ARUSVYX

Collective Ownership of AI Compute Infrastructure

A Network for Building, Certifying, and Collectively Owning Data-Center-Grade AI Compute

ARSX

WHITEPAPER

Version 1.0

TABLE OF CONTENTS

1. INTRODUCTION

2. MARKET CONTEXT

3. THE ARUSVYX THESIS

4. NETWORK ARCHITECTURE

5. TECHNICAL DESIGN

6. THE ASSET LAYER: COMPUTE SHARES

7. ARSX TOKEN ECONOMICS

8. MULTI-CHAIN ISSUANCE AND INTEROPERABILITY

9. GOVERNANCE AND TREASURY

10. ECOSYSTEM AND GO-TO-MARKET

11. ROADMAP

12. TEAM AND ADVISORS

13. RISK FACTORS AND LEGAL CONSIDERATIONS

APPENDIX A: GLOSSARY

APPENDIX B: TECHNICAL PARAMETERS AND REFERENCE VALUES

APPENDIX C: REFERENCES

Abstract

Compute has become the input that sets the price of intelligence. The capacity to train and serve AI models is financed today by a narrow set of balance sheets: hyperscale cloud providers, sovereign funds, private credit, and the debt facilities of a new class of specialised GPU cloud operators. The land, electricity, and demand that make this capacity productive are increasingly located in emerging regions. The ownership of the capacity is not.

Arusvyx is a network for building, certifying, and collectively owning data-center-grade AI compute. Its basic unit is the Compute Consortium: a group of participants who fund, deploy, or operate a cluster of enterprise-class accelerators hosted in a compliant facility, rent that cluster to AI companies under service-level agreements, and receive the rental income according to rules recorded on-chain. Participants enter through capital, through certified hardware, or through operating a Consortium-built cluster. Consumer graphics cards are excluded by design.

The commercial logic is borrowed from the industry itself. Accelerator vendors already deploy hardware to data center operators under revenue-sharing arrangements instead of upfront purchase. That chain stops at the operator. Arusvyx extends it one layer further, to the people and institutions in the regions where the facilities stand.

Two instruments carry the model. ARSX is the network token. It secures certified operators through bonds, governs cluster admission and network parameters, prices network services, and returns a share of protocol service fees to holders through on-chain buybacks. Compute Shares are cluster-specific asset instruments, issued through ring-fenced vehicles, that carry the rental income. It provides priority access to Compute Shares and to the economic machinery of the network. The separation is deliberate, and Sections 6 and 7 describe it in full.

Verification is split into three proofs (hardware existence, uptime, task correctness), each accompanied by a statement of what it cannot establish. ARSX has a fixed supply of one billion units. The largest allocation, community and ecosystem incentives, unlocks only as rented capacity comes online. Issuance is chain-agnostic and anchored by a canonical supply ledger, with the primary chain selected ahead of a token generation event targeted for the fourth quarter of 2026.

1. Introduction

1.1 The Ownership Gap in the Age of Compute

Every prior infrastructure wave created a way for ordinary people to own a piece of it. Railways were financed by public bond issues sold to households. Electricity reached industrial corridors through joint infrastructure syndicates that shared the burden of development. Telecommunications and, later, data center real estate were packaged into listed vehicles that anyone with a brokerage account could hold. The pattern was never perfect, but it existed, and it shaped who benefited when the technology matured.

AI compute has so far skipped this step.

The asset in question is concrete. A training or inference cluster is a set of accelerators in racks, inside a building with a power contract, a cooling system, and a depreciation schedule. It has an address. It produces revenue by the GPU-hour. Its economic life is measured in years rather than decades, which makes the timing of ownership matter more than it did for a railway. And its output, intelligence delivered as a metered service, is now an input to almost every other industry.

The capital that builds these clusters comes from a short list of sources. The largest cloud providers fund capacity from operating cash flow and corporate bonds, at a combined scale now measured in hundreds of billions of dollars a year. Sovereign wealth funds have entered as direct investors and as anchor tenants. Private credit has developed a specialised product in which loans are secured against the accelerators themselves and repaid from take-or-pay contracts. A new class of operators, often called neoclouds, has grown on the back of that credit. Public participation exists only indirectly, through the listed equity of chip vendors and cloud companies, which gives exposure to the supplier or the tenant rather than to the productive asset.

The result is a structure in which the returns on the fastest-growing infrastructure asset of the decade are captured almost entirely by institutions with access to that capital stack. This is not a failure of any single actor. It follows from the financing instruments available. It is also a gap, and gaps of this kind have historically been closed by new ownership structures rather than by waiting for incumbents to change.

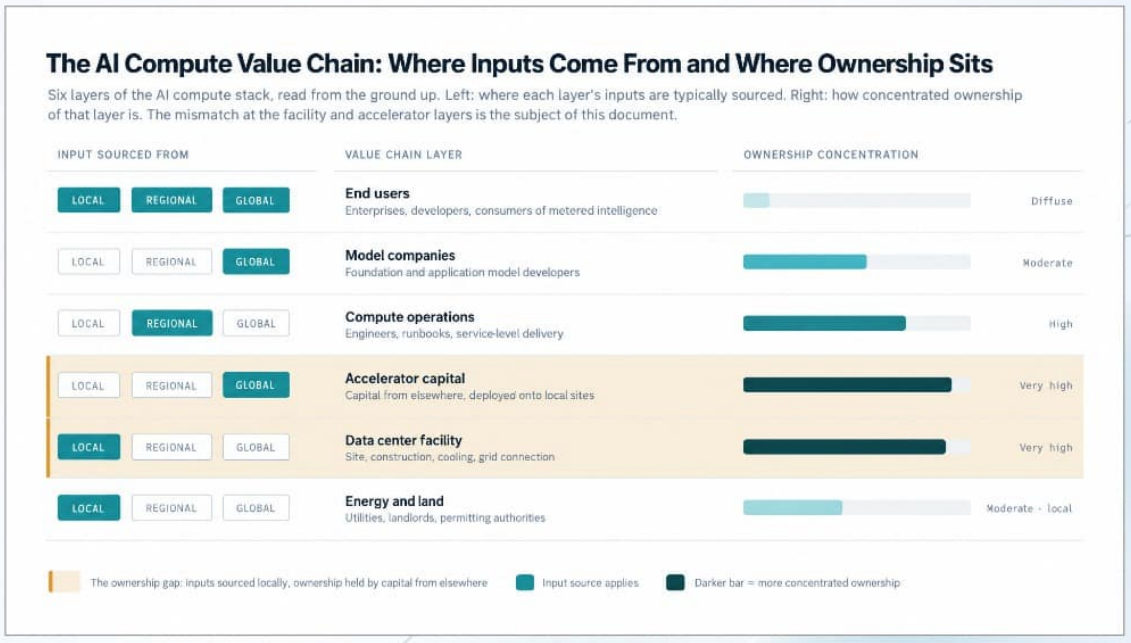
1.2 Who Builds, Who Owns, Who Benefits

The gap is most visible in the regions where new capacity is being built.

Established hubs such as Northern Virginia, Singapore, and Tokyo face binding constraints on power and land. Capacity is therefore migrating toward regions that can offer cheaper sites, available baseload electricity, and governments willing to fast-track permits in exchange for domestic AI capability. Section 2.2 treats these emerging compute corridors in detail. The immediate point concerns ownership.

Consider a single campus of the kind now being announced in Southeast Asia: a 360 MW AI factory designed to house up to 170,000 accelerators, with its first phase scheduled to enter operation in 2027 and multi-year committed offtake publicly reported in the tens of billions of dollars. The site sits on local land. It draws on the local grid. A meaningful share of its future tenants will be regional AI companies. Its construction employs local contractors, and its operation will employ local engineers. Every input except the accelerators and the capital is sourced from the surrounding economy.

The capital is sourced from elsewhere, and so is the ownership. The vendor supplies the hardware under a revenue-share arrangement. The operator holds the lease and the customer contracts. The lenders hold security over the equipment. The host economy participates as electricity supplier, landlord, and customer. It does not participate as owner.

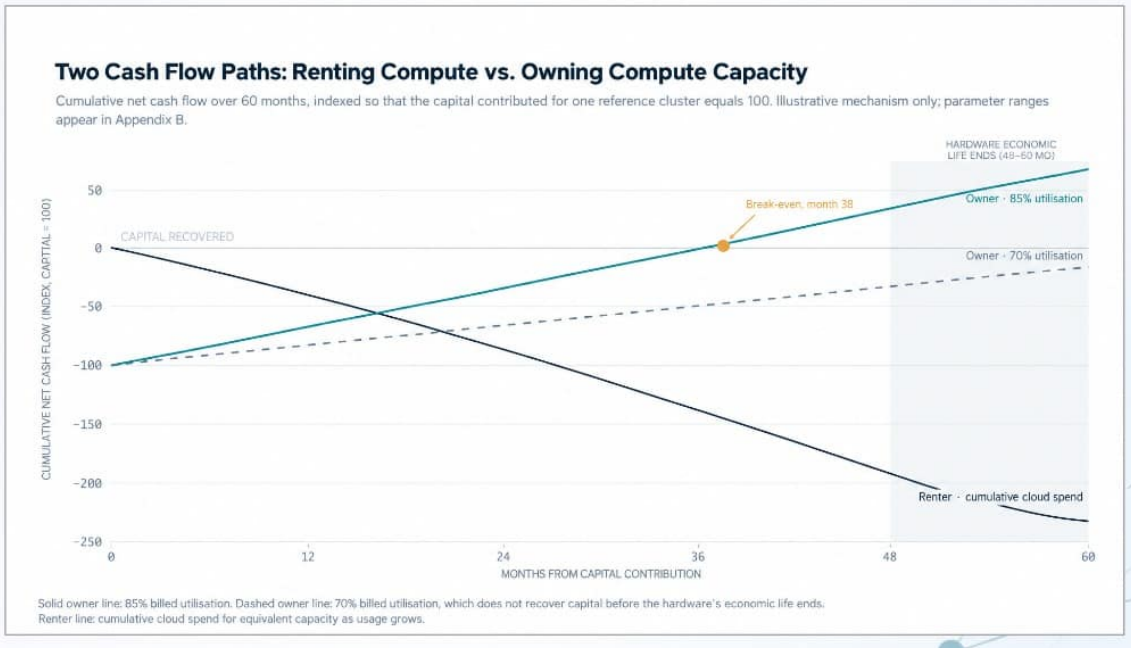


This arrangement is stable for the financiers. It is less obviously stable for the host regions, several of which have begun to describe a preference for participating in the value of AI infrastructure rather than supplying its inputs. Arusvyx is designed for that preference.

1.3 Rent or Own

For an AI company, compute is an operating expense that grows with usage. Cloud invoices scale with inference volume, and a successful product means a larger invoice every month for as long as the product exists. Renting is the rational choice when capital is scarce and demand is uncertain. It is also a permanent transfer of margin to whoever owns the capacity.

For a participant in a Compute Consortium, the same capacity is an asset. Capital goes in once. Rental income comes back over the operating life of the cluster, net of energy, hosting, operator fees, and reserves. Whether the cumulative return crosses into positive territory, and when, depends on two variables above all others: utilisation and the rental rate per GPU-hour. Both can fall. Accelerator generations turn over every few years, and a cluster that is fully rented in its first year may face price pressure in its third.



The figure is not a promise of returns. It states the shape of the choice. Arusvyx exists to make the second path available to people and institutions who have never had access to it, with the risks stated rather than hidden.

1.4 What Arusvyx Is, and What It Is Not

Arusvyx is a network in which participants collectively build, certify, own, and rent out data-center-grade AI compute, organised into Compute Consortia and coordinated by the ARSX token.

It is useful to be precise about the boundaries.

Arusvyx is not a marketplace for idle consumer graphics cards. Clusters admitted to the network must meet data-center-grade hardware, power, and service-level requirements. This excludes a large pool of potential supply, and the exclusion is intentional: enterprise customers do not sign contracts against hardware that cannot be guaranteed.

Arusvyx is not a cloud provider competing with hyperscalers on breadth of services. It provides raw and managed compute capacity for a defined set of workloads and leaves higher-level platform services to its customers and partners.

Arusvyx is not a yield product. Neither ARSX nor any material published by the network promises a rate of return. Rental income flows to holders of Compute Shares in proportion to their interest in a specific cluster, and that income varies with the performance of that cluster.

Arusvyx is not a blockchain. It uses existing chains for identity, bonding, governance, and the recording of distributions, and it is designed to operate across more than one of them.

What remains is a structure with three participation channels (capital, certified hardware, consortium-operated clusters), one basic unit (the Compute Consortium), and two instruments (ARSX for the network, Compute Shares for the asset). The rest of this document describes how those pieces fit together and what each participant can expect from them.

2. Market Context

2.1 The Capital Formation Bottleneck

For much of 2023 and 2024 the binding constraint on AI infrastructure was accelerator supply. Lead times for the most advanced parts stretched past a year, and allocation was rationed by the vendor. That constraint has eased. Manufacturing capacity has expanded, successive product generations have widened the range of viable parts, and the secondary market has deepened.

The constraint has moved. Three things now determine whether a cluster gets built: power, a site, and a financing structure.

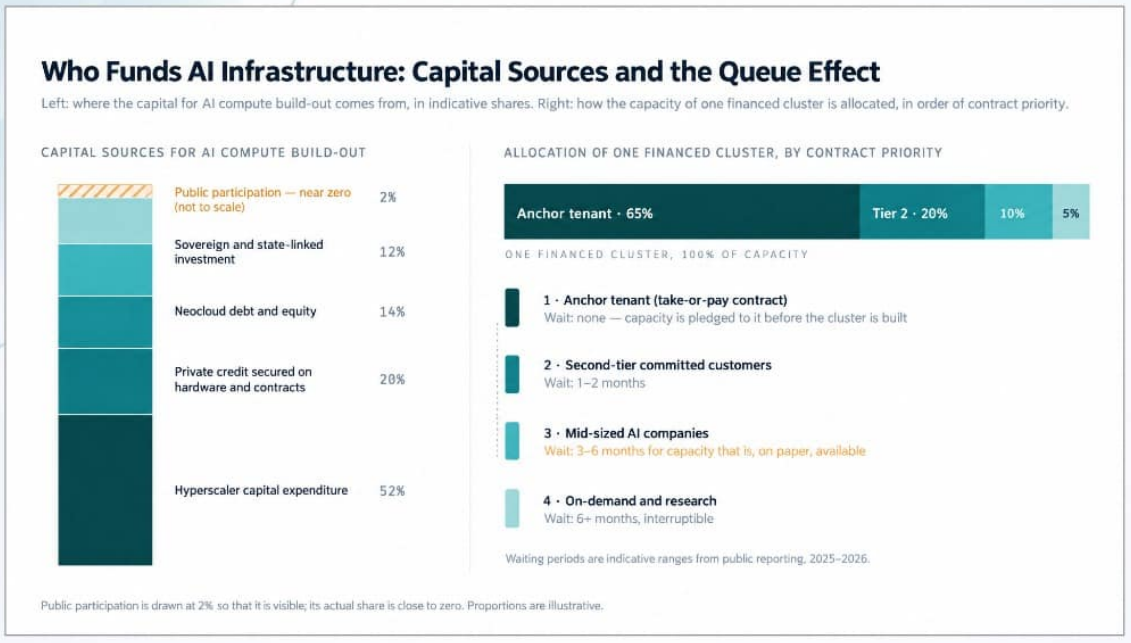
Power and sites are physical problems with physical solutions, and Section 2.2 shows where those solutions are cheapest. Financing is different. The instruments available for funding GPU clusters have been developed for and by large institutions, and they carry that origin in their design.

A typical neocloud financing in 2025 and 2026 works as follows. The operator signs a multi-year take-or-pay contract with an anchor tenant, usually a hyperscaler or a well-funded model company. The contract is pledged, together with the accelerators, as security for a loan from a private credit fund or a bank syndicate. The loan funds the purchase. The anchor tenant receives priority access to the capacity for the life of the contract. Debt raised on this template across the sector has reached the tens of billions of dollars.

The structure has two consequences that matter here.

First, allocation follows the anchor tenant. A cluster financed against a hyperscaler contract serves that hyperscaler. The mid-sized AI company that could use a few hundred accelerators for a fine-tuning run or a regional inference deployment joins a queue behind the anchor, and the queue is long. Public reporting throughout 2025 and 2026 describes model companies and enterprise AI teams waiting months for capacity that was, on paper, available.

Second, the return on the asset is divided between the lender, who takes a fixed coupon, and the operator, who keeps the residual. Neither the host region nor the smaller customer nor any individual has a route into that division.



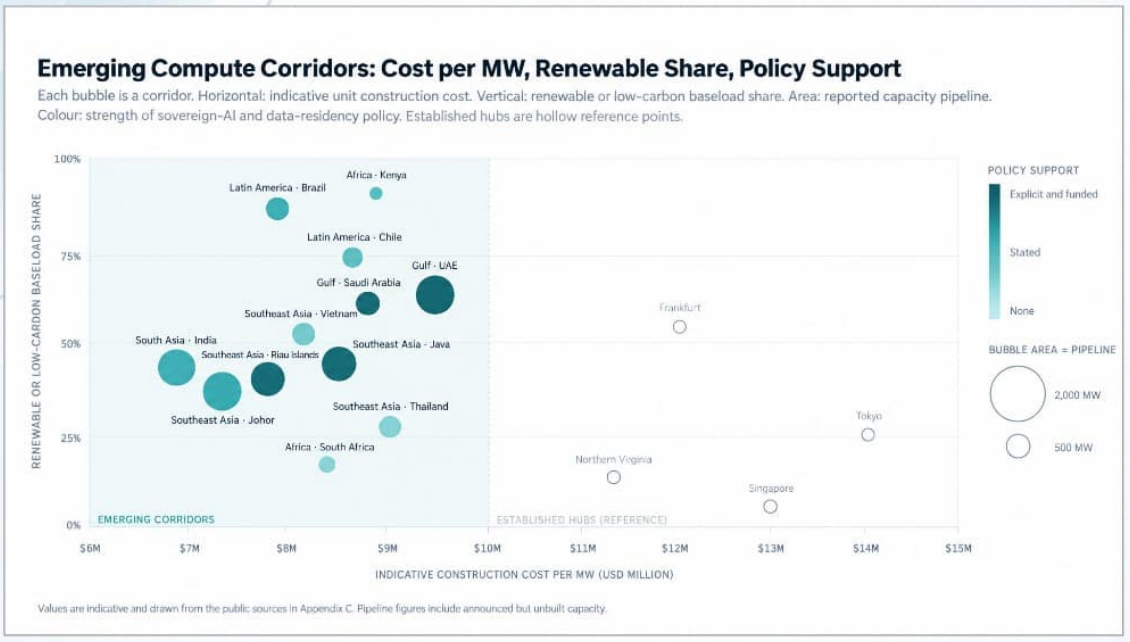
The bottleneck, then, is a bottleneck of entry points. Capital wants to fund compute. Demand wants to rent it. The instruments that connect them are built for a handful of counterparties.

2.2 Emerging Compute Corridors

An emerging compute corridor is a region that combines four conditions: unit construction cost materially below the established hubs, access to renewable or low-carbon baseload electricity, an explicit national policy favouring domestic AI capability and data residency, and growing regional demand for inference and fine-tuning. Several regions meet all four.

In Southeast Asia, capacity is concentrating in Johor and the Riau Islands adjacent to Singapore, with additional pipelines in Vietnam, Thailand, Indonesia, and the Philippines. Publicly reported construction costs in parts of the region sit well below those of Singapore, Tokyo, or Northern Virginia, and several markets combine that cost advantage with geothermal and hydroelectric resources. One market alone has a reported pipeline of roughly 1.3 GW of planned capacity on top of approximately 580 MW in operation, representing announced investment in the range of fifteen to twenty billion dollars.

The Gulf states have committed sovereign capital at scale and offer low-cost solar baseload, with data residency written into national AI strategies. South Asia combines the largest pool of AI engineering talent outside the United States with an accelerating domestic demand curve. Latin American markets, in particular Brazil and Chile, offer hydroelectric and solar power at competitive rates. Several African markets have begun to attract hyperscale investment on the strength of subsea cable landings and renewable generation.



Two features of these corridors matter for the Arusvyx model beyond cost.

The first is that data residency requirements create demand that cannot be served from outside the corridor. A bank, a hospital network, or a government agency that must process data domestically needs domestic compute. That demand is structurally protected from competition by distant hyperscale regions, and it is often mid-sized: hundreds of accelerators rather than tens of thousands. This is the scale at which a Compute Consortium operates.

The second is that new campuses in these corridors are increasingly designed as multi-tenant facilities for AI-native customers rather than single-tenant sites for one hyperscaler. A fragmented customer base means fragmented demand, and fragmented demand is served well by a network of independently owned clusters with shared scheduling and settlement.

Arusvyx will deploy its first Consortium clusters in a Southeast Asian corridor hub selected for construction cost, renewable supply, and regulatory clarity. The replication model for further corridors is described in Section 10.2.

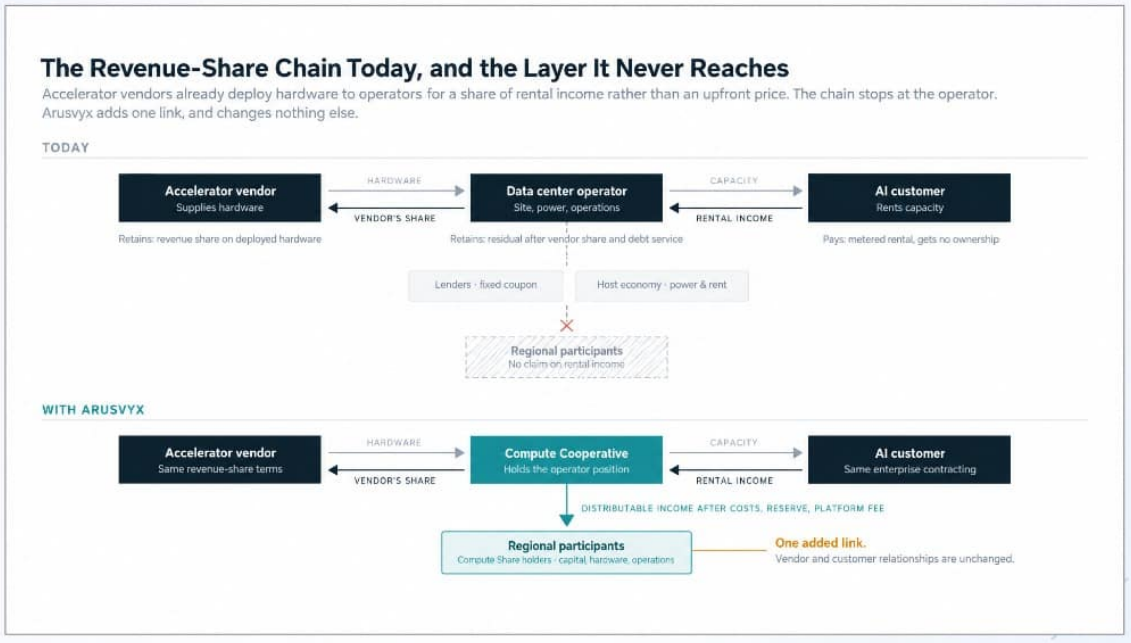
2.3 Revenue Sharing Already Exists. It Stops One Layer Too Early

The idea that participants in a compute asset should be paid from its rental income rather than from a fixed coupon is not new. It is how the top of the industry already operates.

The leading accelerator vendor has, as publicly described, established a program under which GPU infrastructure is deployed to data center operators on a revenue-sharing basis rather than through upfront purchase. The operator contributes site, power, and operations. The vendor contributes hardware and takes a share of the resulting rental income. The largest AI factory campuses now under construction in emerging corridors are financed on this basis.

The arrangement aligns incentives well. The vendor is paid when the hardware is productive. The operator's capital requirement falls. Customers get access to current-generation hardware in regions where it would otherwise arrive late.

Then the chain stops. Below the operator sit the lenders, who are paid a fixed return, and the host economy, which is paid for electricity and rent. Nobody below the operator layer holds a claim on the rental income.



Arusvyx occupies the operator position for its own clusters. A Compute Consortium holds the lease on hosting space, contracts for power, and engages a certified operator to run the hardware. The rental income the Consortium receives is distributed to Compute Share holders after costs and reserves. The vendor relationship, where applicable, follows the same revenue-share terms the industry has already accepted. The customer relationship follows standard enterprise contracting.

What changes is the set of people at the end of the chain.

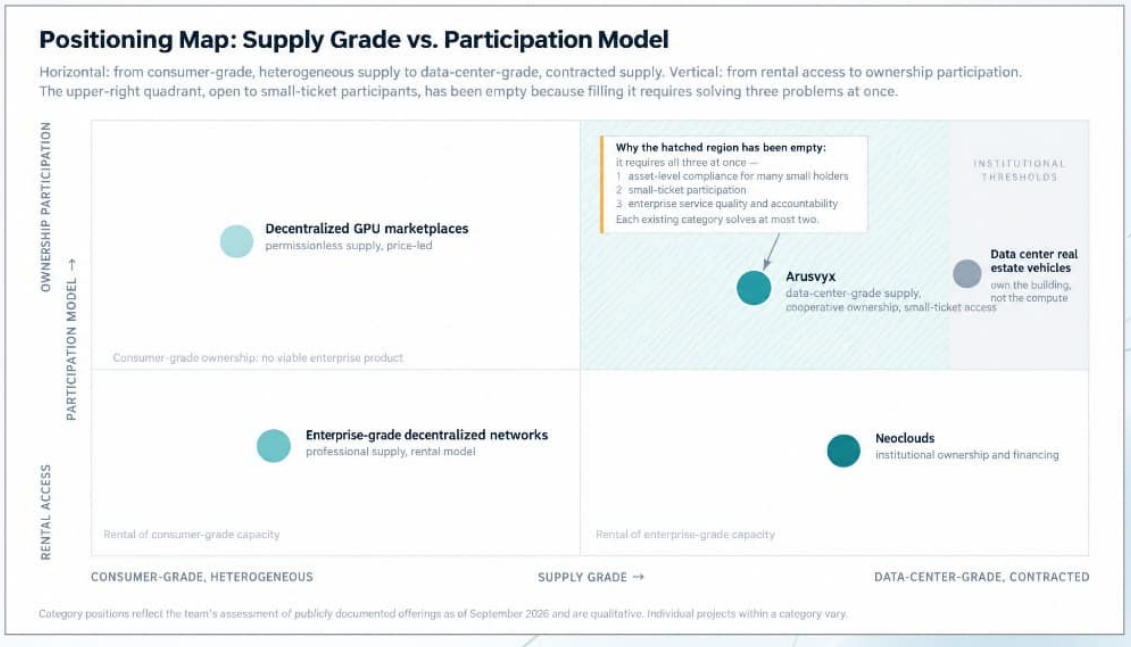
2.4 The Decentralized Compute Landscape

Several projects have set out to open compute to broader participation, and it is worth being accurate about what each has achieved.

Decentralized GPU marketplaces aggregate supply from a large number of independent providers, most of it consumer or prosumer hardware, and match it with price-sensitive workloads. They have demonstrated that permissionless supply can be coordinated at scale, and they offer some of the lowest prices available for batch rendering, certain classes of model training, and experimentation. Their limits are the limits of their supply. Heterogeneous hardware, variable uptime, and the absence of a contractual counterparty make it difficult to offer the service-level guarantees and data-handling commitments that enterprise and regulated customers require. Adoption among those customers has remained modest.

A second group has moved toward enterprise-grade supply, sourcing from data centers and professional operators and offering something closer to a conventional cloud service. These networks address the reliability problem. They remain rental businesses: the customer pays for capacity, the provider earns from it, and the token coordinates the two. Ownership of the underlying hardware stays with the provider.

Neoclouds sit at the enterprise end of the supply spectrum with institutional ownership and institutional financing. Data center real estate vehicles offer ownership, but of the building rather than the compute, at yields that reflect real estate rather than accelerator economics, and with participation thresholds set for institutional investors.



The empty quadrant is not empty by accident. Data-center-grade supply requires capital and professional operations. Ownership participation requires a legal structure that can hold an income-producing asset on behalf of many small holders without becoming unmanageable. Enterprise service quality requires accountability that a purely permissionless network cannot provide. Any two of these can be solved with existing tools. All three together require the Consortium structure, the certified operator model, and the instrument separation that the following sections describe.

3. The Arusvyx Thesis

3.1 Mission and Design Principles

Arusvyx exists to make data-center-grade AI compute an asset that people and institutions in the regions hosting it can build, own, and earn from.

Four principles govern how the network is designed. They are stated here once and applied throughout the sections that follow.

Own, don't rent. Wherever the network faces a choice between structuring participation as a rental relationship or as an ownership interest, it chooses ownership. This is the principle that produces Compute Shares rather than discounted capacity credits, and it is the reason the network accepts the legal complexity of holding income-producing hardware on behalf of many holders.

Collective build. A cluster is financed, deployed, and governed by a group. No single participant is expected to fund a full cluster, and no single participant controls one. The consortium structure was chosen because it has a long record in telecommunications, energy grids, and large-scale industrial infrastructure, coordinating independent contributors around a shared productive asset while keeping governance transparent. The name Compute Consortium reflects this enterprise-grade alignment.

Real yield, real asset. Every distribution in the network traces to rental income from identifiable hardware. The token has no inflationary emission schedule to fund returns, and the incentive pool unlocks only when new capacity begins producing revenue. Where the network cannot show a customer invoice behind a payment, the payment does not happen.

Green and sovereign. Clusters are sited where baseload power is renewable or low-carbon and where data can be processed under the residency rules of the host jurisdiction. This is partly a matter of cost, since renewable baseload in emerging corridors is often the cheapest electricity available, and partly a matter of demand, since residency-constrained workloads cannot be served from anywhere else.

The principles are not free of tension. Ownership requires investor eligibility rules that a pure rental model would not. Collective governance is slower than an operator acting alone. Insisting on invoiced revenue means the network grows only as fast as it can sell capacity. These costs are accepted because the alternative, a faster network that does not deliver ownership, already exists in several forms and does not need another entrant.

3.2 The Compute Ownership Flywheel

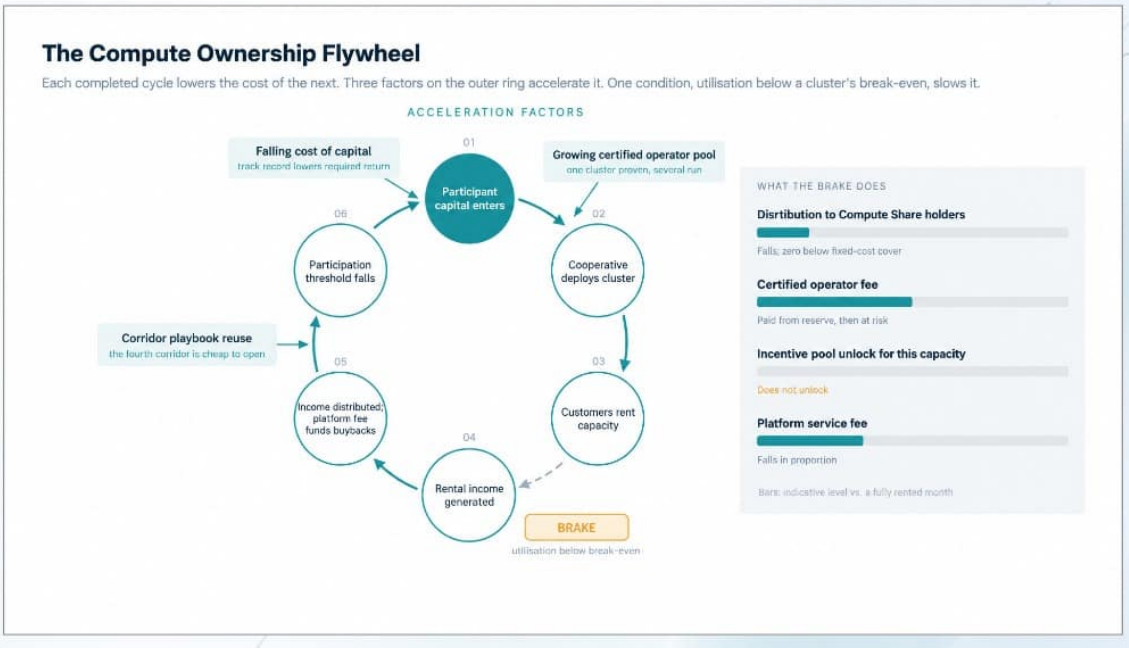
A single Consortium cluster is a modest business: several hundred accelerators, a handful of customers, income that depends on utilisation. The network becomes interesting when clusters compound.

Capital enters through participants who acquire Compute Shares in a new Consortium. The Consortium deploys a cluster with a certified operator in a compliant facility. Customers rent the cluster. Rental income is distributed to Compute Share holders after costs and reserves, and the platform service fee funds ARSX buybacks and network operations. Each completed cycle lowers the cost of the next one in three specific ways.

The network's record of completed clusters lowers the cost of capital for subsequent Consortia. A Consortium with no operating history must offer participants a return that compensates for execution risk. The tenth Consortium, launched by a network with nine clusters producing invoiced income, does not. Lower required returns mean a larger share of rental income is retained by participants or reinvested.

The pool of certified operators grows. Each new corridor initially has few operators able to meet the network's hardware and service standards. Operators who have run one cluster successfully can run several, and their presence makes the next cluster in that corridor faster and cheaper to deploy.

The corridor playbook, covering site selection, hosting contracts, operator certification, Consortium formation, and local regulatory mapping, is reused. The first corridor is expensive to open. The fourth is not.



The brake is part of the design rather than a footnote. A Consortium whose cluster cannot find tenants distributes less, its operator earns less, and the incentive pool for that capacity does not unlock. The flywheel has no mechanism that pays participants for capacity nobody rents. This is what distinguishes it from token networks in which emissions substitute for demand.

3.3 Following a Dollar of Compute Revenue

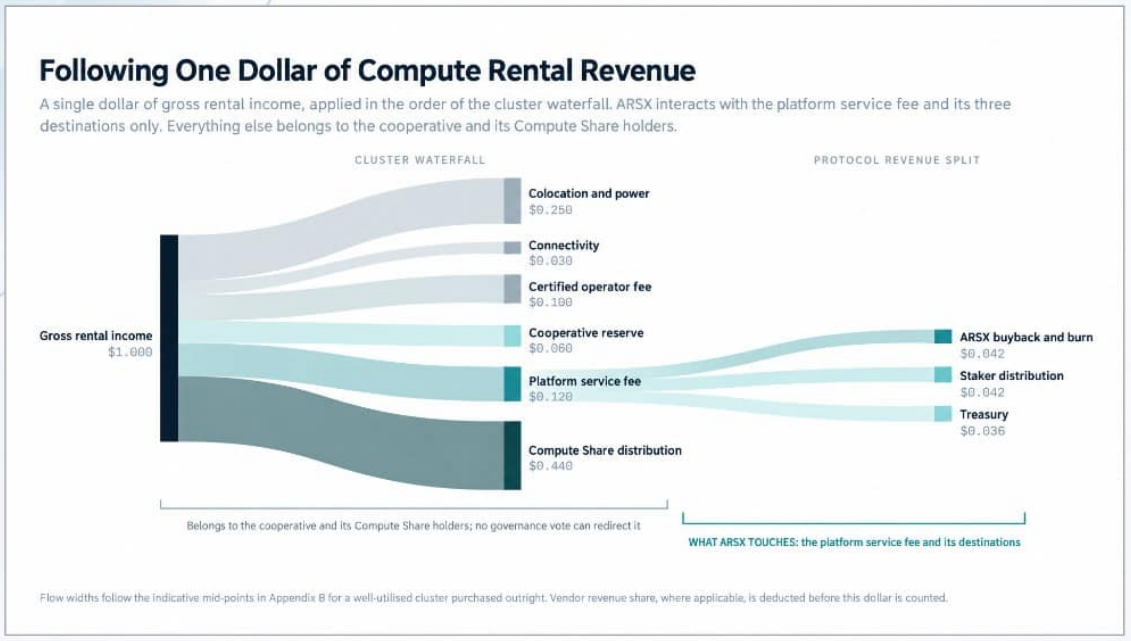
A customer pays the Consortium for a month of capacity. The invoice is settled in stablecoin. What happens to it?

The first deductions are the direct costs of keeping the cluster running: colocation and power under the hosting contract, network connectivity, and the certified operator's fee for managing the hardware and meeting the service-level agreement. Together these typically absorb somewhere between a third and half of gross rental income for a well-utilised cluster, with power the largest single component and the reason corridor selection matters so much.

Next, the Consortium sets aside a reserve. Reserves cover hardware replacement, unplanned downtime, disputes, and, over the life of the cluster, the eventual refresh or decommissioning of the accelerators. The reserve rate is a governance parameter for each Consortium, subject to a network-wide minimum.

The platform service fee is then deducted. This is the network's revenue. Its indicative range is ten to fifteen percent of gross rental income, and it funds three things: on-chain buybacks of ARSX, distributions to ARSX stakers who provide bonding and governance services, and the network treasury. The split among those three is set in Section 7.6.

What remains, after costs, reserves, and the platform fee, is distributed to Compute Share holders in proportion to their interest in the cluster.



Two things are visible in the diagram that are easy to lose in prose. ARSX touches only the platform service fee. The rental income that belongs to Compute Share holders never passes through the token, and no governance vote can redirect it. And the largest single claim on a dollar of revenue is the hardware's own operating cost, which is why the network's siting, hosting, and operator standards are treated as economic questions rather than technical ones.

4. Network Architecture

4.1 Network Overview

The network has six layers. Three are physical or contractual and operate off-chain. Three are coordinative and operate on-chain. The boundary between them is drawn deliberately, and each layer is accountable to a specific source of truth.

Participation layer. Individuals and institutions who contribute capital, hardware, or operations. Their identities, eligibility status, and holdings are recorded on-chain. Their contractual relationships with Consortia and asset vehicles are off-chain legal instruments.

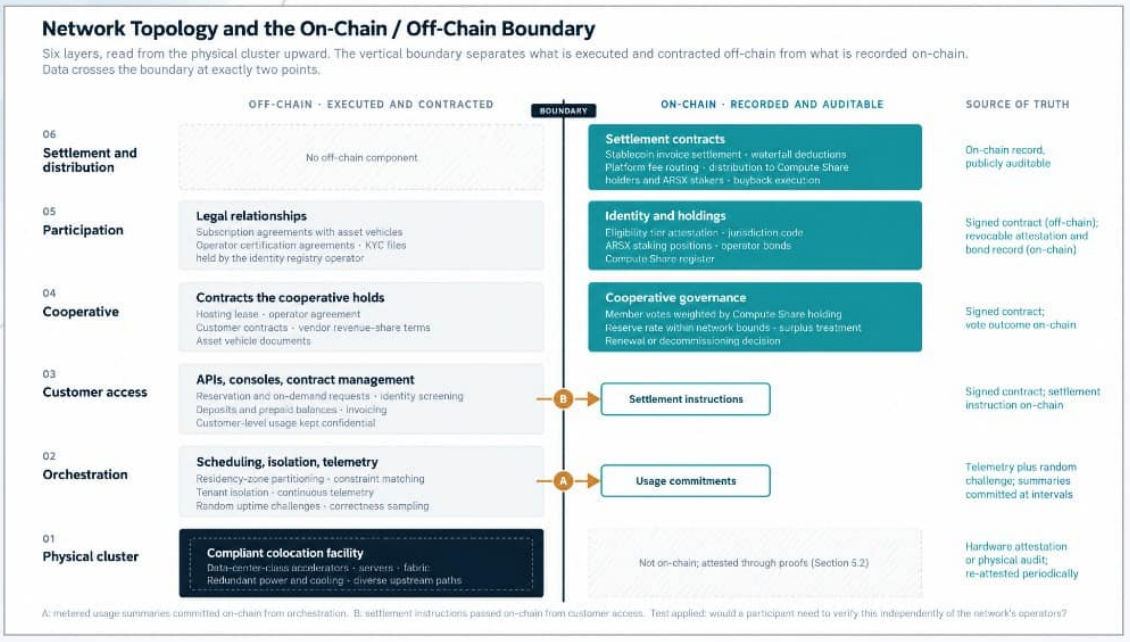
Consortium layer. The Compute Consortia themselves: governance bodies that hold hosting leases, engage operators, admit customers, and receive rental income. Consortium governance decisions are recorded on-chain. The leases, operator agreements, and customer contracts they govern are conventional legal documents.

Physical cluster layer. Accelerators, servers, and networking equipment installed in compliant colocation facilities. Entirely off-chain. Its existence and condition are attested through the verification mechanisms described in Section 5.2.

Orchestration layer. Workload scheduling, tenant isolation, and telemetry collection across clusters. Operated off-chain by the network's software, with scheduling constraints and metered usage summaries committed on-chain at defined intervals.

Customer access layer. APIs, consoles, and contract management through which AI companies reserve and consume capacity. Off-chain, with settlement instructions passed to the on-chain layer.

Settlement and distribution layer. Stablecoin settlement of invoices, execution of the cost and reserve deductions, platform fee routing, and distribution to Compute Share holders and ARSX stakers. On-chain, and publicly auditable.



The choice of what to put on-chain follows one test: is the record something a participant would need to verify independently of the network's operators? Holdings, bonds, governance outcomes, and distributions pass that test. Workload execution and hardware telemetry do not, since they are voluminous, partly confidential to customers, and verifiable only through the attestation and challenge mechanisms described later. Committing summaries rather than raw data keeps the on-chain footprint small and the off-chain systems auditable.

Security of the on-chain layer. Before deployment, every contract is reviewed by at least two independent audit firms, and the reports are published together with each finding and how it was resolved. The non-upgradeable contracts, which hold the supply cap, the vesting schedules, and the Consortium settlement logic, receive a further review after any change to the contracts that call them. A bug bounty funded from the treasury runs continuously, with rewards scaled to severity up to a published ceiling.

The council's signing keys are held on hardware devices by members located in different jurisdictions, and no two members use the same custodian. Key rotation follows a schedule and is triggered immediately on any change in council membership.

When an incident is confirmed, the security council pauses the affected contracts or cross-chain paths within the time bound in Appendix B. Participants are notified within a fixed period, and a post-incident report stating the root cause, the funds affected, and the remediation is published within a further period. Where participants have lost funds, the report is accompanied by a governance proposal on whether and how the treasury reserve is used for restitution. The network does not commit to restitution in advance, and no participant should assume it.

4.2 Three Participation Channels

The network admits three kinds of contribution. They differ in what the participant brings, what the network requires of them, how they are paid, and what they stand to lose.

Capital. A participant acquires Compute Shares in a specific Consortium, subject to the eligibility rules of the asset vehicle issuing them. Their contribution funds hardware, deployment, and initial reserves. They receive distributions from that cluster's rental income for its operating life. They bear the risk that the cluster underperforms, that hardware fails outside warranty, or that the accelerator generation is superseded before the capital is recovered. They exit by transferring their shares where the vehicle permits it, or by holding to the cluster's decommissioning and final distribution.

Certified hardware. A university compute center, a regional data center operator, or an enterprise with an underused enterprise-grade cluster brings that hardware into the network after certification. The operator posts an ARSX bond scaled to the capacity contributed and commits to a service-level agreement. They earn a share of rental income from their own hardware and an operator fee where they also manage Consortium-owned clusters. They risk bond slashing for service-level breaches and bear the ordinary risks of hardware ownership. They exit by giving notice, completing outstanding customer commitments, and withdrawing the hardware and any unslashed bond.

Consortium-operated clusters. A Consortium, rather than acquiring hardware and engaging an external operator, leases colocation space and runs its own cluster, with members or a member-appointed team performing the operator role. This channel suits Consortia anchored by an institution that already has data center staff. The Consortium earns both the Compute Share distributions and the operator fee. It bears the operational risk directly and is subject to the same bond and certification requirements as an external operator.



Nothing prevents a participant from occupying more than one channel. A regional operator may contribute certified hardware to one Consortium and hold Compute Shares in another. The channels describe roles, and roles can be combined.

4.3 Compute Consortia

A Compute Consortium is the unit at which the network holds assets, signs contracts, and makes decisions. Every cluster in the network belongs to exactly one Consortium. A Consortium may in time hold more than one cluster.

Formation. A Consortium is proposed by a sponsor, who may be an institution, a group of prospective participants, or the network foundation. The proposal specifies the target corridor, the intended cluster size and hardware generation, the hosting facility or the criteria for selecting one, the operator or the operator selection process, the reserve policy, and the terms on which Compute Shares will be issued. Network governance reviews the proposal against the standards in this document and admits or declines it.

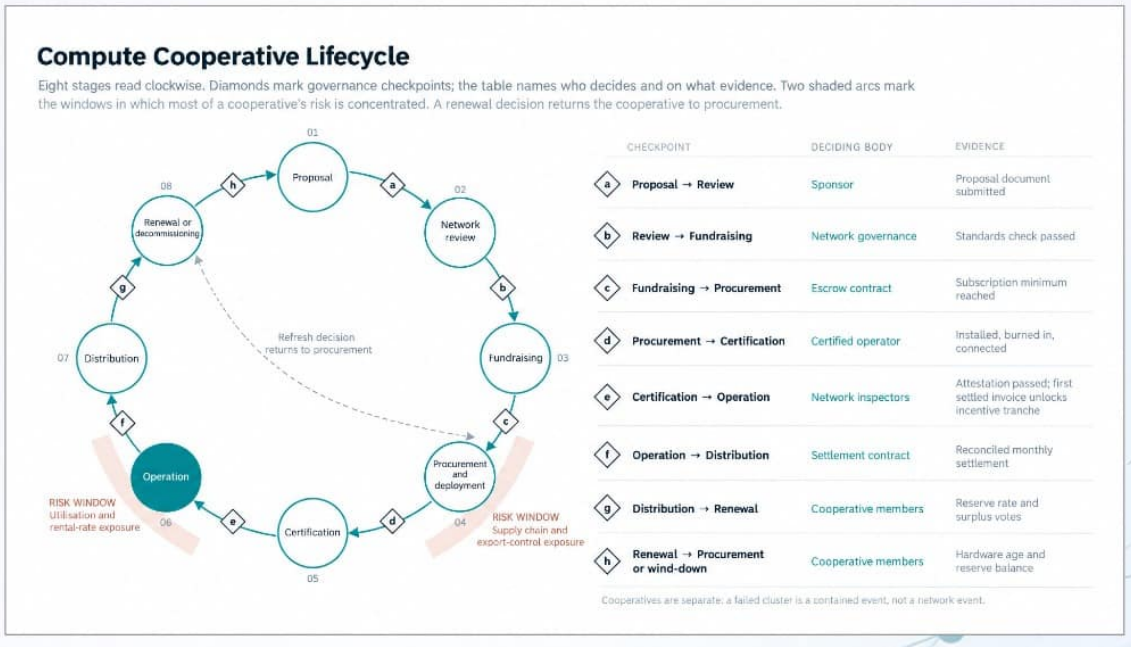
Fundraising. Compute Shares are offered through the Consortium's asset vehicle to eligible participants. The offering has a minimum and a target. Below the minimum, the Consortium does not proceed and contributions are returned. The network's incentive pool does not release anything at this stage; nothing has been built.

Procurement and deployment. Hardware is procured, either purchased outright or, where the Consortium qualifies, taken under a vendor revenue-share arrangement of the kind now common at the top of the industry. The hosting lease is executed. The certified operator installs, burns in, and connects the cluster to the orchestration layer.

Certification. The cluster passes the network's hardware existence and performance attestation. From this point it is visible to the scheduler. The first tranche of the incentive pool attributable to this capacity unlocks when the cluster settles its first customer invoice.

Operation. The cluster serves customers. Income is settled monthly. The Consortium's governance body, composed of Compute Share holders voting in proportion to their interest, decides on customer admission policies where the operator seeks guidance, reserve rate adjustments within network bounds, and the treatment of surpluses.

Renewal or decommissioning. As the hardware approaches the end of its economic life, the Consortium decides whether to refresh the cluster from reserves and new contributions, to sell the hardware and distribute the proceeds, or to wind down. The decision is taken by the Compute Share holders, and the network's role is limited to confirming that the chosen path meets its standards.



Consortia are separate from one another. A Consortium's reserves cover its own cluster. Its Compute Share holders have no claim on other Consortia' income and no liability for their losses. This separation is the reason a failed cluster is a contained event rather than a network event.

4.4 Certified Operators and the Hardware Grade Boundary

The network's commitment to enterprise customers rests on the quality of the hardware it schedules and the accountability of the people who run it. Both are enforced at the point of admission.

Hardware grade. To be admitted, a cluster must meet requirements across five dimensions, with reference values maintained in Appendix B and revised through governance as hardware generations turn over. Accelerators must be data-center-class parts with error-correcting high-bandwidth memory, supported by the vendor for enterprise deployment, and interconnected within the node and across nodes at bandwidths sufficient for distributed training and high-throughput inference. The facility must provide redundant power and cooling at a tier equivalent to what enterprise colocation customers expect. Network connectivity must include redundant upstream paths. The operator must be able to commit to uptime and response-time targets that the network's standard service-level agreement specifies. The hardware must be capable of tenant isolation at the level customers' data-handling commitments require.

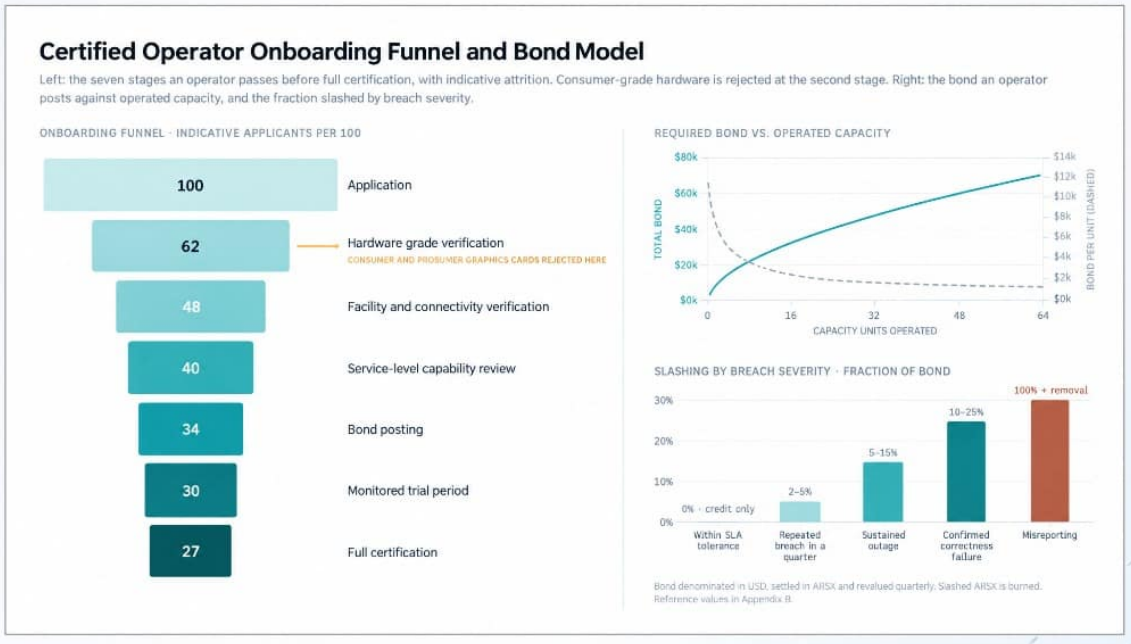
Consumer and prosumer graphics cards fail the first requirement and are not admitted. The exclusion removes a large and enthusiastic pool of potential supply, and the network has considered and rejected proposals to admit it in a lower tier. Enterprise contracts cannot be written against hardware the vendor does not support for that use, and a network that offered two tiers would find its reputation set by the lower one.

Siting and power disclosure. A hosting facility discloses its power source mix at the time a cluster is certified, from its supply contracts and, where the grid operator provides it, hourly matching data. The Consortium publishes the mix alongside the cluster's other attributes, and it is updated annually on the dashboard.

A network minimum share of renewable or low-carbon supply applies to new clusters, with the reference value in Appendix B. A facility below the minimum may host a cluster only where the Consortium's members have voted to accept it with the shortfall stated in the offering documents. The network describes a cluster as renewable-powered only where the disclosed share exceeds the minimum on the basis of contracted or matched supply. Unbundled certificates purchased after the fact do not count toward it.

Operator certification. An operator is admitted after demonstrating that it controls the hardware it proposes to contribute, that the facility and connectivity meet the standards above, that it has the staffing and runbooks to meet the service-level agreement, and that it has posted the required bond. A trial period follows, during which the cluster serves customers under close monitoring before full certification.

The operator bond. Each certified operator posts ARSX in proportion to the capacity it operates. The bond has two purposes. It gives the operator a financial stake in meeting its service commitments, and it provides a source from which customer compensation can be paid when those commitments are missed. Slashing is graduated. A brief unplanned outage within the tolerance of the service-level agreement costs nothing. Repeated breaches, a sustained outage, or a failed correctness challenge cost a defined fraction of the bond, escalating with severity. Evidence of deliberate misreporting costs the whole bond and results in removal from the network. Slashed ARSX is burned.



The bond is denominated in ARSX rather than in stablecoin because it is meant to tie the operator's interest to the network's. An operator whose bond loses value because the network is performing poorly has a reason to help fix that. The network accepts the cost of this choice, which is that bond values fluctuate with the token, and manages it by setting bond requirements in capacity terms and revaluing them at defined intervals rather than continuously.

5. Technical Design

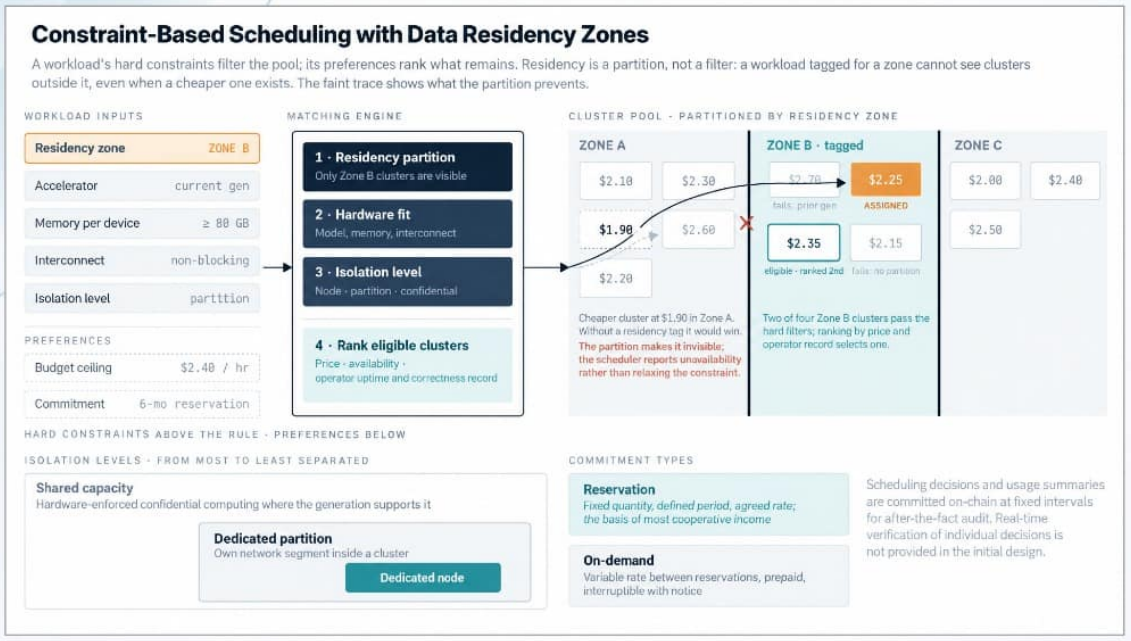
5.1 Workload Scheduling and Data Residency

A customer reaches the network with a workload and a set of constraints. Some constraints are hard: the jurisdiction in which data may be processed, the accelerator generation and memory required, the isolation level demanded by the customer's own compliance obligations. Others are preferences: a budget ceiling, a preferred start time, a tolerance for interruption. The scheduler's job is to find the cluster that satisfies every hard constraint and scores best on the preferences, then to hold that capacity for the duration of the commitment.

Data residency is treated as a partition of the network rather than as a filter applied after the fact. Every cluster belongs to exactly one residency zone, defined by the jurisdiction of its hosting facility and, where relevant, by sub-national rules. A workload tagged with a zone is visible only to clusters in that zone. It cannot spill over to a cheaper cluster elsewhere when its home zone is full, and the scheduler will report unavailability rather than silently relax the constraint. This is a deliberate loss of efficiency. Customers with residency obligations are among the network's most valuable tenants precisely because their demand cannot be served from distant regions, and a scheduler that could be talked into relaxing residency would forfeit their trust in one incident.

Within a zone, matching proceeds on three further dimensions. Hardware fit checks accelerator model, memory per device, and interconnect topology against the workload's declared requirements. Isolation level selects among dedicated nodes, a dedicated partition of a cluster with its own network segment, or shared capacity with hardware-enforced confidential computing where the accelerator generation supports it. Price and availability then rank the eligible clusters, with an adjustment for operator track record drawn from the uptime and correctness history described in the next section.

Two commitment types are supported. Reservations hold a fixed quantity of capacity for a defined period at an agreed rate, and they are the basis of most Consortium income. On-demand capacity fills the gaps between reservations at a variable rate and is interruptible with notice. A cluster's reservation ratio is one of the numbers a Consortium watches most closely, since it determines how much of the cluster's income is contracted and how much depends on spot demand.



The scheduler is operated by the network's software and is, at launch, a coordinated service rather than a fully decentralised one. Its scheduling decisions and the resulting usage summaries are committed on-chain at fixed intervals, which lets any participant audit after the fact whether capacity was allocated as the rules require. Real-time verification of individual scheduling decisions is not provided in the initial design.

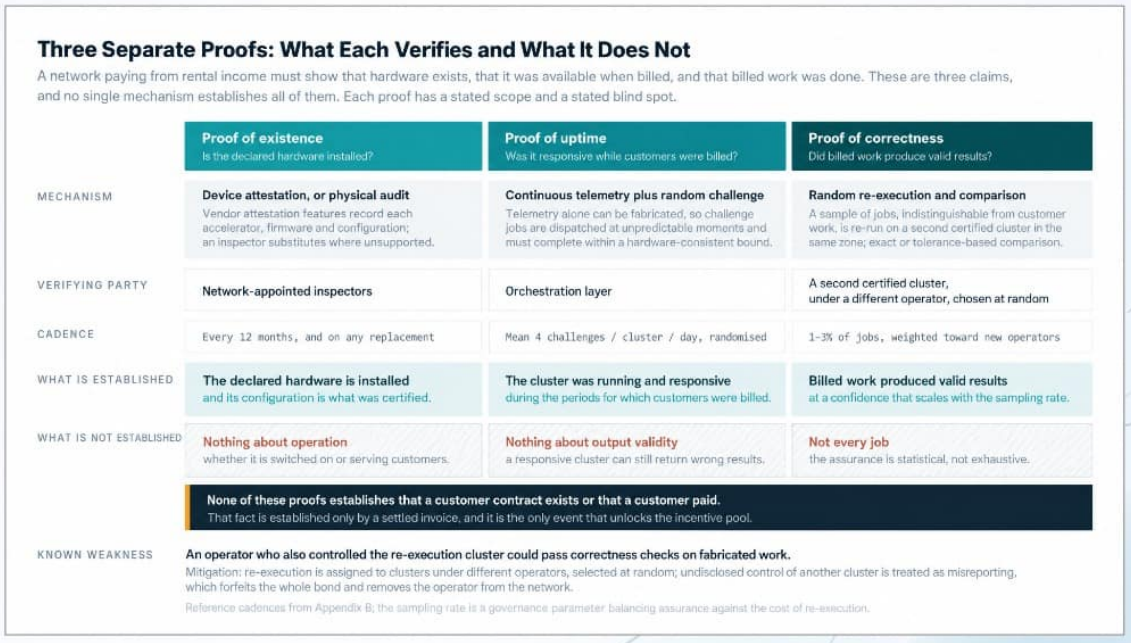
5.2 Verification: Existence, Uptime, Correctness

A network that pays participants from rental income must be able to show that the hardware exists, that it was available when a customer paid for it, and that the work billed was actually done. These are three different claims, and no single mechanism establishes all of them. Arusvyx treats them as three separate proofs, each with a stated scope and a stated blind spot.

Proof of existence establishes that a cluster consists of the hardware its operator declared. At certification, each accelerator is identified through the vendor's device attestation features where available, and its firmware and configuration are recorded. Where a hardware generation does not support cryptographic attestation, a physical audit by a network-appointed inspector substitutes. Re-attestation occurs at defined intervals and whenever hardware is replaced. The proof tells the network what is installed. It says nothing about whether the hardware is switched on or serving customers.

Proof of uptime establishes that the cluster was available during the periods for which customers were billed. Clusters stream telemetry to the orchestration layer continuously. Telemetry alone can be fabricated, so the network supplements it with random challenges: at unpredictable moments, a challenge job is dispatched to the cluster and must complete within a time bound consistent with the declared hardware. A cluster that reports availability but fails challenges is treated as unavailable for the affected period. The proof tells the network that the cluster was running and responsive. It does not confirm that the outputs delivered to customers were correct.

Proof of correctness establishes, statistically, that billed work produced valid results. A fraction of jobs, selected at random and indistinguishable to the operator from ordinary customer work, is re-executed on a separately certified cluster in the same residency zone and compared. Deterministic workloads are compared exactly. Inference workloads with non-deterministic elements are compared within a tolerance defined per model class. A failed comparison triggers investigation and, if confirmed, slashing. The proof gives a probabilistic assurance that scales with the sampling rate, which is a governance parameter balancing assurance against the cost of re-execution. It does not check every job.



Three Separate Proofs: What Each Verifies and What It Does Not

A network paying from rental income must show that hardware exists, that it was available when billed, and that billed work was done. These are three claims, and no single mechanism establishes all of them. Each proof has a stated scope and a stated blind spot.

	Proof of existence Is the declared hardware installed?	Proof of uptime Was it responsive while customers were billed?	Proof of correctness Did billed work produce valid results?
MECHANISM	**Device attestation, or physical audit** Vendor attestation features record each accelerator, firmware and configuration; an inspector substitutes where unsupported.	**Continuous telemetry plus random challenge** Telemetry alone can be fabricated, so challenge jobs are dispatched at unpredictable moments and must complete within a hardware-consistent bound.	**Random re-execution and comparison** A sample of jobs, indistinguishable from customer work, is re-run on a second certified cluster in the same zone; exact or tolerance-based comparison.
VERIFYING PARTY	Network-appointed inspectors	Orchestration layer	A second certified cluster, under a different operator, chosen at random
CADENCE	`Every 12 months, and on any replacement`	`Mean 4 challenges / cluster / day, randomised`	`1–3% of jobs, weighted toward new operators`
WHAT IS ESTABLISHED	**The declared hardware is installed** and its configuration is what was certified.	**The cluster was running and responsive** during the periods for which customers were billed.	**Billed work produced valid results** at a confidence that scales with the sampling rate.
WHAT IS NOT ESTABLISHED	**Nothing about operation** whether it is switched on or serving customers.	**Nothing about output validity** a responsive cluster can still return wrong results.	**Not every job** the assurance is statistical, not exhaustive.

None of these proofs establishes that a customer contract exists or that a customer paid.
That fact is established only by a settled invoice, and it is the only event that unlocks the incentive pool.

KNOWN WEAKNESS

An operator who also controlled the re-execution cluster could pass correctness checks on fabricated work.
Mitigation: re-execution is assigned to clusters under different operators, selected at random; undisclosed control of another cluster is treated as misreporting, which forfeits the whole bond and removes the operator from the network.

Reference cadences from Appendix B; the sampling rate is a governance parameter balancing assurance against the cost of re-execution.

The proofs share a known weakness. An operator who also controlled the cluster used for re-execution could pass correctness checks on fabricated work. The network mitigates this by assigning re-execution to clusters under different operators, selected at random, and by treating any operator found to hold undisclosed control of another cluster as having misreported, with the consequences that carries for its bond.

5.3 Metering and Settlement

Between a customer's job and a Compute Share holder's distribution lies a pipeline with five stages, each running on its own clock.

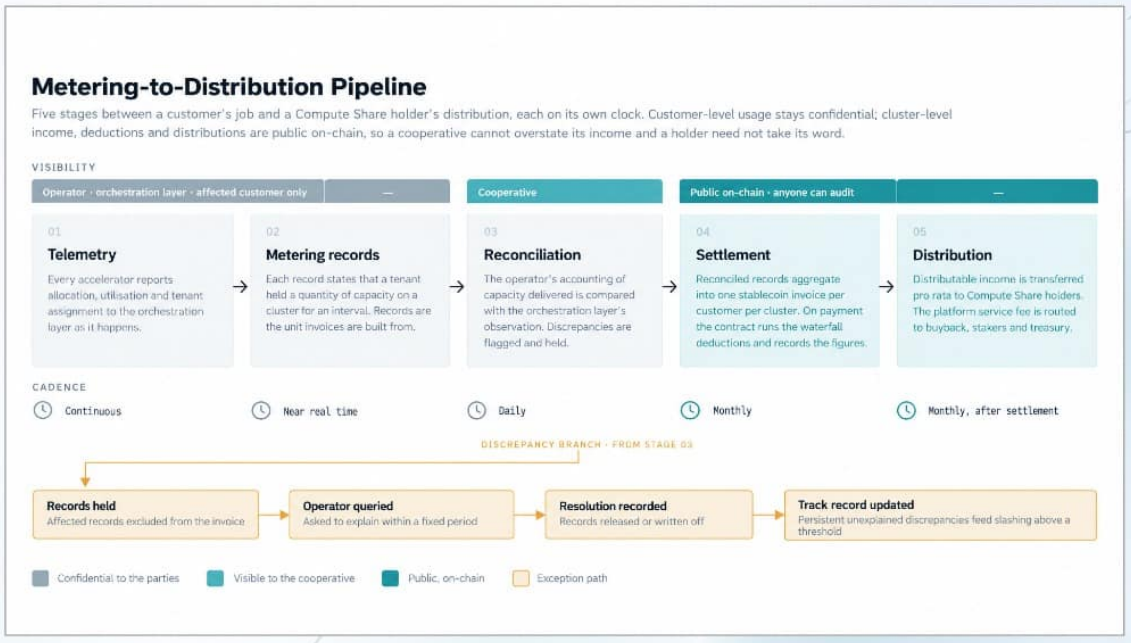
Telemetry is continuous. Every accelerator reports allocation, utilisation, and tenant assignment to the orchestration layer as it happens.

Metering records are produced from telemetry in near real time. A metering record states that a specific tenant held a specific quantity of capacity on a specific cluster for a specific interval, and it is the unit from which invoices are built.

Reconciliation runs daily. The operator's own accounting of capacity delivered is compared against the orchestration layer's observation. The two should agree. When they do not, the discrepancy is flagged, the affected records are held, and the operator is asked to explain. Persistent unexplained discrepancies feed into the operator's track record and, above a threshold, into slashing.

Settlement runs monthly. Reconciled metering records are aggregated into an invoice per customer per cluster, denominated in a stablecoin accepted by the network, and presented for payment. On payment, the settlement contract executes the deductions in the order given in Section 6.2 and records the resulting figures on-chain.

Distribution follows settlement. Distributable income is transferred to Compute Share holders in proportion to their interest. The platform service fee is routed to its three destinations.



Customer contracting and payment. A customer passes identity and sanctions screening before reserving capacity, at a depth proportionate to the size of the commitment. Reservations are invoiced monthly in arrears against a deposit of one month's committed fee, held in the vehicle's account and returned at the end of the term. On-demand capacity is prepaid: the customer funds a balance, metered usage is drawn from it, and the scheduler releases capacity only while the balance covers it.

Invoices are denominated in US dollars and settled in a stablecoin the network accepts. Where a customer's jurisdiction restricts settlement in that form, the licensed intermediary path described in Section 13.2 applies, and the vehicle receives stablecoin in either case. A reservation unpaid beyond the grace period is suspended and the deposit is drawn; the Consortium's governance decides whether to terminate the reservation and release the capacity.

The visibility choices are worth spelling out. Customer-level usage is confidential, because customers regard their compute consumption as commercially sensitive and would not use a network that published it. Cluster-level income, on the other hand, is public. Anyone can see what a Consortium's cluster earned in a given month, what was deducted, and what was distributed. A Consortium cannot overstate its income to prospective participants, and a Compute Share holder does not have to take the Consortium's word for what it received.

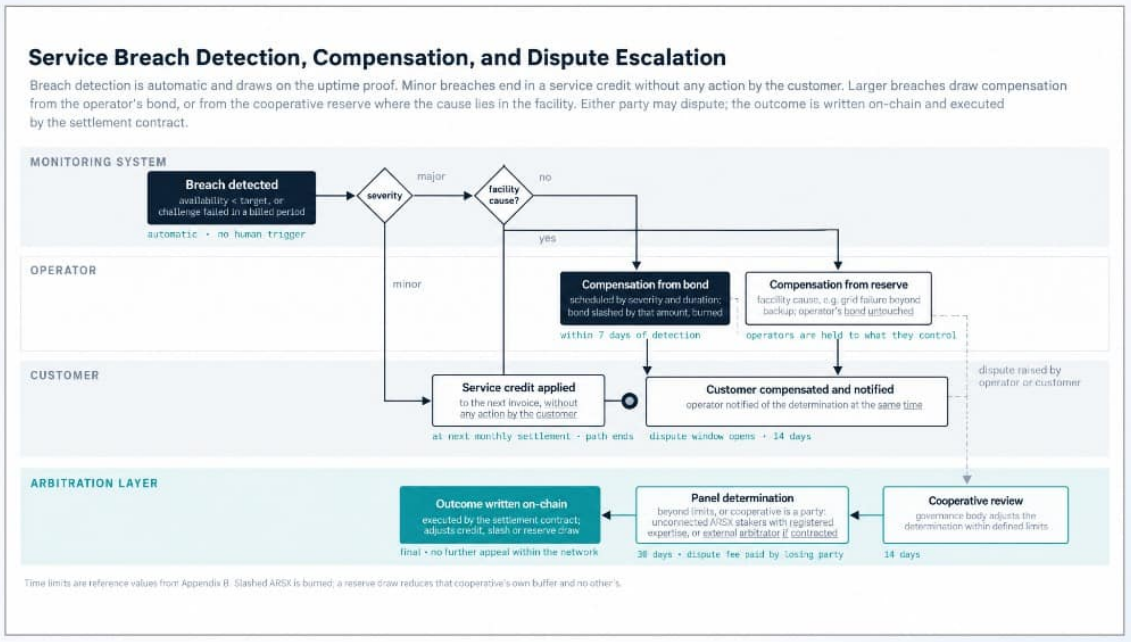
5.4 Service Responsibility and Disputes

The network's standard service-level agreement specifies an uptime target per commitment type, a maximum response time for incidents, and a performance floor expressed in terms of the hardware's declared capability. Consortia may offer stronger terms to particular customers. They may not offer weaker ones.

Breach detection is automatic and draws on the uptime proof. When a cluster's availability over a billing period falls below the target, or when a challenge fails during a period the customer was billed for, the settlement contract applies a service credit to the customer's next invoice without any action by the customer. Minor breaches are settled this way and end there.

Larger breaches move to a second tier. A sustained outage, a confirmed correctness failure, or a breach that caused the customer measurable loss beyond the credit formula triggers compensation from the operator's bond. The amount is set by a schedule that scales with the severity and duration of the breach, and the bond is slashed by that amount. Where the breach originated in the facility rather than in the operator's conduct, for instance a grid failure that exhausted the facility's backup power, the customer is still compensated, but from the Consortium's reserve rather than the operator's bond. Operators are held responsible for what they control.

Either party may dispute an automatic determination. A customer may argue that a breach was more serious than the formula recognised; an operator may argue that a challenge failure was caused by a network fault outside the cluster. Disputes go first to the Consortium's governance body, which can adjust a determination within defined limits. Beyond those limits, or where the Consortium is itself a party, the dispute goes to an arbitration panel. The panel is drawn from ARSX stakers who have registered relevant expertise and are unconnected to the parties, or, where the customer contract so provides, to an external arbitrator. The panel's decision is recorded on-chain and executed by the settlement contract.



An operator may leave the network by giving notice of a length set in its certification agreement. During the notice period it must complete every reservation already committed on its hardware and assist in migrating customers who wish to remain in the network to other clusters in the same residency zone. On completion, any unslashed bond is returned. An operator removed for misreporting receives nothing.

6. The Asset Layer: Compute Shares

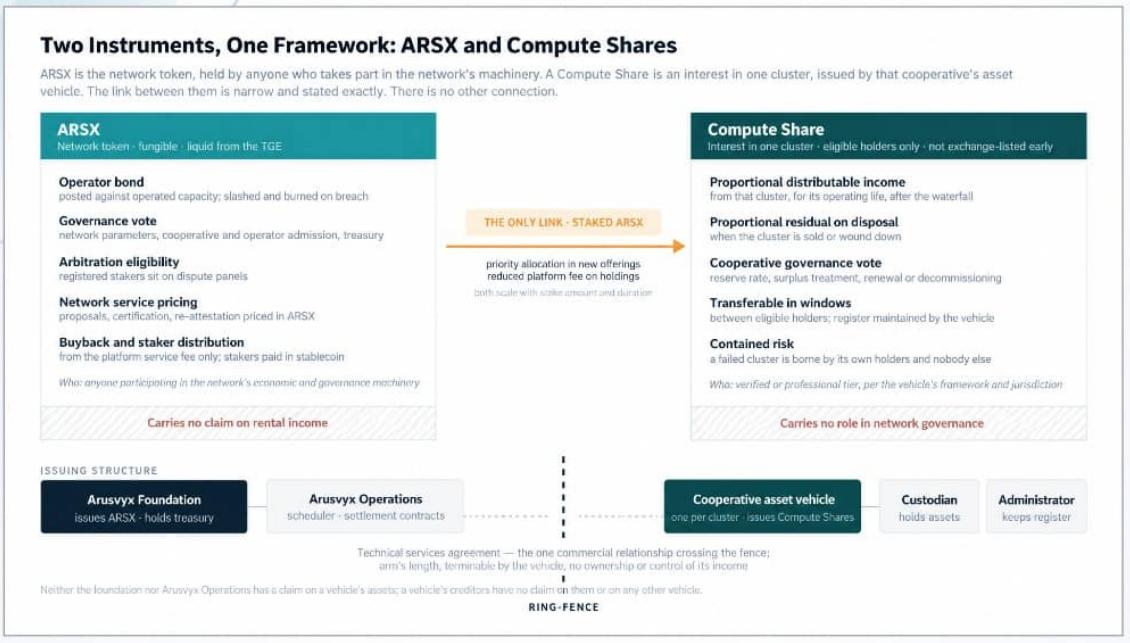
6.1 Two Instruments, One Framework

Arusvyx uses two instruments because the things they represent are different in kind and are best kept apart.

ARSX is the network's token. It is held by anyone who wishes to participate in the network's economic and governance machinery. Certified operators post it as a bond. Stakers use it to vote on network parameters and Consortium admissions and to serve on arbitration panels. It prices network services. A share of the platform service fee is used to buy it back on-chain and burn it, and another share is distributed to stakers. ARSX is fungible across the network, and it is intended to be liquid: it will trade on exchanges from the token generation event onward.

A Compute Share is an interest in a specific cluster. It is issued by the asset vehicle of the Consortium that owns that cluster, and it entitles its holder to a proportional part of the cluster's distributable income for the cluster's operating life, to a proportional part of the residual when the cluster is sold or wound down, and to a vote in the Consortium's governance in proportion to holding. Compute Shares of one Consortium carry no rights in any other. They are transferable only among eligible holders and are not intended for exchange listing in the network's early phases.

Holding ARSX does not entitle anyone to rental income. Holding Compute Shares does not confer any role in network governance beyond the Consortium in question. The link between the two is narrow and stated exactly: a participant who stakes ARSX receives priority allocation when a new Consortium offers Compute Shares, and a reduction in the platform service fee applied to income from Compute Shares they hold. Both benefits scale with the amount and duration of the stake. There is no other connection.



Three reasons drove the separation. The first is legal. An instrument that carries a claim on the income of identifiable hardware is, in most jurisdictions, an asset or investment instrument, and it is offered under rules that govern such instruments. A network token that coordinates bonds, votes, and service pricing is governed differently. Combining them in one instrument would subject the whole network to the stricter regime and would restrict who could hold the token at all.

The second is containment. A cluster can fail. Its hardware can be damaged, its customers can leave, its operator can be slashed out of the network. When that happens, the holders of that cluster's Compute Shares bear the loss. Holders of ARSX, and holders of other Consortia' Compute Shares, do not. A single instrument would spread every cluster's risk across every participant.

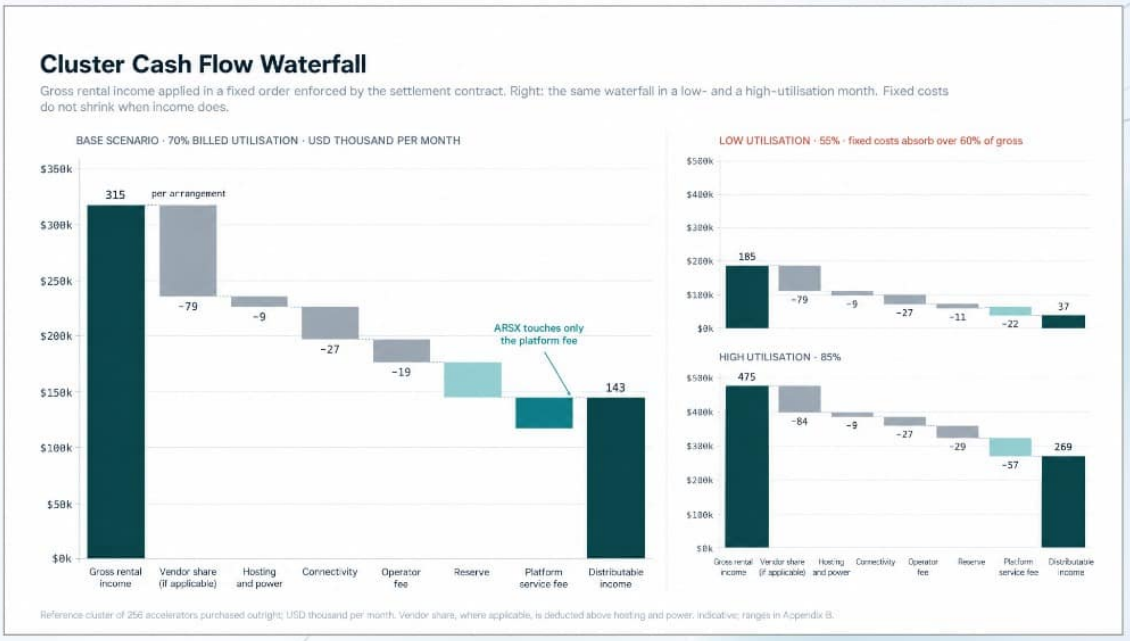
The third is governance. Network-wide votes decide parameters, admissions, and treasury use. They must not be able to reach into a Consortium and redirect its income, and the cleanest way to guarantee that is to make sure the income never passes through anything the network-wide vote controls.

6.2 Cluster Cash Flow Waterfall

Each month, after settlement, a cluster's gross rental income is applied in a fixed order. The order is written into the Consortium's asset vehicle documents and enforced by the settlement contract, and it cannot be changed by a network vote.

1. **Hosting and power.** Colocation fees and metered electricity under the facility lease. Paid first because the facility can disconnect a cluster that does not pay.
2. **Connectivity.** Upstream bandwidth and any dedicated links required by customers.
3. **Certified operator fee.** The operator's compensation for running the cluster to the service-level agreement. Set in the operator agreement, typically as a fixed monthly amount plus a performance component.
4. **Reserve contribution.** A percentage of gross income transferred to the Consortium's reserve until the reserve reaches its target, after which the rate falls to a maintenance level. The target and rates are Consortium parameters within network minimums.
5. **Platform service fee.** The network's revenue, in the indicative range of ten to fifteen percent of gross income, reduced for holders with qualifying ARSX stakes.
6. **Distributable income.** Whatever remains, transferred to Compute Share holders pro rata.

Where the cluster's hardware is held under a vendor revenue-share arrangement rather than purchased outright, the vendor's share is deducted at the top of the waterfall, before hosting and power, in accordance with that arrangement's terms. Compute Shares in such a Consortium represent an interest in the Consortium's residual after the vendor's share, and the offering documents say so.



A month in which gross income does not cover the first three steps is a shortfall. The reserve is drawn to cover it, and no distribution is made. If the reserve is exhausted, the Consortium's members decide among a capital call, a partial sale of hardware, a change of operator, or a wind-down. The network does not fund shortfalls in any Consortium from the treasury or from other Consortia' income. This is the point at which the containment described in the previous section becomes concrete, and it is the reason a prospective Compute Share holder should read a Consortium's reserve policy as carefully as its projected income.

6.3 Eligibility, Transfer and Structuring

Each Consortium's Compute Shares are issued by a dedicated asset vehicle. The vehicle holds title to the cluster's hardware, or the revenue-share contract where hardware is vendor-financed; it is the counterparty to the hosting lease and the operator agreement; and it holds the accounts into which customer payments settle. It has no other business. Its assets are not available to the network foundation, to other Consortia, or to the creditors of either.

Vehicles are domiciled in jurisdictions with an established framework for tokenised asset interests and for the segregation of such vehicles from their sponsors. The choice of jurisdiction for each vehicle is made at Consortium formation and disclosed in the offering documents. Where a corridor's own regulatory framework for tokenised real-world assets matures to the point of supporting such vehicles directly, Consortia in that corridor may be structured domestically.

Who may hold a Consortium's Compute Shares depends on the vehicle's framework and on the holder's own jurisdiction. Some offerings will be confined to professional or accredited investors. Others, where a regulated regime for retail participation in tokenised assets exists, will be open more widely, in some cases with per-person caps. The network's identity registry records each participant's verified eligibility tier, and the settlement contract will not transfer Compute Shares to an address whose tier does not qualify for the Consortium in question.

Transfers are permitted between eligible holders. In the network's early phases they occur in periodic windows, set by each Consortium, rather than continuously, so that the vehicle's register can be maintained and the eligibility of incoming holders confirmed. Compute Shares are not offered for listing on token exchanges during these phases. The intention is that a participant who acquires Compute Shares is acquiring an interest in a cluster's operating life, and that the market for those interests should develop in step with the record of the clusters themselves.

Offering mechanics. Compute Shares are priced at the Consortium's all-in cost of bringing the cluster into service: hardware, installation, connectivity, structuring, and the initial reserve, divided into shares of equal value. Each offering sets a minimum subscription and, where the vehicle's framework caps retail participation, a maximum per holder. Where the hardware is vendor-financed, the offering is sized to the Consortium's residual interest and the vendor's share is shown in the offering documents.

An offering runs for a fixed window. During the first part of the window, only participants with staked ARSX may subscribe, and their subscriptions are filled in order of stake weight until that part closes or the offering is fully taken. The remainder of the window is open to all eligible participants. If total subscriptions exceed the target, allocations within each part are scaled pro rata. Subscriptions are held in escrow until the minimum is reached; if the minimum is not reached by the close of the window, the escrow returns every subscription in full and the Consortium does not proceed.

In-kind contribution. A certified operator that brings its own hardware into a Consortium does so in one of two ways, chosen at Consortium formation and recorded in the vehicle's documents. Under the first, the hardware is transferred to the vehicle at an appraised value and the operator receives Compute Shares of that value, so that its income from the hardware arrives through the waterfall like any other holder's. Under the second, the hardware remains the operator's property and is leased to the vehicle under a term lease, with the lease payment sitting in the operator fee step of the waterfall and no Compute Shares issued against it. The first route suits operators who want a long-term interest in the cluster; the second suits those who expect to redeploy the hardware. Either way, the hardware is certified, the operator posts its bond, and the service-level agreement applies.

Identity registry. The registry is operated by Arusvyx Operations under the data-protection law of the foundation's domicile, using regulated verification providers, and it assigns each participant one of three eligibility tiers. The base tier permits holding ARSX, staking, and receipt of incentive and Corridor Participation Pool allocations. The verified tier permits holding Compute Shares in offerings open to retail participants where the vehicle's framework allows such offerings. The professional tier permits holding Compute Shares in offerings confined to accredited or professional investors. Each tier records the participant's jurisdiction, so that an offering closed to a jurisdiction cannot be subscribed from it. Personal data stays off-chain with the registry operator; the chain carries only the tier, the jurisdiction code, and a revocable attestation that the settlement contract reads on every transfer of Compute Shares.

Insurance. Each vehicle carries property insurance on its hardware at replacement value against physical loss, and business-interruption cover for facility-caused outages that run beyond the period the reserve is designed to absorb. The vehicle is the beneficiary, and premiums are accounted for within the hosting-and-power step of the waterfall. Cover does not extend to utilisation shortfall, decline in rental rates, or obsolescence. Those remain with the Compute Share holders and are the reason the reserve exists.

One further structuring consideration shapes the design. The Consortium's income arises from the lease of a productive physical asset to customers. Returns to Compute Share holders are a share of that income after costs, are not fixed, and fall when the cluster underperforms. Holders share in gains and in losses. These features are the ones that asset-backed, non-interest financing frameworks look for, and the network has designed the waterfall, the reserve mechanism, and the absence of any guaranteed return with compatibility in mind. Formal review under such frameworks is undertaken Consortium by Consortium, and no Consortium will describe its Compute Shares as compliant with any such framework before that review has concluded.

7. ARSX Token Economics

7.1 Token Utility: The Three Rings

ARSX does three kinds of work in the network, and it is easiest to see them as concentric rings ordered by how close each sits to the holder.

The inner ring is entitlement. Staked ARSX gives its holder priority allocation when a new Consortium offers Compute Shares, a reduction in the platform service fee on Compute Shares already held, a vote on network parameters and Consortium admissions, and eligibility to sit on arbitration panels. Priority and the fee reduction both scale with the size and the committed duration of the stake. A holder who does not stake has none of these; ARSX in a wallet is a liquid asset and nothing more.

The middle ring is fuel and collateral. Certified operators post ARSX as a bond against the capacity they run. Network services, from Consortium proposal submission through operator certification to re-attestation of hardware, are priced in ARSX. Bonds are slashed and burned on breach; service fees are routed to the treasury.

The outer ring is value return. The platform service fee, collected in stablecoin from every settled invoice, funds on-chain buybacks that are burned, and stablecoin distributions to stakers. The incentive pool releases ARSX to participants who bring revenue-producing capacity into the network. Neither flow depends on anything but the network's invoiced income and the capacity behind it.



Staking mechanics. ARSX is staked on the primary issuance chain into a position with a committed duration chosen from a set of tiers listed in Appendix B. A position's weight, which governs allocation priority, the platform fee reduction, the staker distribution, and voting, is its amount multiplied by a duration factor that rises with the tier. At the end of its term a position may be closed, renewed, or moved to a different tier.

Early exit is permitted. A position closed before the end of its term enters a cooldown before the ARSX is released, and the staker forfeits the distribution accrued but not yet paid for the current period. Forfeited amounts remain in the staker distribution pool for other stakers. There is no penalty beyond that, and no lock the staker did not choose.

Staking issues no new ARSX. Distributions to stakers come from the stablecoin share of protocol revenue and from nothing else, and a period in which the network earns little pays little.

Stakers are not slashed for how they vote. Stakers who register for arbitration service, by declaring relevant expertise and accepting the panel rules, can be. A panellist found to have colluded with a party, to have concealed a connection to one, or to have disclosed confidential dispute material forfeits a fraction of the position under the schedule in Appendix B and is removed from the panel register. Panel service is paid from a dispute fee charged to the party against whom the determination is made, divided among the panellists who heard the case.

7.2 Why a Token

The question a sceptical reader should ask of any token is whether the network could run without it. For Arusvyx the answer is that most of the network's components could, and the coordination between them could not.

A Consortium could be formed, a cluster deployed, and Compute Shares issued using conventional legal vehicles and bank accounts. What such a structure cannot do is hold thousands of certified operators, across a dozen corridors, to a common service standard with a common penalty, without a central party adjudicating every breach. The operator bond does that. It is posted in an asset every operator can acquire, slashed by a contract every participant can inspect, and burned so that no party benefits from an operator's failure.

Nor can a conventional structure give a participant in one corridor a say over admissions and parameters in another, or a share of network-level revenue that does not depend on which cluster they happen to hold. Governance and staker distribution do that, and both need a unit of account that is fungible across the whole network while Compute Shares, by design, are not.

The token is the layer at which the network is one network rather than a series of unrelated Consortia sharing a software vendor.

7.3 Supply Policy

The total supply of ARSX is fixed at one billion units, minted once at the token generation event. The token contract has no function that can increase this figure, and no governance vote can add one. Supply can only fall, through the burning of bought-back tokens and slashed bonds.

Every distribution of ARSX described in this section, including the incentive pool, the Corridor Participation Pool, and staker rewards, is drawn from allocations set aside at genesis. There is no ongoing emission. When the incentive pool is exhausted, no further ARSX enters circulation from that source, and the network's incentives from that point are the stablecoin distributions and the fee reductions that platform revenue provides.

Fixing supply was a choice against flexibility. A network with a discretionary emission schedule can respond to slow growth by paying more. Arusvyx cannot. The trade was accepted because a fixed supply lets every participant compute the maximum dilution they will ever face, and because the network's growth is meant to be paid for by customers rather than by holders.

7.4 Allocation and Vesting

The genesis supply is divided among seven allocations.

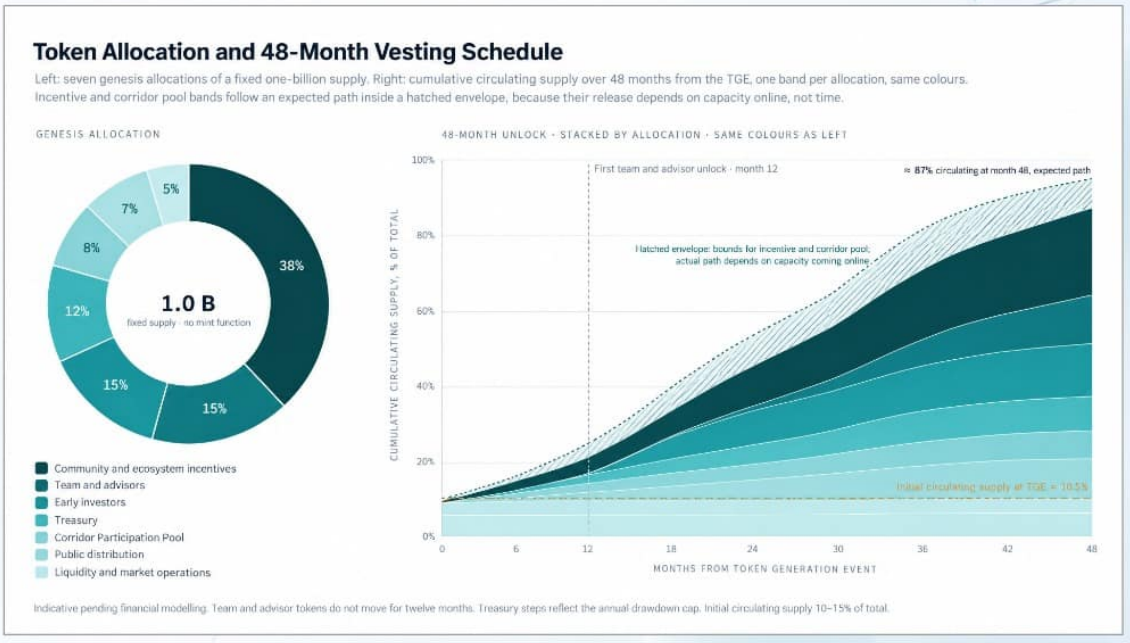
Allocation	Share	Unlock basis
Community and ecosystem incentives	38%	Milestone-driven, tied to revenue-producing capacity (Section 7.5)
Team and advisors	15%	Twelve-month lock from TGE, then linear over thirty-six months
Early investors	15%	Six- to twelve-month lock by round, then linear over twenty-four months
Treasury	12%	Governance-gated, subject to an annual drawdown cap
Corridor Participation Pool	8%	Earned through verified participation; no time-based unlock
Public distribution	7%	Partial release at TGE, remainder linear over six to twelve months
Liquidity and market operations	5%	Released at TGE; liquidity positions locked

The **Corridor Participation Pool** is set aside for participants in the regions where clusters are built: the institutions that anchor the first Consortia, the operators who certify the first hardware, the developer communities that bring the first customers. It is not sold. ARSX from this pool is earned against recorded participation, and an allocation from it vests only after the participant's contribution is on-chain. The pool exists because a network built on regional ownership should not find, a year after launch, that its regional participants hold none of the network token.

Allocations from the pool are earned against four kinds of contribution during a corridor's assessment, anchor, and first-cluster stages: co-sponsoring a Consortium that reaches certification; certifying the corridor's first clusters as an operator; bringing a customer whose reservation settles; and building or localising tooling that the developer programme accepts. The allocation attached to each contribution is set per corridor before the stage opens and published in the corridor brief, so that a prospective sponsor or operator knows what a given contribution earns before making it. No single corridor may draw more than a fixed share of the pool, with the reference value in Appendix B, so that the first corridor cannot exhaust what later corridors will need. An allocation vests when the contribution is recorded on-chain, whether as a certification, a settled invoice, or an accepted grant, and it is not transferable before then.

Team and investor allocations carry the longest locks. No team or advisor tokens move for twelve months after the token generation event, and the full allocation takes four years to vest. Investor locks vary by round; the earliest investors, who took the most risk, also accepted the longest schedule.

The initial circulating supply at the token generation event, comprising the liquidity allocation, the released fraction of the public distribution, and a small first tranche of incentives and Corridor Participation Pool allocations against capacity already contracted, is indicatively between ten and fifteen percent of total supply.



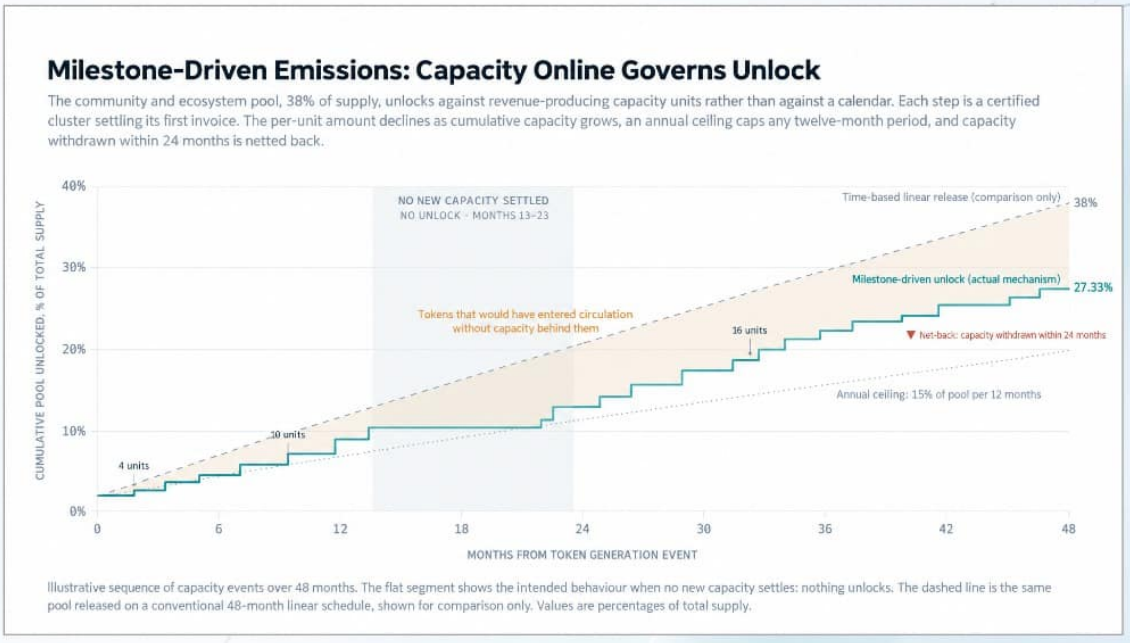
7.5 Milestone-Driven Emissions

The community and ecosystem incentive pool is the largest allocation and the one most exposed to misuse. Most token networks release such pools on a timer, and the release then has to be spent, whether or not there is anything worth spending it on. Arusvyx releases the pool against a physical quantity: certified capacity that has settled a customer invoice.

The unit is a revenue-producing capacity unit, defined in Appendix B in terms of accelerator count and class, normalised across hardware generations. When a cluster passes certification and its first invoice settles, the capacity it represents is counted, and a corresponding tranche of the pool unlocks. The tranche is distributed among the Consortium sponsor, the certified operator, the participants who funded the cluster, and, where applicable, the customer-side partner who brought the anchor tenant, in proportions set by governance.

The amount unlocked per unit declines as cumulative capacity grows. Early capacity is the hardest to bring in and is rewarded most; later capacity, arriving into a network with a track record, is rewarded less. The decline follows a fixed schedule so that any participant can calculate what a given cluster will unlock before deciding to build it.

Two bounds sit on top of the capacity trigger. An annual ceiling caps the amount that can unlock in any twelve-month period regardless of how much capacity arrives, so that an unexpectedly fast build-out does not flood the market. And capacity that later leaves the network, whether through operator withdrawal or cluster decommissioning within a defined period after certification, is netted against future unlocks, so that a cluster brought online only to collect an incentive tranche and then withdrawn gains nothing.



The pool can, in principle, remain largely locked for years if the network fails to bring capacity online. That is the intended behaviour.

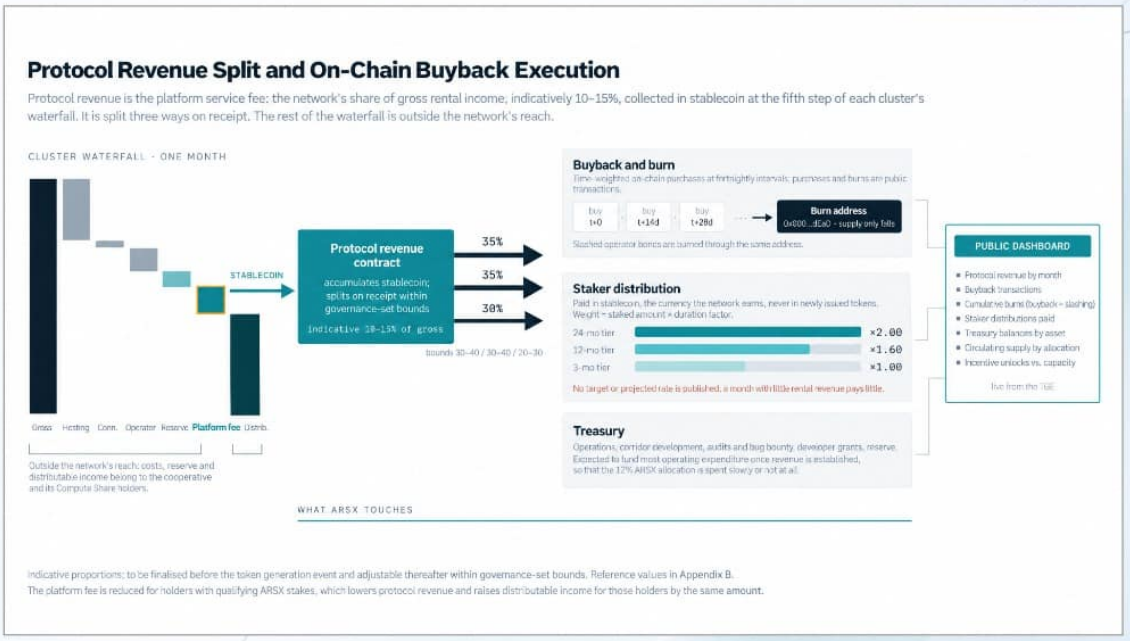
7.6 Protocol Revenue and Value Return

Protocol revenue is the platform service fee: the network's share of gross rental income, indicatively ten to fifteen percent, deducted at the fifth step of each cluster's waterfall. It is collected in stablecoin. It does not include the rental income distributed to Compute Share holders, the operator's fee, or any Consortium's reserve, and the network has no access to those.

The fee is split three ways on receipt. Indicative proportions, to be finalised before the token generation event and adjustable thereafter within governance-set bounds:

- **Buyback and burn**, thirty to forty percent. The revenue contract accumulates stablecoin and executes purchases of ARSX on-chain at scheduled intervals through time-weighted orders, then burns what it buys. Purchase transactions and burn transactions are public.
- **Staker distribution**, thirty to forty percent. Paid in stablecoin to ARSX stakers in proportion to stake and committed duration. Stakers are paid in the currency the network earns, not in newly issued tokens.
- **Treasury**, twenty to thirty percent. Funds network operations, corridor development, audits, and grants, under the governance rules in Section 9.

The network publishes no target or projected rate for staker distributions. They rise and fall with invoiced income, and in a month with little rental revenue they are small. Describing them as a yield would misstate what they are, and the network's documentation will not do so.



7.7 Distribution and the Token Generation Event

ARSX will be generated and first distributed at a token generation event targeted for the fourth quarter of 2026. The event's date, the primary issuance chain, and the final combination of distribution channels will be announced ahead of it.

Prior to the event, ARSX is distributed to early investors in seed and strategic rounds under simple agreements for future tokens, with the lock and vesting terms described above. Participation in these rounds is confined to investors eligible under the rules of the jurisdictions in which they are offered.

Public distribution is expected to use more than one channel. The candidate set includes offerings hosted on established launch platforms, exchange-hosted distributions, community allocation events in which participation is earned rather than purchased, and direct allocations to participants who have contributed capacity or customers before the event. The mix will be chosen to reach participants across the corridors in which the network operates, and it may differ between corridors to reflect local rules on public token distribution. Every channel applies identity verification and jurisdictional restrictions. The network reserves the ability to add, remove, or resize channels up to the event and to conduct further public distributions from the same allocation after it.

The Corridor Participation Pool is distributed on its own schedule, against participation records, and is not part of the public distribution.

No ARSX is offered through this document.

7.8 Liquidity and Market Integrity

The liquidity allocation funds initial trading pairs against stablecoin on decentralised venues on the primary issuance chain and, where agreed, listings on centralised exchanges. Liquidity positions on decentralised venues are locked for a minimum of twenty-four months, and the lock contracts are published.

Where the network engages market makers, the terms of any ARSX lent to them, including quantity, duration, and any options granted, are disclosed before the tokens move. Undisclosed market-making arrangements have been the origin of a number of token market failures in recent years, and the network treats disclosure as a condition of any such engagement rather than a courtesy.

From the token generation event, a public dashboard reports circulating supply by allocation, cumulative burns from buybacks and slashing, treasury balances, protocol revenue by month, incentive pool unlocks against capacity, and outstanding locked positions.

8. Multi-Chain Issuance and Interoperability

8.1 Chain-Agnostic by Design

The network's participants are spread across corridors whose developers, institutions, and exchanges have settled on different chains. Tokenised asset frameworks, identity tooling, and regulated custody are more mature on some chains than on others, and the ranking changes. Liquidity concentrates in different places over time. A token bound to a single chain at launch would either force every corridor onto that chain's tooling or fragment into wrapped versions controlled by third parties.

ARSX is therefore designed to exist on more than one chain from the outset, with one canonical definition of supply and a controlled method of moving between chains.

The primary issuance chain, on which the genesis mint occurs and on which the bonding, governance, and settlement contracts are deployed, will be selected from a candidate set before the token generation event. The selection weighs the maturity of the chain's tokenised-asset and compliance tooling, the cost and finality of settlement at the transaction volumes the network expects, the security record of the available cross-chain paths, and the presence of the exchanges and custodians the network's early participants use. The choice will be published with the reasoning.

Being chain-agnostic has a cost, and it is stated plainly: moving a token across chains requires a bridge or an equivalent mechanism, and bridges have been the single largest source of loss in the industry's history. The design in the following section exists to bound that exposure.

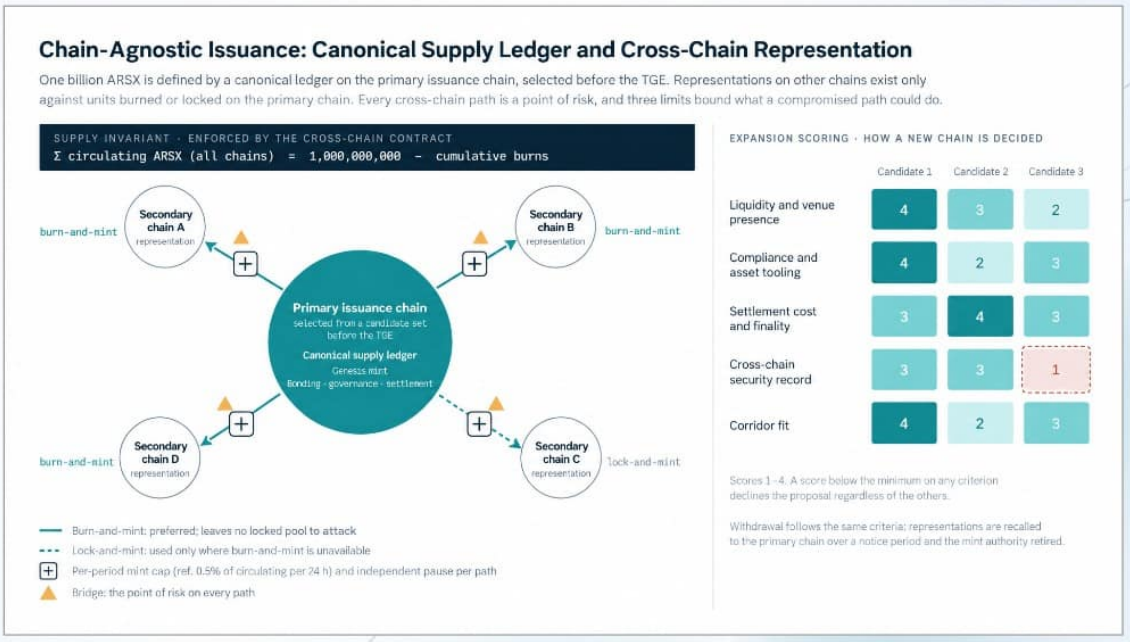
8.2 Canonical Supply and Cross-Chain Representation

The one-billion supply is defined by a canonical supply ledger maintained on the primary issuance chain. Every unit of ARSX on any chain is either a native unit on the primary chain or a representation on a secondary chain that corresponds to a unit removed from circulation on the primary chain.

Transfers between chains use a burn-and-mint model wherever the destination chain's tooling supports it: ARSX leaving the primary chain is burned there and an equal amount is minted on the destination chain by a mint authority controlled by the network's cross-chain contract; the reverse occurs on return. Where burn-and-mint is not available on a destination chain, a lock-and-mint model is used, with the locked ARSX held in a contract on the primary chain and the representation minted against it. The network prefers burn-and-mint because it leaves no pool of locked tokens for an attacker to target.

Whichever model applies, one invariant holds: the sum of circulating ARSX across all chains equals the canonical supply less cumulative burns. The cross-chain contract enforces this by refusing to mint on any chain without a corresponding burn or lock recorded on the primary chain, and the public dashboard reports per-chain circulation against the invariant.

Three limits bound the damage a compromised cross-chain path could do. Each secondary chain has a mint cap per period, so that a forged burn message cannot mint more than a bounded amount before human review. Each path can be paused by the network's security council independently of the others. And the functions that matter most, operator bonds, governance voting, and settlement, run on the primary chain in the network's early phases, so that a failure on a secondary chain affects liquidity there without touching the network's operation. Governance participation from secondary chains, when introduced, will use attested snapshots of staked positions rather than live cross-chain messaging.



8.3 Expansion Criteria

Expansion to a new chain after the token generation event is a governance decision, executed by the foundation in the first governance phase and by token vote thereafter. A proposal to expand is scored on five criteria, each with a minimum threshold below which the proposal is declined regardless of its other merits.

Liquidity and venue presence. Whether the chain hosts exchanges and market makers that the network's participants in at least one corridor already use, and whether liquidity there is likely to be additive rather than a transfer from the primary chain.

Compliance and asset tooling. Whether the chain supports the identity registry, transfer restrictions, and custody integrations the network's Compute Share vehicles require, so that expansion serves the asset layer and not only the token.

Settlement cost and finality. Whether the chain's transaction costs and finality guarantees suit the network's settlement and distribution cadence.

Cross-chain security. Whether a burn-and-mint path with an acceptable security record exists to the chain, and what the mint cap and pause arrangements would be.

Corridor fit. Whether operators, customers, or institutional participants in a corridor the network is entering or expanding within have a concrete reason to hold or use ARSX on that chain.

The same criteria govern withdrawal. A chain on which liquidity has thinned, whose cross-chain path has suffered an incident, or whose tooling has fallen behind can be wound down: representations are recalled to the primary chain over a notice period, the mint authority is retired, and the chain is removed from the dashboard. Adding chains is expected to be more common than removing them, but the network does not treat any secondary deployment as permanent.

9. Governance and Treasury

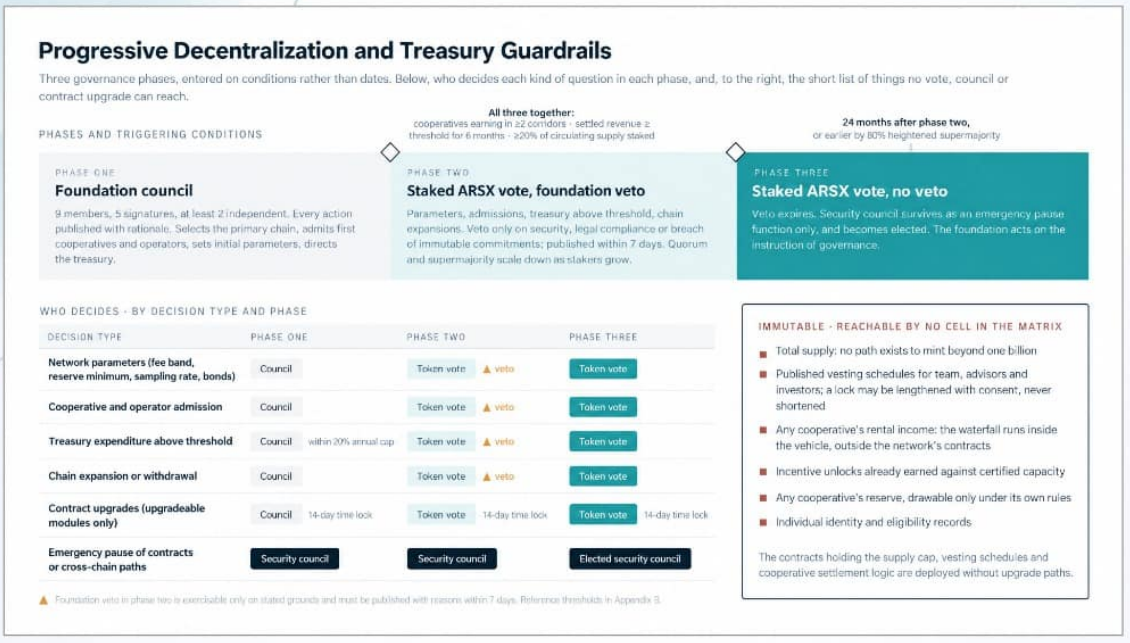
9.1 Progressive Decentralization

Governance of the network moves through three phases. The transitions are triggered by conditions rather than dates, and the conditions are published so that participants can see how far the network is from each one.

In the first phase, from the token generation event until the network has a working record, decisions are taken by the Arusvyx Foundation acting through a multi-signature council (the security council, as it is referred to elsewhere in this document). The council has nine members, of whom at least two are independent of the foundation, its founders, and its investors, and any action requires five signatures. Every action the council takes is published with its rationale. The foundation's remit in this phase is broad: it selects the primary issuance chain, admits the first Consortia and certified operators, sets initial parameters, and directs the treasury. It is broad because a network with no clusters and no revenue has nothing yet for a token vote to weigh.

The second phase begins when the network meets three conditions together: a minimum number of Consortia operating in more than one corridor, a minimum level of settled monthly revenue sustained over a defined period, and a minimum share of the circulating supply held in staked positions. From that point, staked ARSX votes on parameter changes, Consortium admissions, treasury expenditures above a threshold, and chain expansions. The foundation retains a veto, exercisable only on grounds of security, legal compliance, or a breach of the commitments listed in Section 9.2, and every exercise of the veto must be published with its grounds within a fixed period. Proposals require a quorum and a supermajority that both scale down as the number of distinct staking addresses grows.

The third phase begins at a fixed interval after the second, or earlier if a token vote with a heightened supermajority so decides. The foundation's veto expires. The security council survives as an emergency function with the power to pause contracts and cross-chain paths, and nothing else, and its membership becomes elected. The foundation continues to exist as a legal entity for as long as the network needs one to hold intellectual property and sign contracts, but it acts on the instruction of governance.



The network expects the second phase to arrive within its first years of operation and does not intend to hold it back. What it will not do is hand parameter control to a token vote before the network has revenue against which parameters can be judged.

9.2 Scope of Governance and Immutable Commitments

Governance decides what the network has left open, and a short list of things has been closed.

Within scope, and adjustable by vote in the second phase and after, are the platform service fee within a published band, the split of protocol revenue among buyback, staker distribution, and treasury within their bounds, the network minimum reserve rate for Consortia, the hardware grade reference values as generations turn over, the correctness sampling rate, bond requirements per unit of capacity, the per-unit incentive schedule for capacity not yet certified, the admission of Consortia and certified operators, the treasury's annual drawdown, chain expansions and withdrawals, and upgrades to the network's contracts through a time-locked process.

Outside scope, and unreachable by any vote, council, or contract upgrade:

- The total supply. No path exists to mint beyond one billion.
- Vesting schedules already published for team, advisors, and investors. Governance may lengthen a lock with the holder's consent; it may not shorten one.
- The rental income of any Consortium. The waterfall runs inside the Consortium's asset vehicle, and the network's contracts have no function that redirects a Consortium's distributable income.

- The incentive unlocks already earned against certified capacity. A tranche that has unlocked belongs to its recipients.
- The reserve of any Consortium, which may be drawn only under that Consortium's own rules.
- The identity and eligibility records of participants, which governance may not alter for any individual.

Contract upgrades are the route by which a determined majority could, in principle, reach the items above. To close it, the contracts that hold the supply cap, the vesting schedules, and the Consortium settlement logic are deployed without upgrade paths. Upgradeable contracts are confined to the scheduler interface, the governance module itself, the treasury, and the cross-chain module, and any upgrade to those passes through a time lock long enough for participants to exit if they object to it.

9.3 Treasury Management

The treasury holds the twelve percent allocation of ARSX, the treasury share of protocol revenue in stablecoin, and network service fees paid in ARSX. Its purposes are the network's own operating costs, corridor development, security audits and bug bounties, grants to developers building on the network, and a reserve against events the network has not anticipated.

Three rules constrain it.

Drawdown of the ARSX allocation is capped at a percentage of the remaining balance per twelve-month period. The cap applies to the foundation council in the first phase as it applies to token votes later. Stablecoin revenue is not subject to the cap and is expected to fund most operating expenditure once revenue is established, so that the ARSX allocation is spent slowly or not at all.

The treasury does not hold Compute Shares in any Consortium, so that the network's own balance sheet has no interest in favouring one cluster over another in scheduling or admissions. Where the foundation sponsors a Consortium in the first phase, it does so as sponsor and not as holder, and any Compute Shares it must take to meet a vehicle's minimum are distributed to participants at cost.

The treasury does not hold the tokens of other networks beyond operational balances, and it does not lend ARSX. Market-making arrangements draw on the liquidity allocation, under the disclosure terms already described, and never on the treasury.

A dashboard reports treasury balances by asset, expenditure by category, and the remaining annual drawdown at all times. Expenditures above a threshold are itemised with their recipients and purposes. The threshold is low enough that the itemised list covers the great majority of spending by value.

10. Ecosystem and Go-to-Market

10.1 Customer Segments and Workload Fit

The network does not try to serve every workload. Its clusters are hundreds of accelerators rather than tens of thousands, and its advantages are residency, price in the corridors where it operates, and contractual accountability. Those advantages matter to some customers and not to others.

Four segments matter most.

AI-native companies in the corridors, building products on top of open or licensed models, need inference capacity close to their users and fine-tuning capacity for adapting models to local languages and domains. They are price-sensitive, they sign reservations of moderate length, and their workloads are well suited to clusters of the network's size. Several of them are today served from distant regions at latency and cost that a domestic cluster undercuts.

Enterprise AI application developers, building for banks, insurers, telecommunications operators, and government agencies in the corridors, need capacity that meets their end clients' data-residency and isolation requirements. They pay more per unit than AI-native companies and are the segment for which residency zoning was designed. They tend to sign longer reservations and to require dedicated nodes.

Research institutions and universities need reserved capacity for experiments, often in bursts, and need it under contracts their procurement rules can accept. They are less price-sensitive than their budgets suggest, because the alternative is a multi-year wait for institutional hardware. They are also a source of certified operators and Consortium sponsors, which makes them a segment the network cultivates for reasons beyond revenue.

Sovereign and public-sector AI programmes, which several corridor governments have launched or announced, need domestically hosted capacity for model training and public-service deployment. They are the largest single contracts available in a corridor and the slowest to close. The network pursues them through anchor institutions rather than directly in its early phases.



Two workloads are outside the network's ambition. Foundation-model training at the scale of the largest laboratories requires clusters an order of magnitude beyond a Consortium's, and the network does not aggregate Consortia into a single training fabric in its current design. And consumer-facing services that require a global edge footprint are better served by providers with hundreds of points of presence. The network's clusters serve the region they sit in.

10.2 Regional Deployment Model

The network opens a corridor once and then repeats what worked. Each corridor passes through five stages, and the assets produced in each stage are written down so that the next corridor can use them.

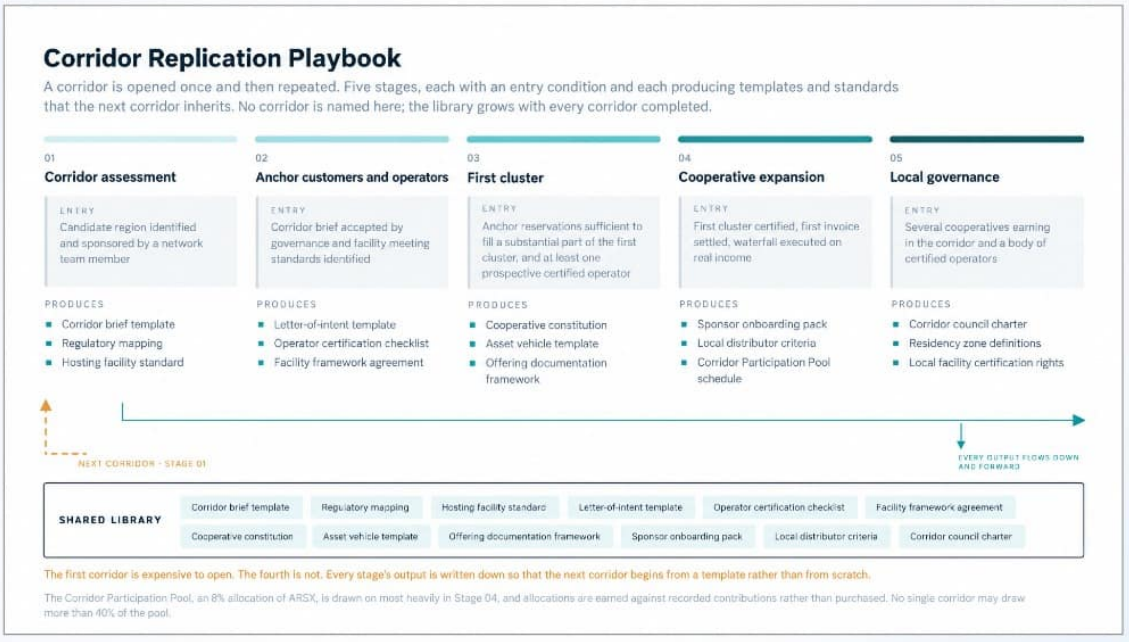
Corridor assessment. The network evaluates a candidate region against the four conditions that define a compute corridor, adds a review of the regulatory position for tokenised asset vehicles and for token distribution, and identifies at least one hosting facility meeting the network's standards. The output is a corridor brief, published to governance, with a recommendation to proceed or not.

Anchor customers and operators. Before any cluster is funded, the network secures letters of intent from customers whose demand would fill a substantial part of the first cluster, and identifies at least one prospective certified operator in the corridor. A first cluster funded without an anchor customer is a cluster whose Compute Share holders bear the full utilisation risk from day one, and the network will not sponsor one.

First cluster. A Consortium is formed, usually with the foundation as sponsor and a regional institution as co-sponsor, funded, deployed, and certified. The first cluster in a corridor is expected to be smaller than later ones, and to serve its anchor customers under close attention from the network's operations team. Its purpose is to prove that the hosting, operator, settlement, and distribution mechanisms work in the corridor's legal and physical conditions.

Consortium expansion. With a first cluster earning, further Consortia are sponsored by regional institutions, operators, or participant groups rather than by the foundation. The network's role narrows to admission, certification, and scheduling. The Corridor Participation Pool is drawn on most heavily in this stage.

Local governance. Once a corridor has several Consortia and a body of certified operators, decisions specific to the corridor, such as residency zone definitions and the certification of local facilities, are delegated to a corridor council of local stakers and Consortium representatives, subject to network-wide standards.



The first corridor is in Southeast Asia. Its assessment is complete, an anchor institution and prospective operators have been identified, and the first cluster is the subject of the roadmap's opening phase. The second and third corridors will be selected by the process above, and the candidates under review span the regions described in Section 2.2.

10.3 Partners and Ecosystem Programs

The network runs three programmes through which outside parties become part of it, and it engages a small number of partner types outside those programmes.

The **certified operator programme** is the route by which regional data center operators, university compute centres, and enterprises with enterprise-grade hardware join the network. It provides the certification process, a reference runbook for meeting the service-level agreement, and access to the Corridor Participation Pool for operators who certify capacity in a corridor's early stages. Operators certified in one corridor may operate in others without re-certification of their organisation, though each cluster is certified separately.

The **developer programme** supports teams building on the network's customer access layer: orchestration tools, model-serving frameworks with residency awareness, cost-management integrations, and the connectors that let an enterprise's existing MLOps stack reserve and consume network capacity. Grants are drawn from the treasury. Developers whose tools bring measurable customer demand are eligible for allocations from the incentive pool alongside the capacity contributors, as described in the unlock rules.

The **Consortium sponsor programme** works with institutions that wish to form a Consortium: it provides the constitution template, the asset vehicle structuring support, the offering documentation framework, and introductions to certified operators and hosting facilities in the corridor. Sponsors are expected to bring an anchor customer or a participant base, or both.

Outside the programmes, the network engages partners of four kinds. Hosting facilities, under a framework agreement that sets the standards a facility must meet and the terms on which Consortia may lease space. Hardware vendors and their financing arms, including the revenue-share arrangements already discussed. Custodians and asset-vehicle administrators in the jurisdictions where Consortia are domiciled. And exchanges and regulated distribution partners in the corridors, through whom participants acquire ARSX and, where the local framework permits, Compute Shares.

The network does not, in its early phases, seek partnerships with hyperscale cloud providers. Their interest in emerging corridors is as operators of their own capacity, and their presence in a corridor is a signal that the corridor is viable rather than an invitation to compete.

11. Roadmap

The roadmap is organised by what has to be true before the next thing can begin, and it is presented without dates for that reason. The one exception is the token generation event, which the network is targeting for the fourth quarter of 2026 and which sits at the boundary between the first two phases.

Phase one: foundation and first cluster. The network's contracts are written, audited, and deployed to a test environment on the candidate primary chains. The first corridor's assessment is complete, its hosting facility is under framework agreement, and its anchor institution has agreed to co-sponsor the first Consortium. The first Consortium's asset vehicle is domiciled, its offering documents are prepared, and Compute Shares are offered to eligible participants. Hardware is procured, installed, and certified. The cluster begins serving its anchor customers under reservation. The phase ends when the first invoice settles and the waterfall executes on real income.

Most of this phase is underway at the date of this document. The corridor assessment, facility agreement, and anchor institution are in place. The contracts are in audit. The offering and procurement are the open items.

Phase two: token generation and network launch. ARSX is minted on the selected primary chain and distributed through the channels chosen for the event. Liquidity positions are established and locked. The public dashboard goes live. The first tranche of the incentive pool unlocks against the capacity certified in phase one. The certified operator programme opens to applicants in the first corridor, and a second Consortium in that corridor is sponsored by a regional institution rather than the foundation. The phase ends when the network has more than one Consortium earning and at least one certified operator not affiliated with the foundation.

Phase three: multi-corridor expansion. The corridor playbook is applied to a second and then a third corridor. Consortia in the first corridor multiply without foundation sponsorship. Secondary-chain deployments of ARSX are proposed and scored. The developer programme funds the first residency-aware tooling. The phase ends when the conditions for the second governance phase are met together: Consortia operating in more than one corridor, settled revenue sustained above a threshold, and a sufficient share of supply staked.

Phase four: governance maturity. Parameter control, admissions, and treasury expenditure pass to staked ARSX. Corridor councils are formed where corridors have the Consortium density to support them. The foundation's veto runs to its expiry. The network's engineering focus shifts from launching mechanisms to hardening them: decentralising the scheduler, raising the correctness sampling rate as re-execution costs fall, and extending governance participation to secondary chains through attested snapshots.

Phased Roadmap with Gating Criteria

Organised by what has to be true before the next thing can begin, and presented without dates for that reason. One exception: the token generation event, targeted for the fourth quarter of 2026, at the boundary between the first two phases. A phase does not begin until every criterion of the preceding phase is met.

PHASE PROGRESS AND EXIT CRITERIA · FILLED = COMPLETE · HALF = IN PROGRESS · EMPTY = NOT STARTED

PHASE 01

Foundation and first cluster

Contracts audited and deployed to test. First corridor assessed, facility under framework agreement, anchor institution agreed.

IN PROGRESS

- ● Contract audits complete, findings resolved
- ● Asset vehicle domiciled, offering documents issued
- ● At least two independent suppliers under quotation
- ○ First invoice settled, waterfall executed
- ◐ Primary issuance chain selected and published
- ◐ Anchor reservations covering most of first-cluster capacity
- ○ First cluster certified
- ○ Unit economics documented without reference to token value

TOKEN GENERATION EVENT · TARGET Q4 2026 held only after the first cluster is certified, and not before the first invoice settles; if procurement slips, the event slips with it

PHASE 02

Token generation and network launch

Minted and distributed through the channels chosen for the event. Liquidity established and locked. Dashboard live.

NOT STARTED

- ○ Token generation event completed
- ○ Dashboard live with all reported series
- ○ Second cooperative in the first corridor earning
- ○ Market-making terms, if any, disclosed
- ○ Liquidity positions established and locks published
- ○ First incentive tranche unlocked against certified capacity
- ○ At least one certified operator unaffiliated with the foundation

◇ PHASE 03 BEGINS ONLY WHEN ALL OF THE ABOVE ARE MET

PHASE 03

Multi-corridor expansion

Playbook applied to a second and third corridor. Cooperatives multiply without foundation sponsorship. Secondary-chain deployments proposed and scored.

NOT STARTED

- ○ Cooperatives earning in at least two corridors
- ○ Staked share of circulating supply above threshold
- ○ Developer programme grants awarded
- ○ Settled monthly revenue above threshold, sustained
- ○ At least one secondary-chain deployment scored and decided
- ○ Corridor library in use for a third-corridor assessment

◇ PHASE 04 BEGINS ONLY WHEN ALL OF THE ABOVE ARE MET

PHASE 04

Governance maturity

Parameter control, admissions and treasury pass to staked ARSX. Corridor councils formed where density allows. Engineering focus shifts from launching mechanisms to hardening them.

NOT STARTED

- ○ Second governance phase triggered
- ○ At least one corridor council constituted
- ○ Scheduler decentralisation milestone reached
- ○ First parameter change enacted by token vote
- ○ Foundation veto expired or waived by heightened supermajority

THREE NETWORK-WIDE METRICS TRACKED ACROSS THE PHASES · NO VALUES SHOWN, SHAPE ONLY



Reference thresholds in Appendix B: cooperatives earning in at least two corridors · settled revenue above threshold for six consecutive months · at least 20% of circulating supply staked. The phase-one criteria are the ones that matter to a reader, because they determine whether the network's economics are real before its token trades.

12. Team and Advisors

12.1 Core Team

The founding team was assembled around the three things the network has to get right: physical infrastructure that meets enterprise standards, an asset structure that can hold it on behalf of many small holders, and a protocol that coordinates both without a central operator. Nobody on the team was hired for a token background alone.



Mattias Kvernberg

Chief Executive Officer

Mattias spent nine years at AirTrunk and then at ST Telemedia Global Data Centres, where he ran site development for hyperscale campuses in Malaysia, Thailand, and Australia. The last of those projects was a multi-tenant AI facility whose anchor tenants were regional model companies, and it was there that he saw the pattern this document describes: local land, local power, local customers, foreign ownership. He left to build a structure that would change the last of those. He holds an MSc in Industrial Economics from the Norwegian University of Science and Technology.



Priyanka Venkataraghavan

Chief Technology Officer

Priyanka led the scheduler and multi-tenancy team for GPU clusters at Oracle Cloud Infrastructure after earlier work on distributed training systems at Microsoft Research India. Her interest in verification predates the network: her doctoral work at IIT Madras concerned attestation for outsourced computation, and the three-proof design in Section 5.2 descends from it. She is the author of the network's scheduling and metering specifications.



Farid Azlan Hashim

Head of Infrastructure and Operations

Farid ran operations for Keppel Data Centres' Johor campus and, before that, was a facilities engineer at Equinix Singapore. He wrote the network's hardware grade standard and the certified operator runbook, and he leads facility selection and operator certification in the first corridor. He trained as an electrical engineer at Universiti Teknologi Malaysia.



Isabel Teixeira Marlow

Chief Financial Officer

Isabel structured infrastructure vehicles at Macquarie Asset Management for eleven years, most recently in the team responsible for tokenised fund interests, and before that worked in project finance at Standard Chartered in Singapore. The Consortium asset vehicle, the waterfall, and the reserve mechanism are hers. She holds an MBA from INSEAD and is a CFA charterholder.



Helena Strandqvist

General Counsel

Helena was a partner in the financial regulation practice at Rajah & Tann, advising digital asset issuers and fund managers across Southeast Asia, and earlier practised at Clifford Chance in London. She designed the separation between ARSX and Compute Shares and leads the network's regulatory mapping in each corridor.



Leonardo Vasconcelos Ribeiro

Head of Protocol

Leonardo built token economic models and on-chain treasury systems at Polygon Labs and, before that, worked on payment infrastructure at Nubank in São Paulo. He is responsible for the supply policy, the milestone-driven emission schedule, and the cross-chain design, and he argued for the fixed supply against a more flexible alternative within the team.



Yusuf Bensaïd Larache

Head of Corridor Development

Yusuf spent seven years in the International Finance Corporation's infrastructure group, working on power and digital projects in North Africa and Southeast Asia, and led the corridor assessment framework the network now uses. He is responsible for the second and third corridors.



Anouk ter Haar

Head of Security and Verification

Anouk led smart-contract security reviews at Trail of Bits and ran the bridge risk programme at Chainalysis. She owns the audit process, the cross-chain mint caps and pause controls, and the correctness sampling design. She studied computer science at TU Delft.

12.2 Advisors



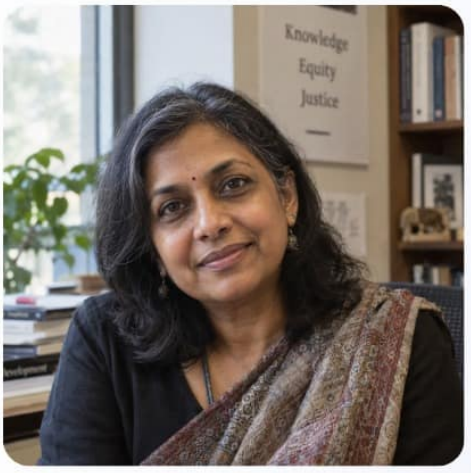
Piet Oosthuizen was chief operating officer of Teraco, the largest data center operator in Africa, through its acquisition by Digital Realty. He advises on facility standards and hosting frameworks.



Mariam Kazim Al-Hashimi heads structured products at an Abu Dhabi Global Market-licensed asset manager and previously served on the Sharia advisory function of a Gulf investment bank. She reviews the Consortium structure for compatibility with asset-backed, non-interest financing frameworks and advises on Gulf corridor entry.



Dr. Sebastian Roeder-Fink is a partner at a Zurich law firm specialising in tokenised securities under the Swiss DLT framework and a lecturer at the University of Zurich. He advises on asset vehicle domicile and transfer restriction design.



Prof. Kavitha Raghunathan holds a chair in distributed systems at Nanyang Technological University and directs a research group on verifiable computation. She advises on the verification architecture and on the research institution segment.



Tomás Echeverría Salgado was a director at the Chilean sovereign wealth fund's infrastructure programme and later led data center investment for a Latin American private equity firm. He advises on Latin American corridor assessment and on institutional co-sponsorship of Consortia.

12.3 Organizational Philosophy

The network is built by a team that lives in the corridors it serves. There is no headquarters city, and the team's compensation does not carry a premium for being in one. Engineering, operations, and corridor development staff are hired in the regions where clusters are deployed, and the foundation's operating entities are established in those regions as corridors open.

Team and advisor allocations of ARSX are locked for twelve months from the token generation event and vest over the following three years, as Section 7.4 sets out. No member of the team holds Compute Shares in any Consortium on preferential terms; where a team member participates in a Consortium, it is on the same terms and through the same eligibility process as any other participant, and the holding is disclosed.

The team expects to be judged on the criteria in Section 11.2 and publishes progress against them for that reason.

13. Risk Factors and Legal Considerations

13.1 Legal Structure

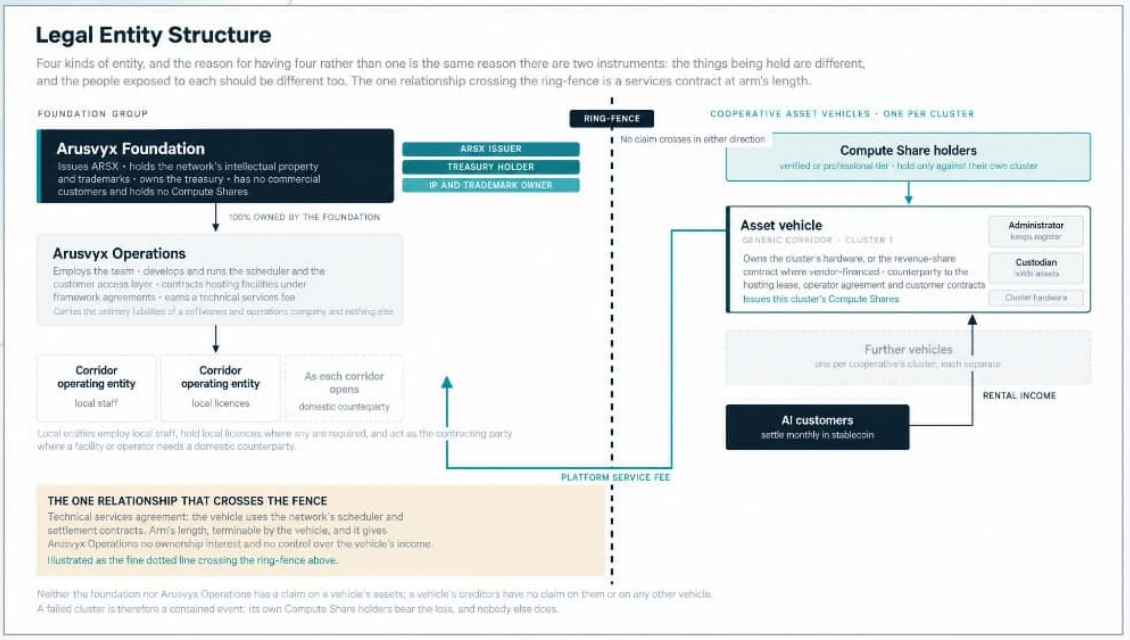
The network operates through four kinds of legal entity, and the reason for having four rather than one is the same reason the network has two instruments: the things being held are different, and the people exposed to each should be different too.

The **Arusvyx Foundation** is a non-profit foundation domiciled in a jurisdiction with an established regime for digital-asset foundations. It is the issuer of ARSX, the holder of the network's intellectual property and trademarks, the owner of the treasury, and the legal person that acts on governance decisions. It has no commercial customers and holds no Compute Shares.

Arusvyx Operations, a company wholly owned by the foundation, employs the team, develops and runs the scheduler and the customer access layer, and contracts with hosting facilities under framework agreements. It earns a service fee from the foundation and, in some corridors, a technical services fee from Consortia. It carries the ordinary liabilities of a software and operations company and nothing else.

Each **Consortium asset vehicle** is a separate entity, domiciled as described in Section 6.3, that owns one cluster, issues that cluster's Compute Shares, and is the counterparty to that cluster's hosting lease, operator agreement, and customer contracts. Vehicles are administered by regulated third-party administrators and their assets are held by regulated custodians. Neither the foundation nor Arusvyx Operations has a claim on a vehicle's assets, and a vehicle's creditors have no claim on the foundation, on Arusvyx Operations, or on any other vehicle.

Corridor operating entities are local companies established in each corridor as it opens. They employ local staff, hold local licences where any are required, and act as the local contracting party for facilities and operators whose own rules require a domestic counterparty. They are owned by Arusvyx Operations.



The one relationship that crosses the ring-fence is the technical services agreement under which a vehicle uses the network's scheduler and settlement contracts. It is a services contract at arm's length, terminable by the vehicle, and it gives Arusvyx Operations no ownership interest and no control over the vehicle's income.

13.2 Regulatory Considerations

The network operates in more than one jurisdiction from the outset, and the regulatory treatment of what it does differs among them. Three questions arise in every jurisdiction and are answered jurisdiction by jurisdiction.

The first is how ARSX is classified. The network has designed ARSX to function as a network utility and governance token, and it takes the position that ARSX confers no right to income, no interest in any entity, and no claim on any asset. Some jurisdictions nonetheless classify tokens by reference to how they are marketed or traded rather than by their functions, and a regulator in such a jurisdiction may take a different view. Where a corridor requires tokens to be registered or admitted before they may be traded by residents, the network will seek admission through licensed intermediaries rather than distribute directly, and it will not market ARSX to residents of a corridor in which it has not done so.

The second is how Compute Shares are classified and to whom they may be offered. The network takes the position that Compute Shares are interests in an asset-holding vehicle and are subject to the rules governing such interests in the vehicle's domicile and in the holder's jurisdiction. This is why offerings are confined to eligible holders, why the eligibility tier is verified on-chain before any transfer, and why Compute Shares are not listed on token exchanges in the network's early phases.

The third is the use of digital assets as payment. Several corridor jurisdictions prohibit or restrict the use of digital assets as a means of payment for goods and services within their borders. The network's customer invoices are settled in stablecoin, and where a corridor's rules do not permit a domestic customer to pay in that form, settlement for that customer runs through a licensed payment intermediary in local currency, with conversion outside the customer's obligation. ARSX is not used as a means of payment by customers in any corridor. The platform service fee, bond, and staking mechanisms are network-internal and are structured so as not to constitute payment for goods or services in the regulated sense, but the network acknowledges that a regulator could disagree and has designed each corridor's settlement path so that a change in local treatment affects that corridor alone.

Data protection law bears on the network at two points. Customer workloads may process personal data, and the residency zoning, tenant isolation, and contractual data-processing terms exist to allow customers to meet their own obligations. The network's identity registry itself holds personal data of participants, and it is operated under the data protection law of the foundation's domicile with transfer safeguards for participants elsewhere.

Tax treatment of distributions to Compute Share holders, of staker distributions, and of ARSX disposals varies by holder jurisdiction, and the network provides no tax reporting beyond what the vehicles' administrators are legally required to produce. Participants are responsible for their own obligations.

The network has obtained legal advice in its foundation domicile, in the domiciles of the first Consortium vehicles, and in the first corridor. It has not obtained advice in every jurisdiction in which a participant might reside, and it does not represent that participation is lawful in any particular jurisdiction.

13.3 Risk Factors

Participation in the network, whether by holding ARSX, holding Compute Shares, operating certified hardware, or sponsoring a Consortium, involves risk. The risks below are those the team considers most significant at the date of this document. They are not exhaustive, and risks that the team has not identified or has judged less significant may prove more so.



1. Accelerator supply and export controls. The network's clusters depend on data-center-class accelerators whose supply is concentrated in a small number of vendors and whose export is subject to controls that have changed several times in recent years and may change again. A Consortium that has raised capital may be unable to procure hardware on the expected terms, or at all, in its corridor. Mitigation: procurement is not begun until the offering has closed, hardware grade standards admit more than one vendor's parts, and a Consortium whose procurement fails returns contributions rather than proceeding with inferior hardware.

2. Utilisation shortfall. Distributable income depends on clusters being rented. Demand in a corridor may be slower to arrive than assessed, anchor customers may not renew, and rental rates may fall as capacity in the corridor grows or as a hardware generation ages. A cluster whose income does not cover its fixed costs makes no distribution and draws on its reserve, and a reserve can be exhausted. Mitigation: no Consortium is sponsored without anchor reservations, reserve policies are set with shortfall periods in mind, and cluster-level income is public so that the position is visible early. The risk is not eliminated by any of these, and it is the principal risk borne by Compute Share holders. The illustrative scenarios in Appendix B show the range of outcomes this risk produces.

3. Regulatory reclassification. A regulator in a jurisdiction relevant to the network may classify ARSX, Compute Shares, or the network's settlement mechanisms differently from the positions set out above. The consequences could include restrictions on trading or holding, requirements to register or license, or, in the extreme, an order to cease activity in that jurisdiction. Mitigation: the two-instrument separation, the eligibility registry, and the per-corridor settlement paths are designed so that an adverse classification in one place affects that place; the network engages regulators in each corridor before entering it; and the treasury holds a reserve for compliance costs.

4. Cross-chain compromise. Bridges and cross-chain messaging systems have been the largest source of loss in the industry. A compromise of a path to a secondary chain could result in unauthorised minting on that chain up to the per-period cap before the path is paused. Mitigation: burn-and-mint is preferred, mint caps and independent pause controls bound the exposure, and the functions on which the network's operation depends remain on the primary chain in the early phases.

5. Operator default. A certified operator may fail to meet its service-level agreement, become insolvent, or withdraw. Customers are compensated from the bond and, where the cause lies in the facility, from the Consortium reserve, but a prolonged failure could lose customers who do not return. Mitigation: bonds are sized to expected compensation, the trial period screens operators before full certification, and clusters can be migrated to another certified operator in the same zone.

6. Energy price and availability. Power is the largest operating cost of a cluster. Prices in a corridor may rise, and grid or facility supply may be interrupted for longer than backup systems cover. Mitigation: hosting agreements fix or cap power rates where the facility will agree to it, and corridor assessment weights the stability of the local supply as heavily as its price.

7. Milestone shortfall. The incentive pool unlocks only against revenue-producing capacity. If the network brings capacity online more slowly than expected, participants who anticipated incentive allocations will receive less, and the network's ability to reward early operators and sponsors is correspondingly reduced. This is the intended behaviour of the mechanism and is nonetheless a risk to participants who rely on it.

8. Liquidity. ARSX may trade thinly, and its price may be volatile, particularly in the period after the token generation event when circulating supply is small. Compute Shares are transferable only within windows and among eligible holders and may have no secondary market at all for an extended period. Participants should assume they may be unable to sell either instrument at a time or price of their choosing.

9. Smart-contract failure. The network's contracts, though audited, may contain defects. A defect in the settlement or distribution contracts could misdirect funds; a defect in the governance or cross-chain modules could be exploited. Mitigation: the contracts that hold the supply cap and the Consortium settlement logic are non-upgradeable and receive the most audit attention, upgradeable contracts pass through time locks, and a bug bounty is funded from the treasury.

10. Corridor policy change. A corridor's government may change its position on data residency, on digital assets, on foreign participation in infrastructure, or on the tax treatment of any of these. A change could reduce the demand that residency zoning captures or raise the cost of operating in the corridor. Mitigation: Consortia are contained within their corridor, and the network's multi-corridor design means that no single corridor's policy determines the network's viability.

11. Competition. Hyperscale providers, neoclouds, and other decentralised networks may enter the corridors the network serves with capacity at scale and pricing the network's clusters cannot match. Mitigation: the network's position rests on ownership participation and residency rather than on price alone, and it does not compete for the workloads on which scale is decisive.

12. Key personnel. The network depends on a small team whose members hold specific knowledge of the verification design, the asset structure, and the corridors. The loss of several of them within a short period would delay the roadmap. Mitigation: specifications are documented to a standard that allows replacement, team allocations vest over four years, and corridor teams are built locally rather than around individuals.

13.4 Disclaimers

This document is a description of the network's design and intentions as of its date. It will be revised as the network develops, and later versions supersede earlier ones.

Nothing in this document constitutes financial, investment, legal, tax, or other professional advice. Nothing in it constitutes an offer or solicitation in any jurisdiction, and in particular nothing in it constitutes an offer of Compute Shares, which are offered, if at all, only under the documentation of the relevant Consortium asset vehicle and only to persons eligible to receive them.

Statements concerning future events, capacity, revenue, partnerships, regulatory outcomes, and the conduct of the token generation event are forward-looking. They are based on the team's current expectations and on assumptions that may prove incorrect. Actual outcomes may differ materially, and the network undertakes no obligation to update forward-looking statements except as required by law.

Figures described as indicative are illustrations of the intended design and will be finalised through financial modelling and legal review. Third-party data cited in this document is drawn from public sources listed in Appendix C, has not been independently verified by the network, and is current only as of the dates given there.

ARSX has no guaranteed value and may lose all of it. Compute Shares carry the risk of the cluster to which they relate and may return less than was contributed, including nothing. Past performance of any cluster, corridor, or comparable asset is not indicative of future results.

The network, the foundation, Arusvyx Operations, the Consortium asset vehicles, and their respective directors, officers, employees, advisors, and agents accept no liability for any loss arising from reliance on this document to the fullest extent permitted by law.

Appendix A: Glossary

Anchor customer. A customer whose signed reservation, secured before a Consortium is funded, covers a substantial portion of the first cluster's capacity.

ARSX. The network token. Used for operator bonds, governance, arbitration eligibility, network service pricing, and the receipt of protocol revenue through buybacks and staker distributions. Carries no claim on cluster rental income.

Canonical supply ledger. The record on the primary issuance chain against which all ARSX on all chains is reconciled, such that circulating supply across chains equals genesis supply less cumulative burns.

Certified operator. An entity admitted to run one or more clusters after verification of hardware, facility, connectivity, and service capability, and after posting an operator bond.

Compute Consortium. The network's basic unit: a governance body of Compute Share holders that owns a cluster through an asset vehicle, holds its contracts, and receives its income.

Compute Share. An interest in a specific Consortium's asset vehicle, entitling the holder to a proportional share of that cluster's distributable income and residual, and to a vote in that Consortium's governance.

Corridor. A region meeting four conditions: construction cost materially below established hubs, renewable or low-carbon baseload, national policy favouring domestic AI capability and data residency, and growing regional demand.

Corridor council. A body of local stakers and Consortium representatives to which corridor-specific decisions are delegated once a corridor has sufficient Consortium density.

Corridor operating entity. A local company owned by Arusvyx Operations, established in a corridor to employ local staff, hold local licences, and act as domestic counterparty where required.

Corridor Participation Pool. An eight-percent allocation of ARSX earned by regional participants against recorded contributions during a corridor's early stages.

Distributable income. What remains of a cluster's gross rental income after the waterfall's deductions, transferred pro rata to Compute Share holders.

Eligibility tier. One of three levels recorded in the identity registry (base, verified, professional) that determines which instruments and offerings a participant may hold or subscribe to.

Hardware grade. The set of requirements a cluster must meet for admission, covering accelerator class, memory, interconnect, facility tier, connectivity, and isolation capability.

Identity registry. The off-chain record of participant verification and jurisdiction, operated by Arusvyx Operations, from which an on-chain tier attestation is derived.

In-kind contribution. The route by which a certified operator brings its own hardware into a Consortium, either by transfer to the vehicle against Compute Shares or by term lease to the vehicle.

Isolation level. The degree of separation between tenants on a cluster: dedicated node, dedicated partition with network segmentation, or shared capacity under confidential computing.

On-demand capacity. Capacity sold at a variable rate between reservations, prepaid, and interruptible with notice.

Operator bond. ARSX posted by a certified operator in proportion to operated capacity, slashed on service breach and burned.

Platform service fee. The network's share of gross rental income, indicatively ten to fifteen percent, deducted at the fifth step of the waterfall and constituting protocol revenue.

Proof of correctness. Statistical assurance, through random re-execution on a second certified cluster, that billed work produced valid results.

Proof of existence. Attestation, through device features or physical audit, that a cluster consists of its declared hardware.

Proof of uptime. Assurance, through telemetry and random challenge, that a cluster was responsive during billed periods.

Reservation. A commitment of a fixed quantity of capacity for a defined period at an agreed rate, invoiced monthly against a deposit.

Reserve. Funds set aside by a Consortium from gross income to cover replacement, downtime, disputes, and end-of-life, drawn under the Consortium's own rules.

Residency zone. A partition of the network by jurisdiction of hosting facility; workloads tagged for a zone are visible only to clusters within it.

Revenue-producing capacity unit. The normalised measure of certified capacity that has settled a customer invoice, against which the incentive pool unlocks.

Security council. The body empowered to pause contracts and cross-chain paths; appointed in the first governance phase and elected thereafter.

Staking position. An amount of ARSX committed on the primary chain for a chosen duration, weighted by a duration factor for allocation priority, fee reduction, distributions, and voting.

Technical services agreement. The arm's-length contract under which a Consortium asset vehicle uses the network's scheduler and settlement contracts; the only commercial relationship crossing the ring-fence.

Token generation event. The event at which the genesis supply of ARSX is minted on the primary issuance chain and first distributed.

Waterfall. The fixed order in which a cluster's gross rental income is applied: vendor share where applicable, hosting and power, connectivity, operator fee, reserve, platform service fee, distributable income.

Appendix B: Technical Parameters and Reference Values

All values are indicative at the date of this document. Final values will be published in the token generation documentation and, where governed, may thereafter be adjusted within the stated bounds. Values marked ▲ are adjustable by governance; values marked ■ are fixed.

Hardware grade reference values ▲

Dimension	Reference value
Accelerator class	Data-center-class parts with ECC high-bandwidth memory, vendor-supported for enterprise deployment; current and one prior generation admitted
Memory per accelerator	Minimum set per generation; reference 80 GB for the current generation
Intra-node interconnect	Vendor high-speed fabric at the generation's standard bandwidth
Inter-node interconnect	Minimum 400 Gb/s per node, non-blocking within cluster
Facility	Redundant power and cooling equivalent to concurrent-maintainable tier; diesel or battery backup rated for a minimum of 24 hours
Connectivity	Two independent upstream providers on diverse physical paths
Isolation	Dedicated node, dedicated partition with network segmentation, and confidential-computing mode where the generation supports it
Excluded	Consumer and prosumer graphics cards of any generation

Revenue-producing capacity unit ■

One unit equals eight accelerators of the current reference generation, certified and with at least one settled customer invoice. Prior-generation accelerators count at a governance-set fraction of a unit, reference 0.6.

Cluster waterfall, indicative ranges for a well-utilised cluster

Step	Share of gross income
Vendor revenue share (where applicable)	Per arrangement
Hosting and power	20–30%
Connectivity	2–4%
Certified operator fee	8–12%
Reserve contribution	5–8% until target, then 2–3% maintenance Δ (network minimum 4% / 1.5%)
Platform service fee	10–15% Δ (band 8–18%)
Distributable income	Remainder

Reserve target ▲: six months of fixed costs, held in stablecoin within the vehicle.

Platform fee reduction for stakers ▲: up to three percentage points, scaling with stake size and committed duration; full reduction at a twelve-month commitment above a governance-set threshold.

Protocol revenue split ▲

Destination	Indicative	Bound
Buyback and burn	35%	30–40%
Staker distribution	35%	30–40%
Treasury	30%	20–30%

Buybacks execute on-chain through time-weighted orders at fortnightly intervals.

Operator bond ▲: reference value per capacity unit denominated in US dollars and settled in ARSX at the prevailing price; revalued quarterly; concave schedule such that bond per unit falls by a governance-set factor for each doubling of operated capacity.

Slashing tiers ▲

Breach	Fraction of bond
Availability below SLA target within tolerance	0% (customer credit only)
Repeated availability breach in a quarter	2–5%
Sustained outage beyond defined duration	5–15%
Confirmed correctness failure	10–25%
Misreporting or undisclosed control of a re-execution cluster	100% and removal

Verification ▲: re-attestation of hardware every twelve months and on any replacement; uptime challenge frequency randomised around a mean of four per cluster per day; correctness sampling rate 1–3% of jobs, weighted toward new operators.

Settlement cadence ■: telemetry continuous; metering near real time; reconciliation daily; settlement and distribution monthly.

Siting ▲: minimum renewable or low-carbon share of contracted supply for new clusters, reference 50%; clusters below the minimum admitted only by Consortium vote with disclosure.

Customer terms ▲: reservation deposit one month's committed fee; on-demand capacity prepaid; grace period 15 days before suspension.

Insurance ■: property cover at replacement value; business-interruption cover for facility-caused outages, reference 90 days.

Offering window ▲: staker-priority part, reference first 48 hours of the window; escrow returned in full if the minimum is not reached.

Staking tiers ▲

Committed duration	Duration factor
3 months	1.00
6 months	1.25
12 months	1.60
24 months	2.00

Early-exit cooldown 14 days; accrued unpaid distribution for the current period forfeited to the staker pool.

Arbitration ▲: dispute fee 1–3% of the amount in dispute with a floor, paid by the party against whom the determination is made; panellist forfeiture for collusion or breach of confidentiality 10–25% of the staked position and removal from the register.

Corridor Participation Pool ▲: no single corridor may draw more than 40% of the pool.

Incident response ■: pause of affected contracts or paths within 1 hour of confirmation; participant notification within 72 hours; post-incident report within 30 days.

Token supply and allocation ■

Total supply 1,000,000,000 ARSX. Allocations as stated in Section 7.4. Initial circulating supply at the token generation event: 10–15% of total.

Vesting ■

Allocation	Lock	Linear vest
Team and advisors	12 months	36 months
Early investors, seed	12 months	24 months
Early investors, strategic	6 months	24 months
Public distribution	0–3 months by channel	6–12 months
Liquidity	None; positions locked 24 months minimum	—
Treasury	Governance-gated; annual drawdown cap 20% of remaining ARSX balance Δ	—
Community and ecosystem incentives	Milestone-driven	—
Corridor Participation Pool	Earned; vests on recorded contribution	—

Incentive schedule ▲: per-unit unlock declines by a fixed factor per tranche of cumulative certified units; annual ceiling of 15% of the pool's initial balance; capacity withdrawn within 24 months of certification netted against future unlocks.

Governance thresholds ▲

Item	Reference
Phase-two trigger	Consortia earning in ≥2 corridors; settled revenue ≥ threshold for 6 consecutive months; ≥20% of circulating supply staked
Quorum	10% of staked supply, declining with staker count to a floor of 5%
Supermajority	66%, declining to 60%
Heightened supermajority (veto waiver, early phase three)	80%
Foundation veto disclosure	Within 7 days of exercise
Contract upgrade time lock	14 days minimum
Phase-three trigger	24 months after phase two, or earlier by heightened supermajority

Cross-chain ▲: per-period mint cap per secondary chain, reference 0.5% of circulating supply per 24 hours; independent pause per path; governance participation from secondary chains by attested snapshot only.

Security council ■: nine members, five signatures required; at least two independent in phase one; elected from phase three.

Appendix C: References

Third-party data cited in this document is drawn from the public sources below. Figures are current as of the dates given and have not been independently verified by the network.

1. Firmus Technologies and DayOne, joint announcement of a 360 MW multi-tenant AI factory campus in Southeast Asia, including planned deployment of up to 170,000 accelerators from the first quarter of 2027 and reported committed offtake, 2026.

2. NVIDIA, public description of the DSX programme under which accelerator infrastructure is deployed to data center operators on a revenue-sharing basis, 2026.

3. CoreWeave, announcement of Southeast Asian facilities totalling 360 MW of contracted IT power, August 2026.

4. Ministry-level statements from a Southeast Asian government on a data center investment pipeline of approximately 1.3 GW above roughly 580 MW in operation, representing announced investment of fifteen to twenty billion US dollars, 2026.

5. Industry cost benchmarks for data center construction per MW across Asia-Pacific markets, 2025–2026 editions.

6. Public reporting on private credit facilities secured against GPU hardware and take-or-pay contracts, 2025–2026.

7. Public reporting on capacity waiting times experienced by mid-sized AI companies, 2025–2026.

8. Publicly available documentation and usage disclosures for decentralised GPU marketplaces and enterprise-grade decentralised compute networks, 2026.

9. Chainalysis and other published analyses of cross-chain bridge losses, cumulative to 2026.

10. National AI strategy documents and data residency regulations for the corridor jurisdictions referenced in Section 2.2, current editions.

11. International Energy Agency and national grid operator publications on renewable baseload share by market, latest editions.

12. Regulatory frameworks for tokenised asset interests in the foundation and vehicle domiciles under consideration, current editions.