



[Introduction](#)

[Core Technology & Platform Architecture](#)

[Core Technology](#)

[Overview](#)

[Substrate Framework](#)

[Integration with Ethereum](#)

[Additional Features](#)

[Economic Model](#)

[System Token](#)

[Payment for Execution](#)

[Storage Fees and Incentives](#)

[Governance](#)

[Platform Architecture](#)

[Platform Components](#)

[DataHaven Runtime](#)

[Consensus and Finality](#)

[Native Bridge to Ethereum](#)

[AVS Support \(EigenLayer\)](#)

[Ethereum Compatibility](#)

[Bridged & Native Assets](#)

[DataHaven Storage](#)

[Storage Runtime Pallets](#)

[Main Storage Providers \(MSPs\)](#)

[Backup Storage Providers \(BSPs\)](#)

[DataHaven Client](#)

[EigenLayer and the DataHaven AVS](#)

[EigenLayer](#)

[DataHaven AVS](#)

[Operators & Rewards](#)

[DataHaven-ETH Bridge](#)

[Bridge Message execution](#)

[DataHaven-ETH Bridge Relayers](#)

[Roadmap & Milestones](#)

[Use Cases](#)

[Authors](#)



Introduction

We live in a world where data isn't just important—it's everything. It's the heartbeat of our digital lives, from the personal identities we protect to the AI models that will define the future. Every interaction, every advancement, is built on data. And yet, the systems that hold it all together—our storage solutions—are failing us.

Today, we are forced to choose between the flaws of the past and the promises of the future. Web2 is centralized and vulnerable, a constant target for censorship, control, and security breaches. Web3, though groundbreaking, often leaves us grappling with insecurity, rigidity, and a lack of true interoperability. We've long known that this is not enough.

It's time for a new approach—one that meets the needs of an AI-first world, where DeFi and the rise of tokenized assets demand storage that is as innovative, secure, and flexible as the technologies they power. One that is tamper-proof, verifiable, and censorship-resistant. One that merges the strengths of decentralization with enterprise-grade security, verifiability, and seamless interoperability for the future of AI, Web3, and tokenized assets. That solution is DataHaven.

DataHaven provides an unshakable foundation for the next wave of digital transformation. Deployed as an Autonomous Verifiable Service (AVS) secured by EigenLayer's re-staking protocol, DataHaven ensures data is:

- **Tamper-proof**, protected by cryptographic proofs that guarantee integrity
- **Verifiable**, allowing AI-generated models and logs to remain auditable and accountable
- **Censorship-resistant**, ensuring no single entity controls access
- **AI-optimized**, designed for federated learning, AI agents, and machine learning workloads

With native EVM compatibility, DataHaven enables seamless smart contract integration, empowering decentralized applications (dApps) to build on a trustless and future-proof storage layer. As AI-driven innovation accelerates, DeFi expands, and tokenized assets rise, the need for sovereign, verifiable, and interoperable storage has never been greater. DataHaven is not just meeting this need—it is driving the revolution.

Why DataHaven?

Blockchains excel at creating public, verifiable data and representing value. However, they are inherently inefficient when it comes to storing large amounts of data, as all information must be replicated across all nodes in the network. This limitation has driven developers to rely on off-chain storage solutions for data-intensive applications, which introduces centralized trust assumptions that compromise the security and censorship resistance that blockchains are



designed to provide. Moreover, off-chain data is vulnerable to unavailability or loss, yet developers and users have largely accepted this weakened security as the norm.

In response to these limitations, projects such as Filecoin and Arweave have emerged to enable decentralized storage while maintaining certain Web3 principles. While these platforms demonstrate the viability of blockchain-based storage, they function as standalone networks. This isolation creates substantial challenges, as the majority of developer activity, asset issuance, and user engagement is concentrated on smart contract platforms like Ethereum, Solana, and Layer 2 blockchains.

For developers building on these smart contract platforms, integrating Web3 storage solutions from separate blockchains results in a fragmented architecture. This not only increases the security attack surface and trust assumptions—as they become the union of both systems—but also leads to a poor user experience due to the complexities of managing accounts and transactions across multiple independent blockchains.

DataHaven was specifically designed to address these challenges by providing a specialized blockchain for efficient decentralized storage while remaining natively integrated within the Ethereum ecosystem. Unlike standalone storage solutions, DataHaven is secured by Ethereum through the EigenLayer re-staking protocol. This approach significantly minimizes additional security and trust assumptions for developers already building within Ethereum, ensuring storage security is directly tied to Ethereum’s economic and cryptographic guarantees.

Moreover, DataHaven features an integral, native, and trust-minimized bridge, a critical distinction absent in competitors, which depend on third-party bridges with their own trust requirements. This native bridge is paramount, as it enables DataHaven storage to be seamlessly accessed and integrated into Ethereum-based applications on both the mainnet and Layer 2 networks.

As a result, DataHaven stands as the first and *only* specialized storage solution that is natively part of the Ethereum ecosystem. This allows developers to incorporate robust, decentralized storage into their applications without the usual trade-offs, ensuring seamless interoperability, high security, and a streamlined user experience.

Core Technology & Platform Architecture

Core Technology

Overview

DataHaven is a decentralized storage solution built on a Proof-of-Stake, layer 1 blockchain secured by Eigenlayer's re-staking protocol. It offers an EVM-compatible environment with fast



finality and integrated data storage. The network consists of “storage providers” managing off-chain data storage and an EVM blockchain coordinating storage requests, tracking ownership, and verifying data. The blockchain secures file storage by managing and verifying proofs from storage providers.

Substrate Framework

DataHaven is built on Substrate, a modular framework in Rust with a rich library of reusable components (pallets). Substrate’s extensibility and active development community enable easy customization. By leveraging Substrate, DataHaven gains access to battle-tested pallets used in networks like Moonbeam and Polkadot, offering features like cross-chain interactions, trustless bridging, governance, and Proof-of-Stake with fast finality. This flexibility, efficiency, and interoperability make Substrate ideal for Web3 innovation.

Integration with Ethereum

DataHaven is fully EVM compatible, supporting popular wallets and tools and providing a familiar environment for EVM developers. Unlike other Web3 storage solutions that operate separately from computation blockchains, causing inefficiencies, DataHaven enables seamless interactions between smart contracts and stored data.

With DataHaven, dApps can interact with the storage layer, managing large datasets, file uploads, access permissions, and storage proofs for data integrity verification. For example, an NFT collection could store digital assets and verify authenticity on-chain. DataHaven also supports secure, collaborative workflows for projects requiring shared access to critical documents, such as Real-World Asset (RWA) initiatives, ensuring decentralized accessibility and control.

Secured by Eigenlayer’s re-staking protocol, DataHaven relies on Ethereum for economic security, unlike standalone storage blockchains that have weaker security. This ensures multi-chain applications are backed by Ethereum’s stake and offers robust protection for assets and data.

Additional Features

- **Interoperability:** DataHaven is EVM-compatible and includes a native bridge to Ethereum and Moonbeam, plus support for General Message Passing (GMP) protocols for secure cross-chain storage interactions.
- **Ease-of-Use:** Specialized pre-compiles enhance user and developer experiences, such as a "CallPermit" pre-compile for gasless transactions and a "Batch" pre-compile to reduce transaction approvals and costs.
- **On-Chain Governance:** Powered by Substrate’s OpenGov pallets, DataHaven enables forkless upgrades and parameter changes via community voting. It also supports decentralized handling of legal requests, ensuring compliance while preserving decentralization.



Economic Model

System Token

DataHaven's economic model is centered around a system token that will be considered native to the network. This token is integral to the network's operation and value accrual and will serve several key functions, as discussed below.

Staking to Secure DataHaven

Although the re-staking of Ethereum will form the bedrock of DataHaven's crypto-economic security, the EigenLayer protocol allows for other assets to be staked/re-staked to help secure an AVS. Holders of the system token will be able to delegate stake to operators in order to participate in securing DataHaven and earn rewards. (It should be noted that Moonbeam's native token GLMR and/or GLMR LSTs will also be able to be staked/re-staked in a similar manner.)

Payment for Execution

The system token will be the gas token for the network used to pay for computational resources and transaction fees. As with a typical EVM environment, the cost of gas will be denominated in the system token. All transactions and smart contract executions will require a fee, which will be paid in the system token. A dynamic fee mechanism will be used to prevent spam and ensure efficient use of network resources. Priority fees will be included in rewards to operators (validators) while the base fee will be split into a portion directed to treasury and a portion to be burned, configurable via governance.

Storage Fees and Incentives

Storage fees may be paid using the system token. Conversely, payments to storage network operators providing storage services may be denominated in the system token. However, a goal of the project is to have predictable storage costs and payments to operators, in dollar terms, per GB per unit of time. In order for storage providers to participate in the network and earn fees for storing data, they will be required to hold collateral denominated in the system token, which may be slashed if they lose user data, ensuring a commitment to data preservation.

Governance

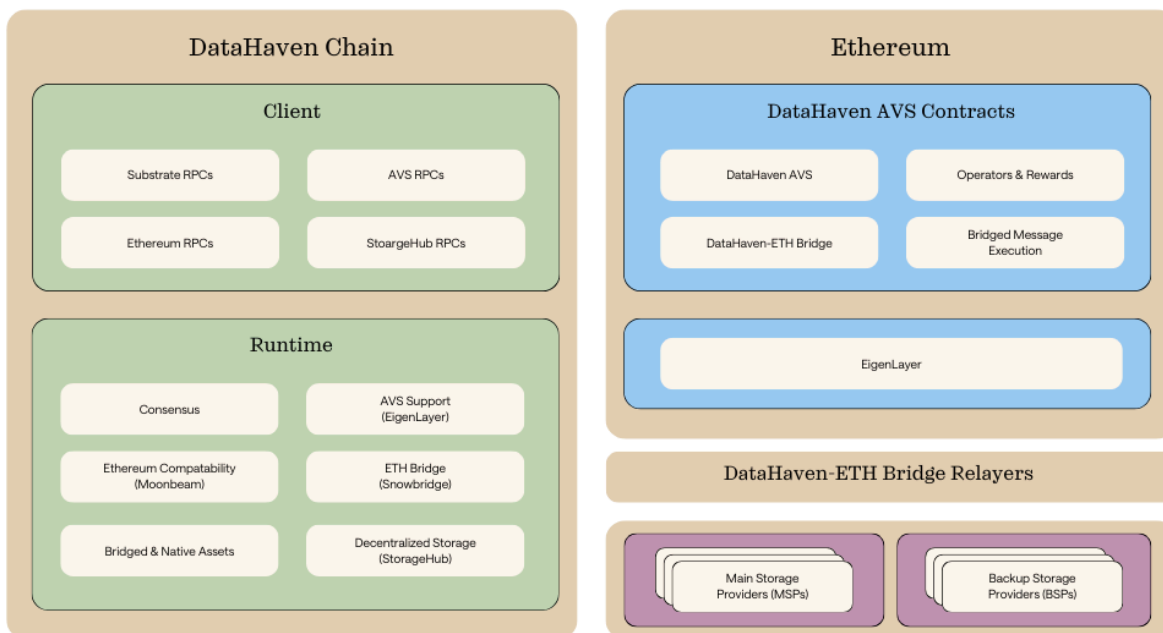
The system token is central to the network's governance, empowering stakeholders to shape decisions. Token holders can vote on proposals and referendums, with voting power based on the amount and lock period of tokens. Token holders govern data storage fees, council composition, and fund allocation through on-chain voting. Additionally, they can delegate voting power for decentralized decision-making.



Platform Architecture

Platform Components

Platform Components



DataHaven's foundational architecture as a Layer 1 blockchain, built using the [Substrate framework](#), enables deep-level customizations at the core protocol level that go beyond the capabilities of smart contracts. These inherent optimizations are crucial for effectively serving DataHaven's intended use cases. The following sections detail the additions made to the runtime and client to add full Ethereum compatibility and decentralized storage capabilities, and for the chain to operate as an EigenLayer AVS and therefore benefit from Ethereum's shared security.

DataHaven Runtime

Aside from managing accounts and balances, the DataHaven runtime will offer a set of powerful capabilities that are commonly included in Substrate-based chains, including support for proxy accounts (proxy pallet), account identities (identity pallet), transaction batching (utility pallet), multi-signature operations (multisig pallet), and scheduled and delayed operations (scheduler pallet).



Consensus and Finality

For consensus, DataHaven will use a typical [hybrid model](#) composed of a block production mechanism and a finality gadget for probabilistic and provable finality. The target block time will be six seconds. DataHaven will use the [BABE](#) slot-based block production mechanism based on a known set of validators with a randomized slot assignment. The block author will be selected pseudo-randomly out of an active set of validators provided to the block authoring mechanism by the EigenLayer AVS integration (based on registered AVS operators on EigenLayer). DataHaven will use the [GRANDPA](#) gadget for block finality.

Native Bridge to Ethereum

DataHaven's chain runtime will feature a trust-minimized native bridge to Ethereum—an extremely rare and highly secure approach in the blockchain ecosystem. Unlike most bridges, which rely on significant trust assumptions, DataHaven's bridge is built directly into the protocol, eliminating the need for additional external trust layers when transferring assets or data between Ethereum and DataHaven. This design is made possible by leveraging a custom Layer 1 with verifiable proofs via [BEEFY](#), ensuring seamless and secure bridging, much like the native bridge of a rollup.

To achieve this, DataHaven will integrate BEEFY alongside GRANDPA finality, allowing its Ethereum-based AVS to track finalized headers efficiently. It will also support [EigenLayer](#) shared security, token transfers, and cross-chain smart contract operations. Built on Snowbridge, a decentralized bridge between Polkadot and Ethereum, the DataHaven Bridge will reuse critical components (e.g., `snowbridge-pallet-ethereum-client`) while customizing others (e.g., system and queue pallets) to fit its needs.

Additionally, DataHaven will enable generic message passing for cross-chain interoperability, utilizing Polkadot SDK components (e.g., [pallet-bridge-messages](#)) and developing new pallets to support arbitrary message formats beyond XCM. The team will fork Snowbridge's pallets for enhanced functionality tailored to DataHaven's architecture.

AVS Support (EigenLayer)

The DataHaven chain runtime will integrate capabilities to make it operate under the shared security model provided by EigenLayer. More concretely, the active set of validators will be provided by the DataHaven AVS, and the validator rewards will be sent to the DataHaven AVS. The active list of DataHaven AVS operators will be cached in a new, special purpose Substrate pallet. The list of validators will be updated through messages transiting from Ethereum to DataHaven via the native bridge. This pallet will also track and report validator rewards based on authored blocks to the DataHaven AVS on Ethereum—again, with messages transiting over the native bridge.



Ethereum Compatibility

DataHaven will provide Ethereum compatibility by integrating core features from [Moonbeam](#), including:

- [EVM compatibility](#) (via [Frontier](#)) allowing Ethereum smart contracts written in Solidity, or other languages targeting the EVM, to run unmodified.
- [Ethereum-style Accounts](#) (H160) thanks to [Moonbeam's unified account system](#).
- Moonbeam [precompiled contracts](#) exposing the chain's runtime core features to the EVM layer for seamless integration in smart contracts.
- Moonbeam [Transaction fees](#) model.

It is worth noting that Frontier provides a high degree of EVM compatibility, with a proven track record deployed on Moonbeam boasting hundreds of deployed applications/protocols and over 56M EVM transactions to date.

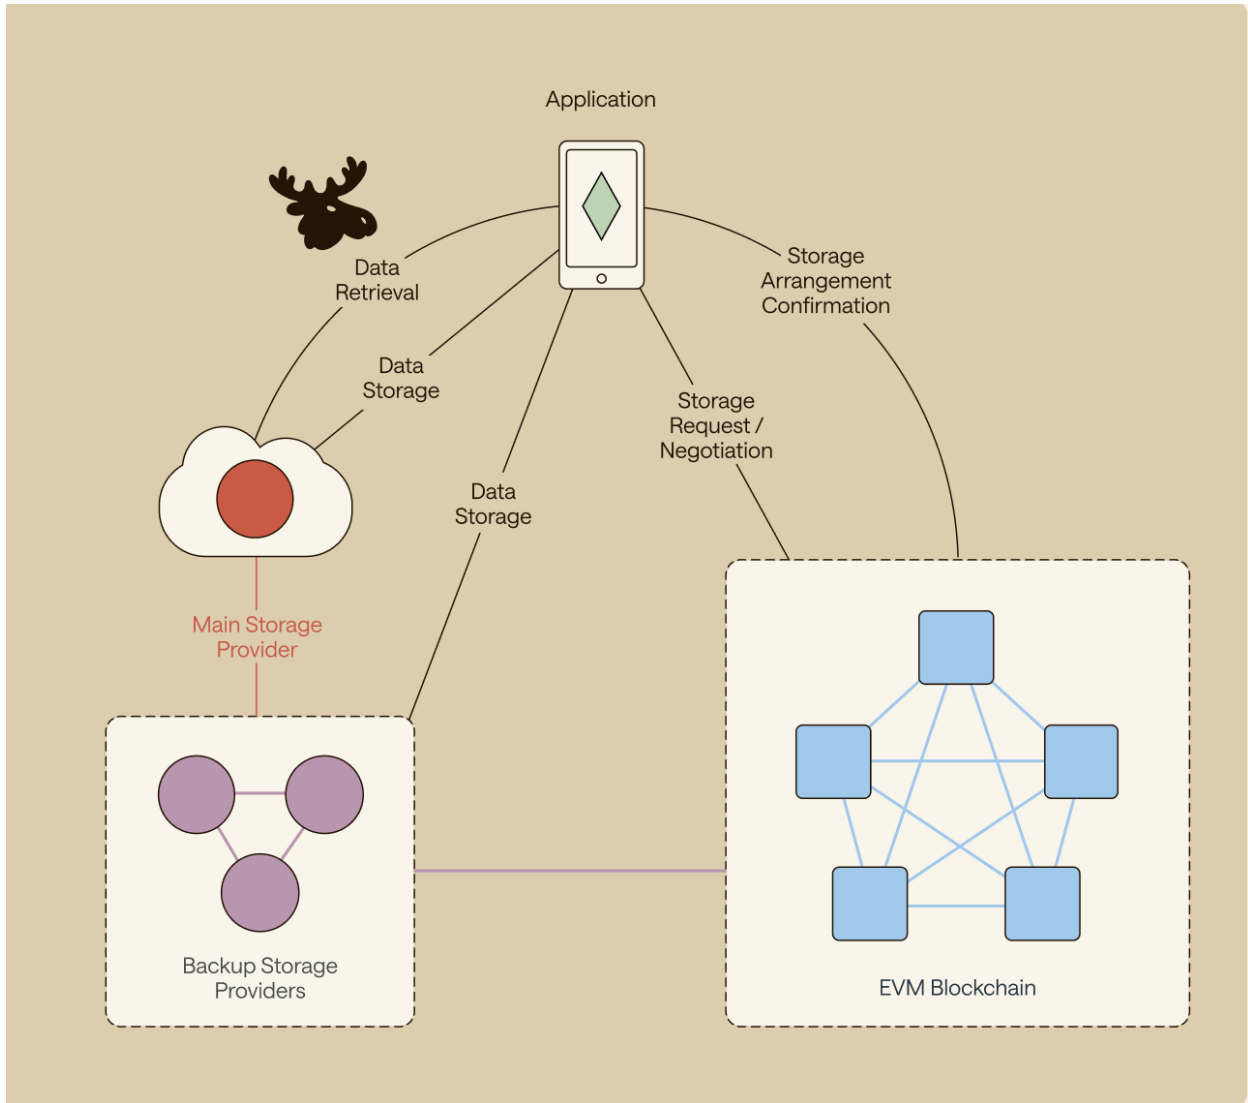
Bridged & Native Assets

DataHaven will support foreign (bridged) assets (i.e. wrapped / derived versions of native ETH or ERC20 tokens on Ethereum) transferred to the DataHaven chain via the native bridge. These will seamlessly interoperate with smart contracts deployed on DataHaven. DataHaven may also support native assets exposed to the EVM layer as ERC20, similarly to assets and foreign assets on Moonbeam.

DataHaven Storage

DataHaven will come built-in with decentralized storage capabilities by integrating features of the [StorageHub](#) project—a storage-optimized collection of Substrate pallets focusing on efficient and decentralized data storage. This addresses the need for scalable, secure, and decentralized storage enabling applications to seamlessly access, utilize, and manage storage even in cross-chain deployments.

The DataHaven execution environment will work in conjunction with Main Storage Providers (MSPs) and Backup Storage Providers (BSPs) to efficiently store files and data in a reliable and decentralized manner. End users and applications will benefit from these features by directly interacting with the DataHaven runtime, RPC APIs and via GMP protocols invoked from smart contracts deployed on other chains.



DataHaven ensures efficient and reliable storage using Merkle cryptographic structures, including **Merkle Patricia Forests** to manage on-chain storage with cryptographic summaries instead of per-file records. It also employs **Merkled Files for Storage Proofs**, challenging providers to prove they are storing files by verifying random chunks.

Storage Runtime Pallets

Several StorageHub Pallets (modules) will form part of the DataHaven runtime in order to provide the coordination and management of storage operations. The **File System Pallet** allocates files to storage providers using Merkle Patricia Forests, managing storage requests and deletions with cryptographic methods to minimize on-chain storage. The **Storage Providers Pallet** handles Main Storage Providers (MSPs) and Backup Storage Providers (BSPs), tracking MSP data retrieval and BSP redundancy, with token collateral and capacity-based price adjustments. The **Storage Proofs Pallet** generates and verifies storage



proof challenges, enabling storage providers to charge for services while distributing challenges based on stored data.

Main Storage Providers (MSPs)

DataHaven's storage enables flexible interactions between users or applications and MSPs. MSPs offer data retrieval services, with multiple customizable implementations at different costs. DataHaven allows AI agents and developers to discover MSPs based on services and value, and ensures transparency by recording agreements. MSPs earn revenue from storage services, with users paying for file storage, thereby incentivizing reliable solutions. To become a provider, MSPs must hold collateral, which can be slashed if data is lost. MSPs manage file "buckets," with each linked to the provider and secured by a Merkle Patricia Forest stored on-chain.

Backup Storage Providers (BSPs)

BSPs ensure data reliability and redundancy in a decentralized network, backing up data to keep it available even if an MSP fails. BSPs must hold collateral that can be slashed if data is lost, and they have lower operational costs since they only perform occasional peer-to-peer retrieval. Each file stored by a BSP is split into chunks, "merkelized," and linked to a file key, with the root stored on-chain as proof of storage. BSPs are periodically challenged to submit proof of storage, with challenge frequency based on storage size. A Fair Distribution mechanism prevents front-running, allowing all BSPs the chance to volunteer for new storage requests.

DataHaven Client

- **Generic capabilities:** The DataHaven client will come built-in with the typical JSON-RPC APIs that can be found in most Substrate-based chains (based on the modules—aka pallets—included in the chain's runtime).
- **Ethereum RPCs:** The DataHaven client will offer full [JSON-RPC compatibility with Ethereum](#), and thus compatibility with existing dApps, wallets, and development tools and libraries.
- **AVS RPCs:** The DataHaven client will expose specific JSON-RPC APIs to facilitate operations related to the EigenLayer AVS operators.
- **StorageHub RPCs:** The DataHaven client will expose the standard StorageHub JSON-RPC APIs to facilitate file, metadata, and storage proof-related operations.

EigenLayer and the DataHaven AVS

EigenLayer

EigenLayer is a protocol built on Ethereum that introduces restaking, a new primitive for Web3 builders that provides a "marketplace for trust," bringing together restakers, operators, and AVSs. It allows users to stake assets such as Native ETH, Liquid Staking Tokens (LSTs), the EIGEN token, or any ERC20 token into EigenLayer smart contracts, thereby extending Ethereum's cryptoeconomic security to additional applications on the network. For more background information on the EigenLayer protocol, consult the official [documentation](#).



DataHaven AVS

DataHaven will function as an [EigenLayer AVS](#), benefiting from Ethereum shared security and providing seamless interoperability between contracts and assets on DataHaven and the Ethereum network. Generally speaking, an AVS is composed of on-chain contracts (on Ethereum) for validation and an off-chain network of operators. Operators execute the service on behalf of the AVS and then post evidence of their execution on-chain to the AVS contracts. If the operators perform the tasks properly, the AVS can distribute rewards; if operators perform poorly or with malicious intent, their delegate stake can be slashed by the AVS, and the operator can be removed from the operator set.

In the case of DataHaven, this implies the following: DataHaven is registered on EigenLayer as an AVS, and operators can opt-in to provide validation services (that is, act as validators by running the DataHaven node binary) for the DataHaven network.

On the Ethereum side, the DataHaven project will deploy several contracts to operate as an AVS:

- A **ServiceManager** contract, implementing the base [ServiceManager](#) interface
- An **OperatorManager** contract, handling registration of operators and maintaining a list of active / selected operators
- A **RewardsProxy** contract, mediating the rewards (or points) calculated on the DataHaven chain and submitting the reward list to the **EigenLayer** [RewardCoordinator](#) contract
- A **ServiceUpgrader** contract, to support upgrades of the DataHaven AVS contracts
- A **BeefyVerifier** contract, to verify BEEFY commitments from the DataHaven chain (using the Snowbridge on-chain Beefy light client), periodically relayed by the DataHaven-ETH bridge relayers
- A root **DataHaven AVS** contract, coordinating the operations of the DataHaven AVS and coordinating the communication with the DataHaven chain via the DataHaven Bridge

Operators & Rewards

The DataHaven AVS contract will maintain a list of selected operators and emit events for updates. This list will be relayed via the DataHaven Bridge to the DataHaven chain to update the validator set. In return, DataHaven will send the current reward shares of active validators through the bridge to the AVS contract and then to the EigenLayer RewardCoordinator to issue staking and delegating rewards.

DataHaven Bridge

DataHaven will integrate [Snowbridge](#) as a DataHaven-Ethereum bridge, reusing existing contracts and infrastructure, modified to fit the needs of the project. The DataHaven Bridge includes an on-chain Beefy light client (Solidity) contract, a Gateway contract for relaying and routing messages, and other supporting contracts (types, primitives, proxy accounts, etc.). These will be reviewed and modified, if needed, to support the DataHaven AVS implementation.



Bridge Message execution

The DataHaven-ETH bridge will need the following types of messages:

- DataHaven-specific operational messages, e.g. update of AVS operators list (ETH → DataHaven); update of the reward shares / list (DataHaven → ETH)
- Well-known message types, e.g. cross-chain asset / token transfer (burn & mint)
- Arbitrary messages: execution of smart contracts functions on DataHaven (call from Ethereum) or on Ethereum (call from DataHaven)

While the DataHaven Bridge is relying on the Polkadot XCM message format for all messages, it will be extended to support additional or arbitrary message formats.

DataHaven Bridge Relayers

As described above, DataHaven will have its own decentralized, trustless bridge solution to Ethereum. The relayers are off-chain, untrusted services watching two Blockchain networks, and relaying messages across them. The DataHaven project will integrate modified versions of the [existing Snowbridge relayers](#), namely:

- DataHaven → Ethereum
 - BEEFY relay: relays signed BEEFY commitments and proofs from the DataHaven chain to the BEEFY light client contract on Ethereum
 - Message relay: relays message commitments and proofs from the DataHaven chain to inbound channel contracts on Ethereum
- Ethereum → DataHaven
 - Header relay: relays Ethereum Beacon chain headers, Execution chain headers and Sync Committees to the Ethereum light client pallet on the DataHaven chain
 - Message relay: relays messages emitted by outbound channel contracts on Ethereum to the DataHaven chain

Roadmap & Milestones

Key achievements to date:

- Storage Hub (core data engine) 1.0 released
- An independent chain based on the Substrate technology instantiated
- Integration with AVS powered by EigenLayer is progressing

2025 goals:

- Private StageNet with selected partners (Q2/2025)
- Public TestNet with more supported apps and available for live testing by users (Q2/2025)
- Mainnet launch with core capabilities enabled (Q3/2025)
 - Storage integrations via APIs and SmartContracts
 - Full EVM support on the DataHaven platform



- Native bridging with both Moonbeam and Ethereum
- Flexible payment models
- Dedicated dApp with a reference implementation of user-facing features (Q3/2025)
 - On-chain governance
 - Storage management UI
 - Account overview including payments
 - Support for common wallets

Future outlook:

- Application-centric features
 - Different data-access models and MSP templates focused on AI, RWA, payments, etc.
 - Advanced permissions and authorization schemes
 - Host-side cryptography
 - Data retention schemes
 - GDPR and privacy-compliant policies for nodes
- Performance improvements
 - Mass file operations/batch processing
 - Systems optimization for data availability, migration paths

Use Cases

DataHaven is AI-first Web3 infrastructure, purpose-built to power the next generation of decentralized AI applications. While it supports a broad range of industries—including real-world assets, web hosting, NFTs, and verifiable backup solutions—its AI use cases stand at the forefront. Below are a few key examples. For a complete list of use cases, visit [this page](#).

- **Storage and Sharing of AI training models:** A developer distributing an open-source AI model, such as LLaMA, GPT, or Stable Diffusion, uploads model weights to DataHaven, ensuring that downstream users can verify they are using tamper-proof version. Enterprises deploying AI models in regulated industries (e.g., finance, healthcare, defense) can guarantee compliance by proving their models have not been altered.
- **Verifiable AI-Generated Code & Software Artifacts:** A self-improving AI agent managing a blockchain protocol autonomously generates a new smart contract module for handling cross-chain transactions. By storing the contract code and verification tests on DataHaven, blockchain validators can cryptographically verify that the deployed contract is identical to its tested version, ensuring trust and security in the network.
- **Guaranteeing privacy of AI Agent Data:** A personal AI assistant stores encrypted conversation history, financial planning data, and medical records on DataHaven. Users selectively grant access to different AI models without revealing raw data, enabling secure, AI-driven health monitoring, investment tracking, or legal document analysis.



- **DeFi and Financial Instruments:** A tokenized bond issuance platform that allows individuals to invest in government-backed bonds previously unavailable to them uses DataHaven to ensure secure access to compliance records and transaction history.
- **Commodity-Backed Tokens:** A gold-backed stablecoin project stores audit reports on DataHaven to verify reserves, ensuring token holders that each unit is backed by real gold.
- **Real Estate Tokenization:** A real estate tokenization platform that allows investors to buy fractional shares of properties uses DataHaven to securely store ownership records and compliance documents, ensuring trust and transparency.
- **Intellectual Property & Patent Tokenization:** A biotech startup tokenizes its patents, enabling researchers or investors to hold fractional stakes in its innovations. DataHaven ensures patent ownership and licensing agreements are accessible, even if the tokenization platform shuts down.
- **Web Hosting:** A DeFi lending platform uses DataHaven to ensure users can always access the service, even if the original web infrastructure is compromised.
- **NFTs:** A high-value NFT collection is stored on DataHaven to guarantee that holders can still access and verify their assets, even if the marketplace or issuer shuts down.

Authors

This litepaper was made possible by the expertise and insights of the following individuals from the Moonbeam Foundation, along with contributions from Moonsong Labs:

- John Bridgewater, Advisor
- Alex DiNunzio, Head of Growth & Strategy
- Piotr Dziubecki, Head of Product
- Aaron Evans, Head of Operations
- Prasanna Ketheeswaran, Chief Financial Officer
- Ryan Levy, Global Head of Business Development
- Sicco Naets, Head of Ecosystem