



Loan Intelligence

Purpose-built loan package intelligence for lending teams and AI agents.

Credit Policy Brief / July 1, 2026 / Loan Intelligence Team

Credit Policy Intelligence Brief

Credit policy review is where loan package work stops being simple document extraction. A reviewer needs to know what the package says, what the lender policy requires, which calculations are supportable, which documents are missing, and what a human should inspect next.

Loan Intelligence treats credit policy as part of the package workflow. The goal is not to let a model approve loans. The goal is to turn policy requirements, source documents, deterministic calculations, and missing-item logic into review-ready decision support.

Why policy review does not scale as a checklist

Manual policy review breaks down because the facts and rules live in different places. DSCR may depend on rent support and debt service. LTV depends on loan amount and collateral value. FICO, property type, entity structure, occupancy, geography, and document requirements each come from different parts of the file.

A checklist can say that someone looked. It does not prove that the current package supports the calculation, that a revised document changed the result, or that the output remains tied to the policy and source context.

The package-native model

The better operating model is package-native. The package is the durable object. Policy context, source documents, calculations, missing items, and generated outputs all hang from that record.

Manual or prompted policy review vs package-native policy intelligence

The policy check should live with the package record, not inside a one-off checklist or prompt session.

Executive question	Manual / prompted review	Loan Intelligence
What is the durable object?	A checklist, spreadsheet row, or chat answer that has to be re-created when the package changes.	A package record connected to source documents, policy facts, rule results, outputs, credits, and audit boundaries.
How are policy rules represented?	A reviewer restates requirements manually or prompts the model with excerpts.	Numeric thresholds, eligibility rules, exclusions, and document requirements become reusable policy context.
How are calculations controlled?	DSCR, LTV, and payment checks are calculated by hand, spreadsheet, or prompted explanation.	Deterministic calculations run when inputs exist; missing or ambiguous inputs remain explicit gaps.
How does the reviewer verify the answer?	The answer often points back to a note, not the exact package and policy basis.	Policy results are designed to stay source-linked so reviewers can inspect the supporting document context.
Can agents use the same workflow?	Agents need bespoke prompts and fragile context management for each policy pass.	MCP/API clients use the same package, credit, policy, status, and result contracts as the workspace.
What stays out of the public contract?	Implementation details vary by tool and can leak into ad hoc workflows.	Queue keys, staged payloads, cache heads, AppData paths, idempotency internals, and provider secrets stay private.

What executives should require

An executive evaluating policy automation should require more than a policy chatbot or a searchable PDF library.

- The system should accept complete or partial packages through workspace, API, or MCP.
- The policy pass should distinguish supported calculations from missing or ambiguous inputs.
- Numeric checks should be deterministic where the package supports them.
- Source and policy references should be inspectable by a human reviewer.
- Credits should be quoted before package work starts and captured only after durable billable outputs exist.
- Public client responses should not expose queue keys, staged payloads, cache heads, AppData paths, or internal idempotency details.

Where agents fit

AI agents can be useful policy operators when they are connected to a stable backend. The agent can collect files, submit a package, poll status, retrieve result manifests, and summarize policy gaps for a human.

The agent should not become the source of truth. Loan Intelligence is designed so the workspace, MCP, and API all use the same package, credit, policy, status, and result contracts.

Decision-support boundary

Credit policy intelligence is decision support. Loan Intelligence can surface how a package appears to align with uploaded policy context, which documents are missing, and which deterministic calculations are supported. It does not approve or decline loans, replace compliance counsel, or remove human credit judgment.

Executive takeaway

The strategic shift is from policy review as a manual restart loop to policy review as a repeatable package operation. Enterprise AI assistants remain valuable at the interface, but the durable policy workflow belongs in the loan package engine.