

SALESFUSION

Governed Outbound Readiness as a State Machine

Technical architecture and implementation report

PUBLIC RELEASE · VERSION 1.0.0 · 2026-07-17

Governed Outbound Readiness as a State Machine

Publication metadata

Field	Value
Author	SalesFusion Research
Publication date	2026-07-17
Version	1.0.0
Canonical slug	governed-outbound-readiness-state-machine
Canonical path	/technical-papers/governed-outbound-readiness-state-machine
Publication class	Technical architecture and implementation report
Approval state	Founder-approved public release

Revenue Cortex connection

Revenue Cortex should not treat outbound as ready because a message sounds persuasive. In the revenue-specific system, outbound readiness is a compound state across fit, source evidence, identity, buyer route, claims, suppression, copy review, expiry, and human approval before any send or launch.

Evidence-status box

SalesFusion uses explicit outbound readiness states as a workflow control pattern. The pattern distinguishes fit, signal evidence, contact verification, route confidence, claim risk, suppression, quality review, approval, expiry, and invalidation. This paper defines the state machine and implementation expectations. It does not claim improved response rates, reduced compliance risk, higher conversion, or better deliverability. Those claims require a false-ready, suppression, and policy-violation benchmark and, separately, a field outcome evaluation.

The design uses risk-management framing consistent with the NIST AI Risk Management Framework, provenance concepts consistent with W3C PROV-O, and LLM-application threat awareness consistent with the OWASP Top 10 for Large Language Model Applications.

Abstract

Outbound readiness is often implemented as a copy-completion flag. If a message exists, the record is marked ready. That model lacks the controls required for governed AI-assisted outbound work. A message can be grammatically complete while the account fit is weak, the signal is stale, the contact is unverified, the route is speculative, the claim risk is high, suppression applies, approval has expired, or new information has invalidated the premise. Governed readiness is a compound state machine.

This paper defines an outbound readiness model in which each dimension has its own state, transitions, blockers, expiry, and invalidation rules. It is the action-gating specialization of the buyer-intelligence pipeline described in Paper 09. The final readiness state is derived from the dimensions; it is not edited directly by a model or operator. The architecture is intended for systems that prepare external messages or recommendations and must block or flag unsupported, stale, suppressed, or unauthorized action. SalesFusion uses this as a pattern; safety and operational effects remain evaluation questions.

The problem

Outbound workflows combine several risks. They use business claims about a target organization, identity claims about a contact, channel and suppression rules, human approvals, and generated language. A single error can create external harm: contacting the wrong person, citing an unsupported fact, using copied source language, ignoring suppression, sending after approval expiry, or acting after a source has changed.

The usual implementation pattern is too coarse. A CRM task may say research complete. A copy generator may say draft ready. A reviewer may say approved. A scheduler may say queued. None of those alone means safe to send. Research can be complete but low confidence. Copy can be polished but unsupported. Approval can be valid for an earlier version. A queued action can become invalid when suppression changes or evidence expires.

The practical answer is not more free-form caution text. It is a machine-readable readiness state with explicit dimensions, transitions, and denial reasons.

Contributions

This paper contributes:

- A compound readiness state machine for outbound action.
- A dimension table covering fit, signal evidence, contact verification, route confidence, claim risk, suppression, QA, approval, expiry, and invalidation.
- Decision rules for deriving `not_ready`, `ready_for_review`, `approved_pending_execution`, `ready_to_execute`, `expired`, and `invalidated`.
- Implementation invariants that block a draft, approval, or queue position from becoming readiness evidence by rule.
- An evaluation design focused on false-ready records, suppression failures, stale evidence, and policy violations.

Definitions

Outbound action means any externally visible message, contact attempt, publication, profile update, CRM mutation, or other external consequential step related to outreach.

Readiness dimension is one independent component of readiness. It has a state, evidence, reviewer or system owner, expiry, and invalidation rules.

Derived readiness is the final state computed from dimensions. It is not manually set.

False-ready means the system marks an outbound action executable even though at least one required dimension is blocking, expired, invalidated, or unknown.

Suppression means a rule, list, policy, preference, legal basis, channel constraint, or operational exclusion that blocks or limits action.

Claim risk is the risk that outbound copy asserts more than the evidence supports, uses restricted source language, implies private knowledge, or makes an unapproved representation.

Compound readiness model

The readiness object has one record per proposed outbound action. The action is versioned and digest-bound. Changing the target, channel, message, evidence, workflow version, or claim set changes the digest and forces re-evaluation.

Dimension	Allowed states	Blocking states
Fit	fit_confirmed, fit_plausible, fit_unknown, fit_rejected	fit_unknown, fit_rejected when fit is required
Signal evidence	current_supported, weak_supported, stale, disputed, missing	stale, disputed, missing
Contact verification	verified, role_possible, unverified, mismatch	unverified, mismatch
Route confidence	high, medium, low, not_applicable	low when route is required
Claim risk	low, moderate, high, prohibited	high without review, prohibited
Suppression	clear, limited, matched, unknown	matched, unknown where policy requires certainty
QA	not_started, passed, failed, needs_changes	not_started, failed, needs_changes
Approval	not_required, requested, approved, rejected, expired, revoked, mismatched	requested, rejected, expired, revoked, mismatched
Expiry	valid, expiring_soon, expired	expired
Invalidation	none, evidence_changed, target_changed, policy_changed, message_changed	any change requiring re-evaluation

The derived states are:

Derived state	Meaning
not_ready	One or more mandatory dimensions are missing, blocking, or unknown.
ready_for_review	Evidence and identity are sufficient for human or policy review, but approval or QA is incomplete.
approved_pending_execution	Approval exists but execution gate has not run final checks.
ready_to_execute	All mandatory dimensions pass at execution time for the exact action digest.
blocked	A dimension fails and requires evidence, correction, or authority.
expired	Time-bound readiness or approval has passed its validity window.
invalidated	A material change requires the state to be rebuilt.
executed	An action attempt occurred and has a receipt or blocker record.

Readiness is computed late. A record can be approved_pending_execution and still fail the final gate because suppression changed, evidence expired, or the contact identity was invalidated.

Transition rules

The state machine starts with a proposed action packet:

```
{
  "readiness_id": "rd_001",
  "action_digest": "sha256:...",
  "target_ref": "account_or_person",
  "channel": "external_message",
  "message_version": "msg_v3",
  "claim_refs": ["claim_1", "claim_2"],
  "workflow_version": "wf_v5",
  "derived_state": "not_ready"
}
```

Each dimension transitions through explicit events:

```
fit.assessed
signal_evidence.verified
contact_identity.verified
route_confidence.scored
claim_risk.reviewed
suppression.checked
qa.completed
approval.decided
readiness.expired
readiness.invalidated
readiness.recomputed
```

The derivation function is deterministic:

```

function derive_readiness(record):
  if record.invalidation != none:
    return invalidated
  if record.expiry == expired:
    return expired
  if suppression in [matched, unknown_when_required]:
    return blocked
  if signal_evidence in [missing, stale, disputed]:
    return blocked
  if contact_verification in [unverified, mismatch]:
    return blocked
  if claim_risk == prohibited:
    return blocked
  if qa in [failed, needs_changes]:
    return blocked
  if approval in [rejected, revoked, mismatched, expired]:
    return blocked
  if qa != passed or approval == requested:
    return ready_for_review
  if approval == approved and final_execution_checks_not_run:
    return approved_pending_execution
  if all required dimensions pass:
    return ready_to_execute

```

This function deliberately treats unknowns as blockers where the consequence class requires certainty. Unknown suppression, unknown contact identity, and missing evidence do not become low-confidence readiness. They become blocked or not_ready.

Implementation

SalesFusion uses this as an explicit state model around outbound preparation. The implementation pattern has three key properties.

First, readiness is derived. Operators and models can update dimension evidence through controlled events, but they cannot directly set ready_to_execute. This blocks a generated draft or checklist note from bypassing the state machine by rule.

Second, readiness is bound to an action digest. The digest includes target, channel, message version, workflow version, evidence references, claim set, suppression check timestamp, QA result, and approval reference. If the copy changes, if the route changes, if the contact changes, or if required evidence changes, the prior readiness becomes invalidated or mismatched.

Third, the execution adapter reruns final checks. Queue time can be long enough for evidence, approval, or suppression state to change. The adapter evaluates the current readiness record immediately before any external action and records the policy decision, idempotency key, and receipt.

The interface can be implemented with endpoints such as:

```
POST /v1/outbound-actions
POST /v1/outbound-actions/{id}/evidence
POST /v1/outbound-actions/{id}/suppression-check
POST /v1/outbound-actions/{id}/qa
POST /v1/outbound-actions/{id}/approval
POST /v1/outbound-actions/{id}/derive-readiness
POST /v1/outbound-actions/{id}/execute
```

The public interface should expose blocker reasons in structured form. A reviewer should see `contact_identity=mismatch` or `claim_risk=prohibited`, not just a red icon.

Evaluation

The near-term evaluation is a state-machine benchmark. It should use fixed scenarios with expected readiness states and expected blockers. Scenarios should include stale signal evidence, copied source language, unverified contact, wrong organization, low route confidence, suppression match, unknown suppression, high claim risk, QA failure, approval for a prior message version, expired approval, and evidence change after queueing.

Metrics:

Metric	Definition
False-ready rate	Proportion of blocked scenarios incorrectly marked executable.
False-block rate	Proportion of executable scenarios incorrectly blocked.
Suppression miss rate	Suppressed or unknown-suppression records marked clear.
Approval mismatch detection	Changed action versions detected before execution.
Expiry enforcement	Expired evidence or approval blocks action.
Invalidation latency	Time between material change and readiness invalidation.
Reviewer traceability	Reviewer can identify the blocking dimension and evidence.

Field evaluation should come later and should not be conflated with this benchmark. Deliverability, reply rate, or revenue outcomes require separate design, baseline, sample, time window, and confounder recording.

Safety and failure modes

The most serious failure is a false-ready external action. Common drivers include treating copy completion as readiness, relying on an old suppression check, reusing approval after message edits, sending to an unverified contact, allowing a model to soften high claim risk, or failing to invalidate queued work when evidence changes.

Another failure is silent overblocking. A system can become unusable if every uncertain but low-consequence action is blocked without clear reason or route to resolution. The remedy is not to weaken consequential gates; it is to define consequence classes and allow lower-risk internal actions, such as drafting or research, under different readiness rules.

Prompt injection in source material is also a risk. Source text should never instruct the outbound workflow. It is evidence input only. Copy generation must use approved claim records, not raw source text.

Limitations

This state machine does not decide whether outbound work is strategically worthwhile. It does not produce better copy by itself. It cannot ensure that a verified contact will welcome contact or that a signal reflects actual need. It also depends on the quality of suppression data, source review, and human approval.

The current evidence supports saying that SalesFusion uses explicit readiness states. It does not support comparative claims about safety, efficiency, conversion, or compliance outcomes.

Practical implications

Teams implementing AI-assisted outbound should replace ready with a derived state. Product surfaces should show dimension states, blocker reasons, expiry, and invalidation history. A draft should be treated as one artifact in the workflow, not proof that action is allowed. Approval should be tied to the exact version that will be executed. Suppression should be checked before execution, not only when research begins.

This model also changes review behavior. Reviewers can approve a claim but reject a message, approve a message but require a fresh suppression check, or mark a route plausible while keeping action blocked until identity is verified.

Conclusion

Outbound readiness is a compound evidence state, not a copy-completion flag. Governed execution requires fit, current evidence, verified identity, route confidence, acceptable claim risk, suppression clearance, QA, valid approval, unexpired state, and no material invalidation. SalesFusion uses this readiness-state pattern and proposes a benchmark to measure false-ready, suppression, expiry, and policy failures. Until such evaluation is reported, the proper claim is architectural and operational: explicit readiness states make the grounds for action inspectable.

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.