BRAIN++ AI FACTORY

Introducing

CHAIN++

A permissioned data governance chain

Proof of Concept · In active development toward production release

Sofia Tech Park · June 2026 · v1.0

Legal Disclaimer

This whitepaper describes the BRAIN++ AI Factory federated chain (CHAIN++), currently at Proof-of-Concept (PoC) stage. Active development is continuing toward a first production release. Nothing in this document describes a production-ready system.

Co-funded by the European Union

EuroHPC Joint Undertaking

Table of Contents

What Type Of Problems CHAIN++ Aims To Solve

CHAIN++ is built to solve one problem: making data sharing between organisations trustworthy enough and acceptable for legal teams. Here are the specific situations where it makes a difference.

Governed clinical data exchange

Healthcare · GDPR · EHDS

Problem: A hospital wants to share patient records with a research consortium for AI training. Its legal team requires provable evidence of what was shared, under what terms, and that it was not used for anything else. Internal database logs are not accepted as evidence.

CHAIN++ provides: Both parties sign an encrypted contract on CHAIN++. The hospital's Data Guardian checks every dataset before transfer. Every transfer is written to a tamper-evident chain. The hospital can produce a solid proof for a regulator at any time, without revealing the commercial terms.

Status: In the PoC.

Multi-party AI model training

Financial Services · AML · Data Sovereignty

Problem: Three banks want to train a shared fraud-detection model. None will share raw transaction data with the others. Each needs proof its data was used only for the agreed purpose and was not retained.

CHAIN++ provides: Each bank signs a separate contract with the AI factory specifying permitted categories, purpose, and expiry. Every access event is independently verifiable by each bank. The final model is shared; no raw data crosses organisational boundaries.

Status: Multi-party contracts with separate terms per participant: in the PoC.

Regulatory compliance audit Regulatory · GDPR Article 30 · AI Act

Problem: A supervisory authority asks a data controller to prove that a dataset was processed only for the stated purpose and within the agreed geographic boundary. The controller's internal systems cannot provide evidence the authority will accept.

CHAIN++ provides: Because every processing event was written to CHAIN++ at the time it occurred, the controller produces a Merkle proof, a short mathematical certificate independently verifiable without trusting the controller's systems. Contract terms remain encrypted.

Status: Tamper-evident audit trail and proof generation: in the PoC.

Cross-factory dataset and model licensing Life Sciences · IP Governance · Cross-Organistion

Problem: A pharmaceutical company at Factory A wants to license a synthetic dataset to a biotech firm at Factory B. The two use different AI factories. Neither factory should be able to read the content or alter the licence record.

CHAIN++ provides: Both companies sign a direct contract on the federated CHAIN++ without either factory as intermediary. Access stops automatically when the licence expires. Both parties verify the contract and transfer record independently.

Status: Cross-factory direct contracts: in the PoC.

AI-driven business automation with governed data Insurance · AI Act · Data Minimisation

Problem: An insurer wants to use an external AI service to assess claims, but cannot give it unlimited access to policyholder data. Legal requires a record of every data item accessed and for which decision.

CHAIN++ provides: The insurer signs a contract specifying exactly which data categories the AI service may access. The Data Guardian checks every dataset before it leaves the insurer's environment. Every access event is permanently logged and auditable.

Status: Contract governance and audit logging: in the PoC. Automated decision-record linking: roadmap.

Who Benefits Most

- **Healthcare and life sciences.** GDPR, EHDS, HIPAA compliance for secondary use of patient data, clinical trial data sharing, pharmaceutical research collaboration.
- **Financial services.** AML data sharing between banks, fraud-detection model training, transaction analytics under data residency constraints.
- **Public sector and research.** Cross-institution research data sharing, government AI procurement under the EU AI Act.
- **AI infrastructure providers.** Any AI factory that wants to serve regulated industries and needs a verifiable governance layer to satisfy enterprise legal and compliance requirements.

The Trust Problem in AI Data Sharing

When a company hands over a training dataset to an AI factory, it signs a contract and trusts. There is no technical mechanism to verify the contract was honoured.

AI factories receive training datasets from multiple customers. Each dataset carries legal constraints: permitted processing purposes, data residency requirements, retention windows, and commercial terms. Today these constraints live in Word documents and database rows. None are cryptographically locked. None are independently verifiable by the customer after the fact.

- The factory has no tamper-evident record that the customer agreed to specific terms.
- The customer has no way to confirm the factory used their data only for agreed purposes.
- Audit logs stored in databases can be altered without detection.
- Contract documents in files can be disputed or repudiated.

CHAIN++ makes all of this cryptographically locked and independently verifiable. The key properties it provides are immediate finality (every record is final within 2 seconds, no possibility of reversal), permissioned membership (only admitted nodes participate), end-to-end privacy (contract terms are encrypted, only named parties can read them), zero-trust key management (signing keys never leave the customer's own environment), and independent audit (every customer independently validates the full chain).

How CHAIN++ Works

Three layers work together. You do not need to understand any of them to use CHAIN++, but knowing what each layer does explains why the guarantees hold.

The Contract Layer

When two parties agree to share data, they sign a contract on the shared chain. Both signatures are stored permanently. The contract terms, what data, for what purpose, for how long, within what geographic boundaries, are encrypted so only the named parties can read them. Everyone else can see that a signed agreement exists between two addresses, but cannot read the terms.

The chain is operated jointly by multiple independent AI factories. No single factory can alter a past record, because doing so would require agreement from the majority of validators, run by different organisations. Every record is linked to every previous record by a cryptographic fingerprint, so any alteration is immediately detectable. This structure is called the Treechain in the BRAIN++ documentation: it is the Merkle Patricia Trie built into every Ethereum-compatible blockchain, not a separate product.

The Privacy Layer

Contract terms travel in an encrypted private channel. The encrypted payload is delivered only to the named parties. Other participants receive a blank response if they try to retrieve it. The keys that encrypt the contracts never leave the parties' own secure key management environments.

Any party can verify at any time that the terms they agreed to have not been altered. They fetch their encrypted copy, decrypt it with their own key, and check it matches the fingerprint on the chain. No intermediary involvement required.

The Network Layer

Every participant connects through an encrypted network tunnel that is configured automatically from the chain. When a customer publishes their connection key on the chain, the factory network configures the connection. When a contract expires, the connection access is removed. No IT configuration tickets, no manual revocation. For direct data exchange between two customers, the chain governs who is allowed to connect to whom. No active contract means no connection.

The Data Guardian

Before any data leaves a customer's environment, a containerised application, the Data Guardian, checks it against the contract they signed. It runs inside the customer's own network, on their own infrastructure, with no inbound ports. It classifies every dataset automatically: personal names, medical identifiers, financial account numbers, IP addresses, and more. If the dataset contains anything outside the agreed scope, the transfer is

blocked with a clear reason. Every decision, approve or block, is written to the chain with a cryptographic fingerprint of the data.

Every transfer must clear two independent gates: the Data Guardian's local policy check and the on-chain contract verification. Both must approve. Failure of one cannot be bypassed by the other.

Trusted Execution Without Confidential Computing

Full hardware-backed confidential computing is the long-term goal. Simultaneously, carefully designed containerisation by a trusted hosting entity already mitigates the most common breach risks with high probability.

Confidential computing, hardware-backed trusted execution environments such as Intel TDX or AMD SEV-SNP, provides the strongest possible isolation: even the host operator cannot read the contents of a running workload. This is the eventual target for sensitive AI processing. However, it is not yet universally available, and its security has been subject to ongoing research disclosure of side-channel vulnerabilities at the hardware level.

In the current PoC, and in the near-term production deployment, most AI factories are operated by trusted entities with strong physical and logical access controls. In this context, a well-designed containerisation and isolation architecture provides meaningful protection against the most likely attack paths.

++ What the Current Architecture Provides

+ **Process isolation**. Every customer workload runs in a dedicated container with its own filesystem namespace, network namespace, and process space. A process in one customer's container cannot directly access memory or files belonging to another customer's container. This eliminates the most common cross-tenant data leakage path.

+ **Network micro-segmentation.** Each customer has a dedicated WireGuard tunnel and a private overlay IP range. Traffic between customers is only permitted when an active signed contract exists on the chain. No customer can reach another customer's network segment without explicit chain-governed authorisation.

+ **Image signing and supply chain integrity.** Container images are signed using Sigstore/cosign before deployment. Only images whose signatures verify against the factory's trusted key are permitted to run. This prevents tampering with the software that processes customer data.

+ **Egress control.** Containers have explicit outbound allowlists. A customer's processing container cannot make arbitrary outbound connections. It can only communicate with the factory's own services and with counterparties that have an active chain-governed contract.

+ **Audit at the host level**. An eBPF-based audit stream logs process execution, file access, and network connections at the host kernel level, below the container layer. These logs are written to CHAIN++ so they are tamper-evident and independently verifiable, not just stored in the factory's own systems.

+ **No persistent data after contract expiry**. When a contract expires, the customer's data is removed from the factory's storage. The chain records the deletion event. The customer can independently verify that the deletion was logged at the correct time.

The Honest Assessment of CHAIN++ Architecture

This architecture substantially raises the bar against the most common attack paths: curious insiders, misconfigured access controls, software supply chain compromise, and cross-tenant data leakage. It does not protect against a determined and privileged host operator with physical access to the hardware. That gap is closed by confidential computing, which is on the CHAIN++ roadmap.

For most enterprise and regulated-industry deployments, the combination of a trusted hosting entity, the containerisation architecture described above, and the chain-based audit trail provides a governance posture that satisfies legal and compliance requirements. The chain provides the provable accountability layer that makes the architecture auditable, not just operationally secure.

In the PoC today	In the PoC today
Process isolation. Dedicated namespace per customer workload.	Egress control. Explicit outbound allowlists per container.
Network segmentation. Chain-governed per-contract connectivity.	Kernel-level audit. eBPF audit stream written to chain.
Image signing. Sigstore/cosign supply chain integrity.	Deletion on expiry. Chain-logged, customer-verifiable.

Confidential computing (Intel TDX / AMD SEV-SNP) is on the roadmap. It will close the remaining gap against a privileged host operator and enable cryptographic attestation of the execution environment.

The Federated Chain in CHAIN++

Multiple AI factories share one chain. No single factory controls it. Customers of any member factory get the same guarantees.

A chain operated by a single company is only as trustworthy as that company. Federation distributes validator control across multiple factory operators. Each factory runs its own validator nodes; confirming a new block requires agreement across factories. No single factory can confirm a block alone or alter a past record.

The identity registry, which companies and customers are authorised participants, is controlled by a multi-signature account requiring approval from multiple governance factories. No single factory can register or expel a participant unilaterally.

What the Federation Adds for Customers

Onboarding	Verification	Contracts
One registration. Recognised by all member factories.	**Same guarantees.** Audit records carry identical tamper-evident guarantees regardless of which factory you use.	**Direct agreements.** Sign contracts with customers of other factories directly, without any factory as intermediary.
One identity. Public keys and connection details published once.	**Cross-factory proofs.** You can verify records related to transactions at partner factories.	**Equal standing.** Both parties' signatures carry equal weight.
No re-negotiation. Existing contracts follow you across the federation.		

Direct Customer Collaboration

Customers of different AI factories can exchange datasets, trained models, and policy rules directly, governed, private, and verifiable. No factory reads the content.

Until now, when a customer at Factory A wanted to share data with a customer at Factory B, the only route was through both factories. Both saw the traffic. Both were intermediaries.

CHAIN++ enables a different model. Two customers sign a direct data exchange agreement on the shared chain. The agreement specifies what data, under what conditions, for how long. Once signed and active, data flows directly between them, over an encrypted connection that the factories route but cannot read.

What Can Be Exchanged

+ **Training datasets.** A data owner agrees to share a dataset with another customer under specific terms: permitted uses, processing purposes, geographic restrictions.
+ **Trained models.** A customer who has trained a model can license it under terms governed by the chain, with automatic expiry.
+ **Policy and compliance rules.** Organisations with developed data governance policies can share those rule sets directly with counterparties.
+ **Audit evidence.** Any customer can share a chain-based proof of their governance compliance with a regulator or auditor, without exposing the underlying contract terms.

How It Works

Two customers agree on terms off-chain. One proposes the contract on CHAIN++, encrypted so only the four named parties can read it: both customers and both factories. The other counter-signs. Either factory activates the contract when the agreed start date arrives. From that moment, every transfer is logged. When the contract expires, access stops automatically. No factory intervention required.

The encryption keys are derived from the chain itself using ECDH key agreement between the two customers' own key management systems. The factory never holds the encryption key. The factory routes the traffic but cannot decrypt the payload.

The Signing Process

- **Step 1, Agree terms.** Both customers negotiate conditions off-chain and agree on a context string for key derivation.
- **Step 2, Encrypt and propose.** The proposer encrypts the terms, computes a commitment hash, and submits the encrypted payload with their signature.
- **Step 3, Counter-sign.** The second customer verifies the terms, decrypts, and checks the hash matches. They add their signature.
- **Step 4, Activate.** Either factory node activates the contract at the agreed start date.
- **Step 5, Transfer and audit.** Data flows over the chain-governed encrypted overlay. Every transfer is logged permanently.
- **Step 6, Expiry.** At the agreed end date, access stops automatically and the contract is recorded as terminated.

For AI Factory Operators

Joining the CHAIN++ federation gives your customers a governance layer that regulated industries require, and gives your platform a differentiation that is difficult to replicate.

Compliance	Operations	Business
Automatic audit trail. Every data access logged permanently. Regulatory audits become mathematical proof, not log reconstruction.	**Zero-touch networking.** Customer connections configure themselves from the chain. Access stops automatically on expiry.	**Verifiable trust.** Customers in regulated industries can verify their own audit records independently. Stronger than any contractual commitment.
Contractual certainty. Every exchange governed by signed agreements that cannot be repudiated.	**Distributed governance.** In the federated configuration, governance is shared. No single factory carries the full compliance burden.	**Network effect.** Every factory that joins extends the addressable market for all existing customers on all member factories.
Data residency enforcement. Contract terms include geographic constraints that are technically enforced.	**Post-quantum readiness.** Regulatory guidance on quantum migration is being issued. CHAIN++ positions your platform ahead.	**New revenue models.** The chain enables customers to monetise data and models directly, with your factory as the trusted infrastructure.

For Customers

You own your records, you verify your own compliance, and you control your own agreements, without depending on the factory for any of it.

Visibility	Control	Exchange
Your own copy. Complete, independently validated copy of the relevant chain data. Nothing taken on trust.	**Your keys stay with you.** Signing keys live in your own environment. The factory never has access.	**Direct agreements.** Sign contracts with customers of other factories. No factory as intermediary.
Mathematical proof. Prove any event occurred, at any time, without factory involvement.	**You define the terms.** Permitted uses, time window, geographic scope, set by you, enforced technically.	**Secure data receipt.** Data arrives over an encrypted connection active only while a valid contract exists.
Contract integrity. Verify your agreement terms have not been altered since you signed them.	**Automatic enforcement.** When your agreement ends, access ends. No revocation request needed.	**One identity.** Your registration is recognised by all member factories. No separate accounts per factory.

Security and Compliance Readiness

What CHAIN++ guarantees, what it is building toward, and where the current PoC limits are.

What CHAIN++ Guarantees Today

- **Non-repudiation.** Both parties' cryptographic signatures are on-chain permanently. Private keys never leave each party's own key management system.
- **Tamper-evident records.** Any alteration of a past record changes the linking fingerprints in all subsequent records. Every participant detects it immediately.
- **Independent verification.** Every customer holds a complete copy of the chain. Any customer verifies any record without asking the factory.
- **Contract confidentiality.** Contract terms are encrypted. Only named parties can decrypt. Non-parties receive nothing.
- **Zero-trust key management.** All signing keys live in OpenBao (Linux Foundation, MPL 2.0). No raw private key is ever exported.

Post-Quantum Readiness

Data governance records may be called into evidence years or decades from now. Quantum computers can break the elliptic curve cryptography protecting most digital records today. CHAIN++ is implementing hybrid post-quantum protection: classical signatures for current compatibility, plus ML-DSA post-quantum signatures (NIST FIPS 204, 2024). Contract encryption adds ML-KEM hybrid key agreement (NIST FIPS 203, 2024). Both algorithms are NIST-standardised. The approach is hybrid: classical and post-quantum run simultaneously, so no infrastructure replacement is needed.

Current Limitations of CHAIN++

The following capabilities are not yet in the PoC. They are in active development or the longer-term roadmap:

- **Formal anonymization.** The Data Guardian classifies sensitive content but does not yet transform data to a validated anonymization threshold.
- **Receiver-side enforcement.** Full or partial contract term enforcement on the recipient's side using a secure contract enforcement container.
- **Contract chaining.** Derivative contracts inheriting parent contract restrictions are planned.
- **Data market discovery.** A marketplace where data owners and consumers find each other is a future direction.
- **On-chain payment settlement.** Token-based settlement is a strategic future vision.
- **Confidential computing.** Hardware-backed TEE isolation is the longer-term roadmap item, closing the final gap against a privileged host operator.

CHAIN++ Roadmap

CHAIN++ is at PoC stage. The features in this whitepaper are built and validated. Active development continues before the first production release.

In the PoC Today

+ **Single-factory chain.** Tamper-evident audit logging, cryptographic contract governance, end-to-end private contract storage.
+ **Data Guardian.** Client-side data classification and contract enforcement before every transfer.
+ **Chain-driven network overlay.** Automatic zero-touch secure connectivity, WireGuard-based.
+ **Containerisation and isolation.** Process isolation, network micro-segmentation, image signing, egress control, eBPF audit.
+ **Delta Lake AI training platform.** Governed datasets at scale.
+ **Federated chain.** Multi-factory shared chain with Gnosis Safe governance, cross-factory contracts, direct customer-to-customer overlay.

Being Added Before Production Release

+ **Formal anonymization**. Validated privacy thresholds in the Data Guardian.
+ **Receiver-side enforcement.** Contract enforcement container at the recipient's side.
+ **Contract chaining.** Derivative contracts inheriting parent restrictions.
+ **Post-quantum signatures and encryption.** ML-DSA and ML-KEM, NIST 2024.
+ **Post-quantum WireGuard**. When upstream implementation stabilises.
+ **Light-client verification**. For participants who cannot run a full observer node.

Strategic Future Vision

+ **Data market discovery.** Marketplace for data owners and consumers to find each other.
+ **On-chain payment settlement.** Token-based coordination for data access transactions.
+ **Confidential computing.** Hardware-backed TEE isolation, closing the privileged-operator gap.
+ **Zero-knowledge compliance proofs**. Prove regulatory compliance without revealing contract terms.

Further Information

Learn more about CHAIN++ and the BRAIN++ AI Factory, access the full technical documentation, and get in touch:

- **Website** https://brainplusplus.bg
- **Technical documentation** https://gitlab.discoverer.bg/vkolev/chainplusplus
- **Contact** office@brainplusplus.bg
- **Location** BRAIN++ AI Factory at Sofia Tech Park, Sofia, Bulgaria