

SALESFUSION

From Signals to Buyer Intelligence

Technical architecture and implementation report

PUBLIC RELEASE · VERSION 1.0.0 · 2026-07-17

From Signals to Buyer Intelligence

Publication metadata

Field	Value
Author	SalesFusion Research
Publication date	2026-07-17
Version	1.0.0
Canonical slug	signals-to-buyer-intelligence
Canonical path	/technical-papers/signals-to-buyer-intelligence
Publication class	Technical architecture and implementation report
Approval state	Founder-approved public release

Revenue Cortex connection

Signals-to-buyer-intelligence is one of the core perception layers of Revenue Cortex. Public signals, forwarded sources, CRM changes, job posts, content, and market events become useful only when the system separates observed facts from hypotheses, route confidence, outreach readiness, and approval state.

Evidence-status box

SalesFusion uses a buyer-intelligence workflow pattern that separates observed signals, sourced facts, inferences, route confidence, recommendations, contact identity, and action readiness. This paper defines that pattern. It does not make reply-rate, conversion, pipeline, or sales-productivity claims. Those claims require analyst-agreement, evidence-precision, and recommendation-utility studies that have not yet been published.

The paper uses provenance concepts consistent with W3C PROV-O, AI risk framing consistent with the NIST AI Risk Management Framework, and threat awareness consistent with the OWASP Top 10 for Large Language Model Applications. These sources support terminology and risk context only.

Abstract

Buyer-intelligence systems often treat a signal as if it were already an outreach reason. A company event, hiring post, product announcement, job change, technology mention, or public statement may be relevant, but relevance is not readiness. The observation must be sourced, converted into a fact with provenance, interpreted as an inference with uncertainty, routed to a plausible business problem, checked against contact identity, assessed for claim risk, and reviewed before an external action. Skipping those steps creates generic outreach, unsupported claims, source-language copying, stale context, and misdirected contact.

We define a stateful architecture for converting signals into buyer intelligence without collapsing evidence into action. Paper 10 specializes the final action-readiness gate as a compound state machine. The architecture has seven separable objects: observed signal, sourced fact, inference, route confidence, recommendation, contact identity, and action readiness. It also includes a strict no-copy rule: source language may inform internal reasoning, but it must not be copied into outreach. SalesFusion uses the design as an internal workflow pattern, with evaluation still needed for analyst agreement and downstream utility.

The problem

Signals are noisy. They are unevenly sourced, often stale, and frequently ambiguous. A public source may say that an organization is hiring for a role, but it does not necessarily prove budget, urgency, buying authority, dissatisfaction, timing, or fit. A webpage may describe a strategic priority, but source language written for a public audience is not safe to paste into a message. A contact profile may point to a person with a similar name, but the identity may not be the right person. A route may look plausible because it matches a familiar selling motion, but the actual evidence may be too thin.

The failure pattern is a single `research_complete` flag. Once set, every downstream system assumes the record is usable: the account is relevant, the contact is correct, the message can cite the event, and outreach can proceed. That is an evidence error. Some records should become questions. Some should become watchlist items. Some should be blocked because the signal is copied from a source that should not be paraphrased into outreach. Some should be suppressed because the contact, company, domain, or channel is not eligible.

Buyer intelligence should therefore be an evidence product, not a scraped-text product. The useful output is not a clever paragraph. It is a dossier of typed claims, uncertainties, routing assumptions, contact checks, action constraints, and reasons to proceed or not proceed.

Contributions

This paper contributes a buyer-intelligence object model, a signal-to-action pipeline, a no-copy outreach rule, decision rules for route confidence and action readiness, and an evaluation plan. It distinguishes seven objects that are often merged:

1. Observed signal
2. Sourced fact
3. Inference

4. Route confidence
5. Recommendation
6. Contact identity
7. Action readiness

The implementation contribution is a SalesFusion pattern for recording each object separately and blocking outbound readiness until all required states pass. The evaluation contribution is a benchmark design that tests whether analysts agree with fact extraction, inference boundaries, confidence labels, route selection, and action gating.

Definitions

Observed signal is the raw detection that something may matter. It can be a public update, structured event, document change, hiring pattern, technology mention, or internal trigger. It is not yet a fact about buying need.

Sourced fact is a claim extracted from an allowed source with provenance, effective time, restriction labels, and source locator. It should preserve what is known without expanding it.

Inference is a bounded interpretation derived from one or more facts. It must be labeled as inference, not fact. For example, a hiring pattern may suggest operational expansion, but it does not prove budget or vendor interest.

Route confidence is the system's confidence that a known business route is relevant to the account or contact, based on evidence fit, recency, specificity, source quality, and missing counterevidence.

Recommendation is a proposed next step, such as collect more evidence, route to analyst review, draft an internal brief, suppress, or prepare for approval.

Contact identity is the verified linkage between person, role, organization, channel, and permitted contact method.

Action readiness is the compound state saying whether the record can support a specific action under current evidence, policy, suppression, QA, and approval conditions.

Signal-to-intelligence architecture

The architecture is a pipeline with explicit gates rather than a text-generation chain.

Layer	Input	Output	Blocking conditions
Signal capture	Raw observation	signal_record	Source not allowed, duplicate signal, tenant mismatch
Fact extraction	Signal and source	sourced_fact	Missing provenance, restricted use, stale source
Inference	Facts and route rubric	inference_record	Fact conflict, unsupported leap, low specificity
Route scoring	Inference and fit model	route_confidence	Missing route evidence, counterevidence, weak fit
Recommendation	Route state and policies	recommendation_record	No action warranted, evidence too thin
Identity check	Person and account records	contact_identity	Unverified person, wrong organization, channel unavailable
Readiness gate	All records	action_readiness	Suppression, claim risk, QA failure, approval missing

The central design rule is that later layers may reference earlier layers but cannot rewrite them. An inference cannot edit a sourced fact. A recommendation cannot silently upgrade route confidence. A contact identity check cannot make weak signal evidence strong. This keeps the intelligence dossier auditable.

The no-copy rule is a separate guardrail:

Source text may be quoted only inside internal evidence review where policy allows it.
 Outreach drafts must not copy source language.
 Outreach claims must use approved, paraphrased, source-backed assertions.
 If the system cannot produce a compliant paraphrase, action readiness is blocked.

The rule exists because source text may carry copyright restrictions, private context, marketing language, prompt-injection text, or claims too strong for the evidence. Even when a source is public, copying its phrasing into outreach can imply more familiarity, permission, or certainty than the workflow has.

Data objects and decision rules

A minimal buyer-intelligence record can be represented as a graph:

```
{
  "signal_record": {
    "signal_id": "sig_001",
    "source_ref": "src_public_version_7",
    "observed_at": "2026-07-16T12:00:00Z",
    "signal_type": "hiring_pattern",
    "raw_use_policy": "internal_evidence_only"
  },
  "sourced_fact": {
    "fact_id": "fact_001",
    "signal_ref": "sig_001",
    "statement": "The source indicates an open role in a relevant function.",
    "effective_at": "2026-07-16T12:00:00Z",
    "epistemic_status": "observed_fact",
    "lifecycle_state": "active",
    "restrictions": ["no_source_language_in_outreach"]
  },
  "inference_record": {
    "inference_id": "inf_001",
    "fact_refs": ["fact_001"],
    "inference": "The organization may be expanding a function related to the route.",
    "confidence": "medium",
    "missing_evidence": ["budget", "owner", "active project"]
  }
}
```

Route confidence is a separate scorecard, not a free-form judgment. A conservative rubric uses evidence dimensions:

Dimension	Question	Example labels
Recency	Is the signal current enough for the route?	current, aging, stale
Specificity	Does the fact identify a concrete business condition?	direct, adjacent, generic
Source quality	Is the source allowed and traceable?	primary, secondary, unknown
Route fit	Does the route map to the observed condition?	strong, plausible, weak
Counterevidence	Is there evidence against the route?	none found, unresolved, blocking
Contact fit	Is the person plausibly connected to the route?	verified, possible, unverified

The recommendation rule should prefer restraint:

```
if source is restricted or stale:
    recommend collect_evidence
elif inference has unresolved conflict:
    recommend analyst_review
elif route_confidence is weak:
    recommend watchlist_or_no_action
elif contact_identity is unverified:
    recommend identity_research
elif suppression applies:
    recommend suppress
else:
    recommend prepare_action_packet
```

Preparing an action packet is not the same as sending. It only means the record may proceed to readiness, QA, and approval checks.

Implementation

SalesFusion uses this pattern as a structured workflow rather than a single research prompt. Signals are recorded with source references and observed time. Facts are extracted into constrained fields with provenance and restrictions. Inferences are labeled, tied to fact references, and prevented from becoming sourced facts. Route confidence is calculated from a rubric that includes recency, specificity, source quality, route fit, counterevidence, and contact fit. Contact identity is handled separately from account intelligence.

The implementation stores null and blocked states. For example, `no_contact_verified`, `source_restricted`, `route_confidence_low`, `claim_risk_high`, and `suppression_match` are valid outcomes of research. They are not treated as failures to be hidden. They provide explicit blocker states when a signal does not support an outreach action.

Interfaces should expose the distinction in review screens. A reviewer should see the observed signal, sourced facts, inferences, confidence rationale, recommendation, identity status, and readiness blockers as different fields. Generated copy, if produced, must be downstream of those fields and cannot directly quote the source. A reviewer should be able to reject the copy while keeping the intelligence record, or reject the inference while preserving the fact.

Evaluation

The first evaluation should measure whether the intelligence objects are correct and useful to reviewers, not whether revenue changes. A benchmark can provide public-safe or synthetic signal packages with ground-truth labels for source eligibility, fact extraction, inference boundary, route fit, contact identity, and readiness state.

Suggested metrics:

Metric	Purpose
Evidence precision	Extracted facts are supported by allowed source material.
Unsupported-inference rate	Inferences exceed the facts or are mislabeled as facts.
Analyst agreement	Independent reviewers agree on route confidence and action readiness.
Source-copy violation rate	Outreach drafts contain prohibited copied source language.
Identity error rate	Recommended contact is not verified for the organization and route.
False-ready rate	Record is marked ready despite suppression, identity, source, or claim-risk blockers.
Recommendation utility	Reviewers judge recommended next step as appropriate under the evidence.

The evaluation should preserve negative labels. A good system may produce many collect_more_evidence, watchlist, or no_action records. That is not necessarily poor performance; it may be correct restraint under weak signals.

Risks and failure modes

The main risk is laundering. A generated dossier can make weak evidence look specific. It can turn a public fact into a private-sounding claim, imply urgency where none exists, or describe an inferred business problem as if the buyer confirmed it. Another risk is identity drift: a signal about an organization is routed to a person whose role, employment, or contact channel is not verified. Source prompt injection is also relevant: a source can contain text that attempts to instruct the agent to ignore rules or contact someone.

Controls include source allowlists, provenance recording, fact/inference separation, confidence rubrics, no-copy enforcement, identity verification, suppression checks, QA, and approval gates. OWASP LLM application risks are useful in reviewing these controls because buyer-intelligence workflows combine retrieval, generated text, external tools, and sensitive business context.

Limitations

The architecture cannot infer buyer intent from thin evidence. It cannot ensure that public signals are accurate, current, or meaningful. It cannot prove that a recommended route is commercially valuable. The SalesFusion pattern supports evidence and confidence separation, but the effect on analyst speed, message quality, or sales outcomes remains unevaluated.

There is also a judgment cost. Separating objects creates more fields to review. Some teams may prefer faster but less governed research for low-consequence internal tasks. The full pattern is better suited where external outreach, reputation, compliance, or data-use restrictions matter.

Practical implications

Teams should treat buyer intelligence as a state record. A useful record says what was observed, what fact was extracted, what inference was made, why a route is plausible, who the contact is, what action is recommended, and why action is or is not ready. It should be normal for a signal to end in no_action.

The no-copy rule should be enforced at the system level, not left as a style preference. Generated outreach should be based on approved claims and constraints, not pasted source language. Reviewers should approve a claim packet and an action packet, not a scraped paragraph.

Conclusion

Signals are not buyer intelligence until they are sourced, interpreted, routed, checked, and gated. SalesFusion uses a workflow pattern that keeps observed signals, sourced facts, inferences, route confidence, recommendations, contact identity, and action readiness separate. The pattern is designed to block or flag unsupported outreach and source-language copying while preserving useful evidence for human review. Its

practical value should be tested through analyst agreement, evidence precision, recommendation utility, and false-ready benchmarks before making business-outcome claims.

References

- W3C, "PROV-O: The PROV Ontology." <https://www.w3.org/TR/prov-o/>
- National Institute of Standards and Technology, "AI Risk Management Framework." <https://www.nist.gov/itl/ai-risk-management-framework>
- OWASP Foundation, "OWASP Top 10 for Large Language Model Applications." <https://owasp.org/www-project-top-10-for-large-language-model-applications/>

Changelog

- 2026-07-17 — Version 1.0.0: founder-approved first public release; publication metadata, evidence-status box, limitations, and verified references retained.