About this document

As part of our ongoing effort to critically evaluate the Nostro project, we approached an independent professional economist and requested an external assessment of the project.

The document presented below is the result of that assessment. We are publishing it in full because we believe that serious projects should be willing to engage with criticism, competing viewpoints, and open professional discussion.

The original analysis was translated into English and made publicly available on Nostro's website so that investors, industry professionals, and other interested readers can review the critique directly and form their own conclusions.

The response prepared by the Nostro team can be downloaded from the same page using the following link:

https://nostro.pro/data_room_downloads/supporting_materials/_response_analysis_of_the_nostro.pdf

We encourage readers to review both documents in full before forming an opinion about the project.

Analysis of the Nostro project

Dear Professor,

This project carries significant risks.

The first category consists of obvious business risks associated with the fact that the team did not start with a verifiable ability to create a more valuable solution for a specific type of customer, but instead began by designing a financial product presented as a private version of financial infrastructure, which contradicts the logic of a venture-backed startup.

The second category consists of more significant risks related to potential fraud, which I will examine in the second part of this analysis.

If all the materials of this project are reduced to the hypothesis of the problem they are trying to solve, it can be formulated as follows:

"*Cryptocurrencies and stablecoins have accumulated a certain market capitalization and a large user base, but have failed to become fully functional money for the real economy due to the absence of regulated infrastructure that would allow businesses to accept crypto assets as easily, with the same legal clarity and accounting transparency as ordinary fiat money.*"

They propose to solve this problem by developing a **private financial ecosystem** with different products for different types of target audiences.

The technology is built around two instruments: N-Token and N-Coin. I will not go into them in detail and will limit myself to noting that both instruments are opaque financial products. The economic nature of these instruments raises serious concerns.

Reasons why the project is inherently doomed to fail and extremely risky:

1. Teams with experience in TradFi that are genuinely trying to introduce new technology into an outdated financial system are easy to distinguish from potential fraudsters with fake projects. Real startups always begin by explaining exactly how they will satisfy an existing customer need (job-to-be-done) in a significantly more valuable way and provide evidence to support it. Such teams identify areas where capital efficiency can be improved within specific financial flows.
   In contrast, this project focuses on a macro-level problem and attempts to justify the development of certain quasi-financial instruments as a solution to that problem. This is immediately concerning. In its current form, I do not believe it has any realistic prospects of successfully attracting venture capital investment. Rather, it may mislead more trusting

individuals among non-professional investors.

2. The economic rationale behind the concept is extremely weak and is built entirely on speculation about the behavior of a non-existent customer in an abstract world.
3. Methodological reasons.
4. The industry itself is moving toward eliminating all kinds of surrogate instruments that “wrap” real value, with the exception of **legally recognized tokenized claims**.

The central logic of such innovations is based on transferring monetary claims and assets from fragmented analog registries into a network of interconnected digital registries (within a blockchain), while the monetary claim, deposit, reserve, security, or property right itself becomes a legally recognized programmable accounting object.

The experiments that are most significant for the market primarily concern central bank wholesale money, tokenized commercial bank deposits, securities, and property rights.
Such initiatives are being implemented through experiments conducted by the largest banks and payment systems, with close involvement from regulators.
These are not the kinds of projects that can be made economically viable within the framework of a venture-backed startup by a team lacking the relevant experience.

This is fundamentally different from the Nostro model, where a team of enthusiasts is attempting to create its own layer of quasi-products and then persuade users to regard it as a new settlement infrastructure.

Thus, they have failed to demonstrate why an early-stage startup, with no indication of experience in serious institutional finance, no understanding of a viable market entry wedge, and no meaningful support, should begin by building some form of infrastructure.

Next, I propose examining the project from several different perspectives:

1. Analysis of the claims made in the address to you and in the startup's materials, and the team's methodological misconceptions.
2. Fraud Risks
3. Analysis of the project through the prism of launching fintech and blockchain startups.

Part I. Analysis of the most significant claims and statements. The team’s methodological misconceptions.

For the purposes of understanding the context of the analysis that follows, I believe it is important to address several claims made in the materials presented to you.

1. **“$3.5 trillion of crypto has been removed from the real economy” The author creates the impression that there is some unused resource and that, if the “right infrastructure” is built, it can be unlocked.**
2. **“The low level of business adoption of crypto is an anomaly”**

The market capitalization of crypto assets is not equivalent to capital that is somehow "sitting outside the real economy" and can supposedly be brought back into economic circulation.

The market capitalization of crypto assets is merely a calculated figure: the market price of an asset multiplied by the number of units in circulation. It does not mean that an equivalent amount of liquid capital exists and is ready to flow into consumption or corporate payments. If all holders attempted to exit the asset simultaneously, its price would fall, liquidity would contract, and market capitalization would evaporate, as has repeatedly occurred during episodes of capital flight. In reality, a $3.5 trillion market capitalization reflects several things: speculative holding, expectations of future price appreciation, the use of assets as collateral, and participation in decentralized yield-generating protocols. It does not imply the existence of corresponding demand to pay for goods and services with those assets.

Stablecoins exist primarily not as standalone payment instruments for purchasing goods and services, but as settlement assets within the crypto ecosystem itself. That was the original purpose for which they were created as a foundational asset. Stablecoins are used for exchange trading, arbitrage, transaction collateralization, liquidity transfers between venues, and similar activities. A high observed volume of stablecoin transactions does not imply a proportionally high level of demand for these instruments in the real economy. A significant share of this transfer volume (approximately 60% of all transfer events according to the BIS paper *The Anatomy of Stablecoin Transactions*) reflects the internal financial mechanics of the crypto market.
While interest in stablecoins as intermediary settlement assets is currently growing, this trend should be discussed with caution given the existing barriers to adoption.

The low level of cryptocurrency adoption in payments is driven primarily by fundamental factors. The vast majority of crypto assets do not meet the requirements businesses have for money. They do not function as a stable unit of account, a convenient means of payment, a legally straightforward asset, or a reliable treasury instrument. In addition, they leave unresolved questions regarding settlement finality, counterparty risk, and the settlement process itself.

Based on my own experience, corporate use cases involve risks related to depegging, as well as the complexity of managing multi-asset treasury operations, including the accounting treatment of tokens, networks, transaction fees, and controls over wallet access. For most corporate treasury functions, maintaining a straightforward fiat-based accounting framework is significantly simpler. This remains the primary barrier to corporate adoption.

In real-world corporate payment use cases, the process typically works as follows:

- In the cryptocurrency market, the primary providers of liquidity are specialized liquidity providers such as Wintermute, B2C2, and others. If a business needs to exchange EUR for USDC, it requires a counterparty that quotes the EUR/USDC pair, provides liquidity, and executes the transaction. This role is performed by liquidity providers. The fiat leg of the settlement between the business and the liquidity provider is typically conducted within the same banking institution.

The principal risk, and the primary concern for decision-makers in this process, is that one party may fulfill its side of the transaction while the other does not, or that the assets involved may become trapped with an unreliable intermediary.

The solution to counterparty risk lies in custody arrangements and the execution architecture built around them: who holds the assets, who confirms the transaction, where settlement takes place, whether a delivery-versus-payment mechanism exists, who bears the risk until settlement finality is achieved, and what happens in the event of a failure.

In many cases, the first question a corporate decision-maker asks when confronted with crypto assets is: "How can they be stored securely? And is there insurance coverage?" When a trusted custodian is involved, it can reserve the assets of both parties on its secure platform during the transaction process, acting as a distant analogue of a central counterparty. Companies such as Fireblocks have built institutional networks, including the Fireblocks Network, which connect institutional participants and enable the secure transfer of assets between institutions while reducing counterparty risk. By doing so, they help overcome some of the barriers to crypto asset adoption faced by banks, financial institutions, and other participants in the payments ecosystem.

The authors of the presentation draw incorrect conclusions from both their market analysis and their assessment of the underlying problem. Liquidity in the form of $3.5 trillion will not materialize as a result of any platform functionality. I consider the reasoning presented to be without merit, and the claim that "the low level of business adoption of crypto is an anomaly" can be explained by the rational behavior of market participants.

3. **"Stablecoins were supposed to turn crypto into money, but they failed to do so"**
4. **"Nostro turns cryptocurrency into money "**

The third claim put forward by the author [Nostro] does not serve as a bridge to the proposed solution. Stablecoins did not become money for the broader real economy, but for entirely different reasons.

If we follow the BIS framework outlined in *The Next-Generation Monetary and Financial System*, money is evaluated against three key criteria:

- **Singleness of money.** Money must be accepted at par without requiring additional questions about whose "euro" or "dollar" it is, because final settlement ultimately occurs through the banking system and settlement assets in central bank money. Stablecoins, by contrast, represent claims on a specific issuer or a specific collateral structure. As a result, issuer risk, collateral risk, and depegging risk arise. In the case of stablecoins, the first principle, the singleness of money, is compromised.
- **Elasticity.** A monetary system remains flexible when it can expand and contract liquidity in line with the needs of the economy. Stablecoins operate differently. They require pre-funding. As a result, such a system does not scale effectively as infrastructure for an economy with more elastic business liquidity requirements.
- **Integrity.** This refers to a system's ability to maintain trust, verify participants, enforce rules, manage risks, support AML/KYC processes, and protect against abuse. In this respect, stablecoins exhibit a number of shortcomings.

Stablecoins have indeed demonstrated the advantages of programmable, tokenized payments.

However, actual market adoption is occurring in the area of tokenizing **existing and trusted monetary claims**, such as central bank reserves, commercial bank money, government bonds, securities, and other forms of property rights.

The conclusions drawn by the Nostro team in this regard also appear unfounded, as Nostro does not possess at least the first two core attributes of money. Instead, it merely creates a surrogate layer on top of existing monetary claims.

The project materials provide no basis for concluding that Nostro possesses the essential monetary

functions of a unit of account, a store of value, or a means of payment.

5. **“On the balance sheet, crypto held through Nostro would be treated like fiat cash”**
6. **“Nostro Token and Nostro Coin jointly perform the functions of money”**

In a conventional fiat payment system, a bank increases the recipient’s deposit balance, while settlement takes place in central bank reserves. In the model described here, however, the recipient of an N-Token would not receive euros. Instead, the recipient would receive a claim on the issuer, which raises a number of questions and risks: who the issuer is, what reserves stand behind it, how redemption is carried out, what rights are attached to the instrument, and so forth.

The reserve asset described as an N-Coin appears to be a complex and opaque instrument, perhaps even a hybrid asset, bearing little resemblance to money. Moreover, it appears vulnerable to significant market manipulation in the interests of large holders. Money, as a store of value, should by definition minimize uncertainty.

Moreover, a monetary instrument cannot simultaneously promise yield and price appreciation. Such characteristics are more consistent with an investment asset than with money itself.

Nor can an asset simply be declared to be money. An asset becomes money when the market uses it as a unit of account, a medium of exchange, a means of payment, and a store of value. In the case of Nostro, there is a declaration that it will serve such a role, but no basis has been presented for concluding that the market will accept it as such.

7. **“The share of stablecoins must increase to 51% in order to stabilize the market”**

The author of the presentation argues that the crypto market is unstable because of the dominance of volatile assets and states that Nostro’s objective is to increase the share of stablecoins to 51%, make financial flows bidirectional, and thereby stabilize the crypto economy.

A larger share of stablecoins cannot, by itself, reduce the volatility of the crypto market or of crypto assets. At most, it represents a form of statistical dilution rather than genuine price stabilization. The claim is methodologically flawed because volatility is influenced by a range of factors, including market liquidity depth and the volume of leveraged trading, which generates liquidation cascades.

These dynamics, in turn, contribute to sharp price declines. In addition, crypto markets exhibit a direct correlation with movements in broader financial markets, as capital flows into crypto assets tend to reflect changes in overall risk appetite across different market cycles.

Accordingly, the claim that stablecoins representing 51% of total market capitalization would stabilize the crypto economy appears to be economically unfounded.

8. **The presentation included the claim that a banking license is required before the project can be tested.**

This is a speculative assertion. Regulated environments provide dedicated regulatory sandboxes that allow startups to test the fundamental viability of their business models before investing in full-scale licensing requirements.

This suggests one of two possibilities: either the team has a poor understanding of how regulated fintech products are actually launched, or the reference to a banking license is being used as a rhetorical device to divert stakeholders’ attention toward supporting the project rather than focusing on the core challenges of customer acquisition, traction, and demonstrated market demand.

Part II. Fraud risks

If one strips away all the window dressing and assumes that, by some extraordinary stretch, developing a new blockchain network makes sense, the first thing that stands out is the absence of a white paper (the project's constitutional document and the primary disclosure document expected of blockchain projects). There is no formal token specification, no definition of token-holder rights, no reserve model, no disclosed tokenomics, no smart contracts, and none of the other foundational elements typically expected of such a project.

I would further ask you to pay attention to what I consider the primary indicator of potential fraud. The entire economic model is built around a proprietary internal token whose price and liquidity could potentially be controlled by the company itself.

For the serious project, this creates a conflict of interest. The company simultaneously designs the token, controls its issuance, organizes the internal market, collects fees on entry and exit, and, most importantly, promises **some form of economic benefit** to token holders.

This is precisely where the risk of price manipulation arises. Mechanically, it resembles the same pump-and-dump schemes that have repeatedly appeared in the market, including recent examples such as Libra or $TRUMP.
In this case, the structure appears to be the following: first, significant marketing hype is created around the asset. Promises are made about future infrastructure, network growth, and future liquidity, while attempts are made to attract influential figures, including individuals such as yourself.

In the early stages, limited supply and internally generated demand may make the token’s price appear stable or rising. At that point, early participants, or the issuer itself, may have an opportunity to realize their gains and exit into cash, while later participants are left holding worthless paper.
It appears to me that this is the true motivation of those who approached you for support, as serious businesses in TradFi are not built in this manner.

Part III. Analysis from the perspective of fintech startup development

If we set aside the first two issues and evaluate this as a startup, we return to the same conclusion: the underlying business logic is flawed.
Distributed ledger technology (DLT) creates economic value when its application to a specific existing financial process eliminates a measurable inefficiency. Such inefficiencies may include lengthy settlement cycles, excessive liquidity lock-up, manual reconciliation, settlement delays, elevated operational risk, insufficient transparency regarding transaction status, costly ancillary processes, or discrepancies between the records maintained by different participants.
In other words, the starting point should be a specific bottleneck that creates costs, slows operations, weakens control mechanisms, or fails to meet the needs of the modern businesses whose financial flows the process serves.

It is only under this logic that a team can legitimately claim a share of the value it creates in the market. Only then does the framework emerge around which future cash-flow generation can be designed and the foundations of future shareholder value can begin to take shape.

.

Therefore, the logic of such a project begins with a fundamental question:

Which specific process today contains a measurable inefficiency? Who bears the cost of that inefficiency? Why are existing solutions inadequate? And why does DLT reduce those costs more effectively than traditional infrastructure? How repeatable and scalable is the opportunity to create value within this particular financial use case and its addressable market?

If these questions are approached correctly, the team should arrive at a limited set of testable hypotheses:

- Which specific customer is experiencing the pain point?
- What job is currently being performed poorly for that customer?
- How is the customer solving that problem today?
- Why is the existing alternative inadequate?
- What early product, prototype, or manual service is capable of changing customer behavior?
- Who is willing to go through the early-stage, imperfect journey alongside the company?
- Does the team have sufficient time, capital, and focus to validate all of the above?

At the Pre-Seed stage, the team must build a sufficient level of evidence-based confidence, gradually replacing its highest-risk assumptions with observable facts.

By commercial validation at the Pre-Seed stage, I do not mean a hypothetical financial model presented by the project. I mean observable customer behavior: a change in how the customer operates through access to a real process, willingness to engage in a paid pilot, persistence of that behavior after the pilot, or any other signal demonstrating that the customer is beginning to alter an existing workflow.

This is how a sequence of evidence should emerge:

> First customer and their job-to-be-done → target segment → a value proposition that solves that job more effectively → behavioral change → revenue potential → future repeatable margins.

To quote you, Professor, most products of this kind possess the characteristics of a **credence good**. A customer must entrust the product with money, transactions, accounting records, payments, data, regulatory obligations, or relationships with counterparties.

As a result, customer discovery and trust-building become particularly demanding, as they require overcoming adoption barriers and addressing the customer's underlying concerns and perceived risks.

If a team is serious about operating in the market, its efforts should be focused on one thing only: gaining access to its first customer, rather than compiling a list of every potential customer it can imagine. The priority should be understanding the customer's process from the inside, acquiring firsthand insight into the customer's operational pain points, and testing a specific workflow within a specific financial flow. This process should result in the identification of a narrowly defined set of participants, a single sufficiently frequent and scalable transaction type with a clear and measurable value metric, a predefined success criterion, and a quantifiable economic impact.

It is equally important for the team to understand the mechanism by which a customer switches from one solution to another. Customers adopt a new solution when the pain points of the current approach and the perceived incremental value of the new product outweigh habit, inertia, and the fears associated with change. In this context, the issuance of a new surrogate token may have the opposite effect, increasing perceived risks and reinforcing customer concerns.

Addressing the challenges described above through the proper lens, while eliminating everything non-essential, should result in the creation of a wedge into the market. Without a market-entry wedge, it is impossible to determine which hypothesis is the primary one. Without a primary hypothesis, the team cannot build a learning roadmap grounded in customer experience and aimed at influencing customer behavior. Without such a roadmap, the team risks building features in isolation rather than reducing uncertainty. And without reducing uncertainty, the project cannot progress toward a validated business, which is a prerequisite for subsequent stages of development.

Only after such validation has been achieved does it make sense to discuss scaling the solution to additional participants, expanding into new markets, or assuming an infrastructure role. Those are matters for entirely different stages of company development (Series A and beyond).

In the public materials and in the pitch presented to you, I did not see this. Instead of a first testable use case and a clear market-entry wedge, what was presented was a high-level vision of a future ecosystem built around questionable economic incentives.

The team's right to enter the market

The next question is whether the team has earned the right to enter this market. By *the right to enter*, I mean a credible entrepreneurial basis for believing that this particular team is capable of securing an advantage within the chosen market opportunity. The team has not explained why it, specifically, has access to the right customer, understands the customer's workflow, can overcome industry barriers, and is capable of generating the first compelling evidence faster than others. This is particularly important given that the market already contains companies and consortia with substantially greater resources, institutional access, and the trust of financial market participants.

Accordingly, a team's right to enter the market should be supported by at least one of several factors: deep domain expertise, existing strategic partners, a unique technology, licensed partners, or privileged access to customers. A startup must demonstrate why it possesses a specific entry point that is either unavailable to, or unattractive for, larger incumbents. Based on what I have seen, the team's right to enter this market remains unproven.

3. The product roadmap has replaced the learning roadmap

Given that every product feature is ultimately a mechanism for capturing and retaining value in the market, having too many features at an early stage is dangerous. The team does not yet know which critical uncertainty each feature is intended to resolve.

At this stage, the startup’s primary artifact should not be a product roadmap but a learning roadmap. The team must understand which high-risk assumptions it is testing, by what minimum viable means, with which customer, and what decision it will make based on the outcome. If, instead, the team predefines an extensive feature set, it begins building a product detached from reality, because it has not yet demonstrated that customers will perform their jobs in that way or achieve the expected outcomes through that approach. This risk runs as a common thread throughout the project's materials. The problem is that it remains unclear which customer pain point is the first mandatory one to validate.

Even a modest excess of features can be dangerous. The more functionality a team commits to before obtaining its first validation, the more capital it burns before the next funding round. That means less runway, greater complexity, and an almost certain likelihood that many of those features will be built with no practical use. Under the proper approach, the team first selects a single customer process and a single improvement metric. Only then does it determine the minimum feature or manual service required to test the underlying hypothesis.

4. The project's economics are not grounded in demonstrated customer value

Nostro's financial model reads more like a top-down fantasy than a description of customer behavior and a mechanism for sustainably capturing a portion of validated customer value. The team must understand what will cause its target users to change their behavior. Only then can it determine, at later stages, where its point of strategic control lies, where the opportunities for nonlinear growth exist, what retention levels can realistically be expected, how cohorts can be expanded, and how unit economics, acquisition channels, and customer acquisition costs should be managed.

At the Pre-Seed stage, the financial focus should be on the cost of validation: how much it costs to test the core hypothesis, whether the company has sufficient runway to obtain that validation, which expenses genuinely reduce uncertainty, and what customer signal would justify moving forward.

- LTV cannot be treated as a fact before customer behavior has been validated.
- CAC cannot be considered a reliable metric before a customer acquisition channel has been proven.
- Retention cannot be assumed before real-world usage has been observed.
- A revenue model cannot be built as though scale has already been achieved.

The logic is currently reversed. The materials present projections that are far removed from reality: billions of euros in profit, millions of users, tens of thousands of businesses, billions in assets within the system, high LTV across multiple user categories, and token issuance. All of this describes a reality that does not yet exist and one that no investor would find credible. In the broader context, it comes across as an attempt to *talk over the issue* and divert attention from the team's true, undisclosed motivations.

Another issue that stands out is the promise of free internal transfers, discounts, or cashback.

We have not yet seen the platform's technical design, but blockchain systems generally follow one of two broad approaches: public networks and permissioned networks.

In public blockchain networks such as Ethereum, transaction costs are expressed through gas fees. Users pay for computational resources, transaction inclusion in a block, and protection of the network against spam. A portion of the fee is burned by the protocol, while another portion serves as an incentive for validators to include the transaction. Therefore, if a product promises users free transactions, it must explain who is actually bearing those network costs: the user, the issuer, or some other party. In economic terms, truly free transactions do not exist, even if they appear free to the end user within the product.

In permissioned DLT networks such as Corda, the logic is somewhat different. There is typically no public gas market or universal validator fee mechanism comparable to Ethereum. Instead, the costs are

embedded in maintaining and operating the underlying infrastructure.

In either case, the question remains the same: if transactions are free, who is funding that economic subsidy? Without a clear answer, the financial model appears to assume that costs simply disappear.

5. The competitive analysis is fundamentally flawed

I will not go into detail regarding what is wrong with the competitive analysis. Everything about it is wrong. I will limit myself to the single most important conclusion.

The team appears to define competitors using a self-created checklist and seems convinced that if a company does not simultaneously offer individual accounts for both consumers and businesses, fiat and crypto payments, and a long list of other features, then it cannot be considered a direct competitor. Perhaps they are unaware that a competitor is not a company that replicates the entire feature set of a project. In the most basic sense, a competitor is any company that accomplishes the same customer job or solves the same infrastructure problem, often through a solution that is more reliable, less expensive, more regulated, or simply more familiar to the customer.

Based on their own positioning, their actual competitive landscape extends far beyond fintech applications. It includes institutional DLT initiatives within traditional finance. This category encompasses projects led by the BIS and central banks, banking and consortium networks, tokenized deposit platforms, wholesale central bank money settlement infrastructure, DLT platforms for the issuance of tokenized equities and repo instruments, and similar initiatives.

For illustration, I will highlight several such initiatives:

- **Project Agorá** — a project led by the BIS and the Institute of International Finance, involving central banks and major regulated financial institutions. Its objective is to explore how the tokenization of commercial bank deposits and central bank reserves can improve cross-border wholesale payments.
- **mBridge** — an initiative of the BIS Innovation Hub and the central banks of Hong Kong, Thailand, the UAE, China, and Saudi Arabia, focused on enabling cross-border settlement through multiple CBDCs.
- **Partior** — an interbank settlement network established with the participation of DBS, J.P. Morgan, Standard Chartered, and Temasek.
- **Kinexys / JPM Coin** — J.P. Morgan’s blockchain infrastructure for programmable payments and tokenized deposits. Unlike Nostro, this is banking infrastructure developed within one of the world’s largest financial institutions.

And many other initiatives typically emerge from the innovation hubs of the world's largest banks and industry consortia. They possess everything required to address challenges of this kind, including institutional influence, the ability to align the perspectives of major market participants on innovation, and the capacity to coordinate efforts across the industry's largest players.

Against this backdrop, the claim that *"we could not find a competitor because nobody has assembled the same set of features"* does not withstand scrutiny. It suggests that the team is comparing itself based on product form rather than the customer job being performed or the underlying institutional problem being addressed. In this market, competitive advantage is built precisely on trust, access to

final settlement infrastructure, legal certainty, and operational reliability.

6. Lessons from previous initiatives

Even if we strip away the fictional construct and assume, for the sake of argument, that we are discussing a real project and a genuine technical initiative, let us further assume that after numerous iterations the economics and underlying logic somehow converge into a coherent model. Even then, another factor must be taken into account.

The market has already seen projects that successfully demonstrated the technological and operational effectiveness of a particular model, yet failed to transform that success into a scalable business. A notable example is **we.trade**, a banking consortium initiative built on Hyperledger Fabric to digitize trade finance for small and medium-sized enterprises.

In this case, the problem was not that the technology failed to work. Hyperledger Fabric proved to be both a reliable technological foundation, and its smart contracts genuinely accelerated settlement processes while reducing the need for manual document handling. The problem emerged at the level of adoption and network economics. Small and medium-sized businesses were not willing to change their operating processes without a complete ecosystem of logistics providers, insurers, customs participants, and banks.

As a result, the project failed to reach **critical mass of participants**, did not achieve the required level of liquidity or fee-based revenue, and faced demand fragmentation across competing initiatives. With less than €2 million in revenue generated between 2020 and 2022 and operating expenses of approximately €1 million per month, the shareholder banks declined to fund another financing round. In May 2022, the consortium entered liquidation.

As experience in this market demonstrates, the technology itself may well be capable of improving a specific process. The key risks lie elsewhere: whether it is possible to assemble a sufficient number of participants around that process, change their behavior, integrate the solution into existing workflows, and establish a repeatable economic model capable of supporting a business with sustainable cash flows.

The primary barrier to scaling is not technology alone, but the fragmentation of networks and standards. When different participants build incompatible ledgers, tokens, and protocols, network effects break down. Banks are then forced either to connect to multiple networks simultaneously or to wait for the emergence of a dominant standard. In this respect, Nostro's private model does not solve the fragmentation problem.

Recommended next steps

If concerns regarding bad faith are set aside and the focus is placed solely on business logic, the core mistake is that the team is attempting to enter this market through its own infrastructure narrative rather than through what should come first: a validated customer job, a narrowly defined financial process, and a measurable economic inefficiency.

The most successful projects focused on a single measurable process: accelerating settlement, reducing trapped liquidity, automating reconciliation, or decreasing manual processing requirements.

Unsuccessful projects, by contrast, often failed due to a lack of critical user mass, weak integration into real-world workflows, and an inability to establish a repeatable economic model.

Accordingly, as part of addressing these shortcomings, the team should be able to answer the fundamental questions any Pre-Seed or Seed investor would ask:

1. Who is the specific customer?
2. What customer job is currently being performed poorly?
3. Why is the existing alternative inadequate?
4. What pilot can realistically be conducted within the applicable regulatory framework?
5. Who is willing to provide workflow access, data access, or a paid commitment?
6. Why has this particular team earned the right to enter the market?
7. Why is a proprietary two-token model necessary?
8. How can the initial use case expand into a large and scalable market?

Until such evidence is produced, the project cannot be considered ready for presentation to investors. As part of a potential pivot, the team should be willing to narrow its focus to a single ICP, a single use case, a single regulated pilot that can be conducted without obtaining a license, within the framework of a regulatory sandbox, and a single measurable economic outcome.