Solvyren Product Whitepaper

v0.4.0 · 2026-11-25 · Mainnet edition

Clearer participation, network records and computing services.

> Under the proposed issuance policy, 80% of each block's reward budget is allocated to AI tasks. Until a task passes review, this share only accumulates as an allowance. No coins from it are issued or enter circulation.

Abstract

Solvyren is a proof-of-work blockchain whose mainnet is live. Blocks are produced with Solvyren MatrixWork SMW-16/2048 v1, a MatMul-class proof-of-work that GPUs compute, and history is confirmed by GRANDPA finality. Anyone with a compatible GPU can mine today through the official pool or their own node, and every block, transfer and reward can be checked in SVYscan.

The longer-term product is a GPU compute network. Solvyren is developing an AI service for users, covering inference, image generation and video generation, together with a platform that aggregates GPU capacity from providers and from company-operated servers. The idea that ties these together is "one card, two uses": the same GPU can mine SVY and, once the AI side opens, also serve paid AI workloads, with a scheduler moving it between the two roles.

This paper explains the live network, the market problem it addresses, the AI and compute platform under development, the economics of SVY and the roadmap. It separates what is running now from what is still being built. AI inference, the production inference API, provider onboarding and cloud server purchases are not open yet.

01 / Vision and user value

Participation should offer more than a rate on a screen. Device owners need to understand work requirements and running costs. Computing customers need to know the resource and delivery scope they select. Both need records they can review.

Solvyren is designed around these needs. Mining is a transparent, rule-based way to secure the network and earn SVY under a published issuance schedule. The AI and compute platform, now in development, is meant to turn the same GPUs into a source of useful work for paying users. Network work, service charges and rewards are each presented in their own context and can be traced through blocks, transactions and account activity.

The core value proposition is one card, two uses (Chapter 03): one GPU that secures the chain and can also supply AI compute. For miners this means a second potential use for hardware they already run. For the network it means a base of GPUs that real AI demand can draw on (Chapters 04 and 05).

02 / Market context: AI compute demand and supply

Demand for AI compute is large and still growing. A conservative way to see its scale is through energy and investment data rather than vendor forecasts. The International Energy Agency estimates that data centres used about 415 TWh of electricity in 2024, around 1.5% of global consumption, and that global investment in data centres reached roughly USD 0.5 trillion in 2024, nearly double the 2022 level. In its Base Case, data-centre electricity use more than doubles to about 945 TWh by 2030, with AI as the most important driver [1]. These are approximate, agency-level figures. They describe the whole data-centre sector, not a market Solvyren can address, and projections may not materialise.

The mix of work is shifting. Training a large model is a concentrated, one-off effort run on tightly connected clusters. Inference — running a trained model to answer a request — happens every time someone uses an AI product, so it grows with usage. Image and especially video generation add heavy per-request workloads. Many inference and generation jobs can run on a single capable GPU and do not need a large training cluster.

Supply is concentrated and costly. Most high-end accelerator capacity sits with a small number of hyperscale cloud providers and large operators, where access is priced for enterprise customers and may require long commitments. The IEA also notes grid-connection and equipment bottlenecks that can delay new data-centre projects [1]. At the same time, many capable prosumer and workstation GPUs (RTX 5090 and RTX PRO 6000 class) are owned by individuals and small operators. These cards are often idle for part of the day, and their owners usually have no direct channel to customers who need inference or generation capacity.

Turning that scattered capacity into a usable service needs more than hardware. It needs a coordination layer: a way to list and identify devices, match jobs to suitable GPUs, measure the work actually delivered, and settle payment in a way both sides can verify. Without that layer, buyers cannot trust distributed supply and owners cannot reach demand.

This is where Solvyren fits. The proof-of-work mainnet already gives GPU owners a reason to run their cards on the network and provides accounts, transactions and finality records that can serve as a settlement base. The AI service in development (Chapter 04) is intended to create direct demand for compute, and the platform (Chapter 05) is designed to route that demand to provider GPUs and company-operated servers, with a planned 20% / 80% split of AI-service revenue between the platform and the GPU providers who serve the work. One card, two uses (Chapter 03) links the two sides: the hardware that secures the chain is the same hardware that can serve AI demand.

This analysis is not a forecast of SVY prices, revenues or returns. Market conditions, competition and technology change, and the AI side of Solvyren is still in development.

[1] International Energy Agency (IEA), Energy and AI, Executive summary, 2025. iea.org/reports/energy-and-ai

03 / One card, two uses

One card, two uses is the central idea of Solvyren participation. A GPU connected to Solvyren has two roles. First, it mines: it computes SMW-16/2048 v1 work, helps secure the mainnet and earns block rewards under the protocol. This role is live today. Second, once the AI side opens, the same GPU can serve AI workloads such as inference, image generation and video generation for paying users, earning rental fees or commissions on top of mining. This role is in development and will roll out in phases.

The two roles are scheduled, not run as one blended job. When an AI job is assigned to a GPU, the scheduler moves that card from mining to the job; when the job ends, fails or loses its heartbeat, the card returns to mining. The design keeps the miner's standing in the pool during an AI job through rolling 60–120 second reservations that stay valid only while heartbeats, progress and job receipts are valid. Mining rewards and AI-service earnings are accounted separately, and AI output is never treated as a block proof.

For miners this means a second potential use for the same hardware, without having to find their own customers. For the network it means GPUs that are already online, measured and tied to on-chain accounts, which the AI service and compute platform can draw on (Chapters 04 and 05). Income from either role depends on demand and network conditions and is not guaranteed.

04 / Solvyren AI service (in development)

Solvyren is developing an AI service for users. The planned capabilities are AI inference (text and chat-style requests), image generation and video generation, offered through Solvyren's own apps and an API. The service is in development; it has not been launched or publicly released, and the production inference API is not open. The models used will be described when the service opens; this paper does not claim that they are trained from scratch in-house, and it makes no performance, quality or model-size claims.

The AI service matters to the network because it is designed to create direct demand for GPU compute. Each request becomes a job that the platform can route to provider GPUs or company-operated servers (Chapter 05), which is what gives one card, two uses (Chapter 03) its practical value.

Access is planned to be paid through subscription tiers covering inference, image generation and video generation. Developers and businesses are planned to call the same capabilities through a metered API. Tier names, prices and quotas have not been set and will be published when the service opens.

The API layer is one of the main ways shared compute is used. GPU providers — including miners who share their card's compute — would serve AI users' requests through it. Subscribed users and API customers pay to use the service, and providers who serve the work earn additional token rewards on top of their mining rewards.

The planned allocation of AI-service revenue is 20% to the platform and 80% to the GPU providers who serve AI users. API usage revenue is included in the 80% provider pool. This split may be

refined before launch. It is separate from block-reward emission and does not change the SVY supply or issuance schedule in Chapters 11 and 12, including the 20% / 80% division of each block's reward budget between mining and the AI reward allowance. Provider earnings depend on demand and are not guaranteed.

Planned allocation of AI-service revenue	Share
Platform	20%
GPU providers who serve AI users (includes API usage revenue)	80%

05 / Server output: cloud computing and GPU compute as a service

Solvyren's compute offering has two parts: a platform that aggregates GPU capacity from independent providers, and GPU server infrastructure operated by Solvyren itself. Both are offered as a service with limited availability; provider onboarding and purchases open in phases.

Part A — GPU aggregation platform (in development). Many owners of high-end GPUs, such as RTX 5090 and RTX PRO 6000 class cards, have compute but no customers and no rental channel. Demand for inference and generation, meanwhile, is fragmented across many small and medium users who cannot easily contract with individual owners. The platform is designed to close that gap. Providers list their GPUs; the platform identifies and checks each device, matches jobs to suitable hardware, meters the work actually delivered, and settles payment with records tied to on-chain accounts. Providers earn rental fees or commissions for accepted work, within the planned 20% / 80% revenue allocation described in Chapter 04. Provider onboarding is not open yet and will open in phases; nobody can list a GPU on Solvyren today. Rental income depends on demand and is not guaranteed.

The aggregation platform adds value to the network in three ways: it gives miners a second use for their hardware, it turns otherwise idle capacity into supply for the AI service, and it keeps metering and settlement auditable through the same accounts and finality records that secure SVY.

Part B — Company-operated GPU servers. Solvyren operates its own GPU server infrastructure and offers its compute output to customers as a service. The current product, shown in the cloud console, is mining hashrate rental: a customer selects a plan by hashrate and term, and rewards are allocated from finalized blocks according to effective hashrate and active time. It provides no server access or control. Official reference prices in USDT are listed below. Each plan has 20 slots, each account can order each plan once, and the 1 MH/s plan is a first-order trial. Purchase is not yet available; plans open in phases with limited availability.

Hashrate, job duration and inference throughput are different measurements and are not combined into one earnings figure. No returns are implied by the plans or prices.

Plan (reference price, USDT)	7 days	15 days	30 days
1 MH/s trial (first order only)	19.9	—	—
3 MH/s	100	190	350

Plan (reference price, USDT)	7 days	15 days	30 days
8 MH/s	240	450	840
20 MH/s	530	999	1680

06 / Technical architecture

Solvyren has two architectural layers. The chain layer is live. The AI and compute platform layer is in development and has not been officially launched or publicly released.

Part A — Live chain layer. The Solvyren node is built on the Substrate framework and runs the mainnet identified by genesis hash 0x678107bc4d21a4e7f9b12ee817af0c93f074cf04fbd7ee716976af42ca92f4c8. Blocks are sealed with SMW-16/2048 v1 proof-of-work and finalized by GRANDPA (Chapter 07). The runtime holds accounts, balances, transfers, the fee burn and the issuance schedule. The current node release is mainnet-retarget-24121, published on GitHub at NYCTEAM/solvyren-svy with checksums; full nodes connect to peers over TCP port 30333.

Miners connect to the official pool at tls://pool.solvyren.com:18444 with a public SVY payout address and a worker name. The pool distributes SMW work, verifies shares by recomputing them, and pays rewards from finalized blocks; GPU and cloud mining rewards are subject to a 0.25% pool fee. Miners can also run their own node. Official miner packages cover Windows (NVIDIA and AMD), Linux and HiveOS; share submission and automatic payouts on AMD hardware are still pending verification.

Read access is public. SVYscan (scan.solvyren.com) indexes blocks, transactions and addresses. A read-only HTTPS JSON-RPC endpoint serves developers, and the official self-custody wallet connects over WSS.

Part B — AI and compute platform architecture (in development). Provider onboarding: a GPU joins under an on-chain account, declares its hardware, and passes device checks before it can receive jobs; self-reported speeds alone do not establish trusted capacity. Scheduling and matching: a scheduler assigns each job to a suitable GPU, binding the job to that provider and to a hash of the request terms, and reserves the maximum payment at assignment. Metering: usage receipts record the model, runtime and measured usage for each job, with one rounding per job. Settlement: an attestor role, kept separate from the scheduler, accepts or rejects results; only accepted work is paid, unused reservations are released, and records are kept so payments can be reconciled. Serving: inference, image and video requests arrive through Solvyren apps and the API and are routed to provider GPUs or company-operated servers. One card, two uses: a GPU moves between mining and AI jobs under the scheduler, with rolling 60–120 second reservations protecting its mining standing during a job (Chapter 03).

These platform components are being designed and tested. None of them is a production service yet, and AI output is not a consensus proof.

07 / Consensus: SMW-16/2048 v1 and GRANDPA finality

Proof-of-work. For each nonce, a miner derives deterministic matrix inputs, multiplies a 16×2048 matrix by a 2048×16 matrix, and binds the result commitment to the chain identity, the expected parent block, the pre-seal header and the beneficiary account. A block is valid when that commitment meets the current difficulty target. Nodes and the pool recompute the same work under the same rules. SMW is matrix proof-of-work for block production; it is not model training or AI inference.

Difficulty. The target block interval is about ten seconds. A rolling difficulty retarget activated at block 24,121, adjusting difficulty from recent block times so that the interval tracks the target as hashrate changes. Full nodes must run mainnet-retarget-24121 or a compatible newer release.

Finality. Proof-of-work produces the best chain; GRANDPA finality confirms history. The best block is the current chain head and can still change in a reorganisation. A finalized block is part of confirmed history. Valid work, inclusion in a block and finalization are three distinct stages, and rewards and settlement rely on finalized records.

Commercial orders have their own delivery and reconciliation process. A finalized network record does not by itself prove cloud service delivery or that an AI answer is correct.

08 / Verifiable AI computation proof: target protocol design

Solvyren names its target mechanism Verifiable AI Computation Proof: participants execute AI computation with defined inputs and acceptance rules, submit reviewable results, and earn rewards for accepted valid work. Research directions include constrained matrix multiplication (GEMM), model operators and inference jobs. Matrix arithmetic is a building block; useful application work requires a real task and delivery need.

A task descriptor should bind network identity, job identifier, block or period context, model digest, input digest, numerical precision, execution budget, deadline and beneficiary account. These bindings identify submissions and help prevent duplicate records and cross-task replay. They do not by themselves prove correct execution.

A GPU executor submits an output commitment, job receipt and required computation records. Verifiers independently recompute supported operators or jobs under agreed rules. Sampled checks and redundant execution are research options whose detection guarantees, costs and numerical tolerances must be specified. An output hash, or another AI rating an answer, is not by itself a complete deterministic consensus proof.

The target reward flow is allocation, computation, correctness and duplicate checks, acceptance, finality and reward accounting. Work weights, verifier incentives and disputes require protocol definition. Issuance rewards and customer-paid inference fees use separate accounting; work cannot be settled twice or create extra issuance because of service demand.

The product family also plans inference jobs. Customers define models, inputs, resources and acceptance criteria; scheduling selects available GPUs; executors return results and run records;

settlement follows order acceptance. Service payment differs from block rewards. A future consensus role needs an explicit protocol rather than automatic equivalence.

Current implementation: the Solvyren mainnet uses Solvyren MatrixWork (SMW-16/2048 v1), a MatMul-class proof-of-work, with GRANDPA finality. Official miners run on NVIDIA GPUs; share submission and automatic payouts on AMD hardware are still pending verification. AI inference uses separate jobs and reward accounting, is in development and is not open. Inference output is not treated as an SMW consensus proof.

09 / Mining and computing compared

Solvyren focuses on integrating products and the participation experience, rather than renaming hash search. Network, Compute and SVYscan address chain records, resource service management and public queries respectively.

Compared with an interface centered only on hash work and mining proceeds, Compute is designed to cover devices, orders, service metering and settlement. Compared with a general GPU computing service, the native network supplies its own accounts, transactions and finality records. Commercial delivery still depends on the relevant order and service records; a chain record does not automatically prove an AI answer correct.

No higher return, lower power consumption or GPU exclusivity is promised. Useful selection criteria are workload support, resource specifications, cost, availability and record transparency.

Dimension	Conventional hash mining	Solvyren product design
Network work	Search hashes against a difficulty target	SMW-16/2048 v1 matrix work ticket (MatMul)
Participation	Solo or pool participation	Owned GPUs and GPU cloud servers merge by valid share
Records	Depend on the network or service	Linked record entry points across Network, Compute and SVYscan
AI computing	Hash search itself is not an AI task	Inference is accepted and rewarded separately, with governed hashrate reservation
Selection criteria	Valid work, cost and competition	Valid shares, service specifications, delivery and transparent records

10 / How to participate today

Mine with your own GPU. Download an official miner from the website and verify its checksum. Enter the pool address tls://pool.solvyren.com:18444, your public SVY payout address and a worker name, then start mining. Equipment, electricity, connectivity and maintenance are your own operating costs. Check the downloads page for the current status of each package; AMD share submission and automatic payouts are still pending verification.

Run a full node. Operators and developers can run a full node from the mainnet-retarget-24121 release to verify the chain independently, and can query public data through SVYscan or the read-only RPC.

AI and compute participation. Becoming a GPU provider for the AI service, subscribing to the AI service and buying cloud server plans are not open yet. They will open in phases and will be announced on the service status page and in release materials.

11 / SVY and asset allocation

Without an accepted AI task, no AI reward is issued or enters circulation. Unused AI capacity accumulates. The operator sets the task reward and reserves capacity on assignment. A verified device receives the accepted amount once, after completing the task and passing review. Reservations do not credit a wallet.

The planned supply cap is 210 million SVY. The team receives 10 million at genesis, leaving a reward budget of 200 million. Issuance measures coins actually issued; circulating supply also depends on how issued coins are held and transferred.

Allocation	Share	Amount (SVY)
Genesis team allocation	≈4.76%	10,000,000
Reward budget	≈95.24%	200,000,000

Transaction fees and burning

The base fee for a standard SVY transfer is 0.000005 SVY. The protocol deducts that fee from the payer and reduces on-chain issuance by the same amount. The fee has no recipient or DEAD address.

For a transfer of 10 SVY in mainnet block #397, TransactionFeePaid records the fee and balances.BurnedDebt records the burn. Both record 0.000005 SVY (5,000,000 atomic units).

The issuance cap is 210,000,000 SVY, including the genesis allocation of 10,000,000 SVY and a reward budget of 200,000,000 SVY. On-chain issuance at a finalized block reflects the SVY actually issued, less all amounts burned.

Of each block’s reward allowance, 20% goes to base mining and 80% accrues as AI reward allowance. Until an AI task passes acceptance, that 80% only accrues: no SVY is issued from it or enters market circulation. Only the approved amount is issued after acceptance.

12 / Issuance driven by block height

Blocks 1 through 1,576,800 each budget 15 SVY. Mining receives 20% (3 SVY); the AI pool accrues 80% (12 SVY). With no accepted AI task at block 1, mining issues 3 SVY, AI issues 0 SVY, and 12 SVY of capacity remains for future tasks.

Every 1,576,800 blocks the per-block budget falls by 11.826% from the previous period. Block 1,576,801 starts the second period at 13.2261 SVY: 2.64522 for mining and 10.58088 in AI capacity. At a ten-second target this is about half a year; block height determines the transition.

13 / SVYscan: readable network records

SVYscan uses blocks, transactions and addresses as entry points to network activity. Users can follow the block containing a transaction, transfers and account-related records. Clear record links make participation outcomes easier to understand and trace.

The explorer presents node-sourced data with finality and indexing progress. When a network connection or record is unavailable, it shows unavailable rather than converting an unknown balance into zero.

Users should verify activity against the correct network identity and public records. SVYscan supports queries; it does not replace consensus or prove commercial delivery.

14 / Wallets and accounts

Solvyren uses native AccountId32 / SS58 accounts. The official self-custody wallet at wallet.solvyren.com creates and imports accounts in the browser; the cloud console uses the same mainnet, so the same recovery phrase produces the same SVY address. Public addresses are for receiving and queries. Private keys and recovery phrases stay with their owner and must never be submitted to the website, support or a query form.

Before a transfer, check the network, address format, amount and fee. Similar names or symbols on other networks do not identify the same asset. Do not choose a transfer network solely because it displays SVY.

A service wallet and a personal self-custody wallet have different responsibilities. Solvyren does not offer custody, and nobody from Solvyren will ask for your keys.

15 / Roadmap

Phase one — Mainnet and mining (live). The proof-of-work mainnet with GRANDPA finality, the official pool, miner packages, node releases, SVYscan and the self-custody wallet.

Phase two — Network hardening. More public nodes and bootnodes, completed AMD verification, miner packaging improvements and further documentation.

Phase three — AI service. Solvyren's AI service for users, covering inference, image generation and video generation, with subscription tiers and a metered API. In development.

Phase four — Compute platform. Provider onboarding for GPU owners, scheduling and matching, metering and settlement with the planned 20% / 80% revenue allocation, and phased opening of company-operated GPU server plans. In development.

Phase five — Verifiable AI computation. Progressive work toward the verifiable AI computation proof described in Chapter 08, subject to explicit protocol decisions.

Phases advance by completion criteria, not calendar dates, and no earnings projections are attached to them.

16 / Current availability and risk notice

Available now: the Solvyren mainnet, public mining through the official pool or your own node, official miner packages and node releases with checksums, SVYscan, the read-only RPC and the self-custody wallet.

Not open yet: AI inference and the production inference API, image and video generation, AI subscriptions, GPU provider onboarding, and cloud server purchases. These are in development and will open in phases. Openings will be announced on the service status page and in release materials.

Risk notice. SVY is the native asset of the network, not an investment product. Mining and any future provider income depend on demand, network conditions, competition, hardware and electricity costs, and are not guaranteed. Forward-looking statements in this paper, including market context and planned features, are not forecasts of SVY prices or returns and may change. Software can contain defects; verify releases against published checksums.

17 / Plain-language explainers and glossary

Blockchain: a shared ledger kept by many computers. Each block lists transactions and points to the previous block, so history cannot be changed quietly.

Proof-of-work (PoW): a way to decide who adds the next block. Miners spend computation to find a valid result; others can check it quickly.

SMW-16/2048 v1: Solvyren's proof-of-work. Instead of hashing alone, each attempt multiplies two matrices (16×2048 and 2048×16), the kind of arithmetic GPUs are built for.

Finality (GRANDPA): the point at which a block is confirmed and will not be reversed. Until then, the newest block (the "best" block) may still change.

Mining pool: a service that combines many miners' work and shares rewards according to valid work submitted. Solvyren's official pool is tls://pool.solvyren.com:18444.

GPU compute rental: paying to use GPU capacity for a period or per job, instead of buying hardware.

Inference: running a trained AI model to produce an answer, an image or a video. It happens every time someone uses an AI product.

API: a programming interface that lets other software send requests to a service and receive results.

Subscription: a recurring payment that gives access to a service at a defined tier.

SVY's role: SVY is the native asset of the Solvyren network. It pays transaction fees (which are burned) and is the unit of block rewards. Holding SVY does not entitle anyone to platform revenue.

One card, two uses: one GPU that mines SVY and, once the AI service opens, can also serve AI jobs, with a scheduler moving it between the two.