On the Radar: SafeLiShare Creates Trusted Bindings between Applications and Data for Safety of AI workloads

Publication Date: July 2024

Rik Turner, Senior Principal Analyst, Omdia

Summary

Catalyst

SafeLiShare provides a technology platform that leverages a hardware-based security feature known as the trusted execution environment (TEE) on modern CPUs. This enables a variety of use cases, namely:

- Confidential Clean room for multi-party access control on data and AI collaboration
- Private LLM confidential access to proprietary LLM data
- Automated data transformation and access policy enforcement throughout data lifecycle
- Location-agnostic and GDPR data sovereignty enforcement

The technology can be deployed in a customer's public, private cloud or on their premises.

Omdia view

Whether it be in medical research or developing more highly targeted financial service offerings, there is immense potential in analyzing data from multiple parties in one place, without any single provider being able to access or see the data from their counterparts. Indeed, entire new business models are likely to emerge, as the owners of particularly good machine learning (ML) or large language models (LLMs) offer AI-as-a-service to groups of customers who want to pool their data for analytical purposes, but without sharing that data with any counterparts.

Leading AI breaches today are often caused by internal and external parties exploiting the lack of comprehensive visibility and controls over large language models (LLMs) and malicious attacks on AI infrastructure. IT in AIOps requires enhanced visibility and control through a centralized governance pane to manage workloads effectively, whether they are in the cloud or on-premises.

SafeLiShare's technology aims to enable that scenario and should be on your watchlist if you are looking for a way to govern your enterprise data and knowledge base securely for a computation, along with data from other parties. It can also be used for training AI models.

Why put SafeLiShare on your radar?

Of particular interest with SafeLiShare is that fact that it is a turnkey secure data proxy any enterprise exploring LLM can use to get full inventory and data governance mesh for enterprise-ready visibility and tamper-proof compliance auditing. It is not only a provider of a confidential clean room, but can also govern LLM and KB data in private AI applications as well as protect run-time model fine-tuning in isolated execution environement, securely. Omdia considers this a distinct advantage vis-à-vis some of the competition in this still emerging market.

Market context

The rise of cloud computing, big data analytics, and now of AI, and in particular of LLMs, has created huge research opportunities for organizations in a variety of vertical markets, including healthcare, insurance, financial services more generally, and pharmaceuticals.

They all gather vast amounts of relevant information about patients, account holders, or clinical trials, and the potential benefits of pooling their data with those of commercial competitors or, in the case of hospitals, of similar institutions, are manifold: analysis of even larger sample groups can improve accuracy and bring insights more clearly into focus.

Such opportunities come with a catch, however, in the form of the privacy concerns that ensue from sharing data in the cloud. It may be unwise from a security standpoint, or indeed altogether illegal, depending on the jurisdiction in which the data custodian is located.

For this reason, the tech industry has embarked on a number of initiatives designed to enable such zero trust model and data collaboration to take place in a secure and compliant manner. There are a range of technological approaches that can be considered, including homomorphic encryption and differential privacy, each with their strengths and weaknesses. Perhaps the most successful to date, however, and certainly the one with the widest support from technology heavyweights, is confidential computing.

This is an approach to securing data, and of guaranteeing its privacy, that relies on the TEE technology delivered by a range of semiconductor manufacturers, most notably Intel, AMD, and ARM. It entails using a secure enclave on the chip, such that encrypted data only exists in plaintext within that enclave, with no possibility of outsiders, not even the parties to the computation, having access to it.

In its latest LLM Secure Data Proxy, SafeLiShare envisions a non-proprietary three-layer confidential secure framework that simplifies confidential computation orchestration, consisting of:

- At the top layer, an integrated IDP connector with distributed Widespread Augmented Verification Environment (WAVE) authorization that provides the strictest confidential data access via airtight data governance
- In the middle, a federated secure fabric for data exchanges regardless of where data is located, hosted, or accessed
- At the bottom, federated secure enclaves, hardware-based and/or software-based secure enclaves, hosting distributed data sources in need of lifetime data protection

SafeLiShare LLM Secure Data Proxy promotes a unified approach of secure services that deliver turnkey data governance and protection measures, federated access controls and compliance protocols to safeguard sensitive enterprise data. Enterprise IT administers can get full LLM inventory and lifecycle on data access.

The increasing prevalence of breaches and hacks against AI systems necessitates a vigilant approach by organizations to safeguard against both external and internal threats. In light of these challenges, it is crucial for organizations to recognize and address security, privacy, and risk concerns driven by generative AI (GenAI). This calls for the establishment of robust AI governance frameworks and the allocation of resources to support these initiatives. Diverse roles must be engaged in managing AI privacy, security, and risk, with AI-mature organizations forming dedicated AI teams to handle these critical tasks. To achieve higher AI maturity and positive outcomes, organizations should implement programs that support AI governance, enforce policies, and create cross-functional AI teams responsible for managing AI-related risks and ensuring compliance.

Product/service overview

While the underlying technology is the same, i.e. a platform that enables customers to make use of the TEEs on the more recent generations of silicon such as the Intel SGX, SafeLiShare articulates its offering as a series of "solutions". These are more aligned with individual use cases for the technology than distinct SKUs, and all of them available via an annual subscription. They are:

Universal AI Security Platform --

ConfidentialAI™ Clean Rooms – The term "clean room" (also written as one word, i.e. "cleanrooms") originated in semiconductor manufacturing, referring to facilities where there was no risk of chamical contamination of the substrates or components used to produce chips. In contexts such as financial engineering and others where data confidentiality is critical, a confidential clean room is a secure space in which data can be viewed and processed without the risk of exfiltration. The Zero Trust Collaboration guarantees no party can see your data or models including the cloud service providers. The SafeLiShare offering is designed to ensure airtight security and privacy control of any assets used within the clean rooms are cryptographically verified. Free trial avaialbe in AWS Marketplace.

LLM Secure Data Proxy with Secure Enclave as a Service – Retrieval-Augmented Generation (RAG) is the process of optimizing the output of an LLM, which it does by referencing an authoritative knowledge base outside of its training data sources before generating a response. The SafeLiShare offering enables RAG workflows to be grounded in the data owned by the organizations taking part in a computation, their vector databases, AI cache or knowledge graphs, without exposing raw data or chunks in prompts, responses, or memory history. The technology can be deployed as a micro service in any cloud or secure service edge environment, or indeed on a customer's premises.

LLM SDP with Secure Enclave as a Service (SEaaS) offers customers:

- Granular control over who has access to private data, such that only authorized individuals can view and interact with sensitive information.
- Auditable data plane logs (who accessed what data) and control plane logs (who changed settings or policy) for purposes of transparency and accountability in data handling.
- An encryption key exchange mechanism to prevent direct access to private data and defend against unauthorized intrusion.
- End-to-end encryption of all queries, responses, RAG history, and login information, with even SafeLiShare being unable to see them, thereby guaranteeing the customer's privacy.

Company information

Background

SafeLiShare was founded in 2021 by CEO Shamim Naqvi, CTO Pramod Koppol, and Chief Architect Goutham Puppala. Naqvi has been CEO (and often co-founder) of a number of tech firms, the most recent one before SafeLiShare being Sensoriant, a provider of technologies in the areas of sensor data, algorithms, and transactions.

In August 2021 SafeLiShare raised a $5m seed round led by Taiwania Capital Management, with participation from FORTH Management andTaya Ventures.

Current position

SafeLiShare is at an early stage of its development, both from a funding and a product offering perspective. It has only raised a seed round, it has a dozen or so employees, and it only unveiled the private beta of its ConfidentialRAG platform in May 2024, at the RSA Conference.

There are a few other players emerging to offer confidential computing platforms, including one that SafeLiShare cites as a competitor, namely the Swiss firm Decentriq. That company is also a provider of data clean room technology based on TEEs, but it currently does not have an offering for securing AI/ML and LLMs.

Future plans

SafeLiShare is now in the process of adding AIOps data access gateway for LLMs integrated with RAG, enabling the platform to secure Private AI and become the “central pane of glass” providing centralized access control, auditing, approvals, model workflow, lineage, and data discovery across enterprise workspaces.

Key facts

Table 1: Data sheet: SafeLiShare

Product/Service name	SafeLiShare ConfidentialAI Platform	**Product classification**	Data access gateway for AIOps for multiparty computational purposes
Version number	**1.0**	**Release date**	June 2024
Industries covered	Finance, insurance, healthcare, and pharma.	**Geographies covered**	**US, Canada**
Relevant company sizes	Enterprise, SMB	**Licensing options**	Annual subscription
URL	https://safelishare.com/	**Routes to market**	**Direct, Channel**
Company headquarters	Morristown, NJ, USA	**Number of employees**	12

Source: Omdia

Analyst comment

350 words.

LLM Secure Data Proxy by SafeLiShare: SafeLiShare has introduced a recent **LLM Secure Data Proxy (SDP)** that offers a unified console. This console allows control over **GenAI**, **LLM**, and **RAG** policy creation and optimization. Enterprise IT, CISOs, and CSOs are actively seeking turnkey solutions to secure AI adoption. Beyond zero-trust collaboration and asset sharing, this solution addresses the need to control location-agnostic AI model access to enterprise data. It enables data loss prevention (DLP) on private data and intelligent responses based on user privileges. The ultimate goal is to upload data to an analytical platform, where it can be decrypted and combined with third-party assets securely. Only the computation results are accessible to participants involved in training, fine-tuning, and research.

Confidential Computing Landscape: Established players like AWS, Google, Nvidia, Apple and Microsoft are actively developing confidential computing capabilities. However, there's also room for cloud-agnostic multi-party collaboration specialists like SafeLiShare. While SafeLiShare is still in its early stages, Omdia recommends monitoring its development. Subscribers interested in multiparty computing and “shadow IT” control in AIOps should explore SafeLiShare's offerings in private AIOps and GenAI secure control use cases. Keep an eye on its progress.

Appendix

On the Radar

On the Radar is a series of research notes about vendors bringing innovative ideas, products, or business models to their markets. On the Radar vendors bear watching for their potential impact on markets as their approach, recent developments, or strategy could prove disruptive and of interest to tech buyers and users.

Further reading

Blockchain Technology and Adoption Trends (December 2019)

"Blockchain is good for more than just Bitcoin," (September 2019)

"CenturyLink goes 'colorless' and takes on the edge cloud" (February 2020)

Service Provider Routers & Switches Market Tracker – 4Q19 (February 2020)

Li You, "Tech-savvy Hangzhou tries out new 'City Brain'," China Daily (retrieved June 17, 2021)

Author

Rik Turner, Senior Principal Analyst, Cybersecurity

askananalyst@omdia.com

Citation Policy

Request external citation and usage of Omdia research and data via citations@omdia.com.

Omdia Consulting

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Omdia's consulting team may be able to help you. For more information about Omdia's consulting capabilities, please contact us directly at consulting@omdia.com.

Copyright notice and disclaimer

The Omdia research, data and information referenced herein (the "Omdia Materials") are the copyrighted property of Informa Tech and its subsidiaries or affiliates (together "Informa Tech") or its third party data providers and represent data, research, opinions, or viewpoints published by Informa Tech, and are not representations of fact.

The Omdia Materials reflect information and opinions from the original publication date and not from the date of this document. The information and opinions expressed in the Omdia Materials are subject to change without notice and Informa Tech does not have any duty or responsibility to update the Omdia Materials or this publication as a result.

Omdia Materials are delivered on an “as-is” and “as-available” basis. No representation or warranty, express or implied, is made as to the fairness, accuracy, completeness, or correctness of the information, opinions, and conclusions contained in Omdia Materials.

To the maximum extent permitted by law, Informa Tech and its affiliates, officers, directors, employees, agents, and third party data providers disclaim any liability (including, without limitation, any liability arising from fault or negligence) as to the accuracy or completeness or use of the Omdia Materials. Informa Tech will not, under any circumstance whatsoever, be liable for any trading, investment, commercial, or other decisions based on or made in reliance of the Omdia Materials.

CONTACT US

omdia.com

askananalyst@omdia.com