




2026

PYTH


THE PRICE OF EVERYTHING

VOLUME 5

PYTH.NETWORK



PYTH



From an idea no one understood to the infrastructure **markets need.**

Over the past five years, Pyth has embarked on an extraordinary journey, fundamentally transforming the financial landscape by reimagining infrastructure from the ground up.

It started as an oracle on Solana, solving a clear need for reliable price data in early on-chain markets.

What began as a crypto-native solution expanded to include equities, commodities, and other real-world assets, long before there was clear demand for them.

That early approach shaped what Pyth is today: a global market data network designed to support a wide range of markets, across different asset classes and use cases.


This report looks at that progression. It covers how the system was built, how it adapted as the market changed, and how it moved from being seen as an oracle to becoming a broader data layer.



Table of Contents

PYTH: A NEW ERA FOR GLOBAL MARKETS	04
THE ORIGIN: BUILDING IN A MARKET THAT DIDN'T EXIST	06
THE FIRST BREAKTHROUGH: THE ORACLE THAT WASN'T	08
A DIFFERENT MODEL: MARKET DATA, BY DESIGN	09
SCALING, NOT PIVOTING	11
THE MOMENT EVERYTHING CLICKED	12
FROM CRYPTO TOOL TO FINANCIAL INFRASTRUCTURE	13
PYTH'S PRODUCT EXPANSION IN 2026	15
THE TEAM BEHIND THE NETWORK	17
MARKET RESET: THE NEXT FIVE YEARS	18

PYTH



Pyth: a new era for global markets

Five years ago, Pyth began with a simple conviction: the world's financial infrastructure was changing faster than its market data systems.

That premise now looks less like an insight and more like a reality.

Markets are becoming more global, more digital, more automated, and more continuous. Trading no longer belongs to one geography, one asset class, or one set of working hours. Capital moves across crypto, equities, commodities, FX, rates, and prediction markets. Institutions increasingly expect infrastructure that is programmable, interoperable, and available on demand.

And a new generation of applications, from exchanges to autonomous agents, does not consume data through static terminals. It consumes data through APIs, live feeds, benchmarks, and machine-readable infrastructure.

Since April 2021, Pyth has grown from an ambitious new network into a broader market data platform built around a different idea: price data should come closer to the source, travel faster, cover more markets, and be accessible through a model designed for the next several decades, not the last several ones.

In 2026, that idea enters a new phase with the launch of the Pyth Terminal, the continued expansion of the Pyth Data Marketplace, and a product suite increasingly built for both institutions and the software systems that will shape tomorrow's markets.

Pyth spent five years building the infrastructure for a new market data model, and 2026 marks an important milestone in advancing that vision through new products, expanded capabilities, and broader access.

PYTH

A data business from day one

“We set out to build something that didn’t exist yet, a data business designed for a different kind of market.

Not just faster or cheaper, but structured in a way that could actually support how financial systems were evolving.

For years, it felt like we had built the right thing for the wrong moment. We had high-quality, licensed data across equities, commodities, and real-world assets, and almost no one needed it.

The market was focused on crypto, where data was abundant and often free, and the infrastructure we were building felt ahead of demand.

That gap forced a constant question: had we misjudged where the market was going, or were we simply early?

The answer became clear as the market changed.

As real-world assets entered the system, as institutions began to participate, and as new financial models emerged, the same data that had been ignored became essential.

What had once required explanation became obvious.

The model didn’t change. The market did.

And when it did, everything aligned.”

Michael Cahill
CEO, Douro Labs



PYTH

The origin: building in a market that didn't exist

Pyth did not begin with a product. It began with a structural gap.

As new financial systems started to emerge, particularly within blockchain environments, a fundamental limitation became clear: these markets lacked access to reliable, licensed, and high-quality data.

Without that foundation, market formation itself becomes constrained.

Price discovery cannot function without trusted inputs, and liquidity cannot scale without credible reference points.

At the time, existing solutions focused primarily on validating data, rather than sourcing and distributing it at scale with institutional-grade data rights and verification.

This raised a more fundamental question: how can markets develop before the underlying data infrastructure exists?

The answer did not lie in liquidity alone, but in price.

Establishing a trusted, real-time price layer made it possible to support early market activity, enabling participation before traditional liquidity thresholds were reached.

In doing so, it created the conditions for entirely new financial systems to form and grow.

This principle became foundational to Pyth's design.

P Y T H

From 2021 onward into the future

2021–2023

In the beginning, Pyth challenged traditional assumptions about market data distribution, demonstrating that first-party, direct-from-source pricing could work at scale.

2024–2025

Pyth expands its coverage beyond crypto, strengthens publisher relationships, enhances performance, and evolves towards universal market infrastructure instead of remaining a niche network.

2026

Pyth is now showcasing a credible alternative model for accessing, distributing, and consuming market data with its Terminal, Data Marketplace, and 24/7 market products, moving beyond just backend development.

“In its early days, Pyth’s focus was on creating a reliable system for delivering high-quality, scalable on-chain data, rather than redefining categories.

Initially, the primary challenge seemed technical, confirming price accuracy, but it soon became clear that sourcing data was more fundamental for real financial use cases.

Pyth was designed to aggregate data from institutions, support various asset classes, and provide on-demand access.

This architecture enabled the development of applications requiring high-frequency, low-latency pricing.

Validation came first from developers learning to use it and later from market recognition, evolving from an alternative oracle model to a broader system for data structure and distribution across markets.”

Marc Tillement

Network Director, Pyth Data Association



P Y T H

The first breakthrough: the oracle that **wasn't**

At the time Pyth was introduced, blockchain oracles were limited in their role, primarily focusing on validating data from a single source without addressing sourcing, licensing, or scalability issues.

This led to constraints in accessing high-quality institutional data, meeting licensing needs, and scaling across various assets and use cases.

Pyth redefined the oracle as a data network by sourcing data directly from institutions, aggregating it at the origin, and distributing independently verifiable data, which significantly enhanced system scalability.

Pyth introduced a new model to enhance oracle systems by adopting a “pull” architecture instead of continuous on-chain updates.

Key features include off-chain signing and data availability, user retrieval only when necessary, and on-chain writing at the point of use.

This approach resulted in higher update frequencies, support for more assets, and reduced operational costs at scale. It enabled the support of hundreds of feeds updating multiple times per second, marking a transition to a scalable market data infrastructure.




“We realized early on that the true opportunity lay not merely in constructing an oracle, but in developing a robust and scalable data system.

This insight led us to shift our entire architectural approach, anchoring it around this transformative idea.

As a result, everything else naturally fell into place, enhanced performance, comprehensive coverage, and the capacity to scale as the network grew.

Jayant Krishnamurthy
CTO, Douro Labs

PYTH



A different model: market data, by design

WHY THIS MODEL WORKS

Pyth's advantage lies in its design approach to market data, not just in incremental improvements in speed or cost.

Traditional market data systems are fragmented, with multiple providers and layers of processing. This creates complexity, increases costs, and makes it harder to maintain consistency across different use cases.

Pyth restructures this model by creating a unified data layer. Instead of relying on multiple intermediaries, it aggregates data directly from trading firms, exchanges, and liquidity venues into a single system.

This simplifies how data is produced and delivered. It reduces fragmentation, removes redundant processes, and allows institutions to access consistent, high-quality data with lower operational overhead.

DESIGN ADVANTAGES

Direct sourcing from institutions

Data is contributed by original producers, improving quality, transparency, and reliability.

Cross-asset coverage

A single system supports equities, FX, commodities, futures, and digital assets, removing the need for separate integrations.

High-frequency delivery

Data is updated at millisecond intervals, supporting latency-sensitive applications.

Programmable access

Designed for modern consumption, data can be integrated directly into applications, exchanges, protocols, and AI systems.

Operational efficiency

By consolidating multiple vendor relationships into one layer, Pyth reduces duplication, reconciliation, and internal complexity.

PYTH

“Pyth’s model emphasizes scalability, performance, and long-term adaptability over short-term optimization, catering to the demands of high-frequency, cross-asset markets.


Initial system limitations in latency, throughput, and reliability led to a fundamental redesign focused on ultra-low latency, high availability, and real-time data processing. This infrastructure now handles hundreds of thousands of price updates per second, allowing for continuous growth without major architectural changes.

A key advantage of Pyth is its direct collaboration with data publishers, which minimizes inefficiencies common in traditional market data systems. This direct sourcing enhances the model’s value as it expands into new asset classes and geographies. Pyth’s early inclusion of real-world assets positioned it uniquely as demand for non-crypto data grew, allowing quick adaptation to market needs.

The architecture supports continuous expansion, accommodating new assets and higher data volumes without increasing complexity. This flexibility, combined with a focus on automation, helps maintain operational efficiency as the network grows. It is designed to adapt to evolving financial systems, delivering high-quality data efficiently and sustainably as new opportunities emerge.”



Ali Behjati
Core Contributor, Pyth Network



Scaling, not pivoting

Pyth's evolution in the past years is marked by technical milestones, but internally, it reflects an increased ability to scale rather than a change in direction.

As the network expanded, it faced new constraints, starting with its reliance on Solana. The introduction of Pythnet addressed control and throughput but introduced performance and maintainability challenges.

Later iterations focused on latency, frequency, and efficiency, refining the system for more sophisticated users. This process continuously extended the original design to meet greater demand and performance expectations, aiming for a scalable, real-time, globally distributed market data layer.

“At the core, it’s the same thing.

The way I see the history of Pyth, especially from an engineering perspective, is not a history of pivoting, but a history of scaling.

At each stage, we reached a bottleneck, whether it was throughput, latency, or reliability, and we had to rethink how the system was built to go further.

The product itself didn’t fundamentally change, it has always been about aggregating data from first party publishers and distributing it in a way that works for modern systems. What changed was how far we could push it.”

Guillermo Alapont
Pyth Pro Core Team



P Y T H

“The model behind Pyth is a completely new paradigm.

We’re bringing in data owners who have never monetized their data before, aggregating it, and distributing it.

People from traditional finance understand the problems immediately, but early crypto users initially faced a learning curve, as those constraints didn’t exist yet.

However, as the ecosystem has evolved, users have quickly adapted to these new concepts.”

Abhimanyu Bansal

Head of Business Development,
Pyth Data Association



The moment everything **clicked**

The shift in Pyth’s model occurred in market understanding rather than product design. Initially, the model was hard to communicate as it didn’t fit traditional categories, requiring extensive explanation of its value. The challenge stemmed from a mismatch with the initial market rather than irrelevance.

Participants from traditional finance quickly grasped the model due to their familiarity with market data complexities, leading to an immediate positive response when they encountered Pyth’s approach.

Early blockchain systems were initially overlooked by traditional market participants due to unrecognized challenges like licensing and data fragmentation.

As the market evolved with real-world assets and centralized exchanges, the need for solutions like Pyth became evident, marking a commercial turning point.

The company’s network growth led to more sophisticated data offerings through advanced analytics and machine learning, establishing Pyth as a key provider for organizations navigating modern financial markets.

PYTH

From crypto tool to financial infrastructure

Pyth's evolution has opened the door to a subscription-based model, allowing clients access to high-performance data tailored to their needs.

It has also fostered deeper engagement with institutional clients, expanding Pyth's reach beyond blockchain-specific products and into a wider array of financial markets.

Unlike traditional market data systems, which are often segmented by asset class, geography, or distribution channel, Pyth breaks free from these constraints by integrating across asset classes.

This integration offers a unified access layer to diverse price feeds suitable for both on-chain and off-chain applications. At the heart of this innovative model is Pyth's publisher network, a collaborative effort of over 138 institutions.

Esteemed firms like Fidelity, Jane Street, Binance, Jump Trading, and Cboe contribute directly to this network, ensuring the delivery of high-quality, real-time market data.

Pyth's journey from a crypto oracle to a cross-market data infrastructure highlights its pivotal role in modern financial markets. It stands not just as a data provider, but as a foundational layer that supports a wide range of financial applications.

 CALADAN

AUIROS

 LMAX
EXCHANGE

 Tradeweb


 TWO SIGMA

Revolut


 SGX
FX


 SUSQUEHANNA


AMBER

 EURONEXT

Fidelity

 WINTERMUTE


 Polymarket

 TOWER
RESEARCH CAPITAL

 Swaap

Tokka Labs

flowdesk™

 imc

Pyth figures

138+ **first-party publishers** providing direct market data

710+ businesses using Pyth feeds globally

60% of the **on-chain perpetuals market** powered by Pyth

3000+ **machine-ready feeds** now live for AI agents in Claude & Cursor

The first **24/7 Oil Index**, and more to come

Thousands of **agents** querying Pyth data daily

P Y T H

Pyth's product expansion in 2026

Pyth's trajectory is increasingly defined by a single unifying idea: modern financial markets require modern data infrastructure.

This strategy is expressed through a set of integrated products, each addressing a different dimension of the market data lifecycle.

Pyth Pro represents the commercial foundation of this model. It introduces subscription-based access to high-performance data, with millisecond-level updates and institutional-grade reliability.

The **Pyth Marketplace** extends this model further by creating a distribution layer for proprietary datasets. Institutions can contribute and monetize their data directly, while consumers gain access to a broader and more diverse set of information through a unified interface.


Beyond pricing individual assets, Pyth is also expanding into the construction of new financial primitives.

Its indices initiative introduces continuous, 24/7 pricing across asset classes such as commodities, equities, and multi-asset baskets, reflecting a view of markets that are no longer constrained by traditional trading hours.


At the same time, the network is adapting to a new category of users. Alongside institutional participants, Pyth is increasingly **serving machine-driven systems** that rely on real-time, high-quality data inputs. This is reflected in its focus on machine-readable formats, simplified licensing models, and direct integration into AI-driven workflows.

Taken together, these developments form a coherent strategy.

Pyth is building infrastructure for continuously operating markets, where data flows directly between producers and consumers, and where participation expands beyond legacy constraints to include both institutional and machine-driven actors.




PYTH



Source: Screenshot of the Pyth app interface

P Y T H



“As the network expanded, it became clear that scaling responsibly meant putting durable operating foundations in place early, clear ownership, disciplined processes, and the controls needed to support institutional standards.

Operational excellence is fundamental to delivering reliable market data infrastructure. As market makers, exchanges, and increasingly automated systems depend on Pyth’s price feeds in real time, consistency and resilience become essential rather than optional.

That means designing for continuous operations, robust counterparty and risk management practices, redundancy by design, and the operational controls needed to support institutional standards.”

Ciaran Cronin
COO of Douro Labs

PYTH

“The trajectory is clear.

Expand coverage.

Increase depth.

Simplify access.

And build a system capable of
supporting any market,
anywhere, in real time.”

Michael Cahill

CEO, Douro Labs

The market reset: the next five years

P Y T H

Pyth is committed to becoming the most comprehensive and scalable source of market data, covering all asset classes and geographies worldwide.

By providing coverage for continuously operating markets alongside traditional financial markets, Pyth aims to broaden the range of data available through a unified infrastructure.

This expansion involves the continued growth of the network's coverage, supporting an increasingly diverse set of assets and use cases. The network is designed to scale with demand while maintaining high-frequency updates and reliable performance.

Over time, Pyth aims to support a broad range of market participants, from institutions to autonomous systems, by providing market data infrastructure that can operate across multiple asset classes and environments.

Through continued expansion and investment in data quality, coverage, and accessibility, Pyth seeks to play an increasingly important role in the evolving market data ecosystem.

In doing so, Pyth aims to make market data more accessible and available to a wider range of users and applications.

2026 marks five years of building Pyth.



It is also the starting point for what comes **next**.



10th of July 2026

Pyth Data Association

Grabenstrasse 25, 6340 Baar, Switzerland

PYTH.NETWORK