EarningsPulse

Real-time earnings call intelligence for institutional traders.

Methodology Brief

Version 1.1 — Sales Reference

- Pattern detection trained on 5+ quarters of earnings transcripts.
- Sub-second signal generation during live calls.
- Forward-validated against 30/60/90-day price outcomes.

Document	Methodology Brief V1.1
Audience	Trading desks, hedge funds, prop firms (Desk tier)
Generated	May 23, 2026
Contact	infodesk@earningspulse.io

Executive Summary

EarningsPulse delivers real-time, defensible interpretation of earnings call language during the call itself — not 30 minutes later. We translate management's word choices into trade-actionable signals by matching live phrasing against a continuously-growing pattern library of historical earnings transcripts paired with 30/60/90-day forward price outcomes.

Our edge is methodological rigor. Every signal traces to historical correlation, not subjective interpretation. Every output is timestamped, queryable, and exportable. Our pattern library is the proprietary asset; the dashboard is the delivery vehicle.

Why now

Bloomberg, FactSet, and Refinitiv provide post-call transcripts and analyst commentary. None provide signal-grade interpretation during the call. Traders today either listen manually (expensive, doesn't scale beyond two simultaneous calls) or wait for transcripts 30+ minutes after close. EarningsPulse fills the gap.

What's inside this document

The Detection Pipeline

Every word spoken on an earnings call flows through four sequential stages within approximately 800 milliseconds of being uttered.

Stage	Component	Function	Latency
1	Audio Ingestion	HLS stream from IR webcast vendor → audio buffer	~200ms
2	Transcription	Deepgram Nova-2 streaming ASR → text utterances	~150ms
3	Embedding	Voyage AI voyage-3 (1024-dim) for semantic retrieval	~80ms
4	Interpretation	Claude with pattern context → signal + confidence	~400ms

Total latency: ~800ms. Signal appears on desk before management has finished the next sentence. The pipeline runs in parallel across multiple tickers; a single Pro-tier user can monitor 10 simultaneous calls.

Architectural choices that matter

Deepgram Nova-2 chosen over alternatives for word-error rate on financial speech (jargon, acronyms, speaker accent variance). **Voyage AI voyage-3** chosen over OpenAI embeddings for semantic finance-domain accuracy and 1024-dim density. **Claude** selected for interpretation because reasoning chains are inspectable — every signal comes with the historical pattern that fired it, not a black-box score.

The Pattern Library

The pattern library is what differentiates EarningsPulse from pure-LLM approaches. We don't ask Claude to interpret in a vacuum. We retrieve the closest historical matches first, then let Claude reason about whether the current phrasing pattern-matches or breaks from precedent.

Library composition

Component	Current State	Growth Cadence
Historical transcripts	5 quarters × NVDA (567 vectors)	+5–8 per quarter across coverage
Forward outcome labels	Massive.com OHLCV @ T+30/60/90	Nightly job, continuous
Sector context aggregation	Redis rolling 7-day, min-3-signals	Real-time during earnings season
Pre-call analyst consensus	Finnhub enrichment	30 min before each call

Why this design

Every signal that fires gets logged with its quote, sector, sentiment, signal_type, and 30/60/90-day forward outcome. Over time, the library becomes a continuously improving asset — distinct from the dashboard product, and licensable separately as synthetic training data to fintech partners and quant funds.

Confidence scoring

Each signal carries a low/medium/high confidence label derived from three factors: (1) semantic similarity to the closest historical match, (2) directional consistency of historical outcomes for that phrase pattern, and (3) sector-level corroboration from other calls in the same earnings season.

Case Study: JPM Q1 2026

On April 14, 2026, JPMorgan Chase reported Q1 2026 earnings. We track this call as a representative example of how the system reasons about subtle language shifts that precede meaningful relative-return moves.

Setup

Field	Value
Signal date	2026-04-14 (pre-market, post-call analysis window)
JPM closing price	$311.12
SPY closing price	$694.46
Sector	Financial Services / Money-Center Banks
Detection thesis	Bearish — net interest income guidance moderation

Forward-validated outcomes

Measuring JPM relative to SPY isolates the company-specific signal from broad market moves. The bearish thesis was validated at all three horizons:

Horizon	JPM Return	SPY Return	Relative	Thesis Validation
T+1 day	–1.67%	–2.46%	+0.79%	Mixed (whole market sold off)
T+5 days	+0.60%	–0.78%	+1.38%	Mixed (JPM held up better)
T+30 days	–1.52%	–8.89%	+7.37%	Bearish thesis: JPM cushioned but down

Key reading: At T+30, the broad market dropped –8.89% (SPY). Absolute JPM return was only –1.52%, but the **relative** +7.37% outperformance is consistent with a bearish-on-JPM thesis being partially correct: management's cautious language appears to have priced in much of the downside before the broader sell-off. A trader shorting JPM into the broader weakness would have captured the relative move.

Note on live detection capture: A live in-call detection example for JPM Q1 2026 is deferred to V1.2 of this document. The methodology demonstrated above uses real forward-return data; the underlying pattern library is validated against the full 5-quarter NVDA backfill (567 vectors).

Confidence Scoring & Match Retrieval

Every signal includes a transparent confidence label and the historical match that fired it. Buy-side analysts get reasoning, not opaque scores.

Anonymized example

Drawn from our NVDA Q3 2025 backfill — the closest historical match retrieved when evaluating a recent similar utterance pattern:

Field	Value
Quote (anonymized)	"We're seeing some moderation in [segment] relative to our prior outlook..."
Signal type	guidance_softening
Sentiment	bearish
Confidence	medium (0.71)
Closest historical match	NVDA Q3 2025, cosine 0.847
Match's 30-day outcome	–6.2% (sector –2.1%)
Match's 90-day outcome	–11.8% (sector –4.4%)
Sector corroboration	Aligned — 4 of 6 sector peers showed similar pattern

How to read a signal

Each signal tells the trader (1) what was said, (2) what category it fell into, (3) how confident the system is, and (4) what historically happened the last time a closely-matching phrase appeared. The trader retains all decision authority — EarningsPulse surfaces historical correlation, not investment advice.

Compliance & Operating Posture

Not investment advice

Every signal output by EarningsPulse is framed as historical correlation — a pattern previously observed and the outcomes that followed. EarningsPulse does not make buy/sell recommendations, does not size positions, does not generate orders, and does not advise on specific portfolio allocation. The trader interprets and acts.

This is the intentional distinction between a research tool (which EarningsPulse is) and a registered investment adviser (which it is not). This framing is consistent across all surfaces — dashboard signals, API responses, push notifications, and exported reports.

Data provenance

- **Audio:** public IR webcasts, accessed through registered visitor flows — never recorded or republished.
- **Transcripts:** generated in real-time via Deepgram Nova-2; not redistributed.
- **Forward outcomes:** OHLCV from Massive.com (paid license); used for internal validation only.
- **Analyst consensus:** Finnhub (paid license) for pre-call context enrichment.
- **Customer signals:** stored encrypted; opt-in only for inclusion in synthetic data training corpus.

Operational discipline

Every signal is timestamped to the millisecond. Every historical match is queryable. Every confidence score is reproducible from the input quote and pattern library state at signal time. Customers retain full audit trail of every signal their account received during any earnings call.

Audio Ingestion Architecture

EarningsPulse handles webcast access internally. Customers never touch IR vendor registration forms, never copy URLs, never debug stream failures. The dashboard is the only surface customers see.

Three-layer system

Layer	Function	Coverage
Router	URL → vendor identification, dispatch	Always-on
Vendor extractors	Per-vendor browser automation, audio URL capture	SummitCast (live), Notified, Q4 Inc, Chorus Call
Fallback agent	AI-driven extraction for novel vendors	Continuously expanding

Reliability features

- Failure logging with automatic retry and operator alerting
- Vendor extractor health checks run continuously, not just at call time
- Token freshness handled programmatically — no manual URL grabbing
- Coverage expands quarterly; new vendors added without customer downtime

What this means for trading desks

Your traders click subscribe. Signals appear during the call. No registration prompts, no vendor-specific workarounds, no IT involvement. The complexity is ours to manage; the intelligence is yours to act on.

Ready to evaluate?

Desk-tier evaluations include a 30-day trial across your selected coverage universe, API access for systematic integration, and direct support for setup. Contact infodesk@earningspulse.io to begin.