PRODUCT REQUIREMENTS DOCUMENT

FinGuard

AI-Powered Loan Eligibility & Risk Assessment

Version 1.0 | Concept Stage | November 2025

Product Manager: Aman Dixit

1. Executive Summary

FinGuard is an AI-powered credit risk and loan eligibility platform that uses alternative data sources and machine learning models to provide more accurate, fair, and explainable lending decisions. Traditional credit scoring relies on a narrow set of variables — payment history, credit utilization, account age — that systematically exclude millions of creditworthy individuals who are 'credit invisible' or have thin credit files.

FinGuard's core thesis: credit risk is better predicted by behavioral patterns, cash flow analysis, and alternative data signals than by backward-looking credit bureau scores. By expanding data inputs and applying modern ML, FinGuard enables lenders to approve more people who deserve credit — while reducing default rates.

> **Stage Note**
> This is a concept-stage PRD. All projections are hypothetical targets. Regulatory review, fair lending legal analysis, and data partnership agreements are required before any production deployment.

2. Problem Statement

2.1 The P0 Problem

45 million Americans are 'credit invisible' — no credit score or a score too thin to qualify for traditional loans. This isn't because they're high-risk; it's because the data the system uses doesn't capture their actual financial behavior. At the same time, lenders face pressure to reduce default rates while growing loan books — a tension that legacy credit models cannot resolve.

2.2 Pain Points by User Segment

For Loan Applicants

- Rejected for loans despite stable income and responsible financial behavior
- No transparency into why they were rejected or what would change the decision
- Alternative lenders charge predatory rates to applicants who are not actually high-risk

For Lenders (B2B Primary User)

- Legacy credit models miss predictive signals available in alternative data
- Manual underwriting is slow, expensive, and inconsistent
- Regulatory pressure for explainability — black-box ML creates compliance risk
- Growing demand for financial inclusion with no tooling to support it responsibly

3. User Personas

3.1 Marcus — The Credit-Invisible Applicant

Attribute	Detail
Age / Role	28 years old, Software Engineer
Background	Recent immigrant, 3 years US residency
Income	$85,000/year
Credit Status	No US credit history, FICO: N/A
Situation	3 years of on-time rent payments, healthy bank cash flows
Pain Point	Banks auto-reject. Car dealership offers 22% APR.
Goal	Get a fair assessment based on actual financial behavior, not absence of a credit file

Marcus's Quote
"I earn good money and pay everything on time. Why won't anyone lend to me?"

3.2 Sarah — The Underwriting Manager

Attribute	Detail
Age / Role	42 years old, VP of Lending at a regional credit union
Experience	15 years in consumer finance
Pressure	Board wants defaults below 2% but also wants loan volume growth
Problem	Legacy model leaving creditworthy applicants untouched; no in-house data science team
Goal	System that surfaces more 'yes' decisions with confidence, with audit-ready explanations

3.3 Priya — The Fintech Founder

Attribute	Detail
Age / Role	35 years old, building a BNPL product for underserved communities
Problem	Available risk APIs are expensive, not tuned for thin-file users, too many false positives

Attribute	Detail
Goal	Accurate, affordable, API-first risk scoring that works for her user segment

4. Alternative Data Sources (Core Differentiator)

Traditional models use: payment history, credit utilization, account age, credit mix, new inquiries. FinGuard uses all of the above PLUS:

Alternative Data Source	Signal It Provides	How Accessed
Bank Account Cash Flow	Income stability, spending patterns, savings behavior	Open Banking APIs (Plaid, MX)
Rent Payment History	On-time payment behavior for largest monthly expense	Experian RentBureau, rental data providers
Utility & Telecom Payments	Long-term payment consistency	Data partnerships, user permission
Employment Verification	Job tenure, income trajectory	Argyle, Work Number API
Mobile Money Patterns	Transaction velocity, merchant diversity	Plaid, Finicity
Gig Income Signals	Income regularity for non-traditional earners	Stripe, PayPal, Venmo patterns

5. Core Features

Feature 1: Multi-Source Risk Scoring Engine

The core ML pipeline ingesting traditional + alternative data, dynamically weighting features by applicant profile type, and outputting a FinGuard Risk Score (0-1000) with confidence interval. Separate model calibration for distinct segments: thin-file, no-file, credit-damaged, traditional. Real-time scoring in under 90 seconds.

Feature 2: Explainability Layer (LIME/SHAP) — Non-Negotiable

Every decision produces a human-readable explanation showing: top 3 positive factors, top 3 negative factors, what would change the decision (actionable guidance), and a regulatory-compliant adverse action notice. This is legally required under FCRA and ECOA.

Feature 3: Cash Flow Analysis Module

Connects to bank accounts via Open Banking APIs to analyze income regularity (24 months), monthly surplus/deficit patterns, recurring obligation detection, savings behavior, and financial stress indicators.

Feature 4: Lender Dashboard

Web console for underwriting teams: application queue with risk scores and explanation cards, one-click approve/decline/manual review workflow, portfolio analytics, model performance monitoring, and a full immutable audit log for regulatory compliance.

Feature 5: API Gateway

REST API for fintech B2B integrations. Score + explanation in single API call. Sub-2-second response at P95. Sandbox environment for testing. SDKs for JavaScript and Python.

6. Feature Prioritization (RICE Framework)

Feature	Reach	Impact	Confidence	Effort	Score
Explainability Layer	4	5	4	3	26.7
Risk Scoring Engine	5	5	4	4	25.0
API Gateway	3	4	3	2	18.0
Cash Flow Analysis	4	4	3	3	16.0
Lender Dashboard	3	4	4	3	16.0
Applicant Portal	4	3	3	3	12.0

Priority Rationale
Explainability ranks #1 because it is legally required (FCRA/ECOA) and is the primary reason lenders can trust and adopt the system. Risk Scoring Engine is the core value proposition. API Gateway enables B2B monetization with minimal engineering effort.

7. Product Goals & Success Metrics

7.1 Primary Objective

Improve loan approval rates for creditworthy thin-file and no-file applicants by 30% while maintaining lender default rates at or below industry benchmarks.

7.2 KPIs

Metric	Definition	Target
Approval Rate Lift	% increase in approvals for thin/no-file applicants vs. legacy model	+25-35%
Default Rate Delta	Change in 12-month default rate vs. control group	≤ +0.5%
Decision Latency	Time from application to decision	< 90 seconds
Explainability Coverage	% of decisions with full factor explanation	100%
False Positive Rate	Approved applicants who default within 6 months	≤ 4%
Lender NPS	Net Promoter Score from lending institution clients	≥ 50

8. Technical Architecture (PM-Level)

8.1 ML Model Stack

- Primary Model: Gradient Boosting (XGBoost/LightGBM) for tabular financial data
- Secondary Model: Neural network for sequential transaction pattern analysis
- Ensemble: Weighted combination with uncertainty quantification
- Explainability: SHAP values for feature attribution; LIME for local explanations
- Fairness Monitoring: Regular bias audits across protected characteristics; disparate impact testing

8.2 Proposed Technology Stack

Layer	Technology
Frontend	React (Lender Dashboard) + React Native (Applicant Mobile)
Backend	Python FastAPI (ML services) + Node.js (API Gateway)
ML	scikit-learn, XGBoost, SHAP, PyTorch
Database	PostgreSQL (decisions) + Redis (real-time scoring cache)
Data Connections	Plaid, MX, Argyle, Experian, Equifax APIs
Infrastructure	AWS with SOC 2 compliant config; Kubernetes for model serving
ML Observability	MLflow (model tracking), Datadog (ops), Arize (ML observability)

9. Compliance Architecture

Critical Requirement — Fair Lending

Regular disparate impact testing across protected characteristics. Fairness constraints baked into model training. Third-party bias audit required before launch. This is the single most important technical requirement in this PRD.

- FCRA Compliance: Adverse action notice generation; user right to dispute every decision
- ECOA Compliance: No use of protected class data; equal credit opportunity monitoring
- SOC 2 Type II: Full audit trail for all data access and decisions
- Data Retention: Configurable policies by jurisdiction; right to deletion

10. Risks & Mitigations

Risk	Impact	Mitigation
Model perpetuates historical lending bias	Critical	Disparate impact testing; fairness constraints; third-party bias audit before launch
Open Banking access restrictions	High	Multi-provider redundancy (Plaid + MX + Finicity); graceful fallback to traditional data
Regulatory reclassification as CRA	High	Legal review in each market; FCRA compliance from day one; licensing agreements
Model accuracy degrades in economic downturn	High	Stress testing on recession-era data; regular model refresh cadence; human-in-loop for edge cases
Lender resistance to replacing legacy underwriting	Medium	Pilot as augmentation, not replacement; side-by-side comparison dashboard to prove lift
Data breach of sensitive financial data	Critical	Zero-trust architecture; encryption at rest and in transit; SOC 2 Type II certification

11. Core User Stories

ID	User Story	Acceptance Criteria	Priority
US-01	As an applicant, I want to connect my bank account to strengthen my application	OAuth bank connection completes in <60 seconds; confirmation shown	Must-Have
US-02	As an applicant, I want to know exactly why I was approved or declined	Decision with top 3 factors in plain English within 5 minutes of submission	Must-Have

ID	User Story	Acceptance Criteria	Priority
US-03	As an underwriter, I want a risk score with explanation for every application	Risk score, confidence interval, and top 5 factors shown on application card	Must-Have
US-04	As an underwriter, I want to override model decisions with documented reason	Override available with required reason field; logged immutably in audit trail	Must-Have
US-05	As a fintech founder, I want a risk score in <2 seconds via API	API returns score + top factors in <2 seconds for 95th percentile requests	Must-Have

12. Revenue Model

Stream	Model	Price Point
API Calls	Per-decision pricing	\$0.50–\$2.00/decision (volume-based)
Lender SaaS	Monthly subscription + per-seat dashboard	\$2,000–\$15,000/month
Referral / Lead Gen	Consumer portal loan brokerage	1–3% of funded loan amount
Model Consulting	Custom model development for large institutions	Project-based

13. Go-to-Market Strategy

Phase 1 — API-First Launch

- Target: 5 fintech startups with BNPL or small-dollar loan products
- Entry: Free tier (500 API calls/month), paid tier for production volume
- Goal: Prove model accuracy on real applicant data

Phase 2 — Credit Union & Community Bank Partnerships

- Target: Regional institutions with financial inclusion mandates and no in-house data science
- Entry: White-label dashboard + model-as-a-service
- Goal: 3 institutional partnerships in Year 1

Phase 3 — Consumer Brand

- Direct applicant portal for personal loan applications
- FinGuard acts as broker routing applications to partner lenders
- Revenue: referral fees from lender partners

14. Open Questions

- Which regulatory framework applies when FinGuard's score is the primary input vs. supplemental data to a lender's existing model?
- What is the minimum viable dataset size to train a reliable alternative-data model?
- How do we handle applicants who refuse to connect financial accounts — fall back to traditional scoring or decline to score?
- Should the applicant portal always be white-labeled through lender partners or does FinGuard build a direct consumer brand?

This PRD is a concept-stage document.
Regulatory review and fair lending legal analysis required before any production deployment.