

THE PRIVATE MARKETS FORUM

The Data Infrastructure Private Credit Needs

Key Takeaways from The Private Markets Forum
webinar held in partnership with AlphaStream AI

POWERED BY



Lorem ipsum dolor sit amet,
consectetur adipiscing elit.
Vivamus ante felis, fermentum

Executive Summary

Private credit has grown into an approximately \$3.5+ trillion asset class¹, on pace to reach \$5 trillion by 2029², yet much of the infrastructure supporting it was built for a significantly smaller market.

Private credit secondaries, for their part, have nearly doubled to \$20 billion in 2025, up from \$10.9 billion in 2024³. It means that in many ways, the industry's next challenge is no longer simply deploying capital. Increasingly, it is building the data infrastructure required to support underwriting, portfolio management, and liquidity at scale.

On June 5, The Private Markets Forum convened senior technology, data, and investment leaders from across the private credit ecosystem, in partnership with AlphaStream AI, to examine the infrastructure required to support this next stage of growth. While the discussion used the rapid growth of private credit secondaries as its starting point, it quickly broadened into a larger conversation about how structured data is reshaping underwriting, portfolio management, and liquidity across the investment lifecycle.

The firms that define this market over the next five years won't be the ones with the most capital. They'll be the ones with the best data."

The discussion's central message: as private credit scales, the firms best positioned for the next phase of growth will be those that can transform fragmented credit information into structured, trusted data, and that increasingly means favoring adaptable, AI-native data models over one-size-fits-all industry standards.

Across the industry, the challenge is the ability to access, connect, and apply that information consistently. Critical data often remains dispersed across loan agreements, financial statements, monitoring reports, fund documentation, and internal systems, making it difficult to evaluate risk, monitor portfolios, and execute transactions with the speed and confidence that an increasingly complex market demands.

¹ Alternative Credit Council (AIMA) and Houlihan Lokey, Financing the Economy 2025. aima.org

² Morgan Stanley, "Private Credit Outlook: Estimated \$5 Trillion Market by 2029." morganstanley.com

³ Evercore, 2025 Credit Secondary Market Survey, as reported in Pensions & Investments, "Private credit secondaries market surges to \$20 billion in 2025, almost double previous record," Jan. 29, 2026. pionline.com

Five themes emerged consistently throughout the discussion:

PRIVATE CREDIT'S NEXT PHASE OF GROWTH DEPENDS AS MUCH ON DATA INFRASTRUCTURE AS CAPITAL.

As the asset class scales toward \$5 trillion, fragmented information is becoming a binding constraint on growth.

SECONDARIES ARE THE INDUSTRY'S FIRST REAL STRESS TEST OF THAT INFRASTRUCTURE.

LP-led deals expose it as a breadth problem (hundreds of borrowers spread across many fund positions), while GP-led deals expose it as a depth problem (concentrated positions where a single borrower can meaningfully move pricing).

THE MARKET IS CHOOSING ADAPTABLE DATA MODELS OVER RIGID STANDARDS.

Rather than waiting on industry-wide data standards, the previous decade's default approach, firms are betting on flexible, AI-native ontologies that absorb inconsistent reporting as-is rather than requiring standardization up front.

BUILD-VERSUS-BUY IS BEING REDEFINED BY AI.

As AI lowers the cost of building, firms are learning to reserve build for the connective data layer that creates genuine edge, and buy or configure everything else - including core systems of record - off the shelf.

DATA STEWARDSHIP IS BECOMING A CONTINUOUS OPERATING DISCIPLINE, NOT A PERIODIC GOVERNANCE EXERCISE.

Automated checks increasingly need to run monthly or weekly, with a human resolving flagged inconsistencies, replacing the old model of occasional, all-at-once data cleanses.

Why This Matters Now

Private credit has reached an inflection point. Over the past decade, the asset class has expanded rapidly, attracting new pools of institutional and private wealth capital while financing an increasingly diverse range of businesses and assets. As strategies have matured and portfolios have grown more complex, the demands placed on firms' data infrastructure have changed just as quickly.

Until recently, fragmented information was largely an operational inconvenience. Investment teams could compensate with institutional knowledge, manual processes, and additional resources. Today, that approach is becoming increasingly difficult to sustain. Larger portfolios, more sophisticated investors, expanding product offerings, and a growing secondary market all require firms to evaluate, monitor, and transfer credit risk with greater speed, consistency, and transparency.

Secondaries illustrate this shift particularly well. Buyers are increasingly underwriting portfolios containing hundreds of loans supported by thousands of pages of documentation, often within compressed diligence timelines. The challenge is rarely a lack of information. More often, it is that critical information remains fragmented across legal agreements, borrower financials, monitoring reports, side letters, and internal systems, making it difficult to assemble a complete and reliable view of risk.

The comparison to more mature credit markets is instructive. The broadly syndicated loan market is a similar size to direct lending, but its secondary market traded roughly \$821 billion in 2024⁴, with trades reported, cleared, and settled within days. Private credit secondaries, by contrast, still run largely on bilateral negotiation and spreadsheets: a gap that is a data and infrastructure problem, not simply a smaller, younger market catching up on its own.

The implications extend well beyond secondaries. The same data challenges influence underwriting, portfolio monitoring, investor reporting, portfolio construction, and liquidity planning across the investment lifecycle. As private credit continues to scale, structured and trusted data is becoming foundational infrastructure that enables better decisions across every one of these functions.

For technology, operations, and data leaders, this represents an important shift in priorities. The objective is no longer simply digitizing documents or automating individual workflows. It is building connected data foundations that allow information to move seamlessly across teams, systems, and investment processes, creating infrastructure that can support both today's operating requirements and tomorrow's AI-enabled capabilities.

⁴ LSTA, "Secondary Trading Volume Soars to \$821 Billion in 2024," Jan. 24, 2025. lsta.org

Five Insights Shaping Private Credit's Next Phase of Growth

1. Private Credit's Next Phase of Growth Depends on Infrastructure.

What's Happening

Private credit has matured into one of the largest and fastest-growing segments of private markets. Institutional adoption has accelerated, new pools of private wealth capital continue to enter the asset class, and investment strategies have expanded well beyond traditional direct lending. As the market scales, however, the operational demands placed on firms have become significantly more complex.

Many of the systems supporting underwriting, portfolio management, and liquidity were built when portfolios were smaller, reporting requirements were simpler, and transactions could rely more heavily on manual processes. Today, those same approaches are being stretched across larger portfolios, more sophisticated products, and increasingly compressed decision timelines.

Why It Matters

The industry's next competitive advantage will be determined less by access to capital and more by the ability to scale decision-making. Firms that invest in modern data infrastructure will be better positioned to evaluate opportunities, manage risk, and support future growth without continuously adding operational complexity.

“That tape is at once the most important document in the deal, but it happens to be the least trustworthy. There's no standard taxonomy. Every GP defines it a bit differently, calculates coverage on different bases. One firm's performance is another's watch list.”

What to Consider

- Infrastructure is becoming a strategic differentiator, not simply an operational necessity.
- Scaling private credit increasingly requires scalable information, not just scalable capital.
- Investments in data infrastructure create long-term advantages that extend well beyond operational efficiency.

2. Secondaries Are a Stress Test of That Infrastructure

What's Happening

The rapid growth of private credit secondaries is exposing infrastructure challenges that have existed throughout the asset class for years. Buyers increasingly need to evaluate portfolios containing hundreds of loans, multiple fund structures, and thousands of pages of supporting documentation within compressed diligence timelines. While the information exists, it often arrives in inconsistent formats with limited standardization, making it difficult to assess risk quickly and confidently.

The nature of that challenge also depends on deal type. LP-led transactions tend to be a breadth problem, diligencing many fund positions, each with hundreds of underlying borrowers. GP-led transactions, which now represent the majority of deal flow, tend to be a depth problem: fewer, more concentrated positions where a single borrower can meaningfully move pricing, but with access to richer disclosure such as credit agreements and valuation models.

Why It Matters

The same infrastructure required to accelerate secondaries also strengthens primary underwriting, portfolio monitoring, investor reporting, and ongoing portfolio management. Firms that solve these challenges for liquidity events will often improve operational performance across the entire investment lifecycle.

What to Consider

- Secondaries are revealing infrastructure limitations that increasingly affect every stage of private credit investing.
- Breadth challenges (LP-led) and depth challenges (GP-led) call for related but distinct data solutions.
- Solving data quality challenges for secondaries creates benefits that extend well beyond liquidity transactions.

3. The Market Is Choosing Adaptable Ontologies Over Rigid Standards

What's Happening

For years, the assumed fix for fragmented private credit data was industry-wide standardization: common templates, shared taxonomies, and agreement across GPs, LPs, and asset servicers on how information should be reported. That view is shifting. Leading practitioners increasingly believe formal standardization efforts move too slowly for a market built on bespoke deal terms and unique fund structures, and that AI has changed the calculus entirely.

Instead of forcing every GP, fund, and borrower to report in the same format, firms are building flexible, dynamic ontologies: a common underlying model that can absorb information as it actually arrives, whether in a 400-page loan agreement, a fund document, a side letter, or a monitoring report, and normalize it after the fact, rather than requiring standardization before the fact.

“I used to think some sort of consistency, some sort of standard was required. I think the technology is there to actually allow us to almost leapfrog that. Is it still a challenge today? Absolutely, because everything isn't 100% there. But just at the pace things have been changing, I think it's going to allow us to leapfrog any need for standardization.”

Why It Matters

This is a meaningfully different posture than the last decade's attempts at market-wide data standards, most of which struggled to gain adoption. Betting on document intelligence and adaptable ontologies rather than waiting for consensus standards means firms can act on the data they can actually get today, rather than the data they wish the market would agree to produce. What matters is a common ontology: one standardization layer that can power document intelligence, workflow intelligence, and ultimately decision intelligence on top of whatever information a counterparty provides.

What to Consider

- Don't wait for the market to agree on a standard; invest in models flexible enough to work with the data you actually receive.
- Greater LP and secondary-buyer demand for transparency is pushing more disclosure into the market regardless of formal standards; firms that can already ingest it will benefit first.

4. Build, Buy, or Both: AI Is Redrawing the Calculus

What's Happening

AI has made it meaningfully cheaper and faster to build proprietary tools, which is reopening a build-versus-buy decision many firms thought was settled. Systems of record (CRMs, investment books of record, core operating systems) are generally not worth rebuilding; the value in AI here is removing the friction of configuring those systems around a firm's own workflows. Where firms are choosing to build is in the connective layer: the data platform that links origination, monitoring, and financial data together, and the agentic tools that orchestrate work on top of it.

“The question is always: where does building allow you to capture some sort of edge or specialization that differentiates you versus your peers? And where is there basically standardization that doesn't add much value? AI is making that decision a little bit harder, because you can build things much more easily than ever before, but you should still only be customizing where it gives you a real edge.”

Why It Matters

Historically, secondaries in particular have been underserved by private-markets technology vendors, most of which were built for traditional LP-side workflows. AI-native, vertical-specific tools are increasingly filling that gap, sitting alongside legacy systems of record rather than replacing them, and interacting bidirectionally with existing platforms rather than requiring a rip-and-replace.

What to Consider

- Reserve internal build capacity for the parts of the stack that create genuine competitive edge.
- Budget for third-party, AI-native tools that solve one vertical problem well, in addition to core SaaS systems of record.
- A central data platform (increasingly a data lakehouse spanning structured and unstructured data) is what makes both build and buy options viable at scale.

5. Data Stewardship Is Becoming an Operating Discipline, Not Just a Governance Checkbox

What's Happening

Firms report a growing emphasis on data stewardship: less about controlling or restricting data, and more about actively owning its quality. In practice this means combining traditional maker-checker controls (increasingly assisted by AI) with a separate, systematic layer of automated checks, borrowed from how public markets manage data at scale, that flags inconsistencies for a human to resolve, rather than relying on periodic, all-at-once data cleanses.

Why It Matters

AI tools amplify whatever they're given. Firms that treat stewardship as a continuous, owned discipline, rather than an occasional cleanup exercise, are the ones able to actually deploy AI safely on top of their data, because the checks that used to run every few years now need to run monthly, weekly, or even daily.

“You can have all the tools, but it basically gets back to garbage in, garbage out. If you don't have someone sitting on top of the data, the human in the loop making sure that it's complete, accurate, and timely, any investment you're putting into technology or AI is worthless.”

What to Consider

- Pair AI-assisted maker-checker workflows with a separate, systematic data-quality check that runs independently of the people entering the data.
- Someone in the organization needs to own data quality explicitly: stewardship works best as a named responsibility, not a shared assumption.
- The payoff compounds: AlphaStream's own team noted its tools could add little value until two years of data organization and cleansing work had been done first.

Practical Takeaways by Function

The infrastructure agenda above plays out differently depending on where you sit. A few starting points for each audience:

CTOs, CDOs, and data leaders at GP credit platforms: Prioritize a connected data platform over point solutions; reserve build budget for the layer that creates genuine differentiation, and buy or partner for the rest.

Operations and technology teams preparing for continuation vehicles or portfolio transfers: Treat loan-tape trust and taxonomy consistency as a pre-transaction workstream, not a diligence-week scramble.

LP technology and data teams managing private credit liquidity: Build the infrastructure to overlay a prospective deal on existing exposure, concentration limits, and forecasted cash flows, before, not after, committing capital.

Secondaries firms evaluating document intelligence: Weigh vendors on how well they normalize inconsistent, non-standard reporting as it exists today, not on how well they support an idealized standardized format that may never arrive.

Conclusion

Private credit's next phase of evolution will be shaped as much by infrastructure as by investment strategy.

As the asset class continues to grow in scale and complexity, competitive advantage will increasingly depend on a firm's ability to transform fragmented information into trusted, connected intelligence that supports better decisions across the investment lifecycle. The organizations that differentiate themselves are unlikely to be those with the largest technology budgets or the most sophisticated AI models. Instead, they will be those that have built the underlying foundations that allow information to move seamlessly across underwriting, portfolio monitoring, portfolio construction, investor reporting, and liquidity.

This represents a broader shift in how technology investments create value. Historically, firms focused on improving individual workflows or solving isolated operational challenges. Increasingly, the greater opportunity lies in establishing common, adaptable data architectures, moving from document intelligence, to workflow intelligence, to decision intelligence, that connect investment processes, reduce operational friction, and create reusable information that becomes more valuable with every transaction, portfolio update, and investor interaction.

For technology, operations, and data leaders, the implication is clear: Preparing for the future of private credit is not simply about implementing AI or digitizing existing processes. It is about building trusted information foundations, and the stewardship discipline to keep them trustworthy, that let the organization adopt new capabilities, respond more quickly to changing market conditions, and scale without continuously adding complexity.

Ultimately, the firms best positioned for the next phase of growth will not be defined by how much information they collect, but by how effectively they organize, connect, and apply it.

As private credit continues to mature, trusted data foundations will increasingly become the operating layer that enables more intelligent investment decisions, more resilient operations, and more efficient liquidity across the market.

POWERED BY

