

Signal Layer Source Taxonomy

A working catalog of signal sources for the Scan phase of Woodring's Loop

A signal layer is a curated, continuously refreshed stream of weak signals an organization watches — five to fifteen at a time, roughly half inside and half outside. This catalog is the source-side reference for building one: where signals come from, what each source is good for, how often it moves, and who owns it. Use it to seed the Scan inventory at Day 1 of an installation — pull from across all five categories, name an owner per source, and bias toward inside sources, which are usually earlier and almost always less contested than the outside news your competitors are reading too.

System data

Your own operational exhaust — the patterns the organization generates about itself before anyone outside can see them. This is the richest inside-signal category and the most consistently ignored, because the people closest to the data are rarely the people who set the planning agenda. System data is cheap, it is already paid for, and it moves earlier than any external feed. The discipline is to wire it deliberately rather than wait for it to surface in a one-on-one.

SOURCE	WHAT IT SURFACES	CADENCE	OWNER
Win/loss notes (CRM)	Reason drift — losses shifting from "price" to "they do a thing we don't"	Per deal / weekly roll-up	Sales ops / RevOps
Customer success ticket queue	New nouns — feature names, workarounds, competitors nobody on the product team coined	Daily	CS leadership
Product / usage telemetry	Cohort breaks — a region or segment whose engagement or churn moves against the portfolio	Daily / weekly	Product analytics
Pipeline & forecast data	Velocity and mix anomalies — deal stages stalling, segment concentration shifting	Weekly	RevOps
Employee referral & recruiting pipeline	Talent flow at your own door — a competitor that stopped feeding you, or started	Monthly	Talent / HR
Support & sales call transcripts	Unprompted repeated questions — three prospects asking the same new thing this month	Daily (AI-clustered)	Enablement
Frontline pattern reports	What reps and field teams see before it reaches a dashboard	Weekly	Field leadership

- Best for: the earliest read on demand shifts, competitive encroachment, and churn risk — usually weeks ahead of any external signal.
- Wire a vocabulary tracker on CS tickets and call transcripts — flag new nouns appearing in the corpus, not just volume.
- Capture win/loss reason drift quarter over quarter as structured data, not a free-text field nobody reads.

Inside signals are usually earlier and almost always less contested — your competitors are reading the same outside news, but they cannot see your ticket queue. Most organizations assume this plumbing is already installed. It almost never is.

Paid feeds

Subscription and licensed data — the outside-signal sources you pay for because they are structured, timely, and harder for a competitor to assemble casually. Paid feeds buy you speed and coverage, not judgment. The trap is treating a paid feed as if its cost makes its signal strong; most of what a feed delivers is strong signal — already-true metrics — and you are paying it to surface the weak signal buried inside. Budget for two or three feeds that cover your highest-conviction categories, not a wall of dashboards.

SOURCE	WHAT IT SURFACES	CADENCE	OWNER
Prediction markets (Kalshi, Polymarket, Manifold)	Skin-in-the-game probability moving days ahead of public reporting — including wallet flows on-chain	Continuous	Strategy / signal steward
Competitive intelligence platforms (Klue, Crayon)	Competitor pricing, positioning, and launch moves at scale	Daily	Product marketing
Hiring & talent data (LinkedIn Talent, Revelio, Live Data)	Where senior engineers are landing, where a competitor is staffing a category	Weekly	Strategy
Patent & IP feeds (PatSnap, Google Patents Pro)	Capability buildup adjacent to your industry, one productization cycle out	Weekly	R&D / strategy
Earnings & filings intelligence (AlphaSense, Sentieo)	Language drift in transcripts and filings — a rulemaking or pivot being scoped	Per filing	Finance / strategy
Credit & risk data (Bloomberg, S&P, Moody's)	Credit spreads in your category, cost of debt for your customers' customers	Daily	Finance
Trade & shipping data (ImportGenius, Panjiva)	Supply-chain second-order effects before they hit your inputs	Weekly	Operations / supply chain

- Best for: structured outside signal in categories where speed and completeness justify the line item.
- Prediction markets earn their place first — a market with capital at risk moves before the journalists reporting the same event.

- Audit every feed quarterly: if a subscription has not contributed a signal to the layer in a quarter, cut it.

A Scan layer that does not ingest from where the price is moving — prediction markets, credit spreads, trade flows — is running on twentieth-century information.

Public data

Free, open outside-signal sources — government data, regulatory text, public filings, open APIs, and the open web. Public data is the most over-relied-upon category precisely because it is easy to read and easy to talk about in a leadership meeting, which is exactly why the signal in it is the most contested: your competitors have the same access. The edge here is not access, it is reading earlier and reading wider — cross-language capture, regulatory text before the announcement, the comment letter before the rule. AI has collapsed the cost of monitoring this category at an order of magnitude more breadth than was affordable five years ago.

SOURCE	WHAT IT SURFACES	CADENCE	OWNER
FRED & central bank data	Macro conditions — rates, employment, credit, the cost environment for your customers	Daily / on release	Finance
Regulatory docket (Regulations.gov, SEC EDGAR, agency comment letters)	The regulatory whisper — rulemaking being scoped before it is announced	Continuous	Legal / policy
Job postings (company career pages, scraped boards)	A competitor naming a category nobody has named yet	Daily	Strategy
Foreign-language press (Mandarin, Korean, German trade media)	Component-ecosystem and policy signal that never reaches the English press	Daily (AI-translated)	Strategy / region leads
Open APIs (Google Trends, GitHub, Wikipedia pageviews, HackerNews)	Attention and developer-velocity shifts adjacent to your category	Daily	Signal steward
Academic & preprint feeds (arXiv, SSRN, PubMed)	Technology shifts one productization cycle from your industry	Weekly	R&D / strategy
Public earnings calls & investor days (competitors, adjacents)	Stated intent and language drift you can read for free	Per event	Strategy

- Best for: regulatory and macro signal, adjacent-technology shifts, and cross-language capture your competitors are not reading.
- Cross-language is the underused edge — the Korean supplier ecosystem talks about itself in Korean, and translation cost has collapsed.
- Wire these as scheduled ingestion (RSS, API, scrape, watchlist), not as manual checks somebody remembers to run.

Public data is where over-scanning starts. The category can ingest the world; let AI handle the volume and keep the leadership team's judgment on what graduates onto the layer.

Human signal

Signal that lives in people, not feeds — and it runs in both directions. Inside: the frontline rep, the CS lead, the engineer who notices something is off before any system logs it. Outside: the customer who tells you what their roadmap looks like, the partner who hears a competitor's plan, the analyst or operator in your network who watches a corner you do not. Human signal is the highest-bandwidth and the least structured source you have. It does not arrive on a schedule and it does not survive unless you build a deliberate path for it to surface — most of it is lost because there was nowhere to put it.

SOURCE	WHAT IT SURFACES	CADENCE	OWNER
Frontline reps & field teams	Unprompted repeated patterns across prospects and accounts	Weekly	Sales / field leadership
Customer & user advisory boards	Where your customers' priorities are moving before they buy on it	Quarterly	Product / CS
Partner & channel conversations	Competitive moves visible to a partner before they're public	Monthly	Partnerships
Expert & advisor network (calls, briefings)	Specialist read on a category nobody internal owns	On demand	Strategy
Industry events & conference floor	Hallway signal — what operators are worried about off the record	Per event	Designated attendees
Internal subject-matter experts	Engineering or domain read on whether an adjacent shift is real	On demand	Function leads
Recruiting & exit conversations	Why talent is choosing a competitor, why it's leaving them	Per conversation	Talent / HR

- Best for: ambiguity resolution — human signal is how you read which of three meanings a weak signal actually carries.
- Give frontline teams a low-friction capture path that does not depend on whether they happen to mention it in a one-on-one.
- Rotate who attends events and debriefs against the signal layer within 48 hours, while the hallway read is still sharp.

Weak signals need a sponsor. Human signal dies fastest in the organizational immune system — name owners who are responsible for carrying it into the weekly review, not just for hearing it.

Derived indicators

Constructed signals — composites, blends, and computed indices built from the four raw categories above. A derived indicator is not a source you subscribe to; it is something you build because no single feed carries the pattern you need to watch. This is where AI earns its keep on the upstream of Scan: semantic clustering, anomaly detection across messy text, and composite indices that recompute as their inputs move. Derived indicators are powerful and seductive — the trap is building a technically interesting index that means nothing to an actual decision. Every derived indicator should trace back to a strategic question the leadership team is already asking.

SOURCE	WHAT IT SURFACES	CADENCE	OWNER
Composite indices (blended signals)	A theme no single source carries — e.g. a category-emergence index across hiring, patents, and search	Recomputed on input change	Signal steward
Semantic clusters (AI-grouped)	Underlying theme across thousands of items grouped by meaning, not keyword	Daily	Signal steward / analyst
Anomaly flags (statistical)	A metric breaking its own pattern across earnings calls, filings, or telemetry	Daily	Analytics
Sentiment & language-drift trackers	Tone shifting in transcripts, tickets, or press toward a new framing	Weekly	Strategy
Cross-source correlation watches	Two unrelated signals moving together — the second-order effect	Weekly	Signal steward
Prediction-market spread monitors	Divergence between market-implied probability and your internal forecast	Continuous	Strategy / finance
Leading-indicator dashboards (curated)	A small fixed set of metrics chosen to lead, not lag — distinct from an operating dashboard	Daily	Strategy

- Best for: watching patterns that span categories and surfacing weak signal buried inside high-volume data.
- Tie every composite to a named strategic question — kill any index that cannot answer "so what decision does this inform?"
- Give the volume to AI and keep the judgment with the leadership team — Principle 4, and it bites first in this category.

A derived indicator is a signal layer instrument, not a dashboard. A dashboard shows what is already true; a derived indicator is engineered to show what might become true. Do not let the two blur.

Seeding the inventory at Day 1

Use this catalog to stand up the first version of your signal layer — the deliverable at the end of week four of the 90-Day Scan. The goal at Day 1 is not a perfect layer; it is a named, owned set of categories with one or two seed sources each and a leadership team that has read them. Work the checklist in order.

- Name six to eight categories, split roughly half inside and half outside — pick categories where you have specific reason to believe early signal lives, not categories that sound comprehensive on a slide.
- Pull one or two concrete seed sources per category from this catalog — bias the inside half toward system data and human signal, which are earlier and less contested.
- Assign a named human owner to every source — automation handles ingestion, but a person owns the layer.
- Confirm at least one prediction-market or price-moving feed is in the mix — the layer should ingest from where the price is moving, not only where journalists report.
- Wire daily ingestion for the feed-based sources (RSS, API, scrape, watchlist) and a low-friction capture path for human and inside signal.
- Stand up the AI-assisted triage layer to cluster and tag incoming items — volume to AI, judgment to the leadership team.
- Set the weekly review on the calendar — same time, same agenda, same artifact updated each week — and commit to running it four consecutive weeks regardless of how thin the layer feels.
- Cap the active layer at five to fifteen signals; if a category is producing nothing after a month, cut it and rebalance toward the categories that are.

The test of whether you have a signal layer is whether your leadership team can name the five to fifteen signals it is carrying without looking anything up. Build the inventory small enough to be carried in working memory, not filed in a document.