# Fraud prevention and risk management vendor RFP template

7/2/2026

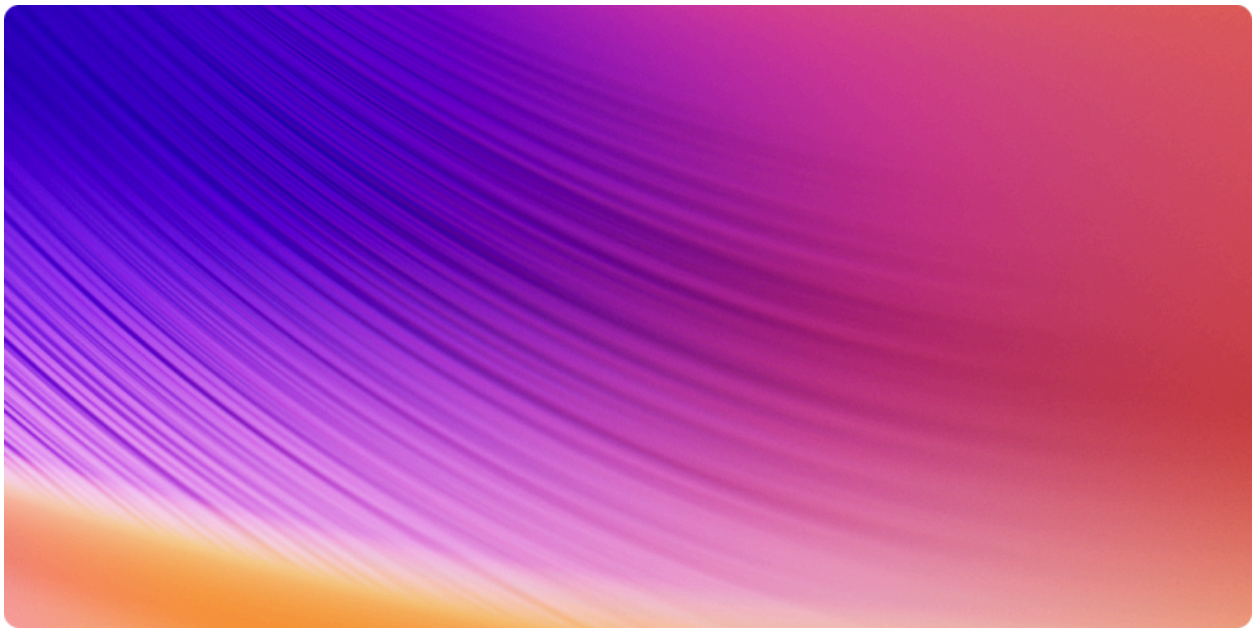

## Table of Contents


1. Cover page
2. Section A: Administrative instructions
3. Section B: Overview and scope of work
4. Section C: Proposal instructions
5. Section D: Evaluation process

6. Section E: Core requirements
7. Section F: Implementation and support
8. Section G: Commercials
9. Section H: Vendor profile
10. Section I: References
11. Section J: Appendices
12. How Stripe Radar can help

This is a structured guide for evaluating fraud prevention vendors. Stripe Radar is included as a reference point throughout; it's a concrete example of what best-in-class fraud prevention infrastructure looks like in 2026.

This guide includes both section outlines and sample copy. You can design your own branded RFP document or use the copy provided.

| | |
|---|---|
| **$1.9T+** | Annual transactions powering Radar's AI |
| **92%** | Chance a card used has been seen before on Stripe's network |
| **32%** | Average fraud reduction for Radar users |
| **$909M** | ACH and SEPA fraud blocked by Radar in 2025 |

# Cover page

The goal of the cover page is to tell vendors exactly what they're looking at and whom to talk to. It also includes key dates up front and information on what the final submission should look like.

### Contact information

| | |
|---|---|
| **RFP manager** | [Full name] |
| **Title** | [Title] |
| **Email** | [email@company.com] |
| **Phone** | [###-###-####] |

## Key dates

| | |
|---|---|
| **Issue date** | [MM/DD/YYYY] |
| **Question due** | [MM/DD/YYYY] |
| **Response due** | [MM/DD/YYYY] |
| **Evaluation period** | [MM/DD/YYYY–MM/DD/YYYY] |
| **Final selection** | [MM/DD/YYYY] |

## Submission format

All responses must be submitted electronically via email in PDF format. Pricing and scoring templates (provided separately in Excel) must be attached in their original formats.

## File naming convention

[Vendor name]–[Project name]–RFP–Response–[Date].pdf

## Purpose of this RFP

[Your company] is seeking a billing partner capable of supporting secure, multicurrency transactions, integrating easily with internal systems via modern APIs, and delivering high reliability, proactive fraud detection, and data transparency across regions.

This document outlines the requirements, evaluation criteria, and process for submitting proposals.

## Short confidentiality notice

This RFP contains confidential and proprietary information belonging to [your company]. It's provided solely for the purpose of preparing a response. Distribution beyond those directly involved in preparing a proposal is prohibited. By accepting this RFP, the recipient agrees to protect this information with at least the same degree of care they apply to protect their own confidential information.

# Section A: Administrative instructions

This section sets the ground rules. Fraud prevention vendor relationships involve deep access to your transaction data and risk infrastructure. Ambiguity here can create legal, operational, and security problems later. Be precise.

## A.1 Statement of confidentiality & nondisclosure

All information in this RFP is confidential and intended solely to enable the vendor to prepare a response. Vendors must not disclose, reproduce, or distribute this document or any portion of it without prior written consent from [your company]. Proprietary information included in proposals should be clearly labeled; [your company] will treat it accordingly.

## A.2 Limitation of financial liability

This RFP isn't an offer to contract. [Your company] is under no obligation to award a contract or reimburse costs incurred in preparing a response. Vendors are solely responsible for their own expenses throughout this process.

## A.3 RFP timeline

| Milestone | Date |
|---|---|
| RFP issued | Q2 2027 |
| Vendor acknowledgment due | [+3 business days] |
| Vendor questions due | [+2 weeks] |
| Q&A distributed to all vendors | [+3 weeks] |

| Milestone | Date |
|---|---|
| Proposal submission deadline | Q3 2027 |
| Evaluation period | Q3 2027 |
| Short list notifications | Q3 2027 |
| Vendor demonstrations | Q3–Q4 2027 |
| Final selection | Q4 2027 |
| Target go-live date | Q1 2028 |

## A.4 Submission guidelines

- All proposals must be submitted by email to [contact email address].
- Vendors must acknowledge receipt within three business days of issue.
- Questions must be submitted in writing by the date listed in A.3.
- All communication must go through the designated RFP manager. Direct contact with other [your company] employees during the evaluation period might result in disqualification.

## A.5 Required submission documents

| Item | Included? | Notes |
|---|---|---|
| Executive summary (PDF) | ☐ Yes ☐ No | |
| Section E requirements response (PDF) | ☐ Yes ☐ No | |
| Completed pricing template (Excel) | ☐ Yes ☐ No | |
| Company profile and financial summary (PDF) | ☐ Yes ☐ No | |
| Three or more references from comparable businesses (PDF) | ☐ Yes ☐ No | |

| Item | Included? | Notes |
|---|---|---|
| PCI DSS v4.0, SOC 2 Type II, and ISO 27001 certifications (PDF) | ☐ Yes ☐ No | |
| Case studies with production fraud reduction metrics (PDF) | ☐ Yes ☐ No | |
| Description of AI model training data and performance benchmarks (PDF) | ☐ Yes ☐ No | |
| API documentation or developer portal link (PDF or URL) | ☐ Yes ☐ No | |
| 12 months of historical uptime data (PDF) | ☐ Yes ☐ No | |

## A.6 Evaluation overview

[Your company] will evaluate proposals on AI model quality, breadth of fraud coverage, false positive rates, agentic commerce readiness, API reliability, and vendor stability. Vendors must demonstrate measurable improvements in fraud reduction, dispute rates, and operational efficiency with production evidence. Claims without data will not be scored.

## A.7 Vendor acknowledgment

Vendors must complete and return the acknowledgment below within three business days of receiving this RFP.

We acknowledge receipt of the RFP titled “[RFP name]” and confirm our intention ☐ to submit / ☐ not to submit a response.

Company name: _____

Authorized representative: _____

Title: _____

Date: _____

# Section B: Overview and scope of work

A vague overview produces generic proposals. Give vendors the specific context they need to respond intelligently: your business model, your current fraud exposure, the attack vectors you've identified, and the outcomes you need to improve. Fraud is developing fast and your RFP should reflect that.

Here are some additional details you can include to customize:

- Headquarters and key markets
- Transaction volume and average order value
- Split between domestic and cross-border transactions
- Primary payment methods in use (e.g., cards, ACH, SEPA, digital wallets)
- Internal teams involved (e.g., Engineering, Finance, Legal and Compliance, Operations)

## B.1 Company background

[Your company] is a [B2B / B2C / C2C] [marketplace / SaaS business / ecommerce business / platform] that's operating in [insert markets]. We process about [X] transactions per month across [X] payment methods and currencies. We currently experience a fraud rate of about [X]% and dispute rate of [X]%, and our fraud costs are about $[X] per month. We're seeking a fraud prevention partner whose platform covers the full stack—transaction fraud, account fraud, and customer abuse—without requiring significant ongoing engineering maintenance.

## B.2 Project's purpose

This RFP exists to identify a fraud prevention partner that protects our business as fraud threats change. Here's our current situation: [describe the gap (e.g., our rule-based system produces too many false positives, we have no coverage for account fraud or abuse, our dispute rate is rising, we lack visibility into emerging fraud vectors)].

## B.3 Scope of work

**Core deliverables:**

- **Transaction fraud prevention:** AI-powered fraud scoring on every payment, with documented false positive rates and fraud reduction figures from production
- **Account fraud prevention:** Automated detection and action on fraudulent accounts across the full lifecycle from onboarding to activity monitoring
- **Abuse prevention:** Detection and blocking of free trial abuse, pay-as-you-go abuse, and first-party fraud vectors
- **Dispute management:** Chargeback automation including evidence submission, win rate tracking, and dispute analytics
- **Custom rules engine:** The ability for our risk team to create and test custom rules without engineering involvement
- **Unified dashboard and reporting:** Real-time visibility into fraud rates, dispute rates, rule performance, and model outcomes in one place
- **API access:** Programmatic integration for teams that need to embed fraud intelligence into custom workflows

**Additional deliverables:**

- **Agentic commerce protection:** Fraud scoring and anomaly detection for AI agent–initiated transactions
- **Support for [insert required payment methods (e.g., ACH, SEPA, wallets, BNPL, stablecoins)]**
- **Platform-level fraud controls for marketplace operators:** Reserves, payout restrictions, and account lifecycle actions

## B.4 Out-of-scope work

Define what's excluded so vendors don't price or assume responsibility for it. Here are some examples:

- Full KYC and AML compliance infrastructure (handled separately by our payment processor)

- Customer identity verification beyond fraud risk signals
- General business intelligence or data warehousing unrelated to fraud

## B.5 Desired outcomes

- Fraud rate reduction of at least [X]% within 90 days of going live (for a benchmark, Radar reduces fraud by 32% on average)
- Dispute rate below [X]% within 90 days
- False positive rate below [X]%—measured by legitimate transactions declined or flagged incorrectly
- Free trial abuse blocked with at least 90% precision (current benchmark)
- Engineering time to maintain fraud infrastructure reduced by [X]%
- Zero-code setup completed within [X] days of contract execution

# Section C: Proposal instructions

Standardize what you receive so you can compare vendors side by side. Fraud prevention proposals tend to lead with product descriptions and ML jargon. Require production outcomes instead.

## C.1 Submission format and structure

Each proposal must follow this structure:

- Executive summary (three pages maximum)
- Responses to all requirements in Section E, numbered to match
- Completed pricing template in Excel

- Vendor profile and financial summary
- Minimum of three client references from comparable businesses
- Supporting documents: compliance certifications, case studies with production metrics, API documentation, description of AI model training data

Submissions that deviate significantly or omit required elements might be deemed noncompliant.

## C.2 Formatting requirements

- Narrative response as PDF; pricing template as Excel file
- Minimum 11pt font, one-inch margins, page numbers
- All monetary figures in USD unless otherwise specified
- File naming: [Vendor name]–Fraud–RFP–[Date].pdf

## C.3 Proposal content guidance

**Executive summary**

- Lead with measured outcomes from comparable deployments: fraud reduction percentages, dispute rate improvements, false positive rates, abuse prevention precision. Don't describe your product. Show what it produces.
- Include your vision for this partnership over three years, specifically how your AI model will continue to outpace fraud development.

**Solution overview and architecture**

- Describe how your platform covers the full fraud stack—transaction fraud, account fraud, abuse prevention, and dispute management—in a single integration.
- Describe the data foundation behind your AI models: the transaction volume, geographic and issuer breadth, and payment method diversity. This foundation determines whether the model performs at scale. The benchmark for leading fraud infrastructure is AI trained on $1.9T+ in annual transactions.

- Describe how your model updates as fraud patterns change and how quickly new fraud vectors are incorporated.

**Agentic commerce**

- Describe your current production capabilities for detecting and preventing fraud in AI agent–initiated transactions. Agentic commerce is a present requirement, not a future road map item. Proposals that position this as forthcoming will be scored accordingly.

**Coverage and payment methods**

- Specify exactly which payment methods you cover (e.g., cards, ACH, SEPA, wallets, BNPL, stablecoins).

**Security and compliance**

- Confirm compliance with PCI DSS v4.0 (effective March 2024) and most recent audit date.
- Provide 12 months of historical uptime data. The expectation for production-grade fraud infrastructure is 99.999%+ uptime.

## C.4 Clarification and questions

Questions must be submitted in writing by [question deadline] to [RFP manager's email]. Answers will be distributed simultaneously to all participants. No informal discussions with other [your company] employees are permitted during the process.

## C.5 Proposal validity

Proposals must remain valid for 90 days from the submission deadline unless extended by written mutual agreement.

## C.6 Right to reject or negotiate

[Your company] reserves the right to reject any proposal, request clarifications, or conduct parallel negotiations with one or more vendors. Participation doesn't constitute a commitment to purchase.

# Section D: Evaluation process

Transparency in scoring pushes vendors to respond with evidence rather than claims. Every criterion maps directly to requirements in Section E.

## D.1 Evaluation methodology

All proposals will be reviewed by a cross-functional team including Engineering, Finance, Legal and Compliance, Risk, and Operations.

**The evaluation runs in three stages:**

1. **Compliance review:** Confirm all required documents are present and meet formatting requirements.
2. **Qualitative assessment:** Score each submission against weighted criteria using a 1–5 scale (5 = exceptional, backed by production evidence; 1 = fails to meet baseline). A score of 5 requires documented production metrics.
3. **Demonstration and final review:** Short-listed vendors present live platform demonstrations. Demos must use a production-parity sandbox, not a scripted walkthrough.

## D.2 Evaluation criteria and weights

| Criterion | Weight | What we're evaluating |
|---|---|---|
| AI model and dataset quality | 25% | Transaction data volume behind the model, geographic and issuer breadth, model update frequency, documented fraud reduction in production |
| Breadth of fraud coverage | 20% | Transaction fraud, account fraud, abuse prevention, dispute management—all in one integration. Coverage across payment methods including ACH, SEPA, wallets, and BNPL. |
| Agentic and emerging threat coverage | 15% | AI agent–initiated transaction protection, first-party abuse detection, model adaptability to new fraud vectors |
| API quality and platform architecture | 15% | Single integration model, API latency, uptime history, sandbox quality, custom rules engine, no-code setup |
| Abuse prevention | 10% | Free trial abuse, pay-as-you-go abuse, first-party fraud—precision and recall from production deployments |
| Operational tooling and reporting | 5% | Unified dashboard, rule performance analytics, dispute reporting, data export capabilities |
| Implementation and support | 5% | Timeline realism, SLAs, support quality |
| Commercials and vendor stability | 5% | Pricing transparency, contract flexibility, financial health |

*Agentic and emerging threat coverage carries a 15% weight because the ability to detect fraud in AI agent–initiated transactions is a structural capability gap between vendors. For example, 65% of business leaders agree fraud is developing too rapidly for their companies to keep up. Vendors that can't demonstrate adaptive coverage today will fall further behind as fraud vectors continue to shift.*

## D.3 Demonstration requirements

Short-listed vendors will demonstrate the following live in a sandbox environment:

- Transaction fraud scoring in real time—trigger a high-risk transaction and show how the model scores it, with an explanation of contributing signals

- Account fraud detection—show how the platform identifies and takes action on a fraudulent account during onboarding and after activation
- Abuse prevention in action—demonstrate free trial abuse detection with the 90% precision benchmark
- Custom rule creation—a risk analyst creates a new rule without engineering involvement, tests it in sandbox, and reviews its projected impact
- Dispute management workflow—show automated evidence submission for a chargeback, with win rate tracking
- Agentic commerce protection—show how the platform applies fraud scoring to an AI agent–initiated transaction and flags anomalous agent behavior
- Unified dashboard—fraud rate, dispute rate, rule performance, and model outcome review in one place

Vendors must provide temporary demo credentials that are valid for at least 10 business days after the demonstration.

## D.4 Negotiation and contract award

[Your company] reserves the right to conduct clarification sessions, request best-and-final offers, and conduct parallel negotiations. No contract is binding until executed by both parties.

⚑ Evaluator notes—remove before sending to vendors

- Score independently before group deliberation. A score of 5 requires documented production metrics—not stated capabilities.
- Ask every vendor to state its AI model's training data footprint in specific terms: transaction volume, years of data, geographic breadth, issuer coverage, and payment method diversity. The answer reveals how much of its fraud detection performance is real vs. claimed. The benchmark is $1.9T+ in annual transactions.
- Probe payment method coverage. For each method on your list (e.g., ACH, SEPA, wallets, BNPL, stablecoins), ask whether the model has been trained on data from that method and what fraud reduction figures look like specifically for it.
- The agentic commerce demonstration is nonnegotiable. Any vendor that cannot show fraud detection on AI agent–initiated transactions in a sandbox today isn't ready for where your platform is going.

- Ask for false positive rates from production, not projections. A model that blocks 40% of fraud but declines 5% of legitimate transactions is a revenue problem, not a fraud solution.
- Ask for 12 months of uptime data, not just an SLA. The difference between 99.900% and 99.999% uptime matters for fraud infrastructure that runs on every transaction.

# Section E: Core requirements

This is the most important section. Require factual, evidence-backed responses. Any vendor worth deploying can point to documented outcomes from real deployments. For each requirement, vendors must indicate its status: Standard (in production today), Configurable (requires setup), Custom (requires development), or N/A.

## E.1 AI model and dataset quality

Fraud prevention AI is only as good as the data it's trained on. A model trained on tens of billions of transactions across markets, issuers, and payment methods produces materially better outcomes than a model built on a narrower dataset. There's a 92% chance that a card used on Stripe's network has been seen before—that level of prior signal is the benchmark. This is a compounding advantage that can't be replicated through engineering effort alone.

| Requirement | Status | Vendor response / evidence |
| --- | --- | --- |
| AI payment model trained on economy-scale transaction data—describe the transaction volume, geographic breadth, issuer coverage, and payment method diversity. Provide specific figures, not a description of your | Standard / Configurable / Custom / N/A | |

| Requirement | Status | Vendor response / evidence |
| --- | --- | --- |
| methodology. The current benchmark is training on $1.9T+ in annual transactions. | | |
| Documented fraud reduction from production deployments on the current model. Provide the average fraud reduction percentage across your customer base. The current benchmark is 32% average fraud reduction. | Standard / Configurable / Custom / N/A | |
| Model update frequency: how often are models retrained and deployed? Are updates continuous or periodic? Describe how quickly new fraud vectors are incorporated after detection. | Standard / Configurable / Custom / N/A | |
| Granular risk score for every transaction (e.g., 0–99 scale), with detailed signal-level insight so our risk team can set their own thresholds for blocking vs. reviewing—not just a binary block or allow output. | Standard / Configurable / Custom / N/A | |
| Per-transaction fraud scoring with a documented false positive rate from production deployments. Provide the number, not a description of how you minimize false positives. | Standard / Configurable / Custom / N/A | |
| Explainability: the ability for our risk team to understand why a specific transaction was flagged or declined at the individual transaction level. | Standard / Configurable / Custom / N/A | |
| Network-level intelligence: the model incorporates signals from across the full transaction network, not only your own historical data. Provide the scale of the network that underpins the model. | Standard / Configurable / Custom / N/A | |

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| Multiprocessor support: risk scores and fraud assessments are available for transactions processed through payment processors other than the vendor's own. Describe how this works and whether the full model capability is available for externally processed payments. | Standard / Configurable / Custom / N/A | |
| Preauthorization risk scores: the ability to receive a fraud risk score before authorization, to inform payment routing decisions. Describe how and at what point in the checkout flow scoring can be applied. | Standard / Configurable / Custom / N/A | |
| Flexible screening points: the ability to screen transactions at different points throughout the checkout flow (e.g., at account creation, preauthorization, post-authorization). Specify all supported screening points. | Standard / Configurable / Custom / N/A | |

## E.2 Transaction fraud prevention

Transaction fraud is rising. The cost of online payment fraud was projected to grow by 15% in 2025. The right infrastructure catches fraud that others can't because its model has seen the pattern before.

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| Real-time fraud scoring on every transaction, with documented median score latency from production. Specify the p50, p95, and p99 scoring latency. | Standard / Configurable / Custom / N/A | |

| Requirement | Status | Vendor response / evidence |
| --- | --- | --- |
| Automatic fraud blocking with configurable thresholds: our risk team sets the threshold and the model acts. No code required to adjust. | Standard / Configurable / Custom / N/A | |
| Dynamic 3DS and SCA application: the platform intelligently triggers 3DS and SCA only for high-risk or mandated transactions, not across all transactions. Applying authentication universally is a conversion killer. Describe the logic and provide documented friction reduction figures from production. Stripe's dynamic application of 3DS reduces fraud by 30%. | Standard / Configurable / Custom / N/A | |
| Chargeback and dispute rate reduction: provide documented dispute rate improvement from production deployments across your customer base. | Standard / Configurable / Custom / N/A | |
| Coverage across all payment methods in use: cards, ACH, SEPA, wallets, BNPL, and stablecoins. In 2025, the benchmark was $909M in ACH and SEPA fraud blocked. Specify which payment methods your model is trained on and the fraud reduction figures for each. | Standard / Configurable / Custom / N/A | |
| Card testing attack detection: automated detection and blocking of multiple payment attempts from the same IP address or email within a defined window of time. Provide documented detection rates from production. Stripe's Payments Foundation Model raised the detection rate from 59% to 97% on large businesses. This is the benchmark for improvement. | Standard / Configurable / Custom / N/A | |

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| PAN portability: the ability to securely handle raw PAN data to perform risk assessments across multiple processors so your fraud history and allow and block lists aren't locked to a single processor. | Standard / Configurable / Custom / N/A | |
| Cross-border fraud detection: describe how your model handles cross-border transactions and what fraud reduction figures look like specifically for cross-border volume. | Standard / Configurable / Custom / N/A | |

## E.3 Account fraud prevention

Account fraud (fraudulent accounts created to exploit your platform) operates across the full account lifecycle. Catching it at onboarding is cheaper than catching it after activation. The benchmark for leading infrastructure is blocking more than 3.5M fraudulent connected accounts per year.

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| Account-level fraud scoring at onboarding: AI-powered risk assessment of new accounts before they're activated, with documented false positive rates. | Standard / Configurable / Custom / N/A | |
| Ongoing account monitoring after activation: continuous risk assessment that detects fraud signals after onboarding, not only at the point of account creation. | Standard / Configurable / Custom / N/A | |
| Automated account actions: the ability to automatically restrict payouts, place reserves, or block suspicious accounts based on risk | Standard / Configurable / Custom / N/A | |

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| signals, without manual review of every case. | | |
| Enhanced verification for borderline accounts (accounts that raise suspicion but aren't conclusively fraudulent). Describe how additional authentication steps are configured and triggered. | Standard / Configurable / Custom / N/A | |
| Flagged accounts view: a list of accounts currently flagged as suspicious, accessible in real time, so our risk team can do triage without running a manual query. | Standard / Configurable / Custom / N/A | |
| Real-time alerts for platform-level merchant fraud: immediate notification when potential merchant or seller fraud is detected, with configurable alert thresholds and delivery channels (e.g., email, webhook, dashboard). | Standard / Configurable / Custom / N/A | |
| Fraudulent connected account blocking: for platform operators, describe how the platform detects and blocks fraudulent seller or merchant accounts. The benchmark is 3.5M+ fraudulent connected accounts blocked per year. | Standard / Configurable / Custom / N/A | |
| Platform-level financial risk controls: automated reserves and payout restrictions triggered by risk signals, configured by the platform operator, without requiring per-account engineering. | Standard / Configurable / Custom / N/A | |
| Radar reduces average fraud exposure by 5.3x by shortening the window between initial detection and platform resolution. Provide your equivalent metric from production. | Standard / Configurable / Custom / N/A | |

## E.4 Abuse prevention

First-party abuse (trial abuse, pay-as-you-go abuse, and friendly fraud) costs businesses $200B annually. In the past year, 94% of businesses have experienced it. Traditional fraud tools aren't designed to catch it. Require production evidence, not a description of a vendor's road map.

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| Free trial abuse prevention: automated detection and blocking of fraudulent free trial sign-ups, with documented precision from production. The current benchmark is 90% precision. | Standard / Configurable / Custom / N/A | |
| Scale of abuse prevention: in the first two months after launch, the benchmark is blocking 715,000 high-risk trials and preventing $6M in losses. Provide equivalent figures from your production deployments. | Standard / Configurable / Custom / N/A | |
| Model adaptability for new abuse vectors: from November 2025–February 2026, the benchmark platform detected 6.2x more abusive free trials across its network. Describe how your model adapts to new abuse patterns and provide equivalent evidence. | Standard / Configurable / Custom / N/A | |
| Pay-as-you-go abuse prevention: detection of customers' exploiting usage-based pricing models. Describe your capabilities and provide production metrics. | Standard / Configurable / Custom / N/A | |
| Customer lifecycle abuse coverage: abuse prevention that spans the full customer lifecycle, not only at the point of payment. Describe how your platform | Standard / Configurable / Custom / N/A | |

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| identifies abuse signals before, during, and after a payment. | | |
| Dispute prevention: proactive identification of first-party fraud before it becomes a dispute. In the past year, 62% of businesses have seen disputes from first-party fraud increase. Provide your dispute reduction metrics from production. | Standard / Configurable / Custom / N/A | |

## E.5 Dispute management

Disputes cost businesses $35 for every $100 in chargebacks, when you factor in operational costs, network fees, and time. The right infrastructure automates evidence submission, shows win rate probabilities per dispute so your team can triage cases effectively, and ideally stops disputes before they're filed.

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| Automated chargeback evidence submission: the platform compiles and submits dispute evidence automatically, using AI to customize the package for each dispute type, without requiring manual data gathering for each case. | Standard / Configurable / Custom / N/A | |
| Win rate probability per dispute: the platform calculates the likelihood of winning each individual dispute so our team can do triage and prioritize responses rather than treat all disputes equally. | Standard / Configurable / Custom / N/A | |
| Evidence recommendations: the platform recommends which | Standard / Configurable / | |

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| specific evidence to submit for each individual dispute, not a generic evidence checklist. | Custom / N/A | |
| Dispute win rate tracking and analytics: real-time visibility into dispute outcomes, win rates by dispute reason, and trends over time. | Standard / Configurable / Custom / N/A | |
| Smart Refunds or equivalent: proactive identification of transactions likely to result in a dispute with the option to refund them before a chargeback is filed, preventing the chargeback fee entirely. | Standard / Configurable / Custom / N/A | |
| Early fraud warnings: the platform notifies you when an issuing bank flags a transaction as fraudulent before it becomes a formal chargeback, giving your team the opportunity to intervene first. | Standard / Configurable / Custom / N/A | |
| Verifi and Ethoca network integration: direct integration with card network dispute prevention solutions, without requiring a separate third-party integration on your side. Specify which networks are supported natively. | Standard / Configurable / Custom / N/A | |
| Dispute management via API: the ability to upload evidence, respond to disputes, and receive dispute events via webhooks programmatically. | Standard / Configurable / Custom / N/A | |
| Dispute prevention signals: proactive risk signals that identify likely disputes before they’re filed, giving our team the opportunity to intervene. | Standard / Configurable / Custom / N/A | |

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| Network-level dispute data: the model incorporates dispute outcomes from across the full network to improve future predictions. Describe the scale of dispute data that informs your model. | Standard / Configurable / Custom / N/A | |
| Fraud monitoring program support: real-time visibility into our standing with card network monitoring programs (VAMP and equivalents), with tools to help our business exit or avoid them. Describe your capabilities and outcomes. | Standard / Configurable / Custom / N/A | |

## E.6 Custom rules and risk team tooling

A fraud prevention platform that requires engineering involvement to adjust risk strategy is a bottleneck. Your risk team needs to move faster than fraud does. That means nontechnical users must be able to write, test, and deploy rules without filing a ticket.

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| No-code rule creation: risk analysts create, test, and deploy custom rules without engineering involvement. Confirm this is available in production today. | Standard / Configurable / Custom / N/A | |
| Plain English rule writing for nontechnical users: rules can be written in natural language (e.g., “Block if the email domain is temporary”) without requiring knowledge of query syntax or rule logic. Describe how this works and provide a live demo. | Standard / Configurable / Custom / N/A | |

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| Block and allow lists: the ability to create and manage lists of specific data points (e.g., suspicious IP addresses, email addresses, card BINs, device fingerprints) and reference them directly in fraud rules. Describe how lists are created, updated, and applied. | Standard / Configurable / Custom / N/A | |
| Rule suggestions based on top fraud indicators: the platform proactively recommends rules based on your specific fraud patterns, not only a generic rule library. Describe how suggestions are generated. | Standard / Configurable / Custom / N/A | |
| Shadow mode or what-if simulation: before a rule goes live, risk analysts can model its projected impact against historical data or run it in shadow mode alongside live traffic to assess impact on fraud catch rate and false positives before deployment. | Standard / Configurable / Custom / N/A | |
| Custom business data in rules: the ability to write rules that reference your own business-specific fields (e.g., loyalty_tier, product_category, shipping_method) beyond standard transaction attributes. Describe how custom metadata is ingested and referenced. | Standard / Configurable / Custom / N/A | |
| Custom rule performance analytics: real-time visibility into how each rule is performing (e.g., transactions affected, fraud blocked, false positives triggered). | Standard / Configurable / Custom / N/A | |
| Configurable risk thresholds by amount, region, account type, or payment method—no code required to adjust. | Standard / Configurable / Custom / N/A | |

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| Manual review queue: the ability for our risk team to flag transactions or accounts for manual review, with workflow tools to manage the queue. | Standard / Configurable / Custom / N/A | |
| Radar for Fraud Teams or equivalent: describe your offering for sophisticated risk teams that need custom rules, detailed analytics, and manual controls. Provide documented outcomes from production deployments using advanced rule tooling. | Standard / Configurable / Custom / N/A | |

## E.7 Agentic commerce and emerging threats

AI agents are already initiating commercial transactions, and fraudulent actors are using the same AI tools that legitimate businesses use. Marketplaces that don't protect AI agent–initiated transactions will face compounding exposure as agentic volume grows. Any vendor that can't demonstrate fraud detection on agent-initiated transactions in a sandbox environment today isn't ready.

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| Fraud scoring on AI agent–initiated transactions: the model applies fraud risk assessment to transactions initiated by AI agents, not only human-initiated ones. Demonstrate in a sandbox. | Standard / Configurable / Custom / N/A | |
| Anomaly detection for agent behavior: the ability to distinguish authorized high-volume automated workflows from fraudulent or compromised agent activity. | Standard / Configurable / Custom / N/A | |

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| First-party abuse coverage for AI-native business models: describe how your platform protects AI companies (a primary Radar target segment) against abuse specific to pay-as-you-go and usage-based pricing models. | Standard / Configurable / Custom / N/A | |
| Model adaptability: describe your process for detecting and incorporating new fraud vectors as they emerge. Provide evidence of how quickly new vectors were incorporated in the last 12 months. | Standard / Configurable / Custom / N/A | |
| 200+ product updates per year or equivalent pace: describe your release cadence and provide evidence of continual improvement in AI fraud prevention. | Standard / Configurable / Custom / N/A | |

## E.8 Analytics, reporting, and real-time intelligence

Fraud analytics and reporting are mandatory. If your risk team can't see what's happening in real time, investigate individual transaction decisions, and receive proactive alerts when an attack is underway, they're always reacting to fraud instead of getting ahead of it.

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| Unified dashboard: fraud insight, payment details, and dispute management in a single view, not split across separate tools. Finance, Risk, and Operations shouldn't need to switch systems to get a complete picture. | Standard / Configurable / Custom / N/A | |
| Real-time dashboard updates: fraud analytics are updated in real time, not on a delay. Specify the latency | Standard / Configurable / Custom / N/A | |

| Requirement | Status | Vendor response / evidence |
| --- | --- | --- |
| between a transaction event and its appearance in the dashboard. | | |
| Per-transaction review: the ability to inspect each individual transaction to understand its risk score, the signals that contributed to it, and the outcome. This is the foundation of any effective risk investigation workflow. | Standard / Configurable / Custom / N/A | |
| Real-time fraud attack alerts: the platform notifies your team immediately when a fraud attack pattern is detected (e.g., a card testing peak), with specific recommendations for how to mitigate the attack. Describe what triggers an alert and what the recommended actions look like. | Standard / Configurable / Custom / N/A | |
| Fraud pattern investigation: the dashboard allows your risk team to investigate emerging fraud vectors and attack patterns, not just view aggregate metrics. | Standard / Configurable / Custom / N/A | |
| Monitoring program standing: real-time visibility into your standing with card networks' fraud monitoring programs (VAMP and equivalents) so your team can act before thresholds are exceeded. | Standard / Configurable / Custom / N/A | |
| Data warehouse sync: the ability to sync fraud data directly to your own data warehouse (e.g., Snowflake, BigQuery, Redshift) for custom analysis. Describe the sync mechanism, schema documentation, and latency. | Standard / Configurable / Custom / N/A | |
| Per-processor analytics filtering: the ability to filter and view analytics by payment processor, for | Standard / Configurable / Custom / N/A | |

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| businesses that run multiple processors in parallel. | | |
| Programmatic dispute management via API: full API access to dispute data, with the ability to upload evidence, respond to disputes, and receive dispute events via webhooks without requiring dashboard-only workflows. | Standard / Configurable / Custom / N/A | |

## E.9 Platform architecture and API quality

A fraud prevention integration that requires ongoing engineering maintenance to stay effective is a hidden tax on your product team. Evaluate the API the way you would evaluate core infrastructure.

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| Zero-code setup: fraud prevention that's active on all transactions with no integration required. Confirm this is available in production today and describe what it covers. | Standard / Configurable / Custom / N/A | |
| API access for programmatic integration—RESTful API with comprehensive, versioned documentation and a public changelog. Specify SDK coverage across Node.js, Python, Ruby, Java, Go, and PHP. | Standard / Configurable / Custom / N/A | |
| Published API latency benchmarks: p50, p95, and p99 response times from | Standard / Configurable / Custom / N/A | |

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| production. Provide real figures, not SLA commitments. | | |
| Uptime: 99.999%+ (under 44 seconds of downtime per year) is the standard for production-grade fraud infrastructure. Provide 12 months of historical uptime data. | Standard / Configurable / Custom / N/A | |
| 100% PCI DSS audit success rate: provide your full PCI audit history. A single failed audit is important information. | Standard / Configurable / Custom / N/A | |
| Webhook support with configurable retry logic, delivery monitoring, and failure alerting. | Standard / Configurable / Custom / N/A | |
| Full sandbox environment with production parity for all fraud flows—transaction scoring, account actions, rule testing, agentic transactions, and dispute management. | Standard / Configurable / Custom / N/A | |
| No-code and low-code tools that allow Risk, Finance, and Operations teams to configure rules, thresholds, and reporting without engineering work. | Standard / Configurable / Custom / N/A | |

## E.10 Security, compliance, and data privacy

Fraud infrastructure sits in the data path of every transaction. A 100% PCI audit success rate is the standard.

| Requirement | Status | Vendor response / evidence |
|---|---|---|
| PCI DSS v4.0 compliance (effective March 2024): specify certification level, most recent QSA audit date, and whether you've maintained a 100% PCI audit success rate across your full audit history. | Standard / Configurable / Custom / N/A | |
| SOC 2 Type II certification: provide most recent audit period and report date. | Standard / Configurable / Custom / N/A | |
| ISO 27001 certification or equivalent. | Standard / Configurable / Custom / N/A | |
| GDPR-compliant data handling with configurable retention, deletion, and portability controls. | Standard / Configurable / Custom / N/A | |
| CCPA compliance for US customer data. | Standard / Configurable / Custom / N/A | |
| Data residency options for markets with localization requirements. | Standard / Configurable / Custom / N/A | |
| Incident response plan with defined client notification timelines: state the contractual commitment. | Standard / Configurable / Custom / N/A | |
| Trusted by 50% of Fortune 100 companies: describe the security and compliance infrastructure that supports enterprise-scale deployments. | Standard / Configurable / Custom / N/A | |

## E.11 Vendor certification statement

I hereby certify that all responses are accurate as of the submission date and that capabilities marked Standard or Configurable are currently available in production environments. Claims not supported by documentation or a live demonstration will not be evaluated.

Authorized representative: _____

Title: _____

Date: _____

⚑ Evaluator notes—remove before sending to vendors

- A score of 5 on any criterion requires documented production metrics. Stating, “We support this,” without evidence is a 3 at best.
- Ask every vendor to state its AI model’s training data footprint in specific terms: number of transactions, years of data, geographic and issuer breadth, and payment method diversity. This is the single biggest differentiator in fraud detection performance.
- Probe payment method coverage. For each method on your list (e.g., ACH, SEPA, wallets, BNPL, stablecoins), ask whether the model has been trained on production data from that method and what fraud reduction looks like specifically for it.
- Ask for false positive rates, not fraud reduction rates alone. A model with a 40% fraud reduction but a 5% false positive rate is costing you revenue.
- Ask to see the plain English rule writing interface live. Nontechnical users should be able to write and simulate a new rule in the demo without assistance.
- Ask vendors to demonstrate shadow mode. Run a new rule against historical data and show the projected impact before it goes live. This is a core workflow for any serious risk team.
- Ask whether block and allow lists can include custom business data fields. A vendor that supports only standard transaction attributes will hit a ceiling quickly.
- Regarding disputes, ask whether Verifi and Ethoca integrations are native or require separate third-party contracts. Native integration matters for speed of intervention.

- Ask to see a real-time fraud attack alert from production. What does the notification look like, how fast does it fire, and what does the recommended mitigation say?
- The agentic commerce demonstration is a hard requirement. Any vendor that cannot show fraud detection on AI agent–initiated transactions in a sandbox today isn't ready.
- Ask for 12 months of historical uptime data and the full PCI audit history. The SLA and current certification status aren't sufficient.

# Section F: Implementation and support

Activating fraud prevention on a live platform carries risk. An incorrectly configured rule or overly aggressive model threshold can block legitimate revenue. This section establishes whether the vendor has the methodology and experience to manage that risk.

## F.1 Implementation approach

Vendors must describe:

- Documented time to first fraud signal for businesses of comparable transaction volume with specific examples, not ranges—the benchmark for zero-code setup is same-day activation
- How they approach initial model calibration—the threshold-setting process that balances fraud reduction against false positive rates for your specific transaction mix
- Their processes for a parallel run—testing the new platform against live traffic before full implementation so you can validate performance before you disable existing controls
- How they manage rule migration, if you're bringing over existing custom rules from another platform

## F.2 Resourcing and governance

Vendors should provide:

- Named account manager and fraud specialist assigned to this engagement
- Escalation hierarchy and decision-making cadence during implementation
- Whether the implementation team is the same team that handles postlaunch support—the handoff is often where service quality drops

## F.3 Training and documentation

Vendors should describe:

- Training available for Risk, Engineering, Finance, and Operations teams
- Quality and currency of documentation, with best-in-class platforms maintaining documentation that risk analysts prefer over asking support questions
- How documentation is updated as the model and product ship new fraud detection capabilities

## F.4 Support model and SLAs

Vendors must specify:

- Support tiers and what's included—a false positive peak or model degradation at 2:00 a.m. is a Severity 1 incident that requires an immediate response
- Response time SLAs by severity, with contractual commitments
- How clients are notified during fraud model incidents or degradation events and what the post-incident review produces
- Historical Severity 1 response time data, not just the SLA

## F.5 Continuous improvement

Describe specifically how your platform uses ML and production analytics to improve fraud outcomes over time. Provide examples with production metrics: fraud reduction improvements delivered to existing clients over 12 months, false positive rate changes, and dispute rate improvements.

## F.6 Vendor attestation

I certify that all implementation and support details are accurate as of the submission date and reflect current production practices.

Authorized representative: _____

Title: _____

Date: _____

⚑ Evaluator notes—remove before sending to vendors

- Ask for specific implementation examples from comparable businesses (e.g., transaction volume, fraud rate at the time, number of markets). Reject ranges.
- Ask about the threshold calibration process specifically. A vendor that sets a single global threshold and moves on isn't treating your fraud profile as unique.
- Request Severity 1 response time actuals from the last 12 months.
- Ask whether the implementation team is the same team that handles postlaunch support.
- Ask how quickly a new fraud vector—one your business hasn't encountered before—would be incorporated into the model after it's detected.

# Section G: Commercials

Fraud prevention pricing varies significantly across vendors and models. Some charge per transaction, others by fraud reduction outcome, and some by tier. Standardize disclosure so you're comparing real economics.

## G.1 Pricing structure overview

Vendors must provide:

- Itemized pricing for every component (base fraud prevention, custom rule access, advanced abuse prevention, dispute management tools, API access, and add-ons)
- A narrative that explains pricing assumptions (transaction volume, average order value, fraud rate, and payment method mix)
- Clear identification of minimum monthly commitments or volume thresholds that affect pricing
- All figures in USD, including conversion logic if other currencies are quoted

## G.2 Pricing components

| Component | Unit | Unit price | Volume assumption | Monthly total (est.) |
|---|---|---|---|---|
| Base fraud prevention (all transactions) | % of transaction or flat fee | | | |
| Custom rules (Radar Plus or Fraud Teams equivalent) | Monthly or per rule | | | |

| Component | Unit | Unit price | Volume assumption | Monthly total (est.) |
|---|---|---|---|---|
| Abuse prevention (Radar Pro equivalent) | Monthly or per event | | | |
| Dispute management tools | Per dispute or monthly | | | |
| Programmatic API access (beyond zero code) | Per call or monthly | | | |
| Platform or connected account coverage | Per account or monthly | | | |
| Advanced analytics and reporting | Monthly or per query | | | |
| Implementation and onboarding | One-time | | | |
| Ongoing support tier | Monthly | | | |
| Add-ons (list individually) | | | | |

## G.3 Volume sensitivity

Provide estimated total cost at the following transaction volumes:

| Transaction volume tier | Estimated monthly cost |
|---|---|
| [Your current GMV] | |
| 2× current GMV | |
| 5× current GMV | |
| 10× current GMV | |

## G.4 Contract terms

- Available contract lengths and pricing incentives for each
- Whether pricing scales down automatically with volume decreases
- Exit clauses and data portability—how fraud signals, rule configurations, and historical data are returned, in what format, and on what timeline
- Minimum spend requirements

## G.5 Assumptions and dependencies

List all commercial assumptions that underpin your pricing. Unstated assumptions discovered after contract execution might be treated as a material misrepresentation.

## G.6 Vendor certification

I certify that all pricing and commercial information is complete and accurate as of the submission date.

Authorized representative: _____

Date: _____

⚑ Evaluator notes—remove before sending to vendors

- Reconcile the narrative against the Excel sheet. Discrepancies are a signal.
- Fraud prevention pricing often buries the cost of custom rules and advanced abuse prevention in higher tiers. Model the total cost at your fraud team's workflow, not just the base tier.

- Ask how pricing changes if your fraud rate improves. Some models penalize success by reducing the product's apparent value as fraud rates fall.
- Data portability is often the real lock-in mechanism. Assess exit terms before you sign, not after.
- Ask vendors to model total cost at 10× your current transaction volume. The pricing curve over growth matters as much as today's rate.

# Section H: Vendor profile

Your fraud prevention infrastructure partner sits in the data path of every transaction. Understand the company as a whole: its financial health, AI development depth, rate of improvement, and track record with businesses that look like yours.

## H.1 Company overview

Provide a two- to three-paragraph summary that explains your history, mission, and market position. Focus on your experience with businesses in [your industry or segment (e.g., ecommerce, SaaS, marketplace, AI companies, platforms)]. Describe your track record of maintaining detection quality as fraud patterns have developed and your history of shipping fraud prevention capabilities ahead of emerging threats.

## H.2 Customer base and track record

Provide specific data on your customer base:

- Number of businesses using your fraud prevention infrastructure

- Aggregate transaction volume protected per year
- Industries and business models represented in your customer base
- Share of businesses in your target segment (e.g., ecommerce, SaaS, platforms) that you protect
- Documented aggregate fraud reduction across your customer base (the benchmark is 32% average fraud reduction)

## H.3 Financial stability

Provide audited financial statements or equivalent evidence of solvency. Private companies should provide a CFO letter that certifies liquidity. Describe your funding structure.

## H.4 Certifications and compliance

| Certification / framework | Status and most recent date |
|---|---|
| PCI DSS v4.0 (effective March 2024) | |
| PCI audit success rate (full history) | |
| SOC 2 Type II | |
| ISO 27001 | |
| GDPR | |
| CCPA | |
| Additional country-specific certifications | |

## H.5 Analyst recognition

Provide independent analyst recognition relevant to fraud prevention and payment security. The current benchmark for a leading fraud prevention provider is recognition as a Leader across payments and fraud categories. Stripe Radar's AI is trusted by 50% of Fortune 100 companies. Describe where your platform stands relative to that benchmark.

## H.6 Pace of improvement

Describe your product release cadence for the past 12 months, including the number of updates shipped and major capabilities launched in AI fraud detection, abuse prevention, and agentic commerce protection. The current benchmark for a leading platform is 200+ product updates per year. Explain how your road map for the next 12–18 months continues to invest in the capabilities that matter for [your segment].

## H.7 Vendor statement of accuracy

I certify that all information in Section H is accurate as of the submission date and that [vendor] has the financial, technical, and operational capacity to perform the described services.

Authorized representative: _____

Date: _____

# Section I: References

References from comparable businesses are more valuable than any demo. Prioritize references that match your business model, transaction mix, and geographic footprint. A vendor that protects a business with a fraud profile similar to yours should have references that can speak to real outcomes.

## I.1 Reference requirements

Vendors must provide a minimum of three references that meet these criteria:

- Comparable business model to [your company]—same industry or transaction type
- Comparable transaction volume or fraud rate at the time of deployment
- At least one reference that's used advanced features: custom rules, abuse prevention, or dispute management
- Active customer in production for at least 12 months

## I.2 Reference Table

| Company name | Contact name and title | Business type | Markets | Tenure | Key use case |
|---|---|---|---|---|---|
| | | | | | |
| | | | | | |
| | | | | | |

## I.3 Reference outcome summary

For each reference, provide documented outcomes: fraud reduction percentage, false positive rate, dispute rate improvement, abuse prevention precision, or operational efficiency gains. Provide specific figures, not ranges.

## I.4 Reference validation

I confirm that each client has consented to serve as a reference and that all information is accurate. [Your company] reserves the right to contact references directly.

Authorized representative: _____

Date: _____

⚑ *Evaluator notes—remove before sending to vendors*

- *Call at least two references by phone. Written summaries are curated by the vendor.*
- *Ask references specifically whether the fraud reduction and false positive rates in the proposal match what they saw in production.*
- *Ask about the implementation experience and threshold calibration process—not just the platform at steady state.*
- *Ask whether the vendor's model adapted to new fraud vectors the reference encountered after going live and how quickly it did so.*
- *Flag references from noncomparable business models. A reference from a physical retail business tells you very little about digital fraud performance.*

# Section J: Appendices

## J.1 Submission checklist (vendor use)

Attach as the first page of your response packet. Incomplete submissions might be excluded from evaluation.

| Item | Included? | Notes |
|---|---|---|
| Executive summary (three-page maximum) | ☐ Yes ☐ No | |
| Section E requirements response | ☐ Yes ☐ No | |
| Completed pricing template (Excel) | ☐ Yes ☐ No | |
| Vendor profile and financial summary | ☐ Yes ☐ No | |
| Three or more marketplace client references | ☐ Yes ☐ No | |
| PCI DSS v4.0 certification and full audit history | ☐ Yes ☐ No | |
| SOC 2 Type II report (most recent period) | ☐ Yes ☐ No | |
| Description of AI model training data and performance benchmarks | ☐ Yes ☐ No | |
| 12-month historical uptime data | ☐ Yes ☐ No | |
| Case studies with production fraud reduction metrics | ☐ Yes ☐ No | |
| False positive rate data from production | ☐ Yes ☐ No | |
| Signed vendor certification statements | ☐ Yes ☐ No | |

## J.2 Glossary of terms

| Term | Definition or usage in this RFP |
|---|---|
| Smart Refunds | Proactive refunding of transactions identified as likely to result in a dispute, preventing the chargeback fee before it's triggered. |
| Early fraud warning | A notification from an issuing bank that flags a transaction as potentially fraudulent before a formal chargeback is filed, giving the business an opportunity to refund and prevent the dispute. |
| Verifi and Ethoca | Card network dispute prevention solutions (Visa and Mastercard, respectively) that enable businesses to resolve disputes before they become chargebacks. Native integration matters. A separate third-party contract adds latency and cost. |
| Shadow mode | Running a new fraud rule against live traffic or historical data without enforcing it so risk analysts can measure projected impact—fraud caught, false positives triggered—before the rule goes live. |
| Block and allow lists | Configurable lists of specific data points (e.g., IP addresses, email domains, card BINs, device fingerprints) that can be referenced directly in fraud rules to block or permit associated transactions. |
| PAN portability | The ability to transfer primary account number (PAN) data and associated fraud history—including block and allow lists—between payment processors so fraud intelligence isn't locked to a single vendor. |
| Preauthorization risk score | A fraud risk score returned before a payment is authorized, enabling payment routing decisions based on risk level instead of blocking only after authorization. |
| AI Payments Foundation Model | An ML model trained on economy-scale transaction data to detect fraud, score risk, and reduce false positives. The benchmark is training on $1.9T+ in annual transactions across millions of businesses. |
| Transaction fraud prevention | Detection and blocking of fraudulent payments—card fraud, card testing, and stolen credentials—at the point of transaction. |
| Account fraud prevention | Detection and blocking of fraudulent accounts across the full lifecycle: onboarding, activation, and ongoing activity monitoring. |
| Abuse prevention | Detection and blocking of customer abuse vectors: free trial abuse, pay-as-you-go abuse, and first-party fraud. The benchmark is 90% precision on free trial abuse detection. |

| Term | Definition or usage in this RFP |
| --- | --- |
| First-party abuse | Fraud committed by real customers—disputing legitimate transactions, exploiting trial offers, or abusing usage-based pricing—causing global losses of $200B annually. |
| Dispute management | Automated chargeback evidence submission, dispute tracking, and win rate analytics. Managing disputes costs $35 for every $100 in chargebacks. |
| False positive rate | The percentage of legitimate transactions incorrectly declined or flagged as fraudulent. This important metric is often omitted from vendor proposals. |
| Custom rules engine | A no-code interface that allows risk analysts to create, test, and deploy custom fraud rules without engineering involvement. |
| Radar Standard | Prebuilt fraud protection with zero integration required, powered by Stripe's AI model. |
| Radar Plus | Custom rules, detailed analytics, and manual controls for risk teams that need to customize their strategies. |
| Radar Pro | Full fraud stack coverage including abuse prevention and adaptive model capabilities for emerging threats. |
| Agentic commerce protection | Fraud scoring and anomaly detection applied to transactions initiated by AI agents. |
| VAMP | Visa's fraud monitoring program. Businesses above VAMP thresholds face fees and potential termination. Dispute management tools should help businesses exit or avoid the program. |
| PCI DSS v4.0 | The current Payment Card Industry Data Security Standard (effective March 2024). A 100% audit success rate is the benchmark. |
| 3DS2 | 3D Secure 2: the authentication protocol for online card payments under PSD2. Dynamic SCA exemption handling minimizes unnecessary friction. |
| Chargeback | A transaction reversed by the card issuer following a dispute. Chargebacks cost the business $35 for every $100 in dispute value. |

# J.3 Requirements quick-reference checklist

For vendor self-assessment before submission.

**AI model and dataset quality**

- AI model's training data footprint described in specific terms—volume, geography, issuers, and payment methods
- Documented fraud reduction from production—32% average is the benchmark
- Granular risk score (0–99) with signal-level explainability per transaction
- Per-transaction fraud scoring with documented false positive rate
- Model update frequency (whether it's continuous or periodic), including incorporation timeline for new vectors
- Network-level intelligence at scale
- Multiprocessor support—risk scores available for externally processed payments
- Preauthorization risk scores for payment routing decisions
- Flexible screening at multiple points in the checkout flow

**Transaction fraud prevention**

- Real-time fraud scoring—p50, p95, and p99 latencies documented
- Configurable blocking thresholds—no code required
- Dynamic 3DS and SCA—applied to high-risk transactions only, not all
- Coverage across cards, ACH, SEPA, wallets, BNPL, and stablecoins
- Card testing detection—raising the detection rate from 59% to 97% is the improvement benchmark
- PAN portability—block and allow lists and fraud history transferable across processors
- Cross-border fraud detection

**Account fraud prevention**

- Account-level fraud scoring at onboarding with documented false positive rates
- Ongoing account monitoring after activation
- Enhanced verification for borderline (not conclusively fraudulent) accounts
- Flagged accounts view—real-time list of suspicious accounts
- Real-time alerts for platform-level merchant fraud
- Automated account actions—reserves, payout restrictions, and blocking
- 3.5M+ fraudulent connected accounts blocked per year—equivalent metric required
- 5.3x fraud exposure reduction—equivalent metric required

**Abuse prevention**

- Free trial abuse prevention—90% precision is the benchmark
- 715,000 high-risk trials blocked and $6M in losses prevented in two months—equivalent required
- 6.2x better detection for new abuse vectors—equivalent required
- Pay-as-you-go abuse prevention
- Customer lifecycle abuse coverage

**Dispute management**

- Automated evidence submission per dispute type
- Win rate probability per individual dispute
- Evidence recommendations per individual dispute
- Smart Refunds or proactive refunding before chargeback
- Early fraud warnings from issuers before formal chargeback
- Native Verifi and Ethoca integration
- Dispute win rate tracking and analytics
- VAMP and monitoring program standing in dashboard

**Custom rules and risk team tooling**

- No-code rule creation in production today
- Plain English rule writing for nontechnical users
- Block and allow lists—IP address, email, card BIN, device fingerprint
- Rule suggestions based on your top fraud indicators
- Shadow mode or what-if simulation against historical data
- Custom business data in rules (e.g., loyalty_tier, product_category)
- Real-time rule performance analytics
- Configurable thresholds by amount, region, account type, or payment method
- Manual review queue with workflow tools

**Analytics and reporting**

- Unified dashboard—fraud, payments, and disputes in one view
- Real-time dashboard updates—not on a delay
- Per-transaction review with risk score and signal breakdown
- Real-time fraud attack alerts with mitigation recommendations
- Fraud pattern investigation tools
- Monitoring program standing in real time
- Data warehouse sync (e.g., Snowflake, BigQuery, Redshift)
- Per-processor analytics filtering
- Programmatic dispute management via API with webhooks

**Agentic and emerging threat coverage**

- Fraud scoring on AI agent–initiated transactions—demonstrate in sandbox
- Anomaly detection for agent behavior
- First-party abuse coverage for AI-native business models
- 200+ product updates per year—improvement cadence evidenced

**Platform architecture and API**

- Zero-code setup—same-day activation
- p99 API latency—production figures required
- 99.999%+ uptime—12-month historical data
- 100% PCI audit success rate—full history
- Full sandbox with production parity including agentic flows
- No-code rules and threshold configuration

**Security and compliance**

- PCI DSS v4.0 (effective March 2024)—100% audit success rate
- SOC 2 Type II
- GDPR and CCPA
- Trusted by 50% of Fortune 100 companies

## J.5 Vendor submission certification

I certify that this submission is complete and that all information provided is accurate to the best of my knowledge. [Your company] reserves the right to verify any claims made in this response.

Company name: _____

Authorized representative: _____

Title: _____

Signature: _____

Date: _____

# How Stripe Connect can help

**Stripe Radar** uses AI models to detect and prevent fraud, trained on data from Stripe's global network. It continuously updates these models based on the latest fraud trends, protecting your business as fraud evolves.

Stripe also offers **Radar for Fraud Teams**, which allows users to add custom rules addressing fraud scenarios specific to their businesses and access advanced fraud insight.

Radar can help your business:

- **Prevent fraud losses:** Stripe processes over $1 trillion in payments annually. This scale uniquely enables Radar to accurately detect and prevent fraud, saving you money.
- **Increase revenue:** Radar's AI models are trained on actual dispute data, customer information, browsing data, and more. This enables Radar to identify risky transactions and reduce false positives, boosting your revenue.
- **Save time:** Radar is built into Stripe and requires zero lines of code to set up. You can also monitor your fraud performance, write rules, and more in a single platform, increasing efficiency.

**Learn more** about Stripe Connect, or **get started** today.