

Financial Sector

Case Study



A Major Private Credit
SaaS Platform

Client Background

A SaaS leader in private lending sought to expand into the private credit sector by modernizing how lenders evaluate complex corporate loans. Traditionally, each opportunity required a team of three underwriters to spend 4–6 weeks manually reviewing a vast mix of borrower, guarantor, and company documents—ranging from personal financials and bank statements to org charts, tax returns, insurance certificates, debt schedules, and pro formas—before assembling a standardized loan memo.

This time-intensive, document-driven process was the primary bottleneck limiting scale, speed, and consistency across the underwriting operation.

The Problem



- Private credit lenders relied on manual review of hundreds of unstructured PDFs, spreadsheets, and legal documents, taking teams of three underwriters 4–6 weeks per deal.
- The process severely limited how many opportunities could be evaluated while making it nearly impossible to maintain consistent diligence and audit standards at scale.
- As deal volume grew, risk increased and speed dropped, with no reliable way to trace data back to its source or ensure underwriting consistency.

Additional Constraints

The client had no prior experience in the private credit sector, creating uncertainty around where to start and how to validate the opportunity. XTAM partnered closely through the discovery and design phases, helping identify and onboard design-partner lenders as part of the go-to-market strategy.

This collaborative approach not only accelerated market learning but also built internal sponsorship and executive confidence, ultimately giving the client the conviction to pursue what became a transformative new business line.

What We Did

Discovery & Sourcing

- Partnered with the client's product and underwriting teams to interview potential lenders and define the right wedge into the private credit market.
- Cataloged diligence questions across finance/accounting, business/management, legal/compliance, underwriting/servicing, and technology/infosec – linking each to its corresponding source documents (e.g., AR/AP agings, bank statements, debt schedules, org charts, insurance certificates, UCC filings, SOC/ISO audits).
- Designed document ingestion and tagging processes so the AI could understand document scope, relationships, and context for each deal.

Architecture Build Out

- Developed a secure, GCP-hosted AI pipeline for document retrieval, OCR, classification, and structured extraction, segmented by diligence category.
- Built REST APIs returning structured answers with confidence scores and source citations down to the page level, flagging low-confidence fields for human review.
- Delivered an embeddable conversational co-pilot API, allowing underwriters to query the AI directly (e.g., “Show AR over 60 days for April” or “Where did we get guarantor net worth?”) and cross-reference across document categories.
- Implemented conversation logging per opportunity to maintain audit trails, support compliance, and enable continuous model improvement.

Pilot & Rollout

- Provided white-glove support through the first 10 loan memos, refining parsing, field mappings, and exception workflows.
- Supplied API documentation, an onboarding run-book, partnered with the Private Credit's front end development team to implement the APIs, and a recommended Salesforce integration architecture, ensuring Salesforce remained the single pane of glass.
- Maintained a clear division of responsibilities – client owned UI, authentication, and approvals, while XTAM focused on AI reliability and performance.
- Trained the client's development team to manage and extend the AI infrastructure long-term, building internal capability and ownership.

“We made a big bet with XTAM and it's worth every penny. This will be transformational for our industry.”
– Private Credit CEO

”

Timeline of Results

Day 30

Day 60

Conducted in-depth customer interviews to validate key underwriting pain points and identify the highest-value automation opportunities.

Built a clickable prototype to showcase the end-to-end experience — from document upload to AI-generated memo preview — helping secure early executive sponsorship.

Recruited design-partner lenders to pilot the solution and co-develop the MVP ahead of broader market rollout.

Launched a **live demo with the first design partner**, enabling real-time feedback on document processing accuracy, workflow integration, and user experience.

Drove **rapid iteration cycles** based on partner input, continuously refining extraction precision, conversational capabilities, and API performance on the path to general availability.

Collaborated with the go-to-market team to expand the design-partner program, formalize onboarding materials, and shape the commercialization strategy for broader rollout.

Companies We've Built

quadpay
a zip company

RENT THE RUNWAY

CVS Health
signifyhealth

care/of

MIRROR

CredSimple
andros*

Handshake

Wunderkind

IDPrimary

SLICE

TEEPUBLIC

splice

The Athletic

grovo
BY COCKSTONE

Sealed

FINAL SITE
A Veritas Capital Company

Legion

reonomy

:Dandi

Orange
Hiring

Why XTAM

Proven Track Record

- Multiple successful transformations and exits across SaaS, EdTech, enterprise software, and consumer tech.

Deep AI Expertise

- 10+ years deploying machine learning solutions—even before AI went mainstream—gives us an edge in knowing what works, what doesn't, and how to execute quickly.

End-to-End Approach

- We combine strategic advisory, engineering execution, and product innovation under one roof—ensuring accountability, speed, and ROI.

Flexible Structures

- We meet partners where they are—whether that means a straightforward consulting engagement or an equity-based deal for deeper, long-term collaboration.

Contact Details

Name	Phone	Email
Brian Williams	973-632-8991	brian@xtam.ai

Client Background


The hundred year old PE backed engineering firm is a long-standing leader in coastal and offshore engineering, known for tackling some of the world's most complex environmental and infrastructure challenges. Its geoscience teams collect and interpret sub-bottom seismic (SEG-Y) and related ocean-floor datasets to map key geologic features that inform permitting, engineering design, environmental compliance, and risk assessment.

The firm provides engineering, construction, environmental, sustainability, resilience, and infrastructure solutions to government agencies as well as commercial, industrial, and energy clients. Over more than a century, it has delivered tens of thousands of projects globally—spanning oil and gas, industrial facilities, coastal infrastructure, environmental restoration, and disaster-response—and has made continuous reinvention a defining part of its DNA: *"After 104 years, we're just getting started."*

With a workforce numbering in the thousands and offices across North America, the firm brings deep expertise in geotechnical and marine engineering, infrastructure asset management, and large-scale subsurface surveys for energy and offshore exploration sectors.

Given the growing complexity and scale of ocean-floor and sub-bottom datasets, encompassing seismic acquisition, geologic interpretation, and risk modeling, the manual process of labeling and classifying geologic features had become a critical throughput bottleneck under private equity ownership, directly impacting speed, margins, and scalability.

The Problem



Manual labeling of complex ocean-floor datasets required teams of PhD geologists up to six months per project.

- Each survey line contained thousands of traces with dozens of possible reflectors and geomorphic features to identify and trace as polylines.
- Consistency and auditability were hard to maintain across regions and projects, slowing approvals and creating rework.
- Throughput was the bottleneck for revenue: every incremental project required more expert headcount.

Additional Constraints

- **Label variability & noise:** Some feature picks were faint or inconsistently visible (e.g., channels that appeared only sporadically), creating ambiguity for both humans and machines.
- **Heterogeneous nomenclature:** Equivalent features carried different names across regions (e.g., *Valley* vs. *Valley Fill*; *Seafloor_ASC* vs. *Seafloor*), complicating training and QA.
- **SEG-Y/CSV alignment:** Feature CSVs and SEG-Y traces used different primary keys. Reliable alignment required matching FFID ↔ Ping (not simple trace indices) and honoring polyline continuity rules across gaps.
- **Time/depth semantics:** Picks were stored as two way travel time; horizontal exaggeration and display scaling made visual QA error prone without normalization.
- **Geodesy differences by region:** Surveys spanned multiple NAD83 State Plane zones in US survey feet, with negative source/group scalars in SEG-Y headers, each requiring precise conversion to WGS84 Lat/Lon for mapping and downstream tools.
- **Operational integration:** The solution had to export predictions into the client's workflow (e.g., SonarWiz polyline CSV) and import cleanly on seismic images, with no off-ramp to "demo-only" artifacts.
- **Speed to value:** As a PE backed business, the firm needed a defensible MVP fast, showing measurable throughput and margin impact with minimal lift from already stretched geoscience teams.

What We Did

Discovery & Label Sourcing

- Partnered with the client's geoscience leads to enumerate high-value features for sub-bottom interpretation: Seafloor, Ravinement Surface (RavSurface), Texas Mud Blanket (TMB), Valley/Valley Fill, Northern/Southern Delta Lobes, Laminated Muds, Faults, Channels, Gas.
- Organized a three-tier feature taxonomy by continuity, aligning model difficulty with geologic expression:
 1. Continuous: Seafloor, RavSurface, TMB (present across entire lines).
 2. Semi-continuous: Valley Fill, Northern/Southern Delta Lobes, Laminated Muds.
 3. Sparse/episodic: Faults, Channels, Gas, etc.
- Codified polyline continuity rules: gaps between clicks should render as a single continuous line unless a new header segment denotes a distinct feature.

Data Standardization & Alignment

- Established a deterministic join between SEG-Y and feature CSVs using FFID ↔ Ping (page-accurate alignment), superseding fragile trace-index heuristics.
- Normalized two-way travel time (TimeFromTX) and corrected horizontal exaggeration to ensure apples-to-apples visualization during QA.
- Built a geodesy layer: honored coordinate_units = feet, applied source/group scalars (e.g., /100), and mapped SOU_X/SOU_Y (or GRP_X/GRP_Y) to WGS84 using the correct NAD83 State Plane

zone per region (e.g., TX South vs. South Central, FL North, LA South, AL West, NC statewide).

- Unified naming conventions (e.g., *Valley* → *Valley Fill*) and reconciled region-specific CSV variants (e.g., *Seafloor_ASC* = *Seafloor*).

Modeling & Architecture

- Built a **specialized AI model** trained to detect and label geologic reflectors and geomorphic features (faults, channels, delta lobes, etc.) directly from sub-bottom profiles.
- Designed feature-aware training guided by the continuity taxonomy: prioritized high-signal continuous reflectors for rapid gains; introduced semi-continuous and sparse classes with targeted augmentation to address class imbalance.
- Implemented polyline reconstruction to convert per-trace predictions into continuous, geologically plausible lines that respect the client's continuity rules.
- Exposed confidence scores and source citations (Ping/FFID + TimeFromTX) to support review, audit, and active-learning loops.

Quality, Evaluation & Human-in-the-Loop

- Built per-feature evaluation harnesses (continuous vs. sparse classes assessed with different metrics and tolerances).
- Flagged low-confidence spans for targeted human review and rapid relabeling—tightening the feedback loop without broad geologist time sinks.
- Created validation checks to catch mis-scaled displays, out-of-CRS coordinates, or mis-joined lines before model training.

Pilot, Integration & Rollout

- Delivered SonarWiz-compatible CSV exports of AI-predicted polylines. Initial imports revealed coordinate/CRS mismatches; we corrected scalars and CRS mapping so predictions plotted precisely on seismic images and in map view.
- Provided a lightweight run-book: CRS selection by region, field mapping, import steps, and a labeling QA checklist.
- Enabled a single-geologist review workflow: accept/edit AI polylines, push back to the training set, and re-export—all within the client's existing tool chain.

“

“The XTAM team weren't geologists, but they learned our world faster than anyone we've ever worked with. They came into a 106-year-old business, understood a hyper-specific domain, and completely changed how we operate.

By the time they walked out of the room, they were operating at the level of the top geologists in the industry.”

— CTO of the 106 year old Engineering Firm

”

The Results

Day 30

Fully onboarded into the Geological domain and understood the context of the problem.

Completed end-to-end data normalization: FFID ↔ Ping alignment, TimeFromTX handling, and horizontal scale normalization.

Finalized the feature taxonomy and unified naming across regions.

Day 60

The team shipped the first working prototype capable of producing AI-generated polylines on the Bay County and Texas datasets for internal validation. Early tests demonstrated that even a 20% labeling accuracy would have been transformational—reducing a six-month manual workflow to just weeks.

Through rapid iteration, model refinement, and targeted augmentation of sparse feature classes such as faults, channels, and gas pockets, the system ultimately achieved **~70% labeling accuracy** with high precision on continuous and semi-continuous reflectors. A game changer in this industry.

This milestone enabled the firm to move from multi-geologist manual labeling to a **single-geologist production workflow**, where experts now focus on rapid review and validation rather than repetitive drawing—effectively turning the primary bottleneck into a scalable quality-assurance process.

Companies We've Built

quadpay
a zip company

RENT THE RUNWAY

CVSHealth
signifyhealth

care/of

MIRROR

CredSimple
andros*

Handshake

Wunderkind

Primary

SLICE

TEEPUBLIC

splice

The Athletic

grovo
BY CORNERSTONE

Sealed

FINALSITE
A Veritas Capital Company

Legion

teonomy
an HBS Group business

Dandi

Orange
Hiring

Why XTAM

Proven Track Record

- Multiple successful transformations and exits across SaaS, EdTech, enterprise software, and consumer tech.

Deep AI Expertise

- 10+ years deploying machine learning solutions—even before AI went mainstream—gives us an edge in knowing what works, what doesn't, and how to execute quickly.

End-to-End Approach

- We combine strategic advisory, engineering execution, and product innovation under one roof—ensuring accountability, speed, and ROI.

Flexible Structures

- We meet partners where they are—whether that means a straightforward consulting engagement or an equity-based deal for deeper, long-term collaboration.

Contact Details

Name	Phone	Email
Brian Williams	973-632-8991	brian@xtam.ai

HR Analytics & Workforce Intelligence Case Study

XTAM




A Complex Data Problem
In Workforce Intelligence

Client Background

Dandi is an AI-native global workforce intelligence platform designed to solve the "last mile" problem in HR analytics. While organizations sit on mountains of people data, extracting actionable insights typically requires technical teams to manually crunch numbers across disparate systems.

Dandi structures and analyzes data across the full HR stack—processing billions of metrics daily—to answer complex workforce questions instantly. By leveraging a modern cloud-native architecture, they have positioned themselves as the "brain" of the HR function, delivering the kind of deep, cross-functional insights that previously required teams of data scientists.

The Problem



- **Time-to-Insight Latency:** Critical questions typically takes weeks of manual data extraction and calculation by technical teams.
- **Failure of Standard AI:** Competitor "Prompt-to-SQL" methods achieved only ~60% accuracy, failing on complex, multi-step metrics.
- **Data Silos:** Disconnected systems (ATS, HRIS) made unified employee lifecycle analysis impossible without heavy engineering lift.

Additional Constraints


- **Strict Security & Privacy:** HR data is highly sensitive. The solution required enterprise-level, granular permissions (e.g., ensuring a manager only sees data for their direct reports), which the AI had to respect implicitly in every answer.
- **Zero-Tolerance for Hallucination:** Unlike creative AI use cases, HR decisions regarding compensation and diversity require 100% mathematical precision.
- **Scale:** The system had to pre-aggregate and query billions of data points in real-time, requiring an architecture far more robust than standard SQL databases.

What We Did

Discovery & Architecture


- Rejected the industry-standard "Prompt-to-SQL" approach due to low reliability.
- Designed a Pre-Aggregation Architecture using Google Cloud Platform (GCP). See below

- Pre-Aggregated Architecture: The platform utilizes Cloud Composer, Cloud Dataproc, and Cloud Bigtable to pre-aggregate billions of metrics upon ingestion, ensuring any data question is answered instantly.
- High-Accuracy Retrieval: This method shifts the LLM's role from complex, multi-step calculation (as required by Prompt-to-SQL methods) to maximum two-step retrieval, achieving a benchmark accuracy of 98%.
- Competitive Edge: By instantly leveraging pre-calculated data, Dandi AI delivers complex trends, forecasts, and correlations that competing systems, limited by raw data processing, find too challenging to address.



AI Modeling & Agent Deployment

- Integrated Google Vertex AI and Gemini 1.5 Pro to drive the "Dandi AI Analytics" agent.
- Built a "Planner-Retriever" system: When a user asks a complex question, the AI breaks it down into sub-steps (e.g., Step 1: Get hiring data; Step 2: Get retention data; Step 3: Compare).
- Instead of writing SQL, the AI acts as an orchestrator, selecting the right "tools" to retrieve pre-verified metrics from Bigtable. This reduced the mathematical burden on the LLM, shifting the task from "calculation" to "retrieval and synthesis".



Security & Standardization

- Built a custom identity-aware middleware using **Google Identity Platform** and **Cloud Endpoints**.
- Ensured that every AI query is filtered through the user's specific role and permissions before accessing the database, guaranteeing that no unauthorized data is ever retrieved or summarized.

“

“Work that used to take weeks now takes seconds. Dandi transforms unstructured HR data into accurate, actionable metrics that power the performance of the people team”

– CTO of Dandi

”

The Results

A Complex HR Query

“Which five departments had the highest hiring percentage last quarter vs the same quarter of last year? And what are the same results for retention percentage for the same segments? What are the 5 takeaways from the difference between the hiring and retention results?”

The Result

Which five departments had the highest hiring percentage last quarter vs the same quarter of last year? And what is the same result for retention percentage for the same segments? What are the 5 takeaways from the difference between the hiring and retention results?



Here are the five departments with the highest hiring percentage last quarter (2024 Q1) compared to the same quarter last year (2023 Q1)

Hiring percentage:

Department	2024 Q1	2023 Q1
AE	70.0%	38.89%
Events	10.0%	0.0%
Lead Gen	10.0%	0.0%
Social	10.0%	0.0%
Frontend	0.0%	33.33%

And here are the five departments with the highest retention percentage for the same periods

Retention Percentage:

Department	2024 Q1	2023 Q1
Accounting	100.0%	100.0%
Demand Gen	100.0%	100.0%
Events	100.0%	85.7%
Mobile	100.0%	92.3%
Recruiter	100.0%	100.0%

Five Takeaways:

AE department hiring surge: The AE department saw significant hiring growth in 2024 Q1 compared to 2023 Q1, suggesting a potential expansion of the team

Increased focus on marketing efforts: Events, Lead Gen, and Social departments, all related to marketing, show increased hiring in 2024 Q1 compared to no hiring in those departments in the same quarter last year. This could indicate a strategic shift towards a more marketing-driven approach.

Frontend hiring slowdown: The Frontend department had significant hiring in 2023 Q1 but none in 2024 Q1. This could point to the team reaching its desired size or a shift in priorities.

Enter a prompt here



Dandi AI understood that it didn't need to calculate any percentages. It simply retrieved the data, sorted it and picked the top 5. Then it generated the 5 takeaways requested.

In most *Prompt-to-SQL* methods, the query has to be generated, the schema needs to be trained on, the model used needs to be able to generate an SQL plan, the planned query has to be executed against raw data, and then the last call will generate the takeaways (assuming the accuracy of the SQL query was sufficient enough to handle the prompt in the first place).

The Impact

The potential impact of Dandi AI on HR business and operations is transformational, shifting the HR function from a slow, administrative cost center to a fast, strategic business partner.

By replacing weeks of manual data extraction and analysis with instant, 98% accurate AI-driven insights, Dandi enables evidence-based HR. This speed allows HR leaders to move from reacting to problems (like high turnover) to proactively managing the workforce and aligning HR activities with core strategic objectives.

Operations become far more efficient, as HR teams can streamline reporting, identify and eliminate pay gaps in real-time, and improve high-stakes functions like recruitment precision and employee retention. Ultimately, this leads to a reduction in time-to-hire, a decrease in turnover (up to 39% decrease), and an overall enhancement of organizational performance and productivity.

Companies We've Built



Why XTAM

Proven Track Record

- Multiple successful transformations and exits across SaaS, EdTech, enterprise software, and consumer tech.

Deep AI Expertise

- 10+ years deploying machine learning solutions—even before AI went mainstream—gives us an edge in knowing what works, what doesn't, and how to execute quickly.

End-to-End Approach

- We combine strategic advisory, engineering execution, and product innovation under one roof—ensuring accountability, speed, and ROI.

Flexible Structures

- We meet partners where they are—whether that means a straightforward consulting engagement or an equity-based deal for deeper, long-term collaboration.

Contact Details

Name	Phone	Email
Brian Williams	973-632-8991	brian@xtam.ai

Case Study

Accelerating Growth & Transformation
Through AI

XTAM



 FINALSITE

Client Background

Finalsite is a leader in K-12 EdTech with a 28-year track record and deep private-equity stewardship. They've captured a significant share of the U.S. district and global independent school market with best-in-class websites, communications, and enrollment products—serving thousands of schools.

The Problem



- Families were asking the same questions across five different channels, 24/7. "When are immunization forms due?" "How do I add lunch money?" "Where's the bus stop for Route 12?"
- Front offices and comms teams were overwhelmed by repeat inquiries and took longer and longer to respond, frustrating parents.
- Meanwhile, web/content owners struggled to keep pages, PDFs, and calendars aligned, making it hard for families to find a single, trustworthy answer.

Additional Constraints

Finalsite needed a **fast, defensible MVP** to prove business value, build internal sponsorship for the project, and scale to their large customer base to make a sizable impact to their revenue growth —while remaining flexible as AI evolves.

They lacked spare bandwidth to architect, ship, and iterate at that speed.

What We Did

We ran **parallel work-streams** across Product, Engineering, and GTM to compress time-to-value.

Discovery & Sourcing

- Mapped high-intent K-12 use cases: enrollment, calendars, transportation, meals, policies.
- Defined **trusted sources** (CMS pages, PDFs, calendar feeds) and the scale of these sources (7M+ files within the first 90 days)

MVP Architecture

- Purpose-built **K-12 AI** platform
- **Customizable** under a district/school hierarchy for appearance, branding, and knowledge base.
- Detailed analytics and insights on what parents are asking and how well the district/school answers parent’s inquiries.
- Finalsight’s first product built end-to-end for free trial -> paid conversion.

Pilot & GTM

- Recruited 10 design-partners for rapid soft-launch.
- Packaged onboarding run-book, ROI model, pricing guidance, demo script, and sales enablement.

“Ask AI has been the most successful product we’ve ever released in the history of Finalsight”
— Jim Calabrese - CEO of Finalsight”

Timeline of Results

Day 30	Day 60	Day 180
Customer interviews to deep dive into the customer pain points. A demo environment was rapidly spun up to demonstrate what was possible. Design clients were recruited.	Live beta was developed and launched on the first 10 partner customers. Rapid iteration and feedback followed for a development towards general availability. Sales was enabled and began selling the product at scale.	The product entered general availability. Within a few months, the product accounted for 12% of the entire year’s ARR growth —the most successful launch in Finalsight’s 28-year history.

Companies We've Built

quadpay
a zip company

RENT THE RUNWAY

CVS Health
signifyhealth

care/of

MIRROR

CredSimple
andros*

Handshake

Wunderkind

Primary

SLICE

TEEPUBLIC

splice

The Athletic

grovo
BY CORNERSTONE

Sealed

FINAL SITE
A Veritas Capital Company

Legion

reonomy
BY BLACK SPARK BUSINESS

Dandi

Orange
Hiring

Why XTAM

Proven Track Record

- Multiple successful transformations and exits across SaaS, EdTech, enterprise software, and consumer tech.

Deep AI Expertise

- 10+ years deploying machine learning solutions—even before AI went mainstream—gives us an edge in knowing what works, what doesn't, and how to execute quickly.

End-to-End Approach

- We combine strategic advisory, engineering execution, and product innovation under one roof—ensuring accountability, speed, and ROI.

Flexible Structures

- We meet partners where they are—whether that means a straightforward consulting engagement or an equity-based deal for deeper, long-term collaboration.

Contact Details

Name	Phone	Email
Brian Williams	973-632-8991	brian@xtam.ai