Hashgraph

BEHIND THE LEDGER TECHNICAL SERIES: PART 2

Network Isolation and the Sovereign Perimeter:

Why independent network design is the foundation of institutional-grade private infrastructure

A technical paper for enterprise decision-makers on how HashSphere's independent network design delivers true data sovereignty, regulatory compliance, and infrastructure control; without vendor lock-in or isolated liquidity.

Satish Ramakrishnan
VP of New Products
Hashgraph

Jeff Barrows
Head of Site Reliability Engineering
Hashgraph

Kash Balhotra
Institutional Digital Assets Lead
Hashgraph

Independent network design is not a configuration option. It is the perimeter itself.

In Part 1 of Behind the Ledger, we established that transaction privacy on HashSphere is enforced cryptographically (through zkSNARKs and encrypted balances) so that no party, including the network operator, can see transaction content. That guarantee, however, only holds if the network boundary itself is secure.

A privacy layer deployed on top of a public or shared base layer does not give users a controlled perimeter. It gives them a private channel inside someone else's infrastructure. For regulated institutions, that distinction is not a matter of preference; it is often a hard regulatory constraint.

Banks and financial institutions operating under frameworks such as GDPR, MiFID II, Basel SCO60, and national data sovereignty legislation are frequently prohibited from allowing client transaction data to be processed or stored on infrastructure they do not control, regardless of whether that data is encrypted in transit. The restriction is not just about confidentiality; it extends to where data physically resides, who has structural access to the underlying infrastructure, and whether the institution can demonstrate, on demand, that those controls were in place and effective. A public or shared base layer fails all three tests by design.

Beyond regulatory obligation, there is a client confidentiality dimension. Institutional clients (sovereign wealth funds, pension managers, corporate treasuries) expect that their transaction activity, counterparty relationships, and position data will never touch infrastructure accessible to third parties. That expectation cannot be satisfied contractually when the underlying network is public. It can only be satisfied architecturally.

This paper focuses on the second of HashSphere's four compounding privacy pillars: independent network design. HashSphere networks have a sovereign perimeter; they are not subnets on a public network. This paper dives into what that means technically, why it matters for compliance and governance, and how HashSphere's stand-alone architecture delivers optionality without the trade-offs that most private network designs require.

Independent network design is not a feature. It’s actually four load-bearing pillars

HashSphere, a private network from Hashgraph powered by Hedera, is built on four compounding privacy properties that must coexist at every layer of the stack:

1. **Transaction confidentiality** — private by cryptographic enforcement, not policy
2. **Independent network design** — a sovereign perimeter, not a subnet on a public network
3. **Flexible validator and governance models** — not delegated to a third party
4. **Mathematically guaranteed fair ordering** — through the hashgraph consensus algorithm

Remove any one of these, and the guarantees of the others weaken.

Part 1 of this series established the foundation: how HashSphere enforces transaction privacy (confidentiality, anonymity, and selective disclosure) across DvP, PvP, and private transfers, without exposing data to any party, including the network itself.

This paper focuses on the second pillar: independent network design. Transaction privacy guarantees are only as strong as the network boundary that contains them. For regulated institutions operating under data residency frameworks, a private channel inside someone else's infrastructure is not a sovereign perimeter; it is a compliance liability.

Subsequent papers will address the remaining two pillars, with a final paper covering how HashSphere's cross-ledger protocol, CLPR, enables interoperability between public and private networks without creating islands of trapped assets.

At Hashgraph, we believe privacy and interoperability on a distributed ledger don't have to be a trade-off.

The problem with shared infrastructure

When regulated financial institutions first evaluate private blockchain networks, the instinct is often to reach for what is familiar: a subnet on an existing public network, a permissioned layer deployed on top of shared base infrastructure, or a network whereby each participant only sees the parts of a transaction they're a party to, or are entitled to see. These approaches vary in design, but share a common characteristic: they offer a version of privacy without offering a sovereign perimeter.

It is also worth being precise about what some of these designs actually are. In many respects, some of these private, permissioned networks are closer to distributed databases with cryptographic audit trails than to a blockchain in any meaningful sense. Consensus is managed by a central coordinator. Privacy is enforced by restricting who receives which data, not by making the data mathematically unreadable to anyone who does receive it. The result is a controlled network, but not a private one; and certainly not a sovereign one.

What a subnet actually gives you

Public networks allow institutions to deploy private or permissioned environments (subnets, appchains, and similar constructs) that appear to be isolated, but the validators underpinning those environments are still communicating over the public network or internet. Data traverses infrastructure the institution does not control, through validators it did not choose, in regions it may not be able to specify.

This also introduces metadata exposure: even where transaction content is encrypted, participant identities, transaction timing, and volume remain visible to outside validators. The risk is concrete: a Tier-1 bank executing a large bond position may have amounts fully encrypted, but the timing pattern and relative size of those transactions is itself a tradeable signal. The content doesn't need to be read for the activity to be inferred.

For institutions subject to data residency requirements, this is not theoretical. If only two validators operate in a given jurisdiction, the institution cannot guarantee its data stays within its regulatory perimeter, because that perimeter is not its to define.

EVM compatibility without the constraints of private QBFT networks

Institutions evaluating private EVM-compatible networks often inherit the architectural constraints of their base layer: fixed validator sets, vendor-coupled governance, and limited control over data residency and network perimeter.

HashSphere is fully EVM-compatible (institutions deploy Solidity smart contracts using familiar tooling without modification, but the underlying architecture is different in the ways that matter. Consensus runs on the hashgraph algorithm with mathematically guaranteed fair ordering. The validator set is institution-defined. The perimeter is sovereign. And open source foundations under LFDT eliminate vendor lock-in. EVM compatibility, without the trade-offs that typically come with it.

WHAT A SUBNET ACTUALLY GIVES YOU

A private environment on top of infrastructure the institution does not control

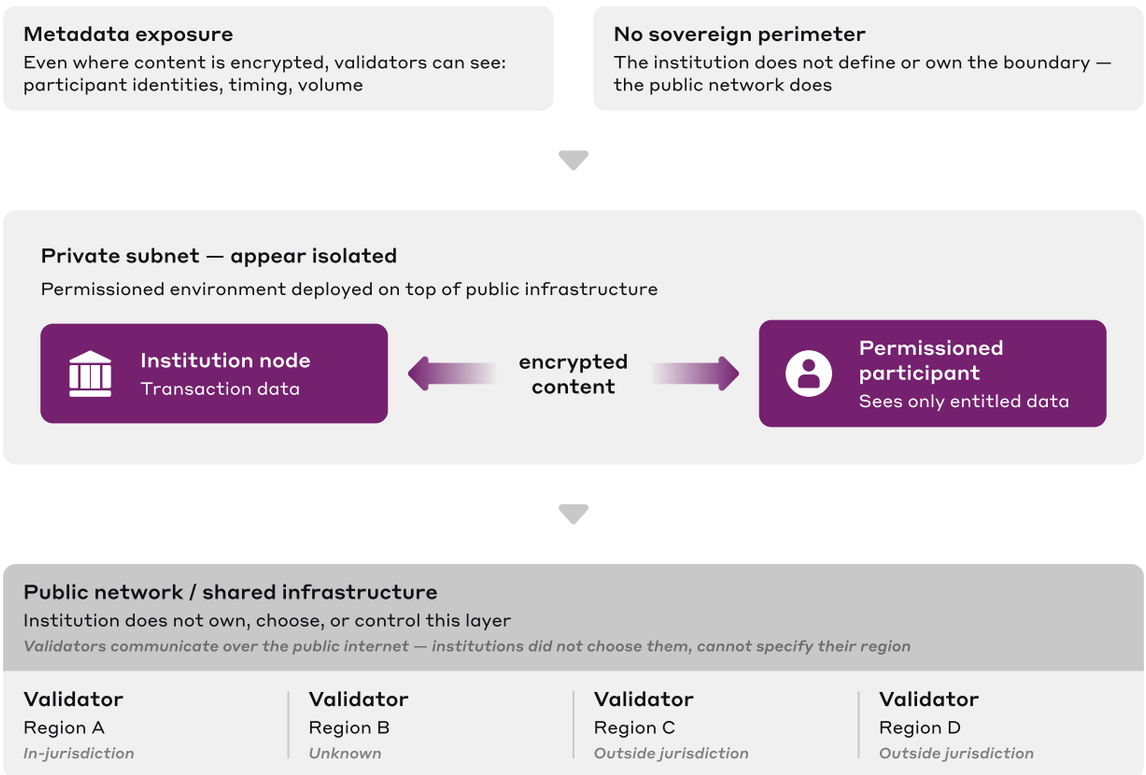

What a central routing party actually does

Many private, permissioned blockchain networks route transactions through a central controller. This coordinator checks for double-spend, sequences settlement, and maintains the integrity of the shared state. To do that, it must see all of the transaction data, including the private kind.

The network's privacy guarantee in this model is not cryptographic; it is contractual. The users are trusting the operator not to look. That is not the same as the operator being mathematically incapable of looking.

There is also an operational risk dimension. A central routing party is a single point of failure. If it goes down, the network goes down. If it is compromised, the data is compromised. If the vendor relationship changes (pricing, acquisition, policy), the institution has limited recourse.

Bob and Alice in practice: Bank A and Bank B are executing cross-border bond settlement on a private permissioned network with a central routing coordinator. The coordinator sees both legs of the transaction, the asset delivery and the cash payment, in plaintext. It sequences them, validates them, and logs them. Bank A and Bank B have a contractual assurance that this data will not be misused. They do not have a cryptographic one.

CENTRAL ROUTING COORDINATOR: WHAT THE OPERATOR SEES

Cross-border bond settlement — Bank A and Bank B

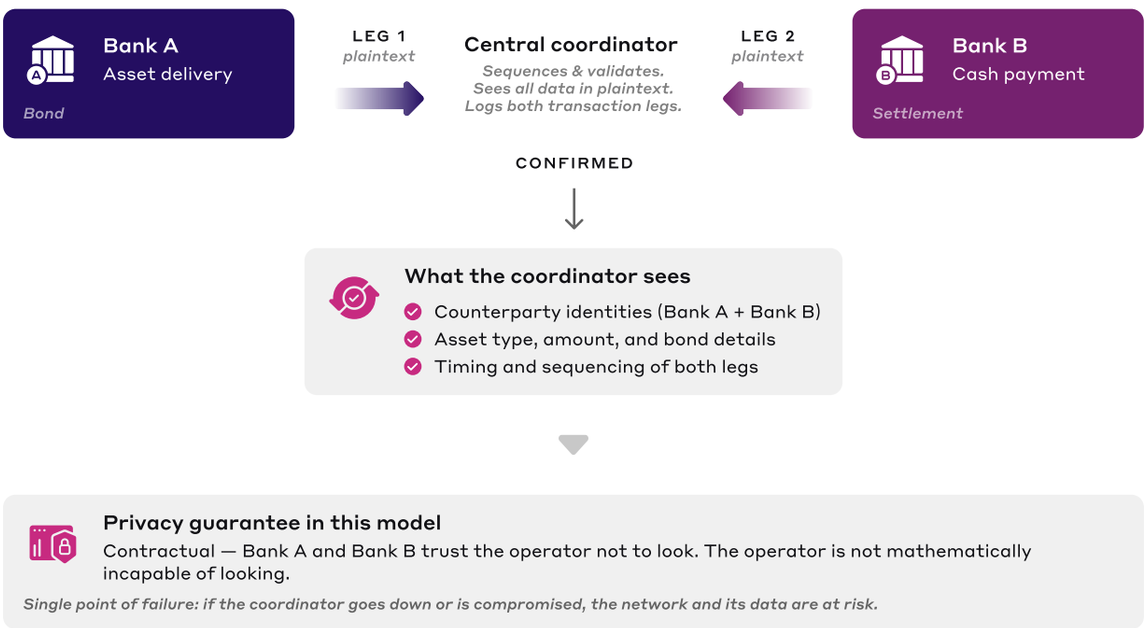

The quantum threat: harvest now, decrypt later

There is a longer-horizon risk that applies specifically to data transiting public infrastructure, even in encrypted form. In a harvest now, decrypt later (HNDL) attack, an adversary captures encrypted data in transit today, with the intention of decrypting it once sufficiently powerful quantum computing becomes available.

If an institution's transaction data never leaves its network boundary, the attack surface for HNDL is substantially reduced. An adversary cannot harvest what they cannot intercept. A breach of the network perimeter itself would still expose data, but it meaningfully reduces the window of exposure compared to data that routinely crosses public infrastructure.

For regulated institutions managing long-dated assets (i.e., sovereign bonds, structured products, pension fund positions), this is not a distant concern. The data being transacted today may still be sensitive in ten years. The infrastructure choice made today determines how much of it is at risk.

HashSphere: Private networks that users can actually own and operate

Developed by Hashgraph, HashSphere networks are fully standalone private networks. They are not subnets on a public network. They do not route transactions through a central coordinator. The network perimeters are defined and controlled by the institution, and everything within these spheres stays within it.

Network architecture: the sovereign perimeter

A HashSphere network, a “Sphere”, is a self-contained environment. The components that make up the network are encapsulated within whatever infrastructure boundary the institution defines: their own data center, their cloud provider, a managed cloud environment operated by Hashgraph, or a combination.

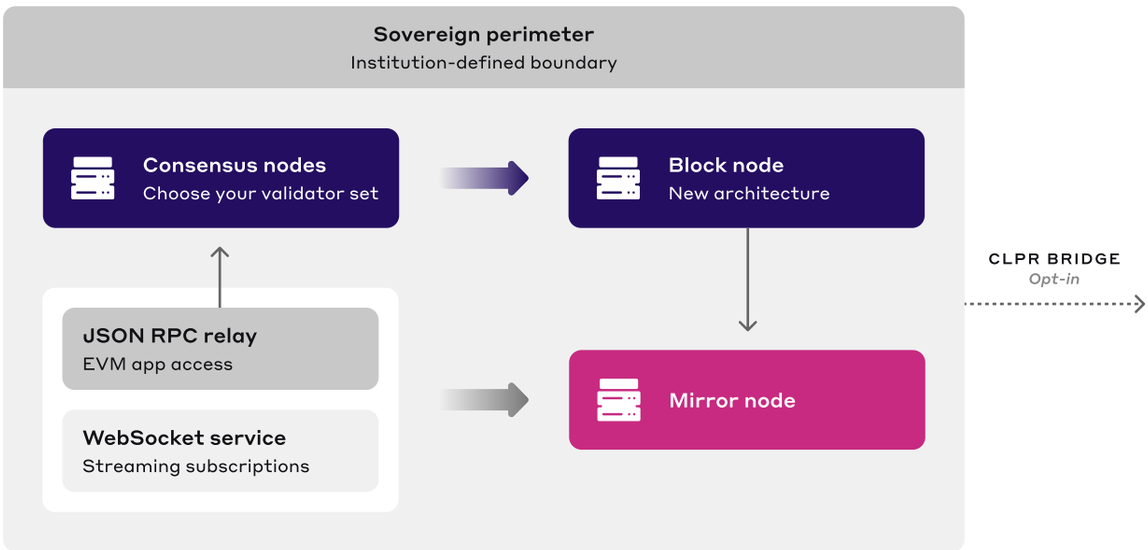

Within this perimeter:

- Consensus nodes process and finalize transactions using hashgraph consensus; mathematically guaranteed fair ordering, without a central sequencer.
- Block nodes store transaction data. The data lives inside the perimeter.
- JSON RPC relay and WebSocket services handle application connectivity; all within the boundary.

Data that never leaves your environment

The core value proposition of independent network design is straightforward: users maintain full control of their data at all times, and it never leaves the chosen environment without an explicit decision to send it out.

This is meaningfully different from what subnet-based designs offer, and the distinction runs deeper than perimeter control. Many private network designs depend on a public L1 for finality, state validation, or data anchoring. That dependency means the network's integrity is structurally tied to external infrastructure the institution does not control. Transaction validity, in those designs, is ultimately contingent on something outside the perimeter. On HashSphere, all data is anchored, validated, and stored entirely within the network itself. There is no dependency on a public L1 or any external system for the network to remain valid and self-consistent. The perimeter is not just a boundary; it is the complete operating environment.

In a HashSphere deployment, the institution defines that perimeter. Data flows inside it. The network operator, whether that is Hashgraph or the institution itself, does not have structural access to data that the institution has not chosen to share.

For institutions operating under data residency frameworks (GDPR in Europe, data sovereignty requirements in APAC, jurisdictional rules in the Middle East), this is the architectural prerequisite. The compliance question is not just "can you encrypt it?" but "can you guarantee it stays within this boundary?" HashSphere answers yes to both questions.

In practice: An Australian central bank requires that all transaction data processed on its network remain within Australian infrastructure. On a subnet-based design, the institution is dependent on validators being available in-region; and if validator coverage in that jurisdiction is thin, the data may route outside it. On HashSphere, the institution deploys consensus nodes in their chosen jurisdiction and defines the perimeter themselves. The data stays.

DATA RESIDENCY IN PRACTICE: AUSTRALIAN CENTRAL BANK

Subnet-based design vs. HashSphere sovereign perimeter

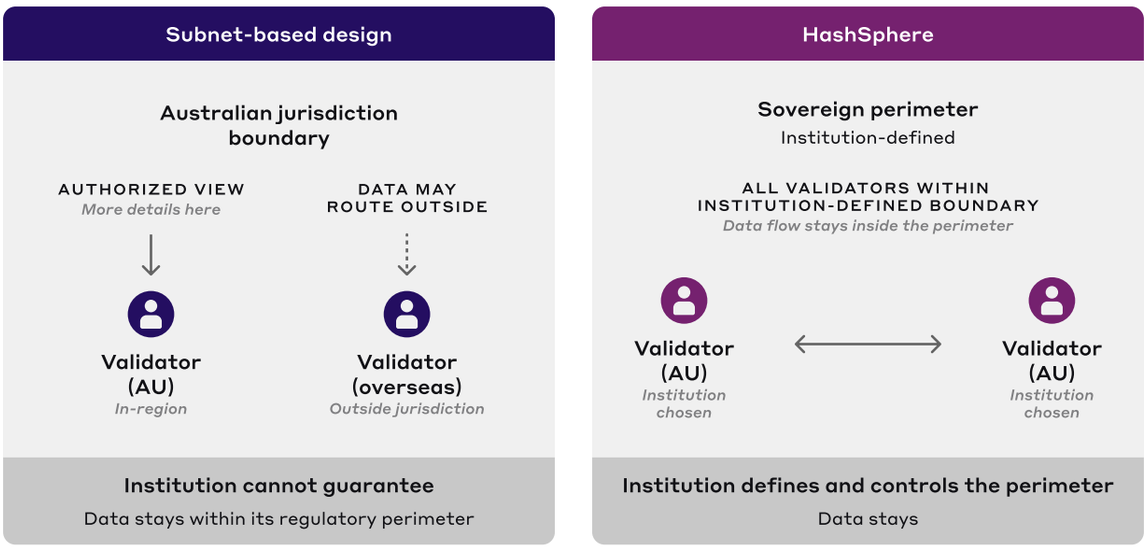

HashSphere consensus nodes are deployed in the institution's chosen jurisdiction
No data leaves without an explicit decision to bridge out.

Open source by design: No vendor lock-in

The Hedera codebase that powers HashSphere has been contributed to Linux Foundation Decentralized Trust (LFDT) as Project Hiero. This means the software underpinning HashSphere is open source, independently maintained by a neutral third party, and not contingent on one vendor.

In practical terms, an institution running a HashSphere network could take the open source codebase and run it entirely independently of Hashgraph. They would not be locked into proprietary infrastructure, a proprietary smart contract language, or a vendor relationship they cannot exit. This is architecturally different from many closed private network designs where the software and the vendor relationship are inseparable.

This is the open source model applied to enterprise infrastructure: the vendor is there to help users run it, and there when they need support. However, if the customer decides to part ways with Hashgraph, the tools and network go with them.

Governance: Configurable by design

One of the more nuanced aspects of private network design is the question of who runs the consensus nodes and what that implies for decentralization. The answer, in HashSphere's case, is deliberately flexible.

Decentralization is a foundational principle of the web3 ethos, but regulated financial institutions relate to it differently. For a Tier-1 bank, a fully decentralized, permissionless network t may pose a liability from a compliance perspective. Regulators want to know who is accountable and need the ability to access private transaction data. HashSphere addresses this through scoped cryptographic verification, allowing regulators to confirm specific transactions without exposing the full ledger. For a deeper look at how selective disclosure works in practice, see the first paper in the Behind the Ledger series: Transaction Privacy on HashSphere. Institutions want to know who to call when something breaks.

HashSphere's governance model is designed around this reality: The network can be as decentralized as the institution's use case and risk appetite require.

Governance models

Three configurations are available, depending on the institution's needs:

Managed service

Hashgraph deploys and manages the network on the institution's behalf, either on Hashgraph's cloud infrastructure or on the institution's own cloud. The institution does not operate consensus nodes directly but retains full data control within its defined perimeter. This is the fastest path to deployment and the model Hashgraph currently recommends for institutions beginning their HashSphere journey.

On-prem

The institution deploys and manages the network itself, using Hashgraph's tooling, documentation, and monitoring endpoints. Hashgraph provides ongoing support: development assistance, deployment guidance, and an SLA, but the institution owns day-to-day operations, including monitoring, backup, restore, and high availability configuration.

Consortium governance

A group of institutions each operates their own consensus node. The network is decentralized across the consortium members. Each party maintains custody of their node, contributes to consensus, and cannot be overruled by any single party. Privacy is maintained within the consortium perimeter, and each member retains full control of their data.

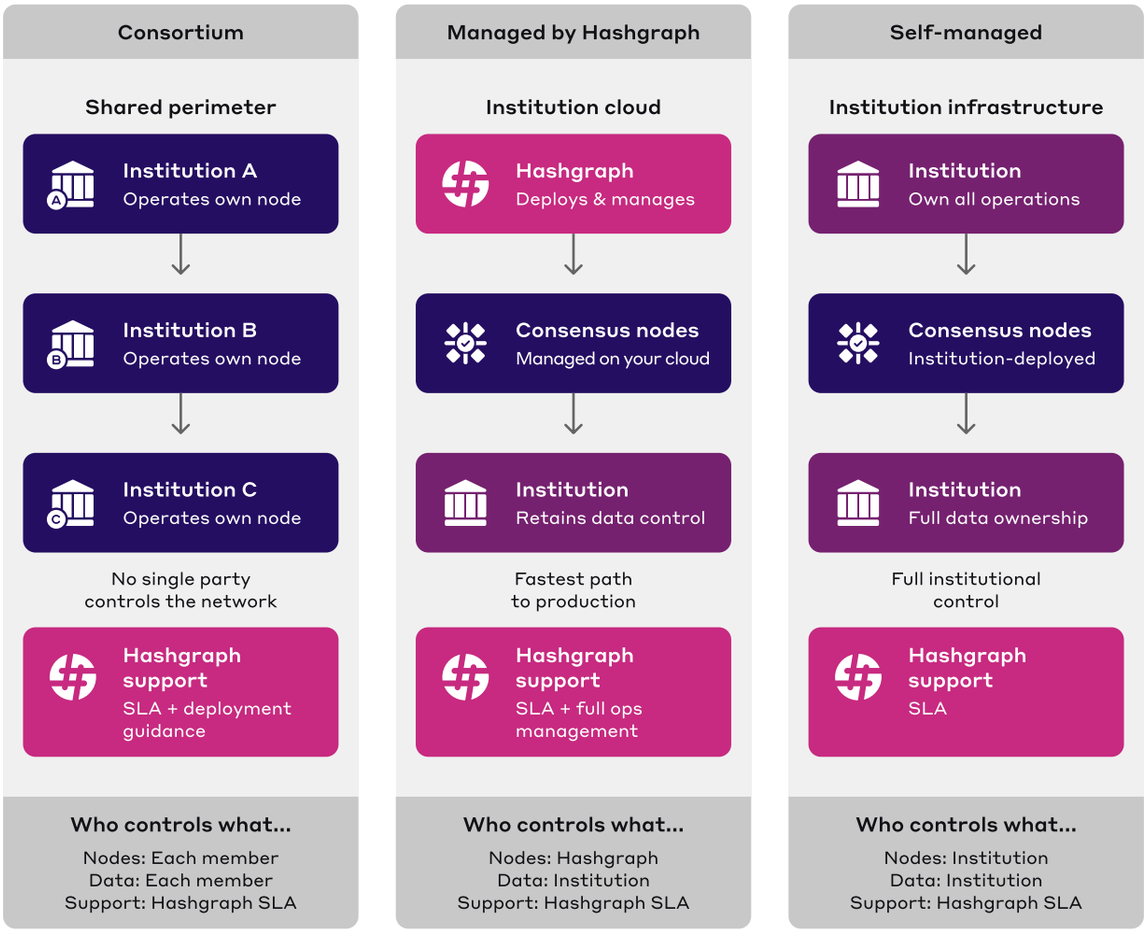

The key principle across all three models: the institution defines who can join the network, who runs the consensus nodes, and what data leaves the perimeter. Hashgraph does not make those decisions by default.

A note on decentralization and institutional adoption

Regulated entities are not primarily motivated by decentralization in the ideological sense. Instead, auditability, accountability, and the absence of single points of failure take precedence to ensure they remain in compliance with the regulators. HashSphere addresses all three without requiring institutions to abandon the governance models they already understand. It can be as decentralized as the user wants it to be.

Regulatory compliance and data sovereignty

Independent network design is not just an architectural preference. For most regulated institutions, it is the architectural prerequisite for compliance.

Data privacy frameworks across major jurisdictions impose strict controls on where transaction data can be processed, stored, and accessed. Basel SCO60, MiFID II, GDPR, and a growing body of national data sovereignty legislation create a regulatory environment in which the question "where does my data go?" is a compliance requirement with material consequences.

Data residency in practice

Consider what data residency actually requires of a private network:

- Transaction data must be processed within defined geographic boundaries.
- Access to that data must be restricted to verified, authorized parties.
- The institution must be able to demonstrate, on demand, that these controls were in place and effective.

A subnet-based design cannot reliably satisfy the first requirement, because the validator infrastructure is not fully under the institution's control. A central-routing model cannot reliably satisfy the second, because the coordinator has structural access to all data. HashSphere satisfies all three by design.

Bob and Alice in practice: A European asset manager is settling tokenized fund positions across subsidiaries. Under GDPR and MiFID II, transaction data (including counterparty identities, amounts, and asset types) must not be processed or stored outside the EU. On a HashSphere private network, the institution deploys its Sphere with consensus nodes located within EU jurisdictions. The network perimeter is defined to match the regulatory boundary. Data does not leave it.

EU DATA RESIDENCY: TOKENIZED FUND SETTLEMENT ON HASHSPHERE

European asset manager — GDPR and MiFID II compliance

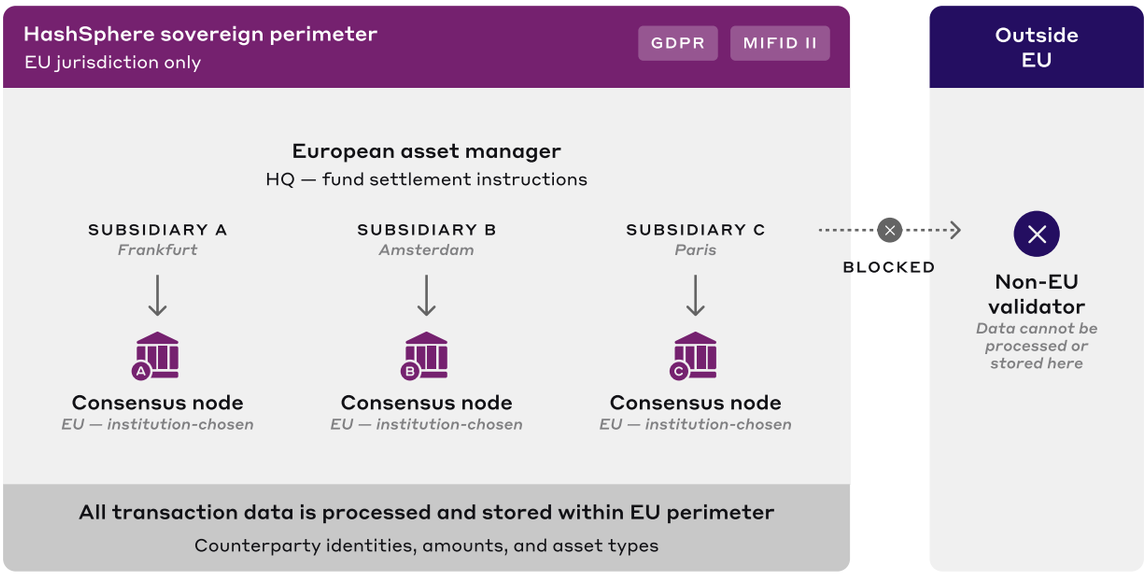

What stays inside the perimeter

- Counterparty identities
- Transaction amounts
- Settlement records
- Fund positions
- Asset types

Network perimeter is defined to match the regulatory boundary. Data does not leave it.

Data pruning: Keeping data exactly as long as required

Financial institutions typically operate under data retention requirements, often seven years, that require transaction records to be ironclad for the duration of the retention period, and then gone. Holding data beyond regulatory requirements is not a neutral choice: it is a liability, both from a discovery perspective and from a security standpoint. Old data that is not needed is data that can be breached.

Most blockchain architectures cannot accommodate this requirement. A conventional blockchain is a cryptographically linked list; truncating it breaks replay and cryptographic verifiability. The data is immutable, permanently on-chain, and cannot be pruned without undermining the integrity of everything that came after.

HashSphere data structure is different. Because transactions are not anchored to the block before them in the same way, it is possible to prune historical data without affecting the cryptographic verifiability of subsequent transactions. An institution can retain exactly what regulation requires, and remove what it does not, without compromising the integrity of the network.

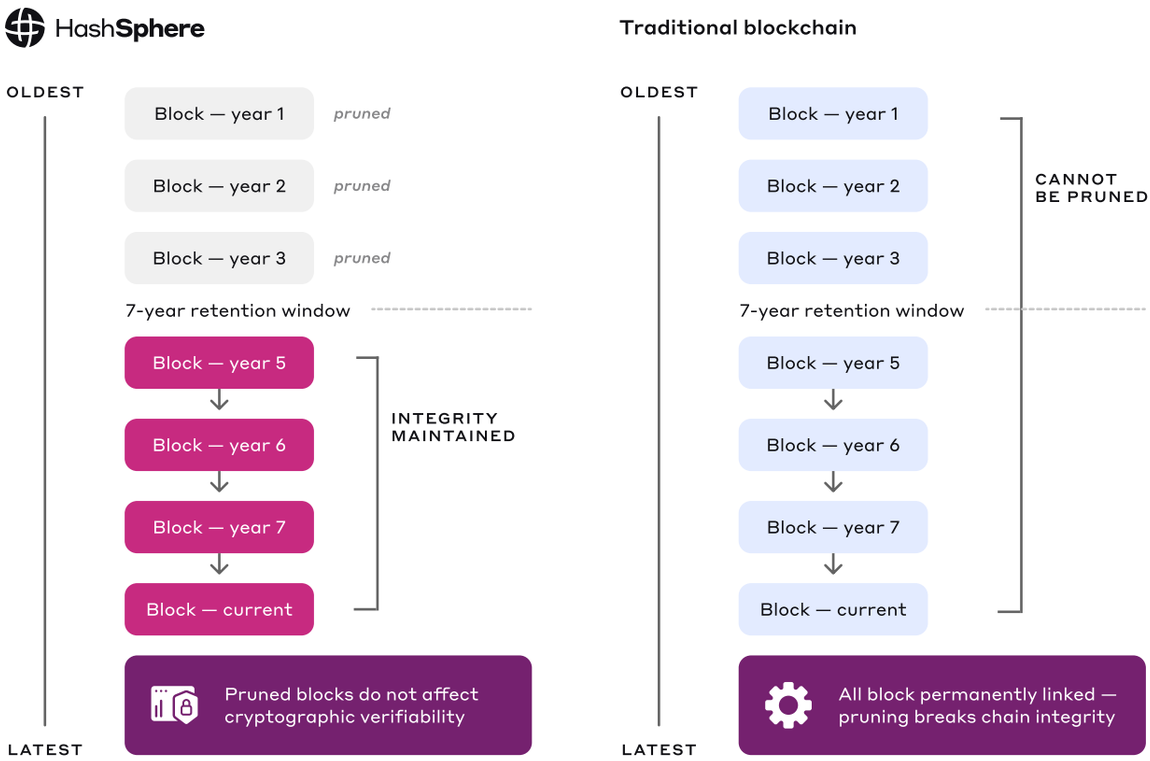

Interoperability without trapped assets: CLPR-Cross-Ledger Protocol

A private network that cannot interoperate with the broader ecosystem is not a network; it is an island. One of the most persistent criticisms of private blockchain infrastructure is the "islands of trapped assets" problem: liquidity locked inside a closed network, unable to move to where it is needed. HashSphere is designed to avoid this problem entirely. The network is private and sovereign by default, but open by choice.

CLPR (pronounced "clipper") is a bridgeless cross-ledger messaging protocol in development by Hashgraph. Transfers between ledgers happen through smart contract calls at each endpoint; no middleman, no custom bridge infrastructure, and no pooled liquidity required. Each state-proof ledger verifies the other's state using its own native proof mechanism, keeping the protocol completely trustless. When an asset issuer has opted in, applications built on CLPR can connect any state-proof public, private, or permissioned network through a single protocol layer.

When used with an application built on top of it, CLPR enables tokens, data, and messages to move between state-proof ledgers (Sphere to Sphere, Sphere to Hedera, or Sphere to other public networks) without trusted relayers or bespoke bridge infrastructure per connection. Cryptographic state proofs and threshold signature schemes (TSS) handle the transfer, preserving each chain's native consensus model and security assumptions.

The practical implication: an institution running a HashSphere network can keep everything within its perimeter by default and, with an application built on CLPR, choose to connect out when cross-network settlement or liquidity access is needed. The data stays private until the institution decides otherwise. The decision to bridge is always theirs.

CLPR ALSO DELIVERS INTEROPERABILITY FOR HASHSPHERE NETWORKS

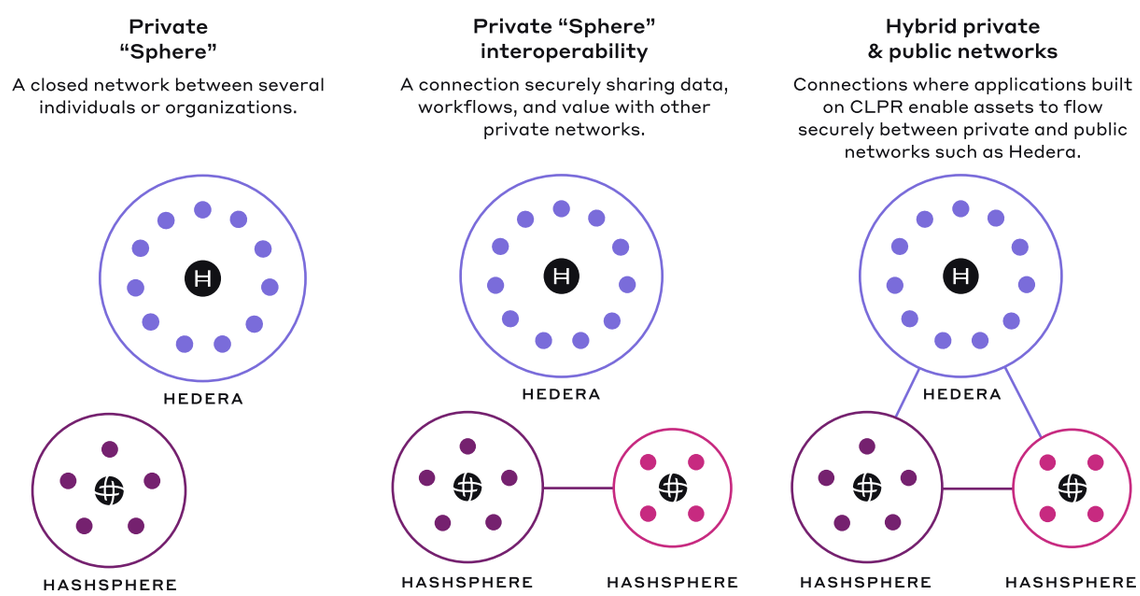

Real institutional use cases

The following deployments represent HashSphere in production. Regulated institutions using sovereign network infrastructure to solve compliance, settlement, and data control challenges that public or shared infrastructure could not address.

Multi-bank tokenized deposits: Cecabank, Spain

Five Spanish financial institutions (ABANCA, Cecabank, Ibercaja, Kutxabank, and Unicaja) completed the first multi-bank proof of concept for tokenized deposits in Spain, orchestrated by Cecabank and powered by Asseto on HashSphere. The trial demonstrated that regulated commercial bank money can be tokenized, transferred instantly between competing institutions, and settled with finality; without issuing cryptocurrency and without changing the legal or accounting nature of the underlying deposit.

HashSphere provided the controlled, permissioned environment that made multi-bank participation credible: each institution retained control of its own issuance and customer relationships, settlement was synchronized with Cecabank's account structure for immediate finality, and sensitive interbank operations never touched public infrastructure. The result was instant, round-the-clock interbank settlement within the existing deposit framework — the efficiency of programmable money without stepping outside prudential regulation.

€350,000,000,000

ASSETS UNDER CUSTODY AT CECABANK

5

PARTICIPATING INSTITUTIONS
ALIGNED WITH EUROSYSTEM TOKENIZATION INITIATIVES

Property risk data infrastructure: The Institutes RiskStream Collaborative

The Institutes RiskStream Collaborative, the insurance industry's largest not-for-profit emerging technology consortium, built an interoperable property risk and resilience portal on a hybrid model combining HashSphere private network infrastructure with the public Hedera network. The solution tokenizes commercial and residential properties, establishing a shared, verifiable source of truth for property risk data across carriers, brokers, and reinsurers.

HashSphere serves as the secure private data layer: sensitive COPE data (Construction, Occupancy, Protection, Exposure) is stored and managed in a permissioned environment, accessible only to authorized participants, while a public Hedera token serves as a universally verifiable property identifier. This architecture eliminates the duplicated data collection and verification that currently runs through five or six organizations on a single placement, replacing fragmented, siloed records with a single trusted reference point.

$1,000,000,000,000+

US PROPERTY AND CASUALTY MARKET

8 of the top 10

US P&C INSURERS IN THE RISKSTREAM ECOSYSTEM
NOT-FOR-PROFIT CONSORTIUM GOVERNANCE

Both deployments share a common architectural principle: a sovereign perimeter that keeps sensitive data under institutional control by design, paired with controlled interoperability when broader connectivity is needed. In Cecabank's case, that means a path to the public Hedera network for cross-border settlement. In RiskStream's case, it means a public trust anchor without public data exposure. The perimeter is real in both. The decision to connect beyond it belongs to the institution.

Deployment options: Getting to production

HashSphere offers flexible deployment paths depending on the institution's technical capacity, cloud infrastructure preferences, and timeline. The fastest path to production for most institutions is the managed service.

Managed service: On a Hashgraph cloud or customer cloud

Hashgraph deploys and manages the HashSphere network on the institution's behalf. Two sub-options are available:

Managed on Hashgraph infrastructure

The network is deployed on Hashgraph's cloud infrastructure. This is the fastest path to a proof of concept and the recommended starting point for institutions evaluating HashSphere. Hashgraph handles deployment, monitoring, backup, restore, and high availability configuration.

Managed on institution infrastructure

he network is deployed on the institution's own cloud infrastructure, managed by Hashgraph. The institution retains full data sovereignty; data never leaves their environment, while Hashgraph handles the operational complexity. This is the preferred model for larger institutions with existing cloud infrastructure and data residency requirements.

In both managed service configurations, Hashgraph operates under an SLA and is available to support both the infrastructure and application development. Professional services are available for institutions building applications on HashSphere.

On-Prem: full institutional control

For institutions that prefer to own all operations (particularly those with established infrastructure teams and existing runbooks for managed services), HashSphere supports self-managed deployment. The institution deploys and manages the network itself using Hashgraph's tooling, deployment guides, and monitoring endpoints. Hashgraph provides reactive support and is available to assist when needed.

HashSphere currently offers the following deployment paths for cloud-hosted infrastructure, where the network runs in a cloud environment with no physical hardware shipped to the customer:

Managed Service (available now)

The network is deployed in the institution's own AWS or GCP account and managed by Hashgraph. The institution retains full data sovereignty while Hashgraph handles operational complexity.

On-Prem (Self-managed) (Alpha, Q4 2026)

The network is deployed in the institution's own AWS or GCP account and managed entirely by the institution, using Hashgraph's tooling and monitoring endpoints with reactive support available.

Physical hardware deployment options (including AWS Outpost, Google Distributed Cloud, and bare metal server deployments in the institution's own data center) are not currently on the roadmap.

	PUBLIC SUBNETS	CENTRAL ROUTING MODELS	SELF MANAGED
Who deploys	Hashgraph	Hashgraph	Institution
Who manages	Hashgraph	Hashgraph	Institution
Data location	Hashgraph cloud	Institution cloud	Institution cloud
SLA	Yes — included	Yes — included	Yes — reactive support
Recommended for	PoC and evaluation; fastest path to production	Institutions with data residency requirements and existing cloud infra	Institutions owning full operations with established infrastructure teams

Bare metal deployment support on roadmap for self-managed.

How HashSphere compares

The private enterprise blockchain market offers several approaches to network design. The table below is not an exhaustive competitive analysis; it is a framework for understanding the structural trade-offs that different designs require.

	MANAGED (HASHGRAPH CLOUD)	MANAGED (INSTITUTION CLOUD)	HashSphere
Data transits public internet	✓	Depends on deployment	Institution
Defined sovereign perimeter	Partial	✗	✓
Privacy mechanism	ZKP token-level confidential transfers	Policy-based access control; contractual	zkSNARKs + encrypted balances; cryptographic
Central coordinator sees all data	✗	✓	✗
In-jurisdiction validator choice	Limited by validator availability	Vendor-dependent	Institution-defined
Open-source No lock-in	Partial	✗ Proprietary	LFDT / Project Hiero
Data pruning capability	✗	Vendor-dependent	✓
Native interoperability	Chain-specific	Limited	✓ Via CLPR
Selective disclosure	Auditor key on mint function (all-or-nothing)	Notary has full plaintext access	Scoped commitment salt verification
Key caveat	Data may route outside jurisdiction; ZK proof vulnerabilities (2025)	Operator sees all data; privacy is contractual, not methematical	Submitted visible; CLPR-bridged data is immutable

Not an exhaustive competitive analysis a framework for understanding structural trade-offs.

Conclusion

The institutions that will define the next era of financial market infrastructure are not looking for a private network in the abstract sense. They are looking for a network they actually control: where data stays within their perimeter by design, where compliance with data residency frameworks is architectural rather than contractual, and where the decision to connect to the broader ecosystem is theirs to make.

That is a harder set of requirements to satisfy than simply deploying a permissioned chain. Most approaches in the market today satisfy some of them. Subnet-based designs give institutions a permissioned environment but leave data traversing infrastructure they do not control. And even where content is encrypted, metadata including participant identities, timing, and transaction volume can leak to validators outside the perimeter. Central-routing models give institutions privacy from each other but require trust in an operator who can see everything. These private, permissioned network models also risk locking users into one vendor and denying access to new sources of liquidity outside of the 'network of networks'

HashSphere's independent network design is different in one fundamental respect: the perimeter is real. Data does not leave it unless the institution decides to send it. The sovereign perimeter is not a configuration option; it is the architecture.

The hashgraph consensus algorithm provides mathematically guaranteed fair ordering within that perimeter. CLPR–cross-ledger protocol when used with an application, eliminates the 'islands of trapped assets' by providing controlled, cryptographic interoperability when the institution chooses to access other private or public networks. Open source foundations under Linux Foundation Decentralized Trust eliminate vendor lock-in. And data pruning capability means institutions can meet retention requirements without carrying data liabilities indefinitely.

Privacy and interoperability on a true DLT-powered network does not have to be a trade-off. Neither does data sovereignty and network openness. HashSphere is built on that premise.

NEXT IN THE BEHIND THE LEDGER SERIES

Flexible Validator and Governance Models: How HashSphere enables institutions to define who runs the network, what they can see, and how consensus is reached, without delegating those decisions to a third party.

Hashgraph

About Hashgraph

Hashgraph is an enterprise software firm powering the next generation of connected digital markets. The company develops solutions for financial services that leverage hashgraph technology and provides ongoing support to the Hedera public network. Through CLPR (cross-ledger interoperability), private network infrastructure, and tokenization platforms, Hashgraph enables institutions to issue, trade, and manage assets seamlessly across systems, helping to drive a more connected digital economy.

For more information, visit hashgraph.com