

THE DISTRIBUTION LAB · INTEGRATION PLAYBOOK

Connecting Agentic Workflows to Guidewire PolicyCenter

The zero-disruption implementation guide for carrier integration engineers, MGA technology leads, and agency technical directors. Implementation patterns, security guardrails, and three production-tested example workflows.

Implementation Patterns

Every connection between a lab module and PolicyCenter uses one of five patterns. Choose per use case — never force one pattern across the board. The anti-corruption layer between them means PolicyCenter is never modified, and the module never learns PolicyCenter's internal data model.

PATTERN 1 · DEFAULT

Cloud REST APIs

PolicyCenter Cloud APIs (Submission, Account, Job, Policy) over HTTPS with OAuth 2.0 client credentials. Reads pull submissions, accounts, and coverage detail; writes land as activities, notes, or documents — never as unapproved transactions.

Use when: the module needs near-real-time reads or structured write-back. **Example:** Cognitive Underwriter scoring a new submission within 20 seconds of creation.

PATTERN 2 · EVENT-DRIVEN

Integration Gateway Events

Subscribe to PolicyCenter's outbound event stream (submission created, policy bound, renewal issued, pre-bind checkpoint). Modules react in seconds without polling, and PolicyCenter carries no query load.

Use when: timing matters — pre-bind checks, triage on creation, renewal triggers. **Example:** Compliance Sentinel flagging an unappointed producer before bind.

PATTERN 3 · BULK DATA

Cloud Data Access (CDA) Extracts

Scheduled extracts from the Guidewire Data Platform into your warehouse, mapped to the canonical schema. Powers analytics-style modules without touching production transaction processing.

Use when: the module scores the whole book or needs history. **Example:** Book Intelligence Warehouse running nightly incremental loads.

PATTERN 4 · LEGACY BRIDGE

ACORD / File-Based Exchange

Where APIs are unavailable or unsanctioned, scheduled SFTP exchange and ACORD XML/JSON message parsing provide a reliable, auditable path. Slower, but honest — and often the only route compliance will approve on day one.

Use when: on-prem PolicyCenter without Cloud API access, or third-party data feeds. **Example:** Loss Run Analyzer attaching normalized analysis to a submission.

PATTERN 5 · IN-CONTEXT

Embedded Sidecar

Module output rendered inside the screens people already use — a score on the submission summary, a draft in the underwriter's activity queue. No new tab, no new login, no adoption campaign.

Use when: the value is in the recommendation, not the automation. **Example:** a triage score with top risk drivers on the underwriter's existing view.

Security Guardrails

These are non-negotiable. Every lab module ships with all eight, and every pilot starts read-only. If a proposed integration can't honor one of these, the integration design changes — the guardrail doesn't.

1 Zero-retention inference

No customer or policyholder data is used to train models. Inference runs in an isolated environment with no persistence of payloads.

2 Your identity provider, your rules

OAuth 2.0 / mTLS against your existing IdP. Service accounts are least-privilege, scoped to exactly the endpoints the module needs, and rotated on your schedule.

3 Read-first deployment

Every engagement begins read-only. Write access is a separate, explicit approval — and even then, writes land as activities, notes, drafts, or documents.

4 The human gate

Nothing binds, sends, or closes on its own. A named human approves every consequential action inside PolicyCenter's existing workflow.

5 Field-level audit trail

Every value read or written is logged with source, timestamp, model version, and confidence score — reconstructable for regulators and reinsurers.

6 Confidence-based routing

Low-confidence extractions and recommendations route to a human queue by default. The module abstains rather than guesses.

7 Replay and rollback

Every sync batch is replayable and every write-back reversible. A failed integration degrades to manual process, never to corrupted records.

8 No production load

Bulk analytics read from CDA extracts or a warehouse replica — never from production PolicyCenter during business hours.

The one-sentence version for your CISO

"Read-only to start, least-privilege service accounts on our own IdP, zero data retention, human approval on every write, and a complete audit trail — PolicyCenter is never modified."

Example Workflows

Three production-tested workflows, end to end. Each names its pattern, its touchpoints, and where the human gate sits.

Workflow A — Submission Triage

COGNITIVE UNDERWRITER · PATTERNS 1 + 2 · REAL-TIME

- 1 **Trigger:** Integration Gateway emits *submission created*; the event lands in the anti-corruption layer.
- 2 **Enrich:** Cloud APIs pull account history, prior policies, and attached documents (ACORD 125/126/130 parsed to the canonical schema).
- 3 **Score:** the module returns an appetite score with top risk drivers and a confidence rating — zero-retention inference.
- 4 **Write-back:** score and rationale post to the submission as a PolicyCenter activity assigned to the underwriter.
- 5 **Human gate:** the underwriter reviews and acts. Nothing is quoted, declined, or referred automatically.

Workflow B — Renewal Communication Drafting

GHOSTWRITER AI · PATTERN 1 · ASYNC, AHEAD OF THE RENEWAL WINDOW

- 1 **Trigger:** nightly job identifies renewals inside the 90/60/30-day windows via the Policy API.
- 2 **Context:** coverage changes, premium movement, and open activities are pulled per account.
- 3 **Draft:** a renewal communication is drafted in the account's tone and templates.
- 4 **Write-back:** the draft is attached to the renewal job as a document, with an activity queued for the account manager.
- 5 **Human gate:** the account manager edits and sends. Nothing leaves without a named approver.

Workflow C — Pre-Bind Compliance Check

COMPLIANCE SENTINEL · PATTERN 2 · REAL-TIME AT BIND

- 1 **Trigger:** a *pre-bind* checkpoint event fires from the Integration Gateway.
- 2 **Verify:** producer license and appointment status are checked against NIPR / state registry data, reconciled with the producer record.
- 3 **Flag:** a mismatch raises an alert activity on the submission with the specific deficiency.
- 4 **Human gate:** the sentinel alerts only — it never blocks the transaction itself. Operations resolves the flag.

The Zero-Disruption Rollout

A typical first module goes from kickoff to production-adjacent pilot in about six weeks, without a single change to PolicyCenter configuration.

WEEK 1

Access & scope

Least-privilege service account provisioned on your IdP; read-only API scopes granted; one use case and one line of business selected. Success metrics agreed in writing.

WEEKS 2–3

Shadow mode

The module reads live submissions and produces output in parallel — visible only to the pilot team. Accuracy is measured against what your underwriters actually did.

WEEKS 4–5

Supervised write-back

Output begins landing as activities and notes in PolicyCenter for the pilot group. Every write is reviewed; the audit trail is validated against your compliance requirements.

WEEK 6

Go / no-go & scale plan

Results reviewed against the agreed metrics. Expand to more lines, add a second module, or stop — the integration layer detaches cleanly with nothing left behind.

Effort at a glance

PATTERN	TYPICAL TIME TO FIRST CONNECTION	POLICYCENTER CHANGES REQUIRED
Cloud REST APIs	Days — read-only pilot in under a week	None — service account only
Integration Gateway events	1–2 weeks	Event subscription configuration
Cloud Data Access extracts	1–2 weeks	None — read from data platform
ACORD / file-based	2–4 weeks	None
Embedded sidecar	Days, after pattern 1 or 2 exists	None

Get the working session

Bring this playbook to a 45-minute working session and leave with a pattern selected, a pilot scoped, and a security review checklist started. Reach Christopher Frankland on LinkedIn ([linkedin.com/in/csfrankland](https://www.linkedin.com/in/csfrankland)) or via bionicagent.com.