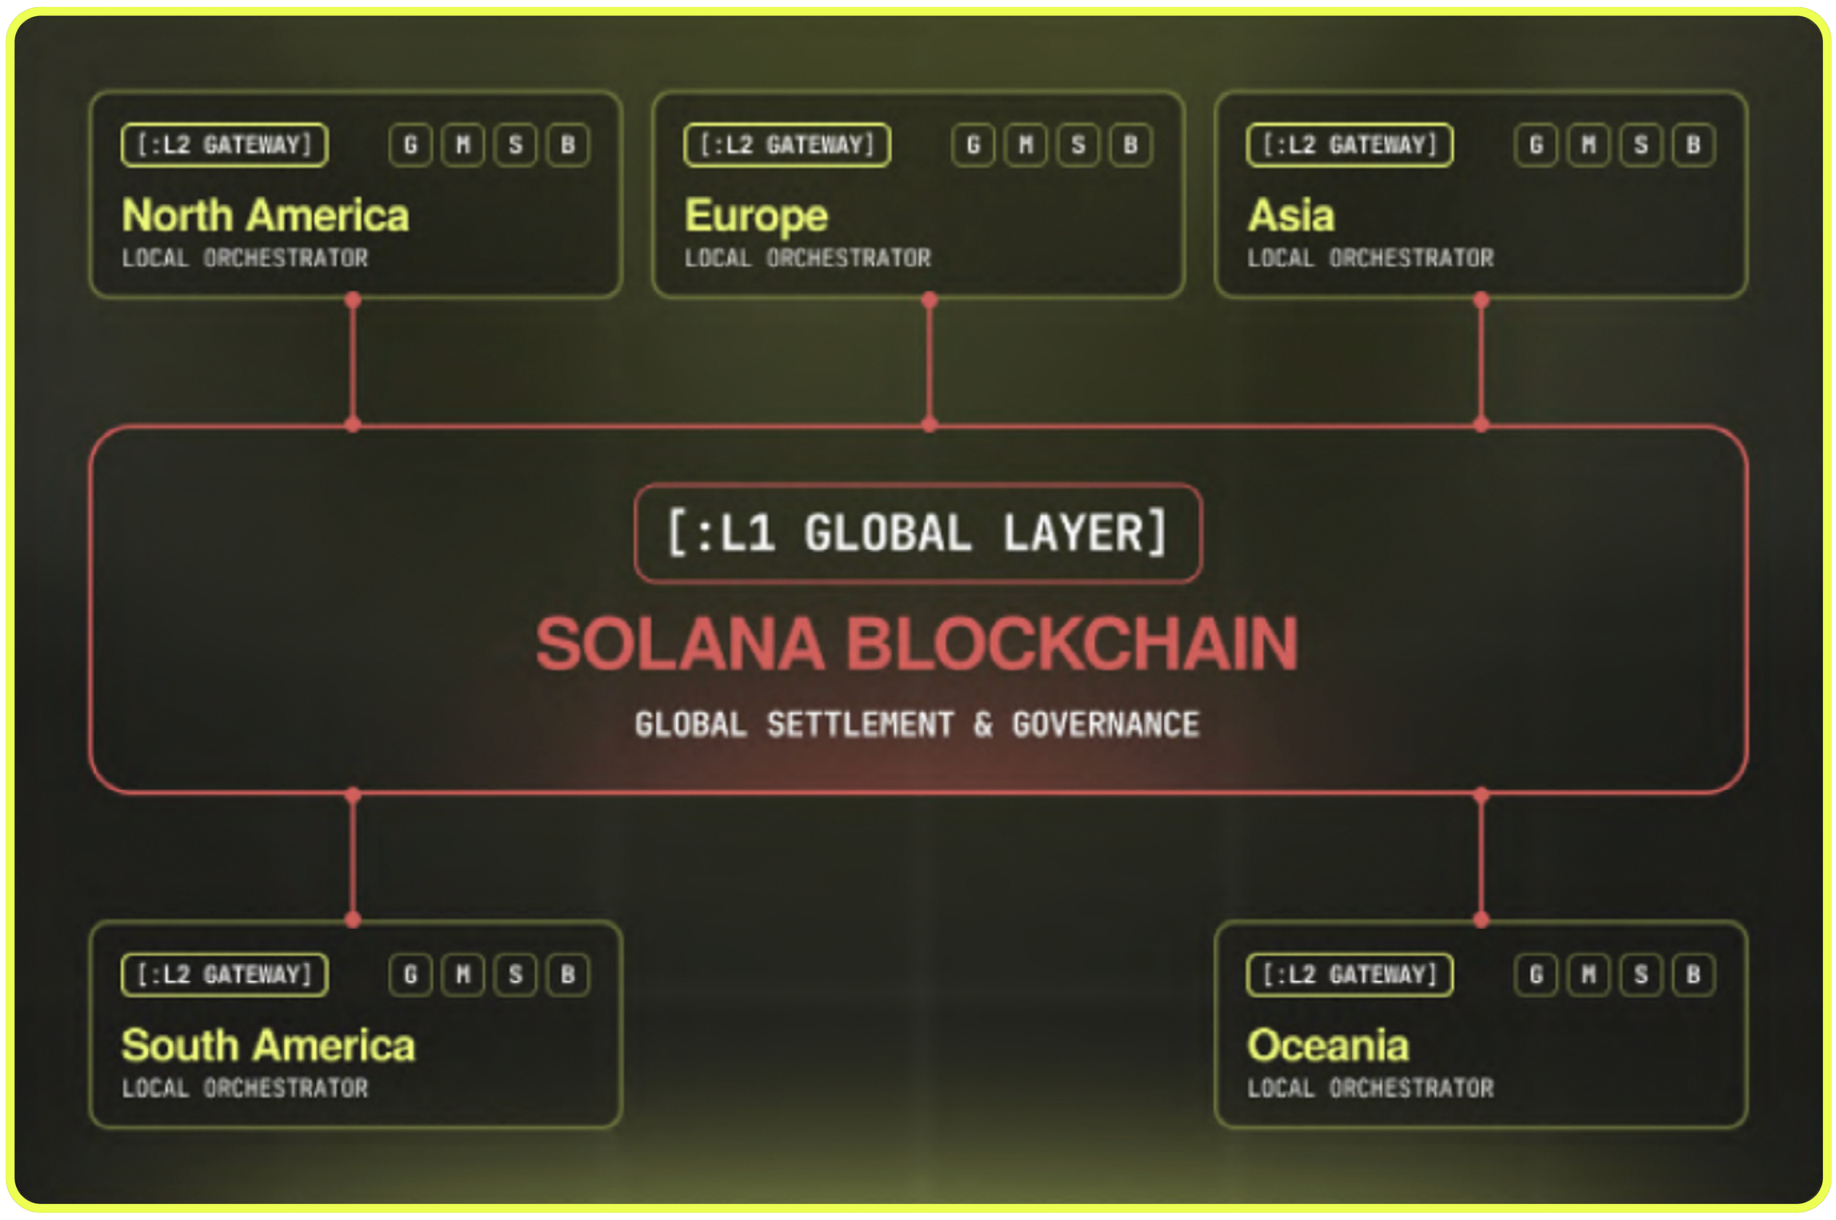

Performance data from Shaga's invite-only phase indicates that this design produces results competitive with local play. In underserved regions, median round-trip times below 40 milliseconds have been recorded, a level sufficient for most real-time genres. For comparison, centralised cloud services often exceed 80-120 milliseconds outside major metropolitan hubs. These findings confirm that geographic distribution of compute is the key determinant of cloud gaming quality, and that consumer PCs, when coordinated properly, can provide superior responsiveness.

Video Encoding Pipeline

If latency-aware routing determines where a game session is processed, the video encoding pipeline determines how it is delivered. Conventional cloud gaming services rely on standard codecs that compress and decompress raw frames, a process that both consumes bandwidth and introduces delay. Bandwidth costs are a significant constraint: Twitch, for example, is estimated to spend between $4 and 6 million per month to deliver high-resolution video streams to viewers.

Shaga addresses this challenge with Neural Game Codecs (NGCs), an experimental but potentially promising R&D, targeting 75% bitrate cut with title-specific neural codec technology. Instead of sending pixels, the system sends causality. Controller inputs are transmitted to the codec, which predicts frame generation locally on the client device. Only when events deviate from the model, such as an unexpected explosion or a user interface change, small delta corrections are transmitted.

The efficiency gains are considerable. Early estimates suggest NGCs can reduce bandwidth requirements by up to 75%, lowering the average stream from 6 Mbps to around 1.5 Mbps. This means that a node with a 100 Mbps connection could host more than 66 concurrent streams, compared to fewer than 20 under conventional methods. Just as importantly, every gaming session generates paired data (frame plus input), which is then fed back into the system to retrain codecs. Over time, this creates a self-reinforcing cycle of hyper-optimised, game-specific codecs that lower costs and improve quality.

The implication is that Shaga could not only deliver gaming more efficiently but also compete with existing content delivery networks (CDNs). If codecs can make streaming economically viable on mobile data connections as low as 0.5 Mbps, the model could extend far beyond gaming into media distribution, turning what is traditionally a cost centre into a potential revenue engine.

Control Bifurcation

Multiplayer environments highlight another architectural weakness of conventional cloud gaming: input relay. Typically, a player's commands are routed first to the streaming node and then forwarded to the game's server, adding unnecessary delay. Shaga's proposed solution, control bifurcation, is to split these data paths. Player inputs would bypass the streaming node entirely and be transmitted directly to the game server, while video continues to flow from the peer-to-peer node.

Although still in development, early modelling suggests that this adjustment could reduce multiplayer latency by 15 to 45 milliseconds. For casual players this may be imperceptible, but for competitive environments where milliseconds determine outcomes, the difference is critical. This feature represents an incremental but important improvement, particularly if Shaga aims to compete in esports-level use cases.

Geospatial-Economic Layer

The cognitive grid's geospatial organisation also serves as the basis for Shaga's economic design. By using H3 indexing, the network can track utilisation and performance at a granular level, identifying areas where supply is under-provisioned or where latency improvements have the highest marginal value. Token incentives can then be calibrated to encourage node operators to join in specific regions, aligning rewards with real-world demand.

This design links emissions to measurable work, primarily electricity consumption. Instead of inflationary token release schedules, Shaga's model emits tokens when nodes perform verifiable compute tasks, such as rendering frames or training Neural Game Codecs. This grounds token value in physical resource use, while creating a natural selection mechanism that favours efficient hardware and cheaper electricity sources.

The result is a system where the same architecture that optimises latency also drives a sustainable incentive structure. Technical performance and economic sustainability are not separate challenges but two sides of the same grid, with geography serving as the bridge between them.

CPU and GPU: A Diverse Set for Efficiency and Resilience

The Shaga network had 217 CPUs and 406 GPUs at the time of analysis but the pool is expanding rapidly as more and more users turn their devices into nodes.

While still early in scale, this dataset provides a representative picture of the types of hardware available within the network. Analysing the spread of both CPUs and GPUs helps in understanding the network's capabilities, its inclusivity, and the foundations of resilience that underpin Shaga's peer-to-peer model for cloud gaming and computation.

CPUs in the Network

CPUs form the backbone of gaming and cloud computing infrastructure. They manage the orchestration of tasks, game logic, physics, and system-level operations that allow GPUs to focus on rendering and parallel workloads. Within Shaga, the CPU landscape demonstrates a clear distribution across performance tiers, reflecting both the diversity of contributors and the breadth of applications the network can support.

The most common CPUs in the dataset include the AMD Ryzen 5 5600 series (19 units), AMD Ryzen 7 5800X (9 units), and Intel Core i5-12400F (6 units), with several nodes contributing modern six- to eight-core processors. Together, these form a strong mainstream tier, well suited to powering high-quality gaming sessions and parallel cloud workloads.

At the upper end of the spectrum, standout contributors include Ryzen 9 series processors and newer Intel i9-14900KF units. These represent high-performance outliers with double-digit core counts and clock speeds suited not only for high-end gaming but also for demanding AI inference, simulations, or multi-instance streaming. On the extreme end of server-grade power, the dataset even includes AMD EPYC processors with 32 cores, highlighting the presence of enterprise-class devices.

Equally important is the long tail of lower-tier CPUs such as Intel i3 laptops and older i5 generations. While individually less powerful, they contribute to Shaga's inclusivity and ensure that less resource-intensive tasks, such as lighter games or auxiliary processes can be efficiently distributed.

The resulting picture is a heterogeneous CPU pool: a network that combines mainstream reliability, high-performance outliers, and low-tier accessibility.

GPUs in the Network

GPUs are the cornerstone of rendering and cloud gaming performance, translating computational power into smooth gameplay and enabling high-throughput workloads such as AI training or 3D rendering. The Shaga dataset of 406 GPUs shows a similarly diverse and layered distribution to that of CPUs, balancing wide accessibility with powerful performance nodes.

In terms of prevalence, the NVIDIA RTX 4060 (39 units), the RTX 3060 series (32 units), the RTX 4070 series (29 units), and the RTX 3050 Laptop GPU (22 units) appear most frequently, together accounting for roughly 30% of all GPUs in the network. Collectively, these cards form a strong mainstream base, offering sufficient performance for modern gaming workloads while maintaining a favourable balance of energy efficiency.

At the high end, the dataset includes RTX 4090 GPUs (11 units), among the most advanced consumer graphics cards currently available. These devices deliver extreme rendering capabilities, real-time ray tracing, and advanced AI workload performance. While limited in number, their presence demonstrates the network's capacity to integrate cutting-edge hardware. Other notable high-performance contributors include the RTX 4080 (9 units), which provide substantial support for high-fidelity, latency-sensitive gaming streams.

Conversely, a portion of the network is composed of older-generation GPUs, including GTX 1050 (10 units) and GTX 1650 cards (15 units), as well as Radeon RX 580 (10 units). While these are no longer optimal for AAA cloud gaming, they remain highly valuable for lighter titles, parallel workloads, or off-peak use cases. Interestingly, the dataset also includes a share of legacy and virtualised devices (e.g., Microsoft Basic Render Driver, Parallels Display Adapter), signalling experimentation with non-standard configurations. The resulting GPU pool mirrors the CPU landscape: a blend of dominant mainstream units, a tail of legacy cards, and high-performance outliers that push the upper limits of capability.

What This Tells Us

Taken together, the CPU and GPU composition of Shaga reflects more than a collection of hardware; it reveals a structural strength grounded in diversity. This diversity enables Shaga to position itself not only as a cloud gaming solution but as a broader decentralised computing fabric. Four strategic themes emerge:

Inclusivity: By accommodating a broad range of CPUs and GPUs, Shaga ensures that virtually any participant can contribute to the network. From high-end RTX 4090s and Ryzen 9 CPUs to everyday laptops with mid-range processors, the system captures untapped potential from across the user spectrum. This inclusivity lowers barriers to entry, fosters community growth, and creates the foundation for long-term adoption.

Resilience: Traditional cloud infrastructure is vulnerable to shocks such as supply shortages, price volatility, or single-vendor dependencies. In contrast, Shaga’s decentralised hardware pool spreads risk across hundreds of devices, manufacturers, and performance levels. If one class of GPU becomes scarce or expensive, the network continues to function through the contribution of other devices. This redundancy makes Shaga resilient in the face of hardware cycles, market fluctuations, or disruptions in global supply chains.

Edge Advantage: Because Shaga’s contributors are distributed across geographies and hardware classes, the network effectively functions as an edge computing topology. Workloads can be routed closer to end users, reducing latency and improving responsiveness. Lower-tier devices play a role here by servicing lighter, localised tasks, while high-end devices handle latency-sensitive, compute-intensive workloads.

Efficiency: The diversity of CPUs and GPUs creates opportunities for task matching. High-performance GPUs can be reserved for graphically demanding games or AI inference, while legacy GPUs and low-power CPUs handle auxiliary or lightweight processes. This distribution ensures that the network extracts maximum utility from every node, minimising waste and allowing Shaga to scale without unnecessary duplication of high-end hardware.

To sum up, the Shaga network shows how a heterogeneous pool of CPUs and GPUs becomes a strategic asset. Mainstream devices provide stability, high-performance outliers push the upper limits, and legacy hardware broadens participation. Together, this mix builds resilience, supports edge distribution, and enables efficient task allocation. As the network grows, diversity will amplify these advantages positioning the network as a resilient infrastructure layer for global computation.

Geospatial Demand-Supply Analysis

One of Shaga's defining strengths lies in the richness of its geospatial dataset, which captures both supply and demand with high accuracy. This is particularly critical for Shaga, where user experience is highly dependent on geographical distance between nodes and gamers. Latency remains a decisive factor in cloud gaming, and Shaga's ability to anchor supply close to demand ensures the lowest possible response times.

To achieve this, Shaga employs the H3 mapping model, a state-of-the-art geospatial framework originally developed by Uber and widely regarded as one of the most precise ways of dividing the earth into uniform zones. H3 works by partitioning the globe into hexagons of varying resolution. At resolution 3, the world is mapped into 12,392 km^2 cells, for a total of 41,162 hexagons. Increasing granularity to resolution 5 divides these further into 2,016,842 unique cells of 252 km^2 each, allowing detailed analysis of supply and demand distribution. This flexibility means that Shaga can study both macro trends across continents and highly localised infrastructure gaps at the city level (source : h3geo).

The dataset, already accessible via snapshot.shaga.xyz and soon to be expanded into a dedicated explorer, enables stakeholders to visualise core network metrics. It maps data centre locations, unserved gamers, globbers, clients, and nodes, offering a detailed view of both supply and demand. Beyond adoption signals, it also incorporates infrastructure variables such as electricity cost, a critical factor for sustainable scaling.

Far more than a visual tool, this mapping framework is at the core of Shaga's governance and token incentives. Incentives can be directed geographically to address imbalances in supply and demand, or to stimulate node growth in underserved but high-potential areas.

This section will examine supply and demand through the lens of this dataset. By reviewing unserved markets, Globber distribution, client conversion, and node infrastructure, we aim to identify areas of opportunity, anticipate constraints, and demonstrate how geospatial data strengthens Shaga's long-term positioning in the cloud gaming industry.

Globbers Concentrated in Underserved Regions

At the centre of Shaga's growth strategy is a precise understanding of unserved gamers. Put simply, these are players who want cloud gaming but cannot currently be matched with the infrastructure needed to support them.

To quantify this, Shaga uses a bottom-up Total Addressable Market (TAM) simulator at H3 resolution 5. Each hexagon estimates two things: how much qualified demand exists today, and how much of that demand could already be covered by ambient capacity, nearby, non-Shaga PCs that could in theory host gaming sessions.

This creates two categories of users:

1. Served gamers are those who could already be allocated to existing capacity within 300 km. They represent the portion of TAM that today's infrastructure could theoretically satisfy, even without Shaga.

2. Unserved gamers are those left behind: players who cannot be connected to enough capacity within reach. They form the residual demand, and therefore the clearest growth opportunity for Shaga.

It is this unserved group that represents Shaga's obvious target market. Current estimates suggest there are roughly 145 million such gamers worldwide, a vast, addressable audience that the network can unlock by directing incentives and deploying infrastructure in the right locations.

UNDESERVED AREAS (GREEN - RESOLUTION 3) **AND GLOBBERS** (YELLOW - RESOLUTION 5)

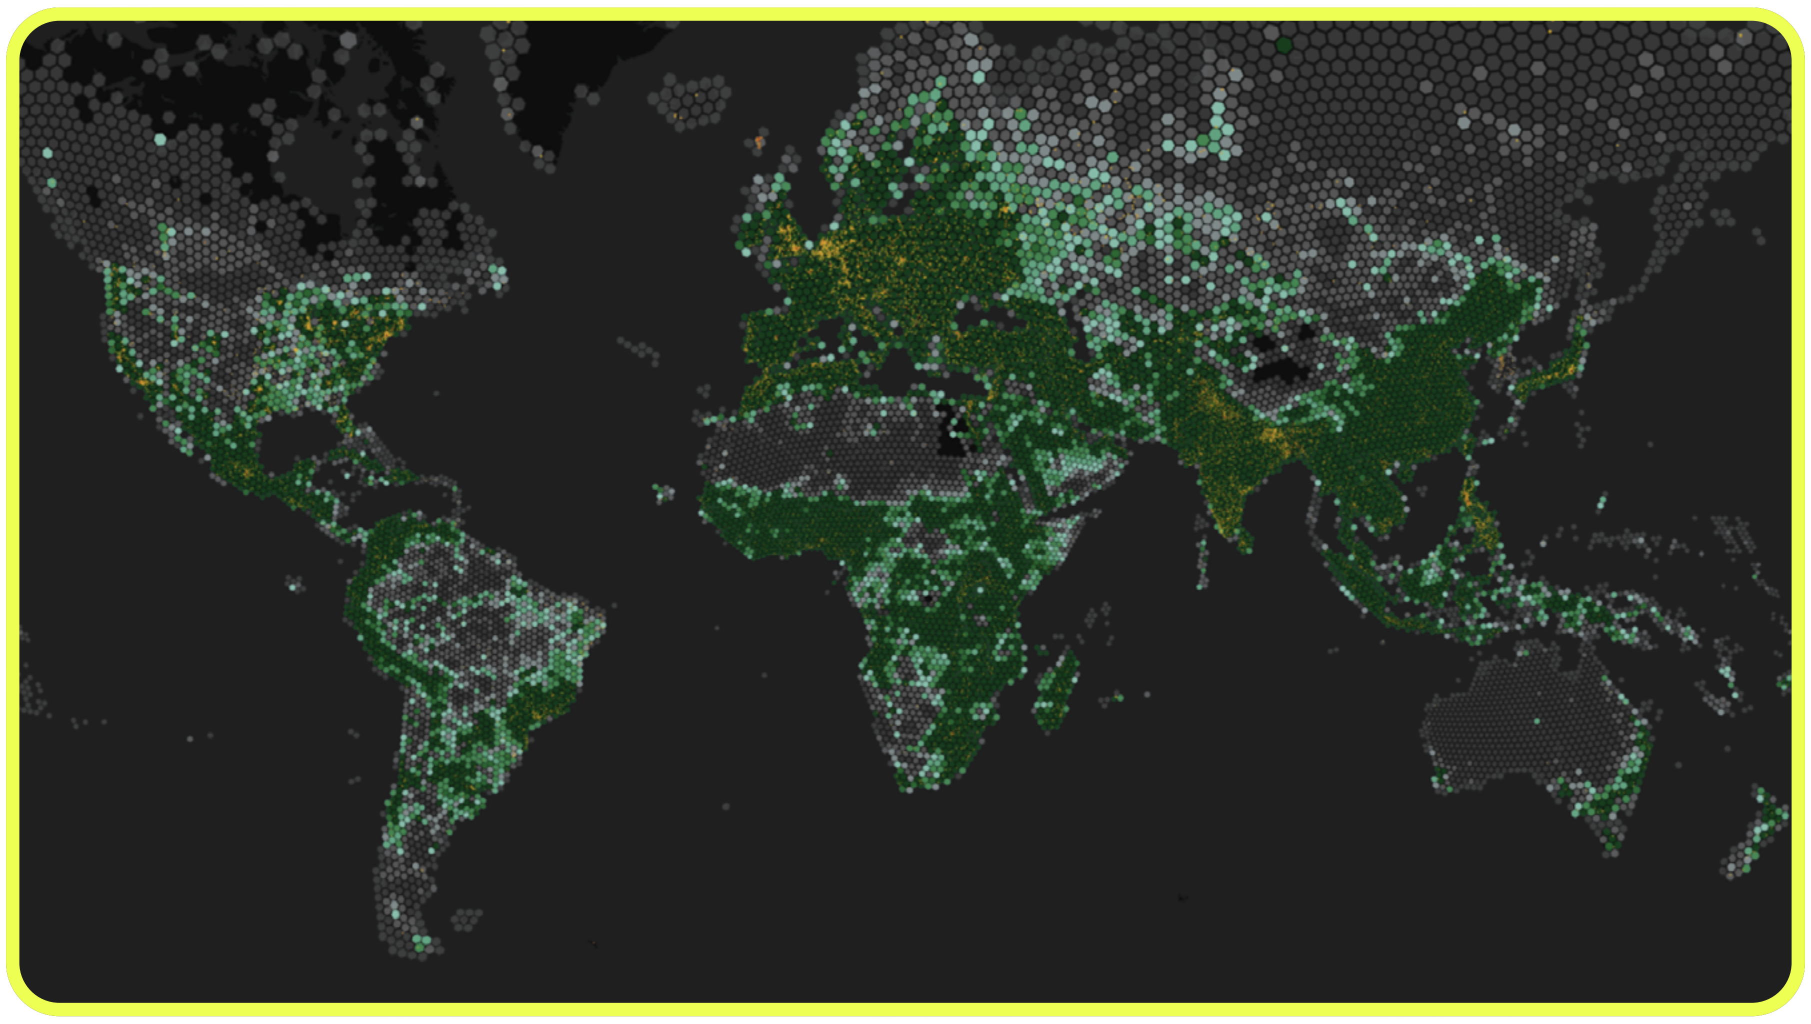

Remarkably, even before any structured incentive programme has been introduced, the network has begun attracting contributors known as Globbers, community members who make their GPU capacity available to Shaga. The most striking feature of this organic growth is its geographic alignment. Globbers have concentrated in the very same areas identified by the TAM simulator as having the highest concentrations of unserved gamers. This convergence between simulated demand and actual contributor behaviour validates the accuracy of the modelling and underscores the urgent need for infrastructure in these underserved regions.

Mapping Interest Through Globbers Distribution

At a global level, the dataset reveals a broad and rapidly expanding contributor base. Uganda leads the chart with 43,352 Globbers, accounting for roughly 19% of the total (217,253 Globbers in the dataset). This is followed by Indonesia with 24,085 Globbers, the United Kingdom with 12,887, China with 11,496, India with 8,039, and Serbia with 7,602. Taken together, these six countries represent around half of the total Globbers base.

NUMBER OF GLOBBERS (TOP 20 — COUNTRY LEVEL)

#	COUNTRY	GLOBBERS
1	Uganda	43352
2	Indonesia	24085
3	United Kingdom	12887
4	China	11496
5	India	8039
6	Serbia	7602
7	United States	6290
8	Australia	6052
9	Maldives	5928
10	Japan	5885

#	COUNTRY	GLOBBERS
11	Bangladesh	5731
12	North Macedonia	5283
13	Vietnam	4568
14	Nigeria	4012
15	Singapore	3908
16	Philippines	3694
17	Germany	3457
18	Luxembourg	2847
19	South Africa	2619
20	Mauritania	2556

Absolute numbers, while sometimes impressive, do not always capture relative traction. Countries vary widely in population, and another insightful metric is Globbers per capita. By adjusting for population size, the analysis highlights pockets of disproportionately strong interest. Smaller regions such as the Faroe Islands, Maldives, Macedonia, Serbia, Cape Verde and North Bangladesh stand out with high Globber density relative to their population. Uganda also remains prominent under this measure, confirming its importance as both an absolute and relative leader.

Further noteworthy concentrations appear across the United States, particularly in the North-East, North-West, and California. In Africa, Nigeria, Morocco, and the metropolitan cluster of Kinshasa stand out. In Asia, traction is visible in Pakistan, the Philippines, and Vietnam. Europe shows meaningful engagement across Germany, Austria, Hungary, and the metropolitan hub of London.

GLOBBERS PER CAPITA (RESOLUTION 3)

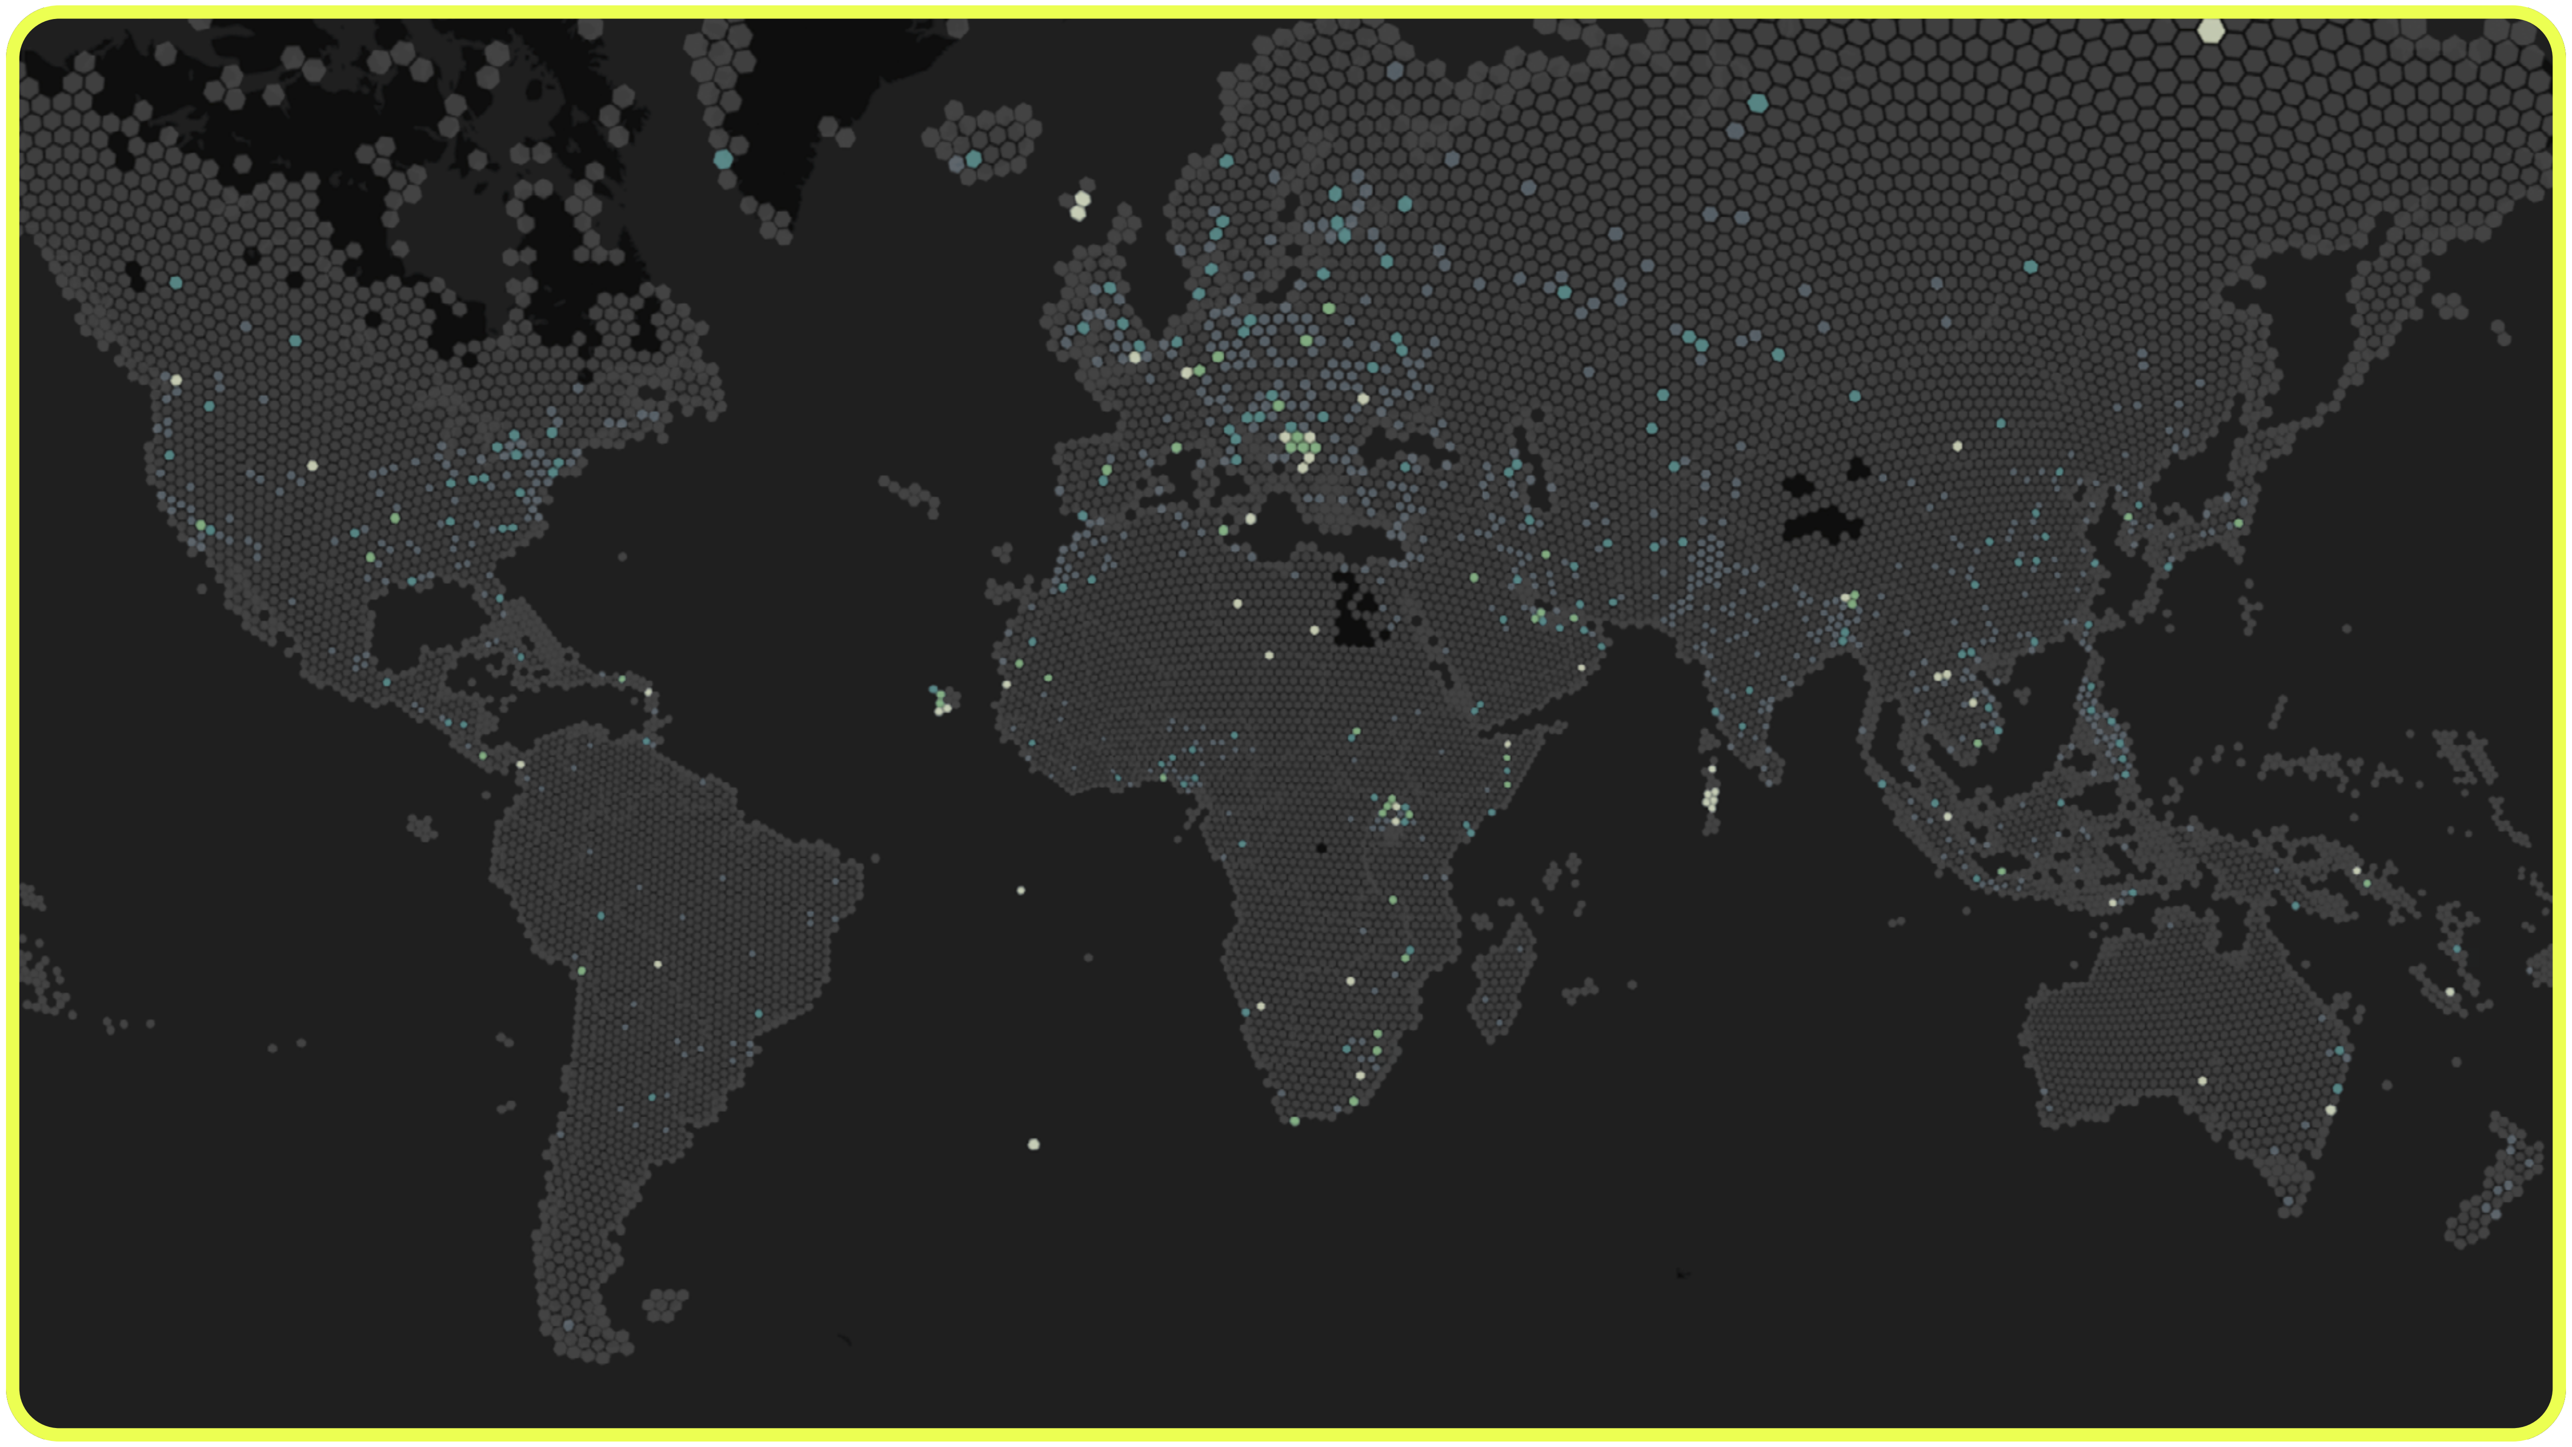

This combination of absolute scale and per-capita paints a layered picture of interest. Some countries, like Uganda, Indonesia and India, demonstrate sheer volume, while smaller regions such as the Faroe Islands or Macedonia signal outsized traction in relative terms.

Both dimensions are strategically relevant for Shaga: absolute numbers highlight where the network can scale to reach mass demand, while per-capita concentrations point to areas of heightened engagement, signalling smaller but highly active pockets of opportunity.

From Interest to Adoption

Identifying interest is only the first step. Globber activity demonstrates willingness to participate in the network, but interest does not automatically translate into active user adoption. For Shaga, the challenge is to convert this early signal into sustainable demand by aligning infrastructure deployment and token incentives with these high-potential geographies. The next layer of analysis therefore turns to client activity and conversion, to assess where traction is most likely to materialise into a stable and growing user base.

When looking at the absolute number of clients, several key areas stand out. The United States leads with 110 clients, distributed across multiple hubs as shown on the map. This is followed by Nigeria (92), India (91), the Philippines (52), and Indonesia (36). Additional notable concentrations appear in specific hubs such as Kenya, north-west Germany, the Netherlands, São Paulo in Brazil, Dhaka in Bangladesh, and the Tokyo metropolitan area.

NUMBER OF CLIENTS (RESOLUTION 3)

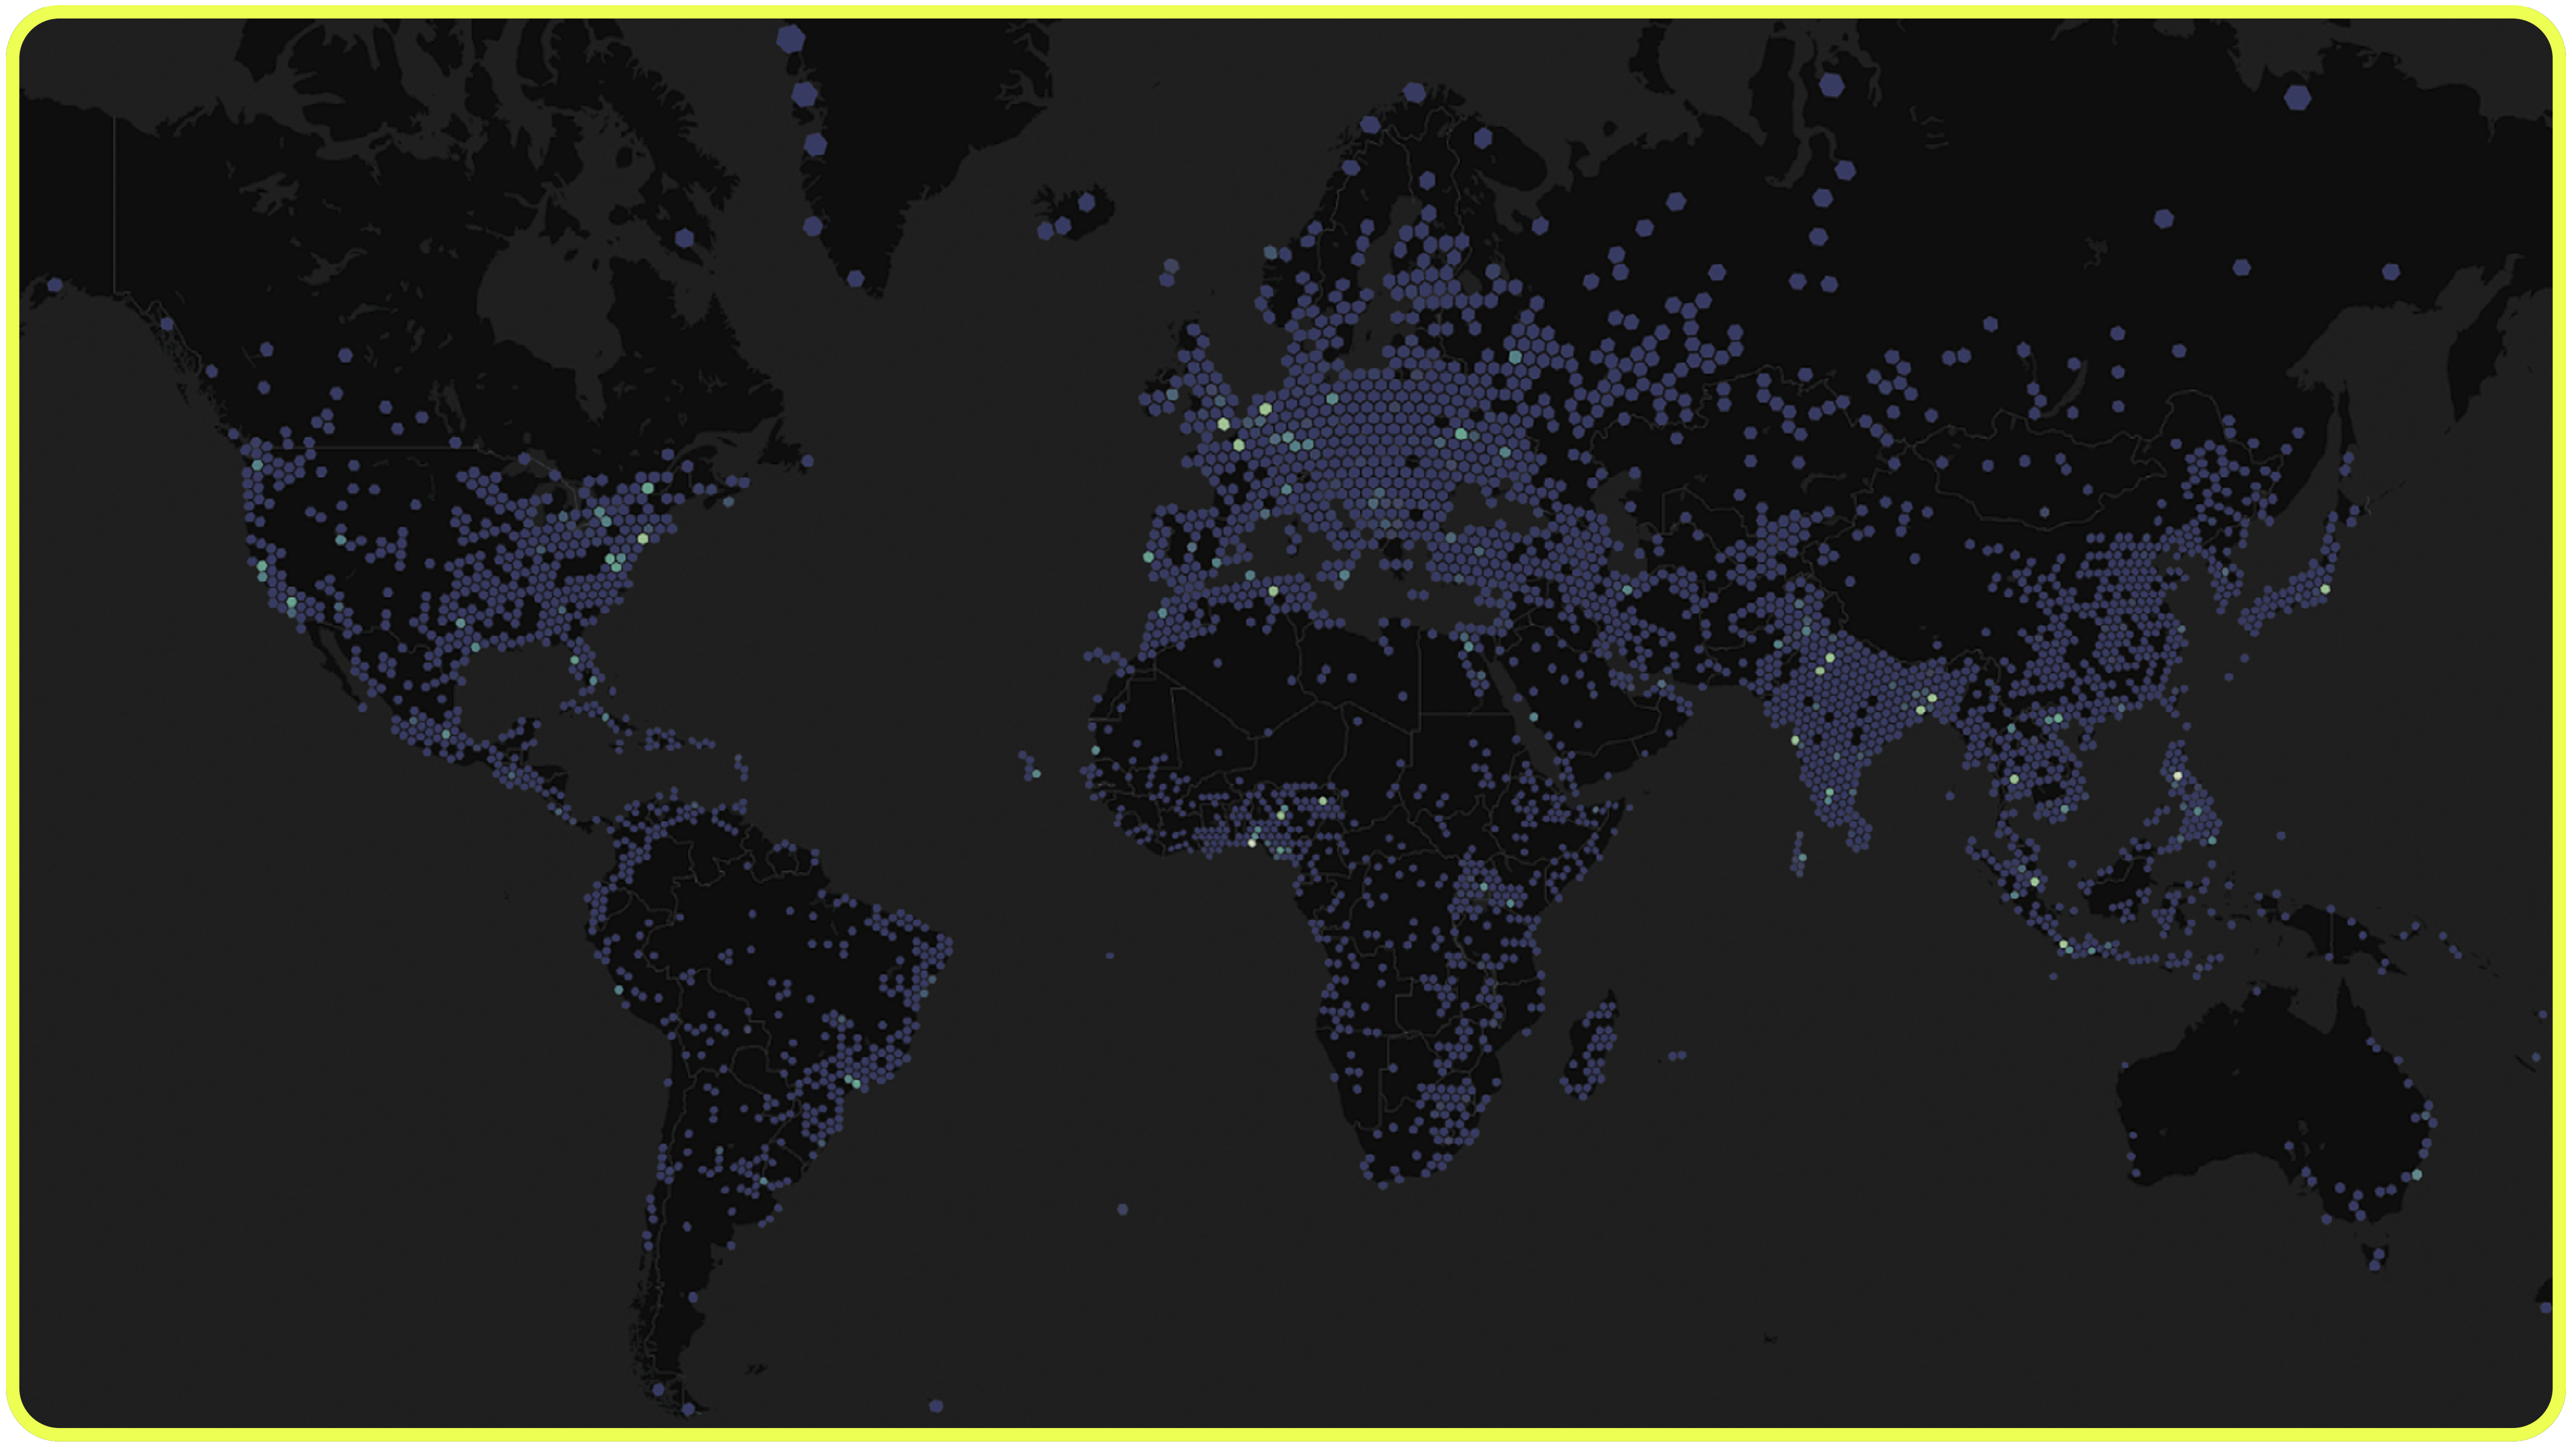

While absolute numbers provide a useful indicator of scale, they are not sufficient to assess adoption on their own. For this reason, we also examine the client-to-globber ratio which can be interpreted as a conversion rate. To ensure the data is representative, countries with fewer than five clients and fewer than fifty globbers are excluded from the analysis, as small bases risk distorting the picture of conversion.

CLIENT TO GLOBBER RATIO (TOP20 — COUNTRY LEVEL)

#	COUNTRY	CLIENTS	GLOBBERS	CLIENT TO GLOBBER
1	Portugal	7	62	11.29%
2	Algeria	16	179	8.94%
3	Argentina	6	73	8.22%
4	Cuba	6	104	5.77%
5	Brazil	24	464	5.17%
6	Ukraine	19	582	3.26%
7	Thailand	21	677	3.10%
8	Egypt	10	346	2.89%
9	Turkey	11	407	2.70%
10	Russia	10	485	2.06%
11	Morocco	7	348	2.01%
12	France	16	900	1.78%
13	United States	110	6290	1.75%
14	Canada	13	765	1.70%
15	Italy	11	699	1.57%
16	Netherlands	10	698	1.43%
17	Philippines	52	3694	1.41%
18	India	91	8039	1.13%
19	Mexico	8	824	0.97%
20	Pakistan	17	2242	0.76%

The top 20 countries by this measure reveal a different distribution than when looking at Globbers in absolute or per-capita terms. Central and Latin America emerges as a particularly advanced region, with Cuba, Brazil, and Argentina all demonstrating conversion rates above 5%. Europe is also strongly represented, with France, Italy, the Netherlands, Portugal, Ukraine, and Russia all appearing in the leading cohort. These markets show that interest has not only been registered but is already translating into measurable user activity.

This contrast is important. Some of the countries that ranked highly in terms of Globber supply or per-capita, such as Uganda (0.012%), United Kingdom (0.17%), China (0.08%), Indonesia (0.15%), Serbia (0.04%), Maldives (0.6%) or the Faroe Islands (0%), are notably absent from the conversion list. In these cases, interest has yet to materialise into active adoption. Conversely, the countries now leading in conversion suggest environments where Shaga's value proposition is resonating more quickly with end-users.

The key takeaway is that while Globbers highlight potential supply, the client-to-Globber ratio reveals which markets are effectively translating interest into adoption. These insights are critical for shaping Shaga's next phase of growth, ensuring that supply-driven enthusiasm is balanced with regions already demonstrating the capacity to convert interest into active participation.

Infrastructure: Anticipating Strain

With demand and potential demand mapped, our focus now shifts to the supply side. Shaga's ability to scale ultimately depends on the capacity of its nodes to support active sessions. At present, the network counts 604 nodes, a strong foundation for a project that was launched only a few months ago without incentives. Nevertheless, when distributed across the globe, the dataset remains relatively limited meaning that any analysis at this stage must be treated as indicative.

When considering node distribution by country, the leading contributors are unsurprisingly those with strong innovation ecosystems, high levels of technology adoption, or deep penetration of crypto activity. The top 10 countries, United States (70), Singapore (59), Japan (32), India (31), and the Philippines (30),China (25), Nigeria (21), Vietnam (19), Canada (18), and more unexpectedly the Maldives (18) account for 55% of all nodes in operation.

NUMBER OF NODES (TOP 10 - COUNTRY LEVEL)

#	COUNTRY	NODES
1	United States	70
2	Singapore	59
3	Japan	32
4	India	31
5	Philippines	30
6	China	25
7	Nigeria	21
8	Vietnam	19
9	Maldives	18
10	Canada	18

Beyond absolute numbers, it is critical to examine the stress on infrastructure at a finer granularity. By analysing the client-to-node ratio at H3 resolution 5 (252 km^2), Shaga can approximate how many clients each node serves. Using a threshold of 50 globbers as a minimum signal of traction, several hotspots emerge. Regions in Bangladesh and Indonesia stand out as areas where nodes are already more heavily solicited. Similarly, parts of Nigeria and India show early signs of “strain”: despite ranking among the top in absolute nodes, their infrastructure may face saturation risks as demand accelerates, particularly when considered alongside their high Globber counts.

This analysis suggests a clear path: these regions should be reinforced with additional nodes to sustain adoption momentum and prevent service bottlenecks.

858C107BFFFFFFF — JAKARTA (INDONESIA) **853CF10FFFFFFFF — DHAKA** (BANGLADESH)

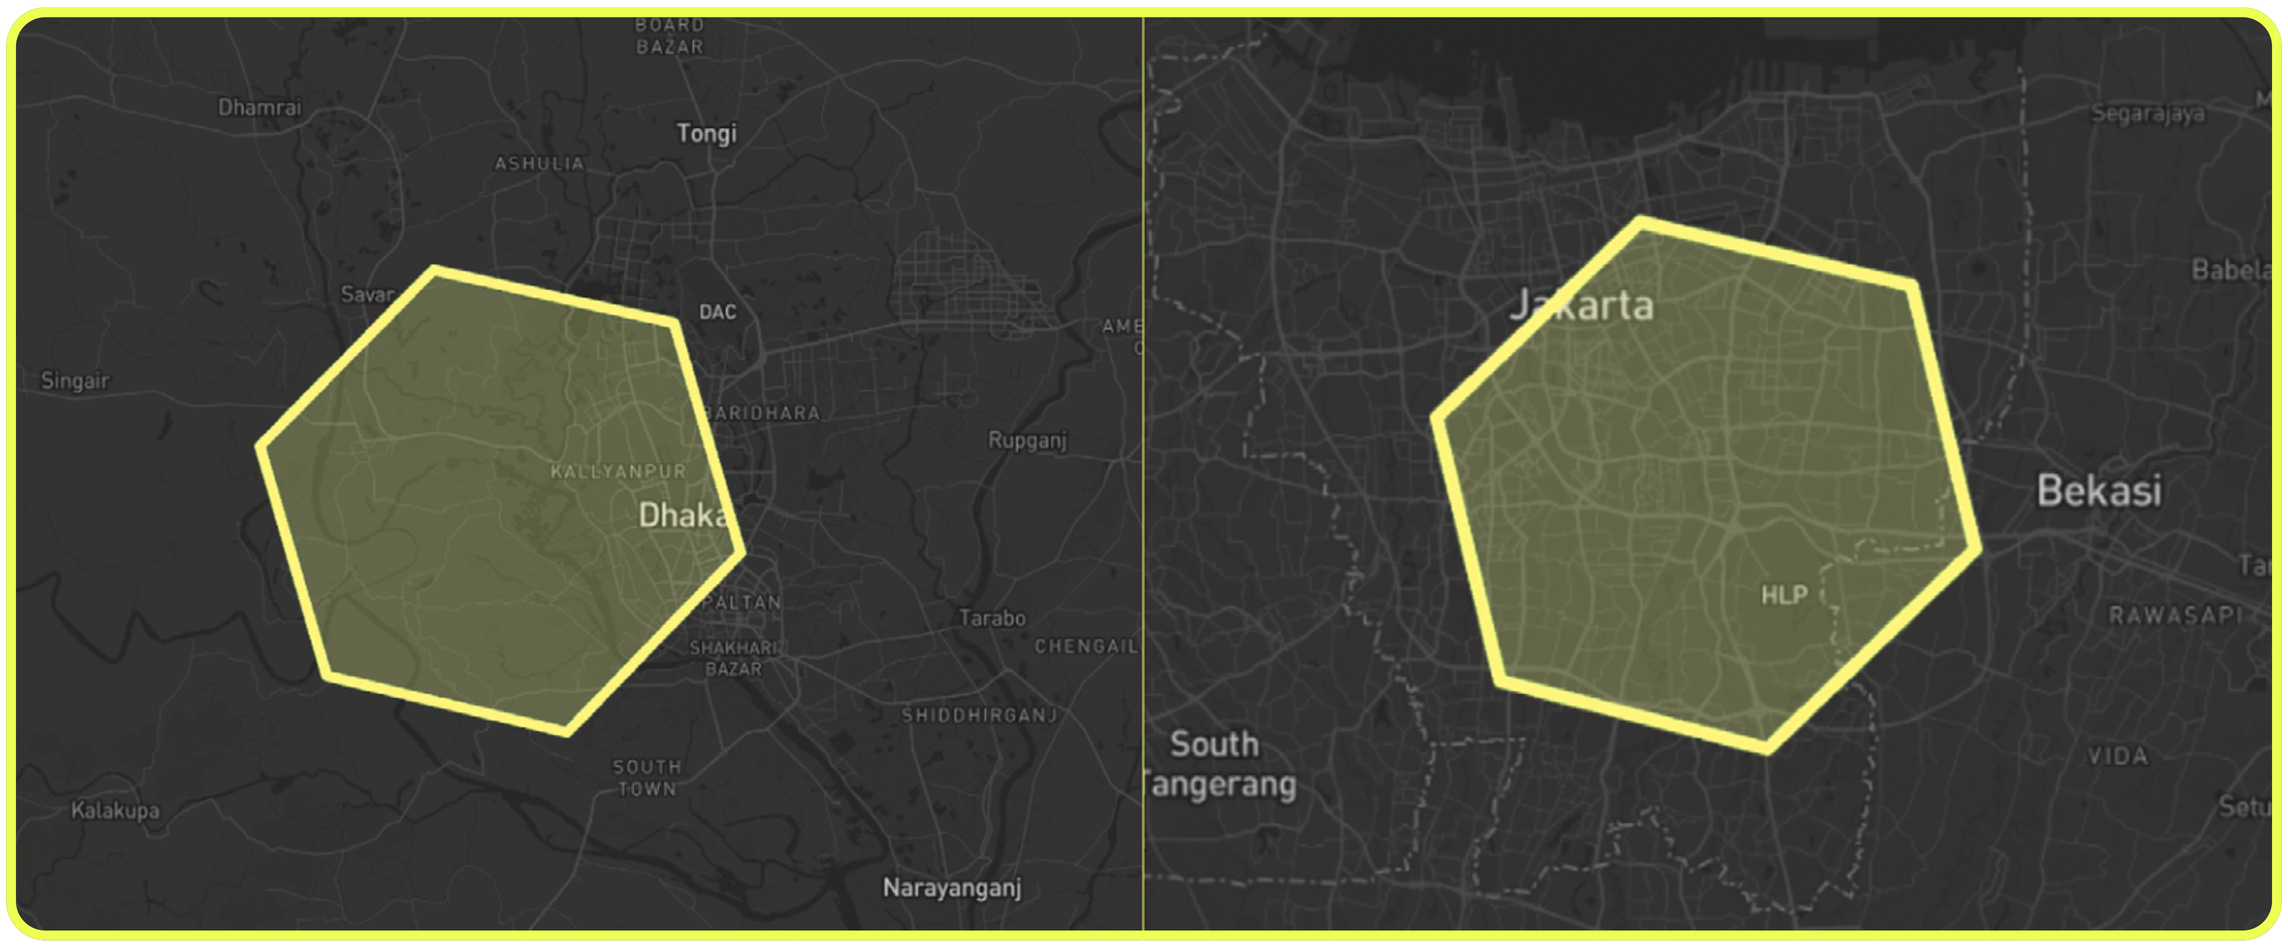

CLIENT TO NODE RATIO (CLIENT TO NODE >3 — RESOLUTION 5)

H3 STRING (RES5)	COUNTRY	CLIENTS	GLOBBERS	NODES	CLIENT TO NODE
853cf10ffffffff	Bangladesh	13	2884	1	13.0
858c107bfffffff	Indonesia	13	989	1	13.0
8558f46bfffffff	Nigeria	11	397	2	5.5
8560145bfffffff	India	5	566	1	5.0
852a100ffffffff	USA	16	1496	4	4.0
853cf2c7fffffff	India	8	638	2	4.0
85588383fffffff	Nigeria	4	121	1	4.0
852cf303fffffff	Iran	4	84	1	4.0
85581bcffffffff	Nigeria	4	81	1	4.0
857a6e43fffffff	Kenya	4	69	1	4.0
85196953fffffff	Netherlands	7	127	2	3.5
85608b57fffffff	India	10	411	3	3.3
85694ec3fffffff	Philippines	29	916	9	3.2
854995bbfffffff	Mexico	3	413	1	3.0
85446c33fffffff	USA	3	57	1	3.0

A further layer of insight comes from analysing the globber-to-node ratio. In regions where more than 500 Globbers correspond to a single node, infrastructure strain is likely to become pronounced if those Globbers transition into active users.

The map highlights several such high-ratio regions, including Bangladesh, Nigeria, Indonesia, Uganda, Pakistan, and the Balkans, as well as dense European clusters in Germany, London, and the Netherlands. Additional isolated hubs also stand out such as São Paulo, California, Tokyo, Sydney, the Maldives, Bahrain, Dubai, the Faroe Islands, and Cape Verde, each of which may require reinforcement to capture future conversion efficiently.

Taken together, these patterns underscore a dual challenge for Shaga: strengthening infrastructure in regions where demand pressure could intensify rapidly, while strategically pre-positioning nodes in areas where interest is high but conversion is still emerging.

GLOBBER TO NODE RATIO (RESOLUTION3)

Toward Incentive Design

This first geospatial analysis highlights where Shaga's incentives and infrastructure could be most effectively directed.

On the demand side, Globbers are joining in large numbers across Uganda, Indonesia, the United Kingdom, China, India, and Serbia, with smaller countries like the Faroe Islands, Maldives, and Cape Verde showing outsized participation on a per-capita basis. These regions are where the peak interest is.

However, client adoption paints a different picture. The United States, Nigeria, India, and the Philippines lead in absolute numbers, while Latin/Central America (Brazil, Cuba, Argentina) and parts of Europe (Portugal, France, Italy, Netherlands, Ukraine, Russia) show particularly strong client-to-globber conversion ratios. These countries provide some of the clearest signals of growing usage.

From an infrastructure perspective, node distribution is strongest in tech- or crypto-friendly countries such as the United States, Singapore, Japan, India, and the Philippines. However, when examining the Globbert-to-Node ratio, data hints at potential pressure points in Bangladesh, Nigeria, and Indonesia, where infrastructure may struggle to keep pace with rising interest. Looking further ahead, more speculative hotspots also emerge in areas like the Balkans, central Europe, and global hubs such as São Paulo, California, Tokyo, and Sydney, regions where demand could quickly outstrip supply once incentives accelerate conversion.

Taken together, these findings suggest a layered strategy. High-conversion countries represent the first wave of durable adoption and should be reinforced with scaling incentives. Meanwhile, high-interest but low-conversion regions form a critical second wave, requiring targeted rewards and infrastructure to translate enthusiasm into usage.

This dual approach provides a roadmap for how Shaga can turn early signals of interest into globally balanced, sustainable growth.

Economics of Compute & Cost Advantage

Session Cost Model

The fundamental cost driver of Shaga is electricity. Unlike centralised cloud providers, where operating expenses include data centre real estate, cooling systems, staffing, and hardware depreciation, Shaga's network is powered by ordinary gaming PCs operated by individuals. This creates an unusually transparent cost structure in which the variable cost of provision is almost entirely determined by electricity consumption.

The session cost model can be expressed as:

$$\textbf{\textit{Cost per player}} = \frac{\textit{GPU watt draw * local electricity cost* session minutes}}{\textit{concurrent players}}$$

The two key variables in this model are local electricity tariffs and the level of concurrency achieved. A rig serving a single player in a high-cost energy market such as Switzerland may approach $0.20 per player-hour, while the same rig serving four players in a low-cost market such as India reduces effective costs to less than one cent per hour.

Across geographies, the cost spectrum therefore ranges from fractions of a cent to roughly twenty cents per player-hour. As a base case for the network at scale, an average of around $0.025 per hour is a reasonable benchmark.

BENCHMARK ECONOMICS OF COMPUTE

		Number of players							
		1	2	4	6	8	10	12	Countries
Electricity Cost ($/kWh)	0.05	$0.025	$0.013	$0.006	$0.004	$0.003	$0.003	$0.002	China / India / Kazhakstan
	0.1	$0.050	$0.025	$0.013	$0.008	$0.006	$0.005	$0.004	Brazil / Turkey / Thailand
	0.15	$0.075	$0.038	$0.019	$0.013	$0.009	$0.008	$0.006	Spain / Poland / Malaysia
	0.2	$0.100	$0.050	$0.025	$0.017	$0.013	$0.010	$0.008	France / USA / Japan / UK
	0.25	$0.125	$0.063	$0.031	$0.021	$0.016	$0.013	$0.010	Argentina / Portugal / Philippines
	0.3	$0.150	$0.075	$0.038	$0.025	$0.019	$0.015	$0.013	Ireland / Belgium / Singapore
	0.35	$0.175	$0.088	$0.044	$0.029	$0.022	$0.018	$0.015	Germany / Demark / Sweden
	0.4	$0.200	$0.100	$0.050	$0.033	$0.025	$0.020	$0.017	Switzerland / Netherlands

When compared to industry incumbents, the efficiency advantage becomes clear. Amazon Web Services' G4dn GPU instances are priced at approximately $0.52 per hour to the client. Microsoft's Xbox Cloud Gaming charges $20 per month for a single-user subscription, which, assuming an average of 40 hours of gameplay per month, equates to an effective rate of about $0.50 per hour (source: kotaku).

Against these baselines, Shaga's underlying cost of approximately $0.025 per hour per player creates a twenty-six-fold gap between the cost base of Shaga and the prices charged by centralised providers.

This gap represents a strategic choice. Shaga could seek to capture the spread as direct margin by pricing closer to industry norms, or it could use its structural efficiency to reduce costs for end users. In practice, Shaga pursues the latter path by generating revenue from an entirely different source: the monetisation of synchronised frames and control data for AI applications.

Revenue and Surplus Allocation

Shaga benefits from the emergence of a new data asset class: synchronised frames and control inputs generated during gameplay.This data, which is essential for AI training in interactive environments, cannot be scraped from existing sources and must be mined directly from live play sessions.

Market quotations for this gameplay data vary widely. Data brokers typically advertise prices in the range of $10 to $30 per hour, though these figures incorporate their own margins. Internal floor estimates for direct sales or licensing are closer to $1 per hour. Some specific scenarios have seen brokers quoting $50 to $100 per hour. These ranges illustrate the significant upside potential, with Shaga remaining economically sustainable even at the conservative $1 per hour level.

Assuming the cost of operating a rig is approximately $0.10 per hour, the $1 per hour data floor yields a $0.90 gross margin. At $5 per hour, margins rise to $4.90 per rig-hour. In the event of $50 per hour broker demand, the margin approaches $50 per rig-hour.

The implication is clear: even at the lowest conceivable uptake, Shaga's revenue comfortably covers the cost of gameplay for users. Put differently, the sale of a single hour of AI data is sufficient to pay for approximately nine hours of gaming.

The surplus generated ensures that all stakeholders benefit:

- **Nodes are rewarded** for their uptime and capacity provision, ensuring a stable and resilient supply base.
- **Players benefit directly**, as their gameplay costs can be reduced or eliminated, creating a powerful retention driver.
- **The treasury captures residual value** to fund growth initiatives and maintain long-term stability.

This dynamic creates a virtuous circle: lower costs attract more users, more users generate more valuable data, and more data creates higher revenues that can be recycled into further cost reductions and network expansion. Each participant strengthens the system while directly benefiting from its growth.

Data Monetisation Loop & Partnerships

Why Partnerships Matter

The foundation of Shaga's model lies in its hardware base: hundreds of millions of gaming PCs that can be converted into cloud nodes. However, hardware alone does not create revenue. To translate this capacity into income, the network requires demand anchors in the form of popular games, token ecosystems, and distribution partners that can guarantee consistent session hours and predictable monetisation.

The agreement signed on March 12 with Star Atlas illustrates this principle. The deal combined three components: an equity stake in Shaga, integration of game assets, and alignment with the ATLAS token economy. Star Atlas, a Solana-based grand strategy game built in Unreal Engine 5, has heavy graphical demands and a committed global player base. Running such a title through Shaga demonstrates that AAA blockchain games can be streamed to low-spec devices without compromising quality.

For Star Atlas, the benefit is access to new users in regions such as Southeast Asia and Latin America, where mobile penetration is strong but high-performance PCs are limited. For Shaga, the partnership validates its streaming layer with a flagship content partner and establishes a pipeline of sessions that can be monetised directly through gameplay and indirectly through associated data.

The same logic applies beyond Star Atlas. Any publisher with a hardware-intensive title faces adoption limits when distribution depends on costly rigs. By integrating with Shaga, that barrier is removed, enabling broader reach. In turn, Shaga secures the steady traffic and session time that drive its economic loop.

From Gameplay to Data Assets

If partnerships with publishers secure Shaga's demand base, partnerships with data buyers define its margin structure. What distinguishes Shaga from other cloud gaming models is that each gameplay session produces more than entertainment. Every stream generates a synchronised pair of outputs: rendered video and player control inputs. This combination is significantly more useful than video alone, providing a training resource for interactive AI and robotics that captures not only outcomes but the decisions leading to them.

This transforms gameplay into a secondary asset class. Broker quotes already suggest ranges of $10-30 per annotated hour, with higher rates for exclusive or richly tagged datasets. Even under a conservative floor of $1 per hour, the economics are favourable. Operating costs on consumer hardware average around $0.10 per session hour, meaning data resale alone can subsidise the bulk of user play. At mid-range benchmarks of $5-10, margins expand dramatically, producing surplus that can be recycled into lower access costs for players, higher rewards for node operators, and fiat-denominated stability for the treasury.

This dual monetisation loop depends on building relationships in both directions. On one side, publishers like Star Atlas, TRIB3, Wanderers, LowLife Forms, and Aurory are already leveraging Shaga to bring graphically intensive titles to broader audiences, effectively extending their reach to mobile and lower-spec devices. On the other, emerging partnerships with AI labs and large model developers ensure that the by-product of these sessions is not wasted but resold into markets where demand for behavioural datasets is accelerating.

The Monetisation Loop

The final piece is how partnerships connect into a system that compounds over time. Shaga's economics are built to improve with scale. Each new partner, whether a game publisher or an AI buyer, adds not only extra revenue but also greater leverage across the network. Gameplay sessions are monetised once through access fees and again through data resale.

These revenues cycle back into the system to lower costs for players and increase rewards for node operators. The effect is circular: lower costs attract more users, more users generate more data, and more data strengthens both revenues and the appeal of Shaga to future partners.

Partnerships also shape where this loop grows. By using its dataset to map where unserved gamers and potential node operators are concentrated, Shaga can focus on regions with the strongest adoption potential. In places such as Southeast Asia or Latin America, this creates a three-way benefit. Publishers expand into new audiences, local PC owners earn by running nodes, and Shaga gains both traffic and gameplay data. This targeted approach makes growth more efficient and ensures momentum builds where it matters most.

For Shaga to succeed, however, partnerships must go beyond just publishers. Collaborations with large language model developers and companies like Qwen are equally critical, turning gameplay data into a valuable input for AI training. Shaga's future depends on building both sides of this ecosystem. It needs strong relationships with games to bring in users and with LLM buyers to turn their activity into revenue.

Tokenomics & Incentive Design

The SHAG token is at the heart of Shaga's economic model. It is used as a real tool to stimulate all stakeholders, from gamers and streamers to publishers and nodes. SHAG is an incentive token, but unlike typical designs, its flows have been carefully engineered.

Unlike most DePIN projects that emit tokens blindly, Shaga treats incentives as a system of regulation powered by data. Its tokenomics are built on the ethos of "incentives go to who adds value and where it matters." ensuring that every distribution builds lasting value rather than unchecked inflation.

Nodes Rewarded Based on Data

Nodes are not rewarded randomly, which is one of Shaga's key differentiators. Rewards unlock only in line with measurable network outputs. This ensures that supply expands only when genuine usage justifies it.

To achieve this, Shaga relies on two variables: **geolocation and electricity cost.**

Using the H3 map to measure demand and infrastructure maturity, Shaga directs emissions to regions where node development is most needed, ensuring that quality of service reaches the gamers who are actually there.

Electricity cost then defines how much a node must earn to stay online. A node must cover electricity, bandwidth, and ideally hardware depreciation. This simple logic anchors the emissions system: only nodes that provide value and operate efficiently receive meaningful rewards.

TOKEN EMISSION

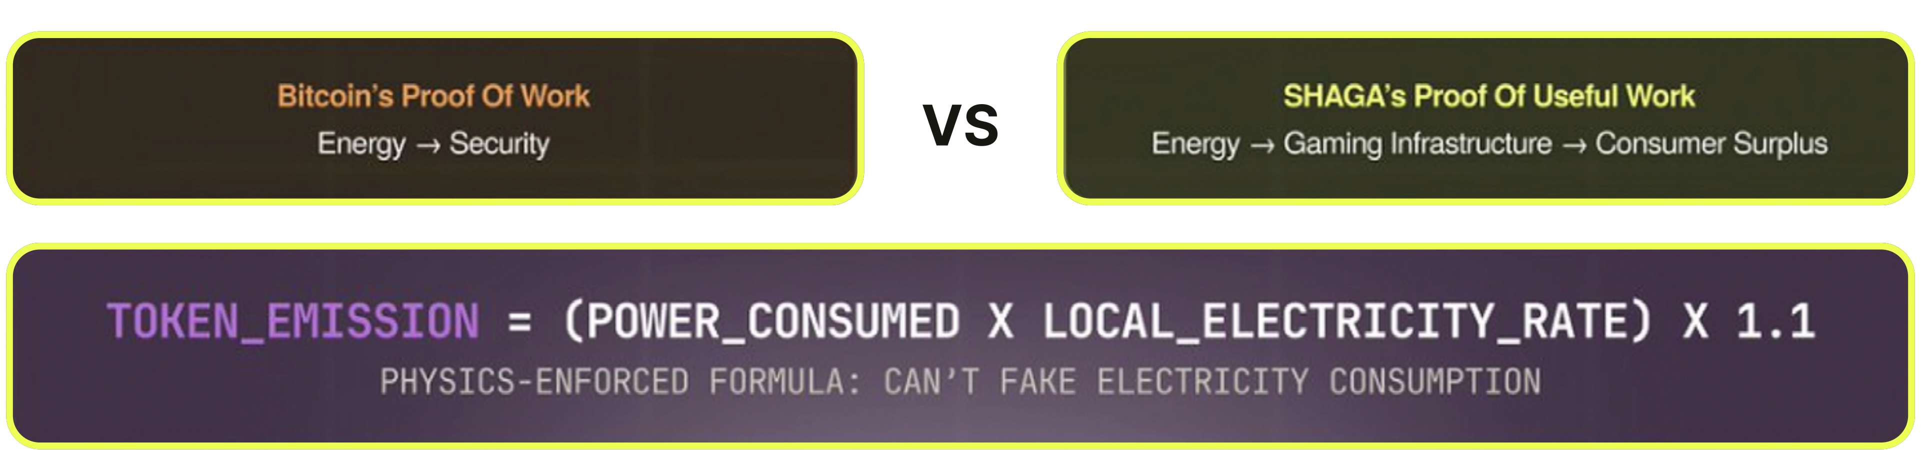

By using data and maths to drive emissions, Shaga solves a core weakness of many DePIN projects. Nodes are neither over-incentivised nor left under-compensated. They are paid fairly, based on real demand and actual activity. This recalls the early years of Bitcoin, when only those willing to stay online and contribute meaningfully persisted.

NODE SUPPLY ANALYSIS: COST THRESHOLD FOR PARTICIPATION

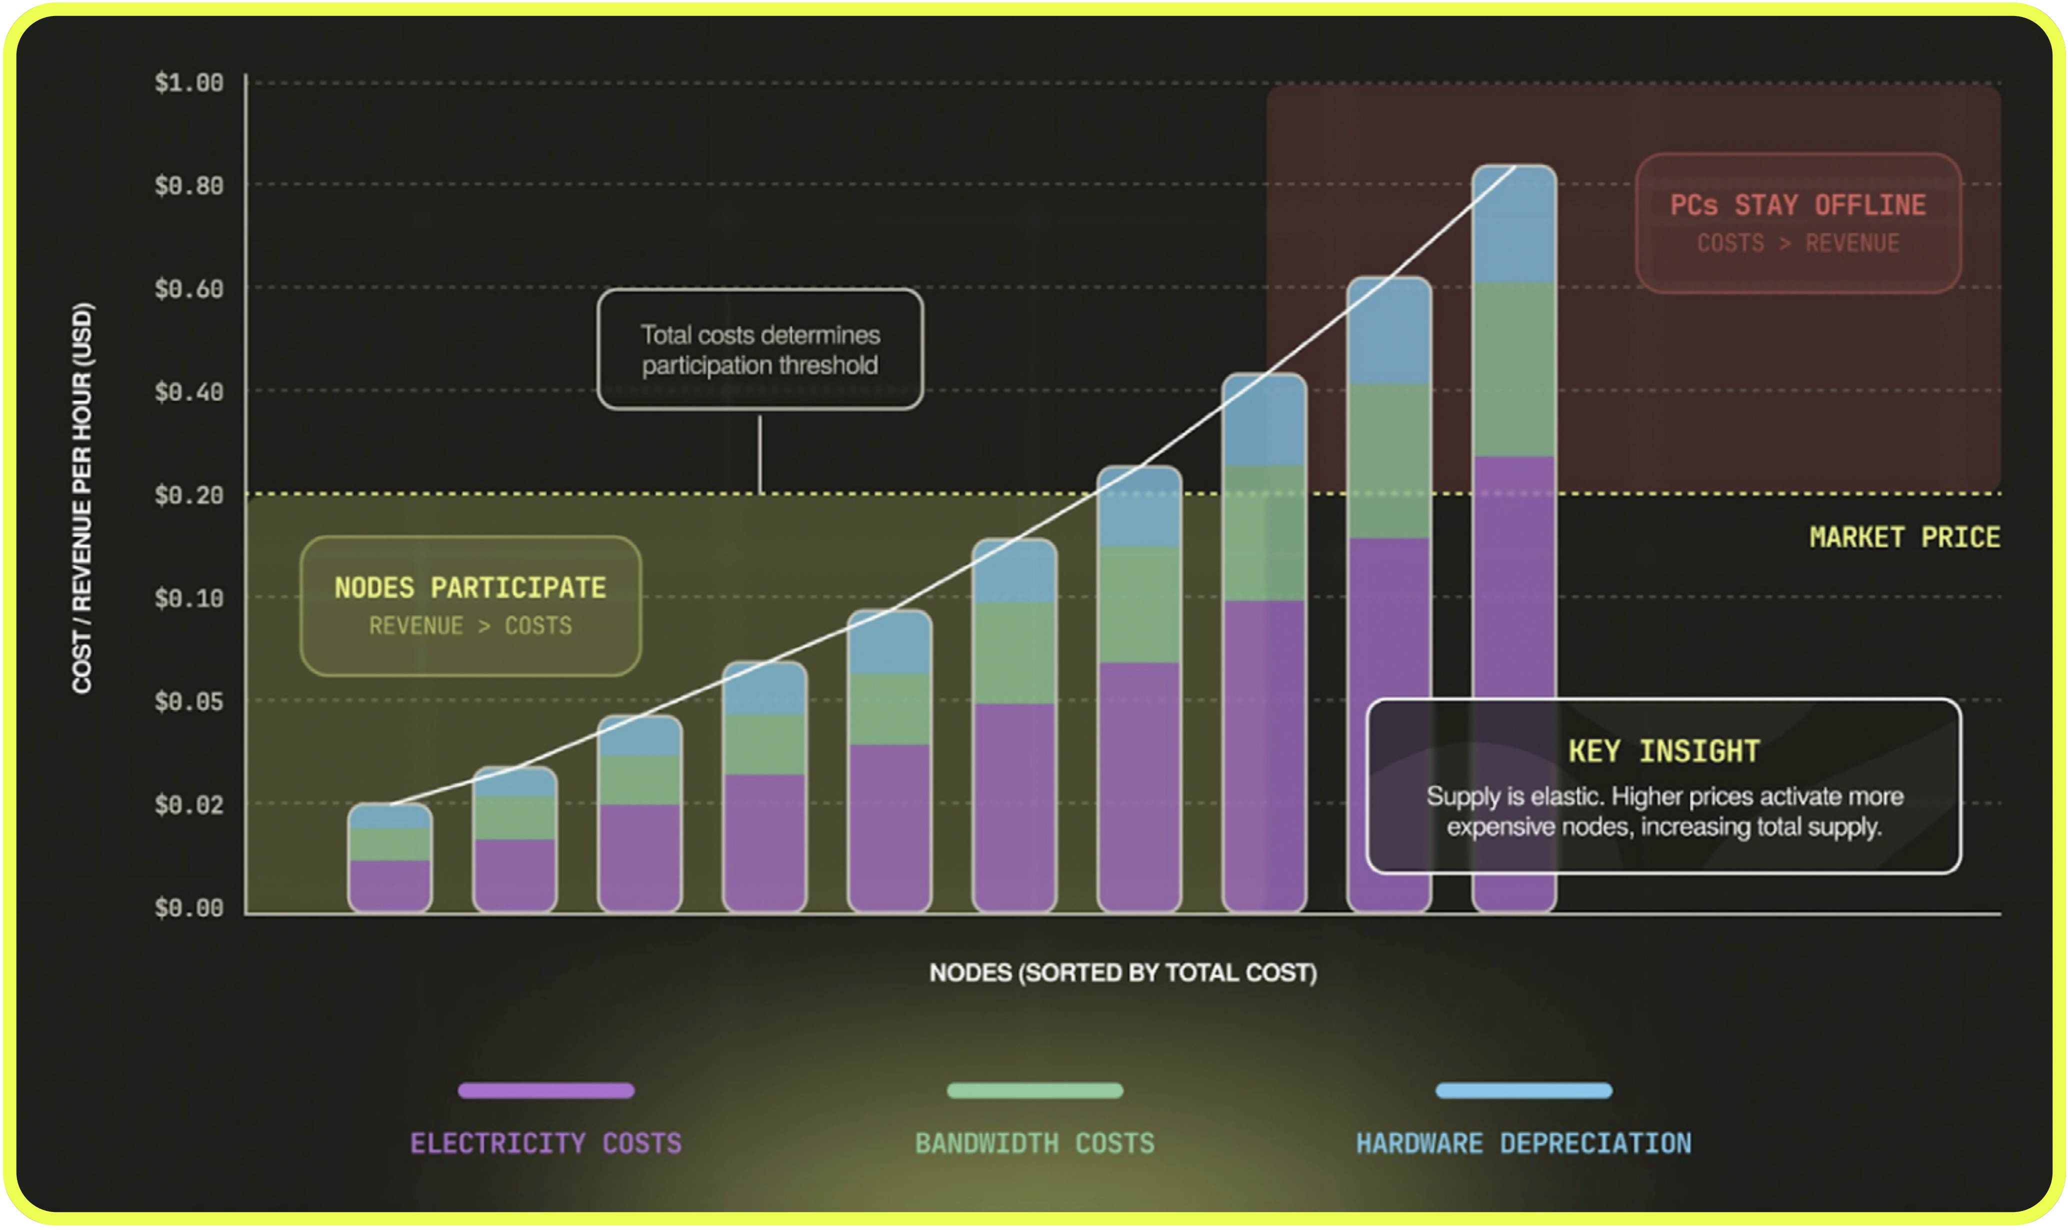

SHAG Stimulating Demand

Covering infrastructure costs is only half the equation. For Shaga to thrive, it must also attract players. Where traditional companies spend heavily on marketing, Shaga uses SHAG as a driver to convert new users into active clients.

Streamers are the first engine of this growth. By receiving incentives, they can run contests, referral campaigns, and community events, bringing their audiences directly into the Shaga ecosystem. This is a cornerstone of Shaga's adoption strategy, tapping into streamer communities to expand reach.

Gamers can also be incentivised directly. Sign-up rewards, achievement campaigns, and targeted promotions encourage engagement and retention.

Crucially, all of this remains data-driven. Emissions can be directed to regions where demand is highest (Globbers), to players who are most active, or to streamers who deliver measurable results. As with nodes, the principle is the same: the more value you bring to the network, the more you earn.

Revenue Flow

Yet, aligned incentives do not change the fact that stakeholders will probably sell tokens to cover costs. For SHAG to sustain long-term value, it must generate ongoing demand. This is achieved through utility and revenue-backed mechanisms.

On the utility side, SHAG can be required for premium services. Accessing exclusive streams or communities may require a subscription fee in SHAG. Gamers may also bid in SHAG for priority access to high-performance nodes. Streamers and publishers will need SHAG to boost visibility, sponsor events, or promote content. Governance functions are also tied to SHAG, giving holders influence over emissions and treasury decisions. In this way, tokens continually recycle between stakeholders and the treasury, which acts as the central mechanism to redistribute value and support long-term growth.

Beyond direct utility, external revenue plays a decisive role in value capture. The network produces frames and control data that can be sold, generating significant income for the treasury. From there, Shaga periodically buys back SHAG from the market and burns it, offsetting inflation from emissions and reinforcing long-term value.

REVENUE FLOW

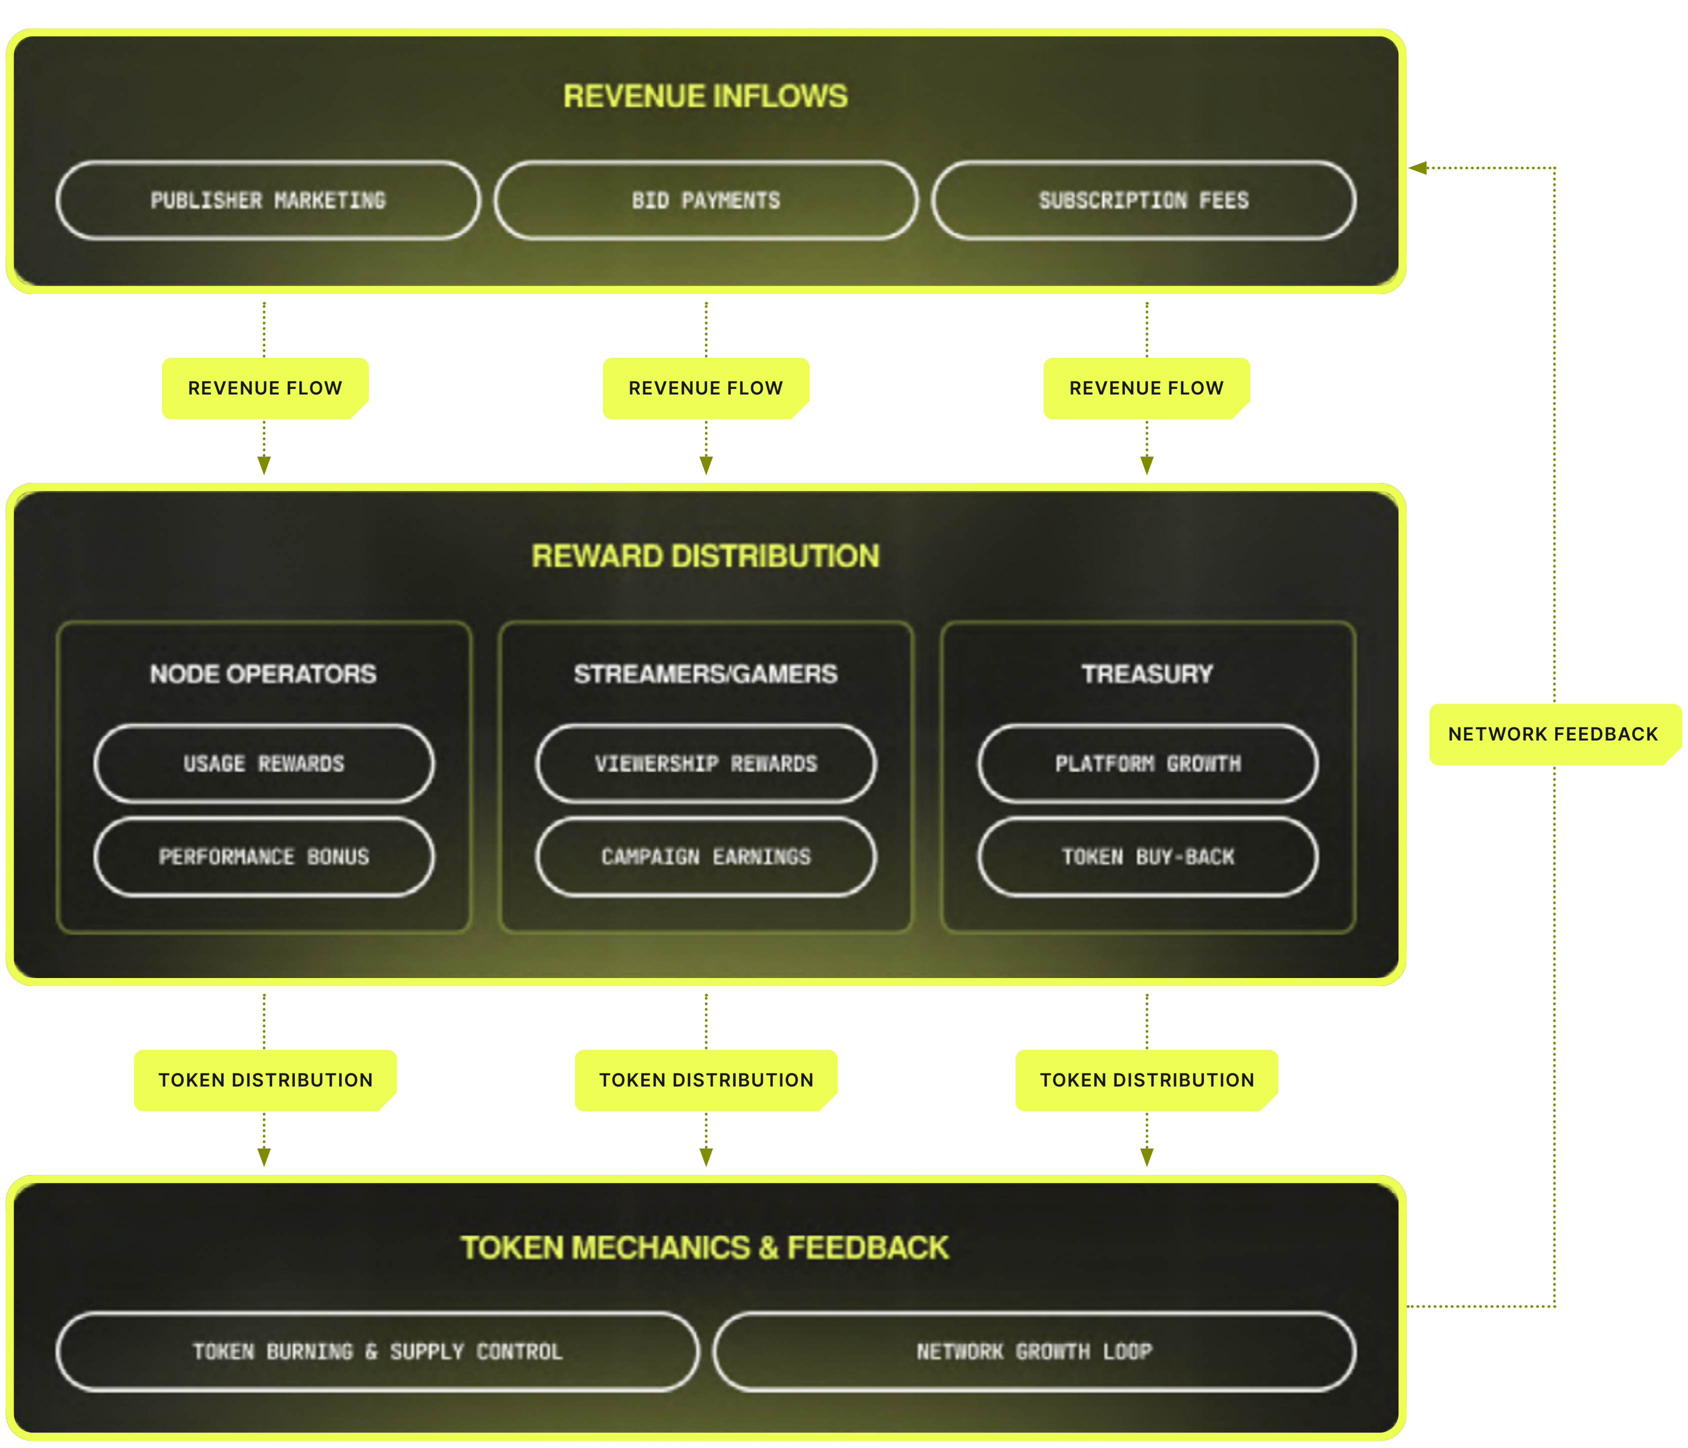

In addition to AI data, other revenue streams reinforce this model. Game advertisements, premium subscriptions, and event sponsorships all contribute recurring external value, creating a diversified foundation for sustainability.

Sustainability by Design

What sets Shaga apart is that its tokenomics are firmly tied to reality. Emissions are driven by data and aligned with demand. Stakeholders only earn when they deliver value. Revenues stem from real sources such as AI data sales, advertising, and subscriptions.

By treating SHAG as a currency flowing between stakeholders rather than a marketing budget to be spent, Shaga avoids the pitfalls that have undermined other DePIN projects. It transforms tokenomics from a speculative mechanism into a system where growth is controlled, sustainable, and lasting.

Assessing Shaga's Valuation

Valuing Shaga requires an approach that bridges both institutional finance and decentralised infrastructure dynamics. Traditional decentralised networks have often anchored their valuations in speculative token supply models or comparisons with loosely related peers. This tends to either overstate optionality or obscure the fundamentals of cost and revenue generation.

A more robust framework is to ground Shaga's economics in measurable usage patterns, map earnings across conservative and optimistic scenarios, and then cross-check with existing industry comparables. This hybrid method creates a valuation narrative that is both credible to investors and compelling to participants.

Comparable Multiples

A brief analysis of sector peers highlights the dispersion of valuations in decentralised and centralised compute:

- **CoreWeave**, one of the fastest-growing GPU cloud providers in Web2, reported annualised revenues of around $5 billion in 2024, supporting a valuation in excess of $23 billion, or roughly 4.6x revenue (investors.com and investors.coreweave.com)
- **Rescale**, an HPC-focused compute platform, generated $35 million in annual revenues with a $1 billion valuation, equivalent to an almost 30x revenue multiple (source: getlaka)
- **Render** generated under $1 million in revenue during 2024 yet traded at a $2.3 billion fully diluted valuation, implying a multiple in the range of 2,500x revenue (source: getlaka)
- **Akash** is currently on a $450 million fully diluted basis with projected $4 million in annual revenues, or around 112x revenue. (source: stakecito)

Taken together, these cases show how nascent and speculative the space remains. The breadth of revenue-to-valuation multiples, ranging from single digits to several thousand, signals that the market has yet to converge on reliable pricing benchmarks. For Shaga, this means direct comparison offers little precision.

Instead, the more instructive path is to position its valuation on the basis of user-side economics, AI-driven revenue potential, and adoption TAM, while using peer multiples only as a contextual reference.

Earnings-Based Scenarios

Assumptions

Shaga's cost base begins with the electricity required to run a gaming PC. A typical 500W rig consumes around $0.10 per hour at an average tariff of $0.20/kWh. Unlike centralised data centres, this excludes overheads such as cooling, real estate, or staffing. This lean structure enables Shaga to operate at a fraction of the cost of AWS or Xbox equivalents.

On the revenue side, the breakthrough lies in data monetisation. Gameplay Action Pairs (frames plus control data) are of intrinsic value to AI training and cannot be replicated without live gameplay due to computational irreducibility. Each hour of captured gameplay therefore represents a potential asset in the AI training market. Broker quotes for such data have ranged as high as $30-$100 per hour, but to remain conservative our modelling considers only $1 per hour as a floor and $5 per hour as a mid-case.

The scale of this opportunity depends on adoption. The underserved gaming TAM is estimated at 150 million users globally. Based on Xbox's reported average of 40 hours of online play per month (480 annually), we model Shaga's network under adoption scenarios of 1%, 5%, and 10%.

In addition, we apply a conservative filter in which only 10% of gameplay hours are monetised as AI data (below broker assumptions) to preserve prudence.

Modelling

The table below presents the potential surplus generated by the network under the assumptions outlined above.

Shaga reaches break-even in the conservative case where only 10% of gameplay hours are monetised at $1 per hour (assuming electricity costs of $0.20/kWh). In this scenario, AI revenues effectively subsidise gaming, making it cost-free for nodes to operate and for users to play. However, this level of monetisation does not provide sufficient headroom for growth, meaning higher pricing is required to generate meaningful returns.

Such an outcome underlines the resilience of the model, with sustainability achievable even at the most conservative input levels.

At $5 per hour (which is more reasonable), three scenarios can be outlined depending on adoption. In the bear case, where 1% of the TAM is converted, the network generates a surplus of approximately

$288 million. In the base case of 5% adoption, the surplus expands to $1.4 billion. In the bull case of 10% adoption, annual surplus reaches $2.88 billion.

Such surplus creates a virtuous cycle: more players generate more data, more data drives higher revenues, and those revenues sustain growth and attract further adoption.

FROM IDLE GPUS TO GLOBAL INFRASTRUCTURE

PENETRATION IN TAM	GAMERS	GAMING HRS / YEAR	COST AT $0.2/KWH ($)	AI HOURS / YEAR	REVENUE AT $1/H ($)	SURPLUS AT $1/H ($)	REVENUE AT $5/H ($)	SURPLUS AT $5/H ($)
1%	1,500,000	720,000,000	72,000,000	72,000,000	72,000,000	0	360,000,000	288,000,000
5%	7,500,000	3,600,000,000	360,000,000	360,000,000	360,000,000	0	1,800,000,000	1,440,000,000
10%	15,000,000	7,200,000,000	720,000,000	720,000,000	720,000,000	0	3,600,000,000	2,880,000,000

Valuation Framework

From the results above, we can now apply the valuation multiples seen at the beginning of the section. To avoid reliance on highly speculative figures, we restrict the analysis to 4.6x and 30x. Shaga's valuation can therefore be modelled by applying these multiples to projected network surplus (at a selling price of $5 per hour) across different adoption scenarios.

At 1% adoption, Shaga generates $288 million in annual surplus. Applying multiples of 4.6x and 30x results in valuations of approximately $1.3 billion and $8.6 billion respectively. In the base case of 5% adoption, surplus expands to $1.44 billion, corresponding to valuations of $6.6 billion at 4.6x and $43 billion at 30x. Under a bullish 10% adoption scenario, surplus reaches $2.88 billion, implying valuations of $13.2 billion and $86.4 billion respectively.

It is important to highlight three considerations. First, the analysis relies on a number of assumptions such as TAM estimates, pricing, average gaming hours, and cost structures, which, while conservative, inevitably carry some margin of inaccuracy. Second, the adoption percentages, though reasonable, represent long-term outcomes: the network growth will be gradual and cannot be achieved overnight. Third, valuation multiples should be treated as speculative. As the project matures and demonstrates a sustainable model, these multiples are expected to compress toward a more traditional range of 4.6x to 10x, rather than the upper bound of 30x.

Opportunities for Stakeholders

The strength of Shaga's model lies in aligning incentives across multiple groups. Each stakeholder both benefits from and contributes to the loop, expanding the network as adoption grows.

FROM IDLE GPUS TO GLOBAL INFRASTRUCTURE

STAKEHOLDER	OPPORTUNITY	CONTRIBUTION TO FLYWHEEL	HOW TO PARTICIPATE
Node Operators	Earn tokens and fiat royalties from data resale	Provide distributed GPU capacity, powering low-latency sessions	Install the Shaga client on a compatible gaming PC to turn idle hardware into a cloud node. Use the profitability calculator to plan returns.
Players	Access free or subsidised gameplay by consenting to data sharing	Generate synchronised gameplay and control data during sessions	Play games through the Shaga platform and toggle data-sharing consent. Participation reduces or eliminates costs of gameplay.
Game Publishers	Expand audience reach by streaming hardware-intensive titles	Bring high-demand games onto Shaga, driving session volume	Integrate via Shaga's SDKs and APIs. Titles with high GPU requirements or Unreal Engine 5 builds benefit most.
AI Labs / Robotics Firms	Purchase annotated gameplay-control datasets for training	Convert gameplay into a revenue stream by paying for data	Subscribe to Shaga's data catalogue, select annotated session types, and negotiate pricing tiers for access.
Investors	Exposure to real-use-case compute and data markets	Provide capital to expand the network and validate Shaga's valuation	Acquire and stake Shaga tokens to gain exposure to adoption growth and network rewards.

Shaga works because every stakeholder has a clear way to join and a direct reason to contribute. Players gain cheaper access to premium games. Node operators turn idle GPUs into income. Publishers unlock new audiences in regions where expensive hardware is out of reach. AI labs secure a rare source of synchronised gameplay and control data to train their models. Investors have two paths, either backing Shaga early through equity or participating later as token holders once the network is live.

Together these roles form the engine of the flywheel. More publishers bring more players. More players create more gameplay sessions. More sessions generate more data. More data attracts AI buyers. The revenue from data sales and gameplay flows back into the treasury, which lowers costs for players and increases rewards for node operators. Each group makes the system stronger, and as adoption grows the network becomes more valuable and more difficult to displace.

Risks & Mitigations

Shaga faces a mix of external and project-specific risks that together define the conditions under which the network can scale.

Regulatory uncertainty is the most immediate. Gameplay data enriched with control signals falls into a grey area across jurisdictions, intersecting with privacy regimes such as GDPR in Europe and CCPA in the United States. Shaga's opt-in model and anonymisation standards are important safeguards, but compliance will need to be adapted region by region as regulation evolves.

Reliability is a second area of concern. A peer-to-peer grid benefits from geographic distribution but is also exposed to hardware churn, uneven connectivity, and strain in areas where demand quickly exceeds available supply.

Earlier analysis highlighted imbalances such as high globber-to-node ratios in Bangladesh, Nigeria, and Indonesia, where infrastructure could become saturated if adoption accelerates. Latency-aware routing, targeted incentives, and the planned bug-bounty programme reduce these risks, yet maintaining consistency across heterogeneous devices remains an ongoing challenge.

Economic sustainability also depends on inputs that Shaga cannot control. Electricity costs form the base of the session model; if prices rise in key geographies, the cost advantage over centralised providers could narrow. At the same time, GPUs remain contested resources. Resale markets and AI miners can draw operators away if hardware values climb. By tying emissions directly to verifiable compute, Shaga creates a steady return path for contributors, but hardware cycles will always influence supply at the margin.

The largest tail risk remains demand for synchronised gameplay data. Broker quotes and early partner discussions indicate appetite, yet this is an emerging market without stable pricing benchmarks. If buyer demand proves shallower or more volatile than expected, revenues may settle closer to conservative floor estimates, limiting the surplus available to subsidise sessions and fund growth.

Finally, Shaga is not insulated from the wider crypto environment. Token volatility, shifting investor sentiment, and liquidity shocks have historically shaped the trajectory of decentralised networks regardless of their fundamentals. Although Shaga's design grounds token issuance in measurable work, broader market downturns could still affect treasury value, operator incentives, and user acquisition.

None of these risks are unique in isolation, but together they outline the boundaries of Shaga's execution space. They underscore the importance of measured scaling, careful monitoring of external conditions, and continuous adjustment as the network expands from beta to full commercialisation.

Roadmap & Catalysts

Shaga's roadmap follows a phased structure that addresses technical reliability, incentive design, monetisation, and scale. The closed beta, launched in the first quarter of 2025, provided the first opportunity to test the network in practice.

Results from this phase showed median round-trip times below 40 milliseconds in underserved regions, a level of responsiveness that is competitive with local play. With hundreds of active nodes already contributing, the beta confirmed that consumer hardware organised through latency-aware routing can provide a service comparable to centralised infrastructure.

The next stage is the token generation event in the third quarter of 2025. The introduction of the SHAGA token formalises the incentive system that has been described throughout the report. Token emissions are linked to verifiable compute tasks, such as rendering frames or contributing to codec training, grounding supply expansion in measurable work rather than idle staking. This step establishes the economic framework required for the network to grow beyond the beta cohort.

Attention then shifts to the commercial layer. In the fourth quarter of 2025, Shaga plans to open a data marketplace where gameplay streams enriched with control signals can be purchased by AI developers, robotics teams, and generative media firms. Earlier sections highlighted the value of these datasets, with broker quotes ranging from $10 to $30 per hour, and higher in specific cases. The marketplace provides the mechanism to translate these signals into external revenue, testing whether buyer demand is sufficient to subsidise gaming sessions and support the broader cost model.

The final near-term milestone is the release of a mobile SDK in the first quarter of 2026. Extending to mobile devices expands both access for users and the diversity of data generated for resale. As shown in the network composition analysis, heterogeneity of hardware has already proven to be an asset, with mainstream devices ensuring reliability and high-end outliers contributing capacity. Mobile integration continues this pattern by enlarging the potential base of participants while preserving compatibility with the dataset model.

The sequence from beta to token launch, marketplace, and mobile integration provides a clear structure for evaluating progress. Each stage addresses a distinct requirement: technical validation, incentive alignment, monetisation, expansion, and together they define the criteria against which Shaga's execution can be measured.

Conclusion

Shaga demonstrates that a decentralised model can overcome the inefficiencies of cloud gaming while opening an entirely new layer of value in AI-ready data. By aggregating idle GPUs into a distributed mesh, the network delivers compute at the true edge, with latency and responsiveness already proven competitive. This technical foundation is reinforced by innovations such as the neural game codec, which reduces bandwidth needs dramatically, and control bifurcation, which lowers multiplayer latency by bypassing unnecessary relays. What strengthens the network even further is inclusivity: any kind of CPU or GPU can contribute, from high-end RTX 4090s to everyday laptops. This diversity ensures resilience, efficiency, and accessibility across markets, making Shaga a truly global infrastructure from the ground up.

To scale efficiently and rapidly, Shaga's economic model is central. Costs are reduced to electricity consumption, gameplay is monetised as AI training data, and everything is tracked on an H3 geospatial framework. When comparing Shaga's costs with those of centralised incumbents, the difference is striking: a twenty-six-fold gap separates Shaga's underlying expenses from the rates competitors charge. Yet instead of capturing this spread as margin, Shaga chooses to subsidise gaming by selling AI-ready data. Even at the lower bound of the price spectrum, around five dollars per hour, revenues comfortably cover costs and leave a surplus. This surplus is the foundation of the flywheel effect: more users generate more gameplay, more gameplay produces more data, more data creates more revenue, and more revenue sustains growth. The model compounds naturally, turning scale into self-reinforcement.

The geospatial map is the instrument that keeps this growth targeted and efficient. Built on Uber's H3 indexing framework, it directs incentives precisely to where they are needed to stimulate interest, drive conversion, or reinforce infrastructure. For example, in markets such as Nigeria and India, adoption signals are already strong, suggesting momentum that incentives can accelerate further. In regions like Bangladesh, where interest is high but infrastructure lags behind, the map ensures that rewards are channelled into strengthening local capacity. By linking adoption signals and node strain to concrete geographic cells, Shaga avoids the common inefficiency of incentives distributed "blindly".

This brings us to tokenomics, where Shaga differs sharply from other DePIN projects. Incentives are not distributed randomly or based on idle staking. Instead, emissions are tied to two concrete anchors: geography and measurable outputs. Through H3 indexing, rewards can be distributed spatially, directing incentives to underserved but high-potential areas. For infrastructure, rewards are tied to verifiable outputs such as electricity consumption, creating a clear link between real-world contribution and rewards. On the adoption side, incentives are directed to areas with potential, and growth can be measured easily through the activity of Globbers and clients.

With these foundations, the growth of the network and the sustainability of its revenues appear highly likely. Still, risks remain. The resale of gameplay data sits in a regulatory grey area. Frameworks such as GDPR in Europe and CCPA in the United States will require close compliance, and rules may evolve quickly as AI training data becomes more politicised. Infrastructure strain is another challenge. Regions such as Bangladesh or Nigeria, where interest is high, risk outpacing available nodes if reinforcements do not arrive in time, potentially undermining user experience. Energy prices and hardware cycles also pose external risks: electricity costs could narrow Shaga's cost advantage in certain markets, while GPU shortages or rising resale values could reduce supply. Finally, the AI data market itself, while growing rapidly, is still in its early stages. Broker quotes point to strong demand, but long-term pricing remains uncertain, and volatility could affect revenues. These risks underline the importance of measured scaling, adaptive incentives, and ongoing monitoring of both technical and market conditions.

Taken together, Shaga is technically validated, economically grounded, and strategically positioned. If execution continues in line with its roadmap, from the launch of the SHAGA token to the opening of a gameplay data marketplace and eventual mobile integration, it has the potential to evolve from a promising experiment into a global infrastructure layer. In that scenario, idle hardware ceases to be wasted capacity and instead becomes the foundation for the future of both gaming and AI.

DLResearch × SHAGA

From Idle GPUs to Global Infrastructure

The Shaga Model