FinSeminarIQ™

Product Features

Turn Retirement Dinner Attendees Into Clients

Prepared by Axon Tech Labs

axontechlabs.com

2026

FinSeminarIQ™ - Executive Brief

Axon Tech Labs | axontechlabs.com | 2026

Introduction

FinSeminarIQ is a purpose-built platform for financial advisory firms running retirement dinner seminar campaigns. This document provides a detailed feature-by-feature overview of the platform covering all five functional areas — campaign management, prospect intelligence, event operations, data and analytics, and platform infrastructure.

Campaign Management

Full Campaign Lifecycle

FinSeminarIQ manages the entire process from prospect data import to post-event analytics — no spreadsheets, no manual tracking, no disconnected tools.

FinSeminarIQ covers the following lifecycle stages of a campaign:

- Importing and management of venues for events
- Importing and management of prospect pools
- Generating mail house export packages
- Creating campaigns and materializing prospects
- Managing registration and check-in for campaign events
- Collecting questionnaire answers
- Conducting post-event prospect scoring, segmentation, and analytics
- Managing the invite suppression information

Campaign Structure Flexibility

A campaign can have up to 4 events — on the same day or different days, at the same venue or different venues:

Format	Example use case
Single event, single venue	Small market, one night, one location
Two events, same venue	Back-to-back nights at the same restaurant
Two events, different venues	Two nights in different parts of the metro area
Four events, multiple venues	High-density market, maximize reach across the region

Automatic Mailer Count Calculation

The platform automatically calculates the mailer count from the total seated capacity across all campaign events, accounting for spousal attendance rate, response rate, and show rate. The campaign manager can override any of these values — the platform recalculates instantly. No manual math required.

Venue proximity ranking

Venue proximity ranking — when a prospect pool is selected, venues are automatically re-sorted by distance from the pool's geographic centroid. The centroid is computed from the pool's ZIP codes using US Census Bureau data — each ZIP code is weighted by its total household count, so a ZIP with more households pulls the centroid toward it proportionally. This means the centroid gravitates toward where prospects are actually concentrated, not just the geometric middle of the ZIP code boundaries. The closest, most convenient venues for the prospect population appear at the top of the list automatically.

Prospect Selection Radius

The platform calculates the prospect selection radius automatically — the smallest radius around the campaign center point that captures exactly enough prospects from the pool to reach the mailer target.

For a single-event campaign the center point is the venue location. For a multi-event campaign, the center point is the capacity-weighted centroid of all venue locations — venues with more seats pull the center toward them proportionally.

The radius is determined by sorting all prospects in the pool by distance from the center point. The platform then finds the distance to the Nth closest prospect — where N is the mailer target. In other words, the radius is exactly wide enough to capture the required number of prospects, and no wider. A dense pool produces a tight radius. A sparse pool produces a wide one.

The calculation works as follows:

Total seats
- → household target (÷ spousal attendance rate)
  - → mailer target (÷ response rate × show rate)
    - → sort pool prospects by distance from center point
      - → radius = distance to the Nth closest prospect

The result is a market-specific radius — not a fixed rule of thumb. The campaign manager can adjust the spousal attendance rate, response rate, and show rate at any time, and can override the calculated radius directly. The platform's calculation is a starting point, not a constraint.

Campaign ROI Model

The economics of dinner seminar campaign can be diagrammed as follows:

10,000 mailers sent
- → ~100 register (1% response rate)
  - → ~80 attend (80% show rate)
    - → ~40 per night across 2 nights
      - → ~8–12 book a consultation (10–15% conversion)
        - → ~3–5 become clients

Metric	Value
Typical campaign cost	$17,000-$36,500
Average AUM per new client	$750,000
Annual advisory fee (1%)	$7,500
Average client tenure	8 years
Lifetime client value	~$60,000

Break-even requires less than one new client. Every additional client is pure margin.

Prospect Intelligence

Prospect Wealth Scoring

Every prospect is scored at campaign materialization time using a five-component composite index. The score is a percentile rank within the specific campaign pool — not an absolute dollar threshold. A score of 85 means this prospect is wealthier than 85% of the candidates pulled for that specific campaign.

Component	Weight
Net worth	35%
Investable assets	25%
Household income	20%
Home value estimate	10%
Discretionary income rating	10%

The score resets per campaign and automatically adapts to local market conditions — a high score in Roseville and a high score in Scottsdale both mean the same thing: top of the local field.

FinSeminarIQ uses wealth score to categorize prospects into tiers:

Tier	Score
VIP Private Wealth Targets	Greater or equal 75
Core Growth Opportunities	Between 50 and 74
Nurture & Mass Market	Below 50

Prospect Selection Strategies

Three selection strategies are available per campaign, giving the advisor control over how the mailing list is constructed.

Strategy	How it works	Best for
Top N	Take the highest-scoring prospects within the radius	Maximum wealth concentration
Random Draw	Random sample from prospects above the minimum score threshold	Avoiding repeat patterns in dense markets
Stratified	Divides pool of prospects into three strata and fills each to a target allocation	Balanced mix of high-value and emerging prospects

For the Top-N strategy, prospects with equal wealth scores are ranked by distance from the campaign center point — the closer prospect is selected first.

Stratified Strata Breakdown

Stratum	Definition	Default allocation
A – The Anchors	Top 10% by wealth score	20% of mailers
B – The Core Growth	Score 60th-89th percentile AND business owner	50% of mailers
C – The Next-Gen	Age 50-60, score below 60th percentile, high proxy indicators	30% of mailers

Planning Clusters

Prospects are automatically assigned to one of seven planning clusters based on their self-scored planning need responses. The cluster tells the advisor exactly what to lead with in the follow-up call — no guessing, no reading through questionnaire responses.

Planning Cluster Definitions

Cluster	Name	Pattern
A	Holistic Wealth	All five planning dimensions scored consistently at 7 or above with no single dominant area
B	Preservation & Protection	Tax and healthcare planning scores combined exceed income and investment scores by more than 3 points
C	Accumulation & Income	Income and investment planning scores both ≥ 8, regardless of scores in tax, healthcare, and legacy dimensions
D	Legacy & Multi-Gen	Legacy planning score ≥ 8 while all other four dimensions — income, investment, tax, and healthcare — score below 8
E	Single Dominant Need	One planning dimension scores ≥ 8 and stands at least 3 points above the average of the remaining four
F	Multi-Need Complexity	Two or more planning dimensions score ≥ 7 with no single dimension scoring ≥ 8 and standing 3 points above the rest
G	Emerging Needs	All five planning dimensions score below 7 with no dominant pattern — prospect has not yet identified a clear planning priority

Advisor Pitch by Planning Cluster

Cluster	Name	Pattern
A	Holistic Wealth	Comprehensive integrated plan
B	Preservation & Protection	Lead with LTC riders and tax-loss harvesting
C	Accumulation & Income	Lead with distribution guardrails and annuities structures
D	Legacy & Multi-Gen	Lead with trusts and generational wealth transfer
E	Single Dominant Need	Laser-focus on the one dominant planning area
F	Multi-Need Complexity	Phased planning roadmap across multiple areas
G	Emerging Needs	Start with a broad retirement readiness conversation

The cluster tells the advisor exactly what to lead with in the follow-up call — no guessing, no reading through questionnaire responses.

Risk Profile

Two questionnaire questions drive the risk profile — investment style (conservative, moderate, aggressive) and how the prospect would react to a 20% portfolio drop (sell, hold, buy more). The combination produces one of nine risk profiles, including behavioral mismatch flags.

Behavioral Mismatch

A prospect who describes themselves as aggressive but would sell on a 20% portfolio drop is flagged as Behavioral Mismatch — Overconfident. A prospect who describes themselves as conservative but would buy more on a 20% drop is flagged as Behavioral Mismatch — Underestimating Risk Tolerance. These flags are critical signals for the advisor going into the first conversation — the prospect's stated style does not match their actual behavior under pressure.

Risk Profile Matrix

Profile	Self-Declared Investment Style	Reaction to 20% Drop	Flag
Conservative Steady	Conservative	Hold	-
Conservative Cautious	Conservative	Sell	-
Conservative Opportunist	Conservative	Buy more	Mismatch
Moderate Balanced	Moderate	Hold	-
Moderate Cautious	Moderate	Sell	-
Moderate Opportunist	Moderate	Buy more	-
Aggressive Growth	Aggressive	Buy more	-
Aggressive Balanced	Aggressive	Hold	-
Aggressive Overconfident	Aggressive	Sell	Mismatch

Event Operations

QR Code Registration

Each prospect receives a personalized mailer with a unique QR code. When they RSVP, they scan the QR code on their smartphone — no phone call, no website to navigate to, no form to fill out manually.

The QR code registration process is fully automated — no staff involvement required at any stage:

- Instant confirmation — seat reserved, registration cap updated in real time
- Automatic waitlist management — if the event is full, the prospect is placed on the waitlist automatically
- Immediate admin visibility — the office manager sees the registration the moment it happens
- Timestamped attendance log — records the exact UTC time each prospect was activated at the door, providing a durable per-event attendance record stored in DynamoDB.

Questionnaire on Smartphone

Prospects receive a place card at their seat with a QR code linking to the questionnaire. During the presentation, they scan and answer 25 questions on their own smartphone.

The questionnaire process is designed to maximize completion rates while minimizing disruption to the event:

- No paper trail — digital capture at the table, no scanning or manual entry
- Instant storage — answers written immediately to the client's own database
- Client data ownership — data lives in the client's dedicated AWS account, not a shared vendor database
- Paper fallback — for prospects who prefer not to use a smartphone, a paper questionnaire is available. The office administrator enters paper responses directly into the platform after the event using the admin application — ensuring every prospect is scored and segmented regardless of how their answers were collected

25 Questions Cover

- Retirement timeline and savings
- Financial situation and debt
- Investment style and risk tolerance
- Primary retirement goal and biggest concern
- Advisor relationship status and consultation interest
- Self-scored planning needs across five dimensions — income, investment, tax, healthcare, and legacy

Invite Suppression

When a campaign is closed, the platform automatically classifies every prospect based on their behavior and assigns a suppression tier. Most prospects rotate back into the eligible pool — suppression is time-limited for all but the most disengaged. This prevents duplicate invitations while ensuring no high-value prospect is permanently excluded based on a single non-response.

Tier	Behavior	Suppression
Permanent	Attended, completed questionnaire, explicitly not interested	Never invited again
Long Cool-Down	No-show after registering, or attended but refused questionnaire	Skip next 4-6 campaigns
Short Rotation	Mailed but did not register, or registered and cancelled	Skip next 1 campaign
No Suppression	Waitlisted or attended and wants consultation	Invite immediately

Automatic Escalation

- Two no-shows within 12 months — permanent suppression
- Three consecutive non-responses — extended cool-down

All suppression records are visible and manageable on the Suppression screen. Manual override is available when needed.

Data & Analytics

Operational to Analytical Data Push

Transactional data — prospects, registrations, check-ins, and questionnaire answers — is stored in real time in DynamoDB. An AWS Step Functions state machine orchestrates the full pipeline, extracting from DynamoDB, transforming, and loading into Redshift Serverless.

The following attributes make the data push process efficient and reliable:

- No timeout ceiling — Step Functions handles any data volume without time constraints
- Single transaction — on any failure the pipeline rolls back automatically, leaving tables in their pre-ingest state
- Two modes — on demand triggered from the admin app, or automated schedule so dashboards are always up to date

AI-Powered Analytics

All nine dashboards are connected to Amazon Q generative AI. Any team member — advisor, office manager, or regional director — can ask questions in plain English and get instant answers from their own data.

Examples of questions that can be answered on the spot:

- "Which planning cluster is most common in Roseville?"
- "What percentage of prospects in each tier want a consultation?"
- "Which office has the highest percentage of VIP prospects?"
- "What if the VIP tier threshold was lowered from 75 to 65 — how many prospects would move up?"

No SQL. No analyst. No waiting. The answer appears in seconds, sourced from the firm's own prospect data. Included in the platform — not a separate add-on.

Nine Dashboards

All nine dashboards are embedded directly inside the admin application. Advisors and managers access their analytics in the same tool they use to run campaigns — no separate login, no context switching, no training on a second platform.

The nine dashboards cover every aspect of the campaign lifecycle:

- Executive Overview — KPIs, office performance, tier distribution, and score analysis across all offices
- Advisor Follow-Up Dashboard — prospect priority rankings, planning clusters, tiers, and geographic map
- Questionnaire Insights — aggregate view of questionnaire responses across all campaigns
- Campaign Analytics — campaign activity and funnel performance across all offices and time periods
- Pool Intelligence Dashboard — pool composition, wealth profile, demographics, and geographic spread
- Prospect Directory — searchable, filterable directory of all prospects across all events
- Venue Analytics — venue performance, attendance rates, consultation rates, and geographic distribution
- Event Analytics — individual event performance, attendance, and questionnaire completion rates
- Invite Suppression Analytics — suppression pool visibility, tier distribution, and office-level suppression health

Platform & Infrastructure

Multi-Office with Role-Based Access

The platform ships ready for all offices at no additional per-office cost. Multi-office is built in from day one.

Two roles control access across the platform:

Role	What They See	What They Can Do
Office User	Only their office's data	Manage campaigns, events, and prospects for their office
Superuser	All offices	Browse, compare, and manage data for any office including importing prospect pools

Headquarters purchases the data, imports it into the right office pool, and the local office manager sees it ready to use — no file transfers, no emails, no manual handoff.

Serverless Infrastructure

Every component is serverless — no servers running 24/7, no idle compute costs. The client pays only for actual usage.

Each component bills only for what is actually used:

Component	What You Pay For
DynamoDB	Per read and write requests
Lambda	Per function invocation
Redshift Serverless	Per query
S3	Per GB stored
API Gateway / CloudFront	Per request

Most platforms charge $2,000–$5,000 per month for a firm this size. FinSeminarIQ's serverless architecture means infrastructure costs scale with actual usage — not a flat monthly fee regardless of activity.

Client Owns Everything

The entire system runs in the client's own AWS account. The client owns the infrastructure and all their data — permanently.

This model provides several important guarantees:

- No SaaS vendor agreement to negotiate
- No data processing addendum
- No vendor lock-in clause
- No shared infrastructure with other firms

Exit Is Instant

Delete one IAM role and Axon Tech Labs access is gone in seconds. Your data stays in your account. Your dashboards keep working. Your prospect history is intact. No migration project, no data export negotiation, no 90-day wind-down.

For a FINRA/SEC-regulated environment, this model gives the firm direct control over its data residency, access logs, and compliance posture — without relying on a vendor's compliance certifications or shared infrastructure.

Enterprise Security Architecture

Security is designed in from day one — not bolted on after.

The following security controls are built into every deployment:

Layer	Security Control
Network	Dedicated VPC with private subnets — no internet access to data or compute
Operational database	DynamoDB in private subnet — accessible only from within the VPC
Analytical database	Redshift Serverless in private subnet — no public endpoint
Inter-component traffic	Security group rules scoped to exact source, destination, and port
Authentication	Amazon Cognito — battle-tested, AWS-managed identity service
Data ownership	All data in client's own AWS account — no shared infrastructure
Vendor access	Axon Tech Labs operates via a scoped IAM role — revocable instantly
Audit trail	All actions logged in AWS CloudTrail — visible to client at any time

No passwords stored anywhere in the system. All authentication via AWS IAM and Amazon Cognito.

Data Provider Independence

The platform is designed to be data provider agnostic. A mapping profile translates the vendor's column names to the platform's internal field names — configured once per vendor, reused for every future import.

Key benefits of the mapping profile approach:

- Vendor independence — switch data providers without changing anything in the platform
- Multi-vendor support — different offices can use different data providers simultaneously
- Schema resilience — if a vendor changes their column names, update the mapping profile once — no code changes

- Validation built in — mandatory fields enforced at import time — incomplete records rejected before entering the system

The initial deployment includes a pre-built mapping profile for Aidentified. Adding a new data provider is a configuration task, not a development project.

Preloaded Venue Database

The platform ships with a curated venue database for each office's geographic region. Venues can be added and edited at any time through the admin application — the preloaded data is a starting point, not a constraint.

Each venue record includes:

- Name, address, city, state, ZIP, and county
- Category — restaurant, hotel, country club, winery, golf course, event venue
- Cuisine type, star rating, and price tier
- Distance from office, website, phone, and GPS coordinates

Venue Analytics

Management gets a full view of venue performance across all campaigns — which venues produce the highest attendance rates, which produce the highest consultation rates, and the average wealth score of prospects by venue. Helps campaign planners choose venues based on past performance data, not gut feel.

Prospect Pool Derive and Merge

Two pool operations give precise targeting control without purchasing new data.

Two non-destructive operations are available:

- Derive — create a focused subset from an existing pool by applying filters such as ZIP codes, wealth score range, age range, or net worth bracket
- Merge — combine two or more existing pools into a single combined pool for campaigns that draw from multiple geographic clusters

Original pools are never modified. One large data purchase can power multiple focused campaigns with different targeting strategies.