

Buy vs. Build: A Strategic Framework for Loan Analytics Software at Credit Unions

Executive Summary

Credit unions face an increasingly complex lending environment. Interest rate volatility, evolving member behavior, expanding origination channels, and mounting regulatory expectations have made sophisticated loan analytics not just a competitive advantage — but an operational necessity. At the center of this challenge lies a fundamental technology decision: should a credit union buy a purpose-built loan analytics solution, or invest in building one internally?

While building a proprietary solution offers theoretical benefits in customization and flexibility, a thorough analysis reveals that for the vast majority of credit unions, purchasing a proven solution delivers superior outcomes — faster time to value, lower total cost, reduced operational burden, and meaningfully lower technology risk. Even in an era of rapid advances in artificial intelligence and low-code development tools, the buy case remains compelling and, for most institutions, decisive.

The Loan Analytics Challenge

Before evaluating the build-versus-buy question, it is worth appreciating why loan analytics is difficult, and why the difficulty scales non-linearly with portfolio size.

Dirty Data In; Dirty Analysis Out

Cleaning, normalizing and structuring loan data is the first hurdle to building a meaningful loan analytics solution. Moreover, how that data is maintained and stored is not only a privacy and security concern but also a performance issue. Loan analytics beyond simple aggregations, such as weighted average coupons over time, require high volume calculations. These depend upon well-architected databases and intelligent algorithms that use server memory to run the requisite number of base and stress scenarios on the short timelines to be useful for credit unions of even modest scale.

Internal Rate of Return at Scale

Internal Rate of Return (IRR) is among the most important metrics a lender can track. It reflects the true return on a loan, accounting for the timing and magnitude of all cash flows — origination fees, scheduled payments, prepayments, and charge-offs. For a single loan upon maturity, calculating IRR is straightforward. However, for a portfolio of tens of thousands or hundreds of thousands of loans that are at various stages of repayment, calculating accurate IRR ranges at the individual loan level, and then aggregating it meaningfully across segments, vintages, products, and risk tiers, becomes a significant computational and data engineering challenge. Notably, IRR requires knowing all of the cashflows of each loan until maturity — which until maturity or charge-off is unknown. Therefore, IRR requires a lender to project forward all the potential outcomes to the end of the loan life across multiple charge-off and prepayment scenarios and then provide an IRR range, which narrows over time as additional loan behavior becomes known.

That IRR range for each cohort of analysis must be recalculated each month as actual cash flows deviate from projected ones. A solution that cannot dynamically recalculate IRR as prepayments occur, delinquencies emerge, or rate environments shift provides an incomplete historical snapshot rather than actionable intelligence.

Compounding the difficulty is that meaningful investigation requires the ability to change and modify loan cohorts over time. When evaluating a loan programs with dynamically changing underwriting criteria, lenders need to compare and contrast different loan cohorts among multiple key attributes like credit score, DTI, LTV, geography and time.

Prepayment Analysis

Prepayment behavior is equally complex and consequential. For example, while many credit unions have had de minimis losses in their mortgage portfolios for nearly twenty years, nonetheless, they may have lost money on those portfolios due to prepayments depending on the costs of origination. In many cases, credit unions would have earned higher returns investing in government guaranteed securities versus even a no-loss residential loan program. When a member pays off some or all of a loan early, the credit union loses anticipated interest income and must redeploy that capital — often in a different rate environment. Understanding prepayment patterns across loan types, member demographics, geographic markets, and interest rate cycles requires both historical data depth and sophisticated statistical modeling.

At scale, manually tracking prepayment speeds (not just tracking repayments in full), including constructing conditional prepayment rate (CPR) curves, and modeling the impact of prepayment on portfolio yield and duration is not practically achievable in spreadsheets. A credit union managing a consumer loan portfolio may be processing thousands of prepayment events per month — some full and some partial — each of which affects IRR calculations, income projections, and asset-liability positioning.

Analyzing “Building”

In certain circumstances, building proprietary loan analytics infrastructure is the right call for some institutions.

Total Customization is Possible

A self-built solution can be engineered to reflect the exact pricing models, risk frameworks, and reporting structures that a credit union uses internally. If an institution has developed a member scoring methodology, or prices loans based on a unique combination of behavioral and financial variables, a commercial off-the-shelf solution may not accommodate that logic without significant configuration — or at all.

Protecting Proprietary Pricing Models

The largest and most analytically sophisticated lenders — major banks and specialty finance companies — sometimes possess genuine, proprietary pricing models that they want to protect and operationalize in software. These institutions often can afford teams of engineers that structure and maintain the database, develop and maintain back-end code, design and update the user-interface, develop and maintain the front-end code, and data engineers to continually monitor and clean incoming data. With those labor intensive and costly resources, building ensures that the institution's intellectual property is not inadvertently shared with a vendor who also serves competitors, and that the pricing engine can evolve precisely in step with the institution's thinking and the feedback loop provided by the analytics solution.

AI and Modern Development Tooling Can Help

The recent emergence of generative AI coding assistants, low-code platforms, and modern data stack tools (e.g., Power BI, Tableau, Looker, and even Excel) have in many instances lowered the cost and timeline of building internal software. A well-resourced engineering team today can accomplish in months what used to take years. This development has given new life to the internal build argument, and credit unions evaluating their options today should not dismiss it on the grounds of traditional development economics alone.

Internal Development Expertise Required

However, these advantages still require the internal database, back-end, front-end and design expertise even if a firm may need fewer individuals to develop and maintain an effective web application. These conditions do not exist for the overwhelming majority of credit unions. And for those institutions, the build path carries risks that often go under-appreciated in initial planning conversations or prototyping.

Key Personnel Risk

Internal development carries material key personnel risk. Maintaining an unfamiliar code base is difficult for even the most talented and experienced engineering teams. As development professionals transition to new roles out of the firm, new people will be hired to maintain the analytics framework. This personnel churn is costly and frequently can result in static applications that cannot be updated at commercially reasonable expense.

A Case for Buying

1. Total Cost of Ownership

The instinct to build is often driven by a perception that buying is expensive — that vendor pricing is a perpetual tax on the institution. This framing ignores the true total cost of ownership on the build side.

Building a loan analytics platform requires not just initial development, but ongoing investment in:

- **Data engineering:** pipelines to ingest, clean and normalize loan data from core systems, payment processors, and servicing platforms, including all third-party serviced loans.
- **Computation infrastructure:** cloud or on-premise resources capable of recalculating IRR across a multitude of scenarios, prepayment speeds and all other meaningful aggregations across hundreds of thousands of loans on a regular cadence.
- **Software maintenance:** bug fixes, dependency updates, core system integration maintenance as vendors release new versions.
- **Feature development:** ongoing enhancements as business requirements evolve, regulatory standards change, or new loan products are introduced.
- **Talent:** experienced data engineers, back-end and front-end engineers and user experience design professionals to maintain a complex codebase — all of whom need managers and career advancement — a scarce and expensive resource in the current labor market.

When all of these costs are fully loaded, the total cost of ownership for a built solution typically far exceeds the cost of a commercially available purpose-built one.

On the other hand, commercially available solutions distribute their development and maintenance costs across a broad customer base, allowing vendors to offer functionality at a fraction of what it would cost any single institution to build.

2. Time to Value

Speed matters. A credit union experiencing margin compression, entering new markets, evaluating existing products or introducing new ones, preparing for a regulatory examination, or seeking to improve pricing discipline on a growing loan portfolio cannot afford an even a 12-month build cycle — or the more likely multi-year effort it takes to build a fulsome analytics platform. Commercially available loan analytics solutions can typically be implemented and deliver value within weeks to a few months, depending on data complexity and integration requirements.

Time to value is not merely a convenience consideration — it has direct financial implications. Every month that a credit union operates without accurate IRR attribution or prepayment analysis is one in which pricing decisions may be suboptimal, concentration risks may be underappreciated, and income projections may be materially inaccurate.

3. Reduced Internal Maintenance Burden

Credit union technology teams are stretched. Core system administration, cybersecurity operations, digital banking support, and regulatory compliance technology demands consume the bandwidth of most IT departments. Adding a complex, internally-built analytics platform to the maintenance queue along with the need to continually clean and maintain data ingestion pipelines, introduces both operational risk and opportunity cost — pulling technical resources away from other strategic priorities.

Commercially purchased solutions shift the maintenance burden to the vendor. Software updates, infrastructure patches, performance optimization, and integration updates become the vendor's responsibility, contractually defined and commercially incentivized.

This is a non-trivial benefit. In a sector where technology talent is difficult to recruit and retain, preserving internal capacity for mission-critical systems is a strategic imperative. While not all vendors also perform data cleaning, selecting vendors that maintain, clean and normalize the data ingestion process is a critical differentiation to review.

4. No Risk of Building Into Obsolete Technology

Perhaps the most underappreciated risk in the build argument is the danger of technology obsolescence. The data and analytics technology landscape has shifted dramatically over the past decade — from on-premise data warehouses to cloud-native architectures, from batch processing to streaming computation, from traditional BI tools to AI-augmented analytics. Any institution that invested heavily in building an on-premise analytics platform five years ago is now managing the cost and risk of migrating or rebuilding on modern infrastructure.

This cycle will continue. Generative AI, vector databases, real-time computation frameworks, and new data lakehouse architectures will reshape what best-in-class analytics looks like over the next five years. A credit union that builds a solution today risks locking into a technology stack that becomes a liability before the development investment has been fully amortized.

The best commercial vendors, by contrast, are continuously investing in modernizing their platforms — because their survival depends on it. Credit unions that buy effectively outsource the risk of technology obsolescence to a vendor whose entire business model is built around staying current.

5. Proven Methodologies and Regulatory Alignment

Commercial loan analytics vendors have invested years in developing and validating the underlying methodologies their platforms use — IRR calculation approaches, prepayment modeling frameworks, credit loss estimation methodologies. These methodologies have been stress-tested across thousands of institutions, refined in response to regulatory feedback, and validated against real-world outcomes.

A credit union building from scratch must develop and validate these methodologies internally — a process that requires both domain expertise and significant time. Errors in IRR calculation logic or prepayment modeling assumptions are not merely inconveniences; they can result in materially misstated financial performance reporting or flawed strategic decisions.

The AI Factor: Does Generative AI Change the Calculus?

Given the rapid maturation of AI-assisted software development, it is worth directly addressing whether these tools meaningfully shift the build-versus-buy balance for credit unions.

Generative AI tools — including large language model-based coding assistants — can accelerate the initial development and prototyping of a loan analytics codebase. They can help generate data pipeline logic, scaffold reporting dashboards, and assist with documentation. These are real productivity gains that lower the barrier to entry for building.

However, several important limitations remain:

- **AI tools accelerate development, but do not eliminate it.** A credit union still needs experienced technical staff to direct, review, validate, and maintain AI-generated code. The reduction in cost and timeline is meaningful but not transformative for most community-scale institutions. Simply put, prototypes handling a handful of loans are a far cry from production platforms handling millions of computations in real-time.
- **The hardest parts of loan analytics are not coding problems.** Continually ingesting and maintaining clean data, designing a robust IRR calculation framework that handles edge cases correctly across a multitude of scenarios, building prepayment models that are statistically valid across loan types, and ensuring data integrity across integration points are problems of domain expertise and data engineering — areas where AI tools provide limited assistance.
- **AI-generated code still requires ongoing maintenance.** Code written by an AI assistant today will need to be maintained, updated, and extended by human engineers tomorrow.

In short, AI lowers the floor on building but does not lower it enough to change the fundamental recommendation for most credit unions. The complexities and costs of building and maintaining meaningful loan analytics still strongly weigh in favor of buying over building even after taking into account the gains from AI-supported coding.

Conclusion

The decision to buy or build loan analytics software is not primarily a technology question. It is a question of strategic priority and risk tolerance.

For credit unions navigating the complexity of IRR calculation at scale, prepayment analysis across diversified portfolios, and the need for timely, accurate performance reporting, the market offers proven, purpose-built analytic solutions that can deliver value in a fraction of the time and at a fraction of the cost of internal development.

The allure of the build path — total control, custom pricing logic, AI-assisted development — is real but largely theoretical. When weighed against the practical realities of cost, time to market, maintenance burden, and technology obsolescence risk, buying is not merely the pragmatic choice. For the vast majority of credit unions, it is the strategically correct one.

The most sophisticated lenders in the world do not build their own cores. They buy, configure, and integrate best-in-class tools — and direct their internal talent toward the decision making that only humans can do with the analysis and information that commercially available platforms provide.

This article is intended for informational purposes and general strategic guidance. Credit unions should conduct thorough due diligence, including consultation with qualified technology and financial advisors, before making platform investment decisions.

About LoanStreet

Contact us today to learn more about our platform and products. With over 1,300 financial institutions registered and hundreds already receiving our reports, we are ready to help you build a more efficient loan participation model.

 sales@loan-street.com

 www.loan-street.com