Loan Intelligence

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

Executive Brief / July 1, 2026 / Loan Intelligence Team

Loan Intelligence Executive Brief

Loan Intelligence is a purpose-built loan package intelligence engine for lending teams and AI agents. It reads complete or partial loan packages, extracts structured facts, calculates deterministic metrics where inputs support them, identifies missing items, and returns source-linked outputs for human review.

The product is not a chatbot and not an autonomous underwriting engine. It is the operating layer that lets a lending team or AI agent process package work through a repeatable system instead of rebuilding context in every prompt session.

Why this matters now

Lending teams are experimenting with enterprise AI assistants, but real loan packages are a poor fit for session-by-session prompting. A package can contain 10 to 30 source documents, late-arriving revisions, policy exceptions, borrower/entity details, property facts, rent rolls, income records, and collateral materials.

General AI can help explain a document or draft a follow-up. It does not, by itself, create a durable package record, quote work before processing, preserve source-document boundaries, capture credits only after outputs exist, or expose a stable MCP/API contract for agent operation.

Purpose-built package engine vs enterprise AI assistant

General AI assistants remain useful at the edge of the workflow. Loan Intelligence is the operating layer that makes package review repeatable.

Executive question	Enterprise AI assistant	Loan Intelligence
What is the durable object?	A chat session, prompt chain, or uploaded context window.	A loan package record with status, source references, outputs, credits, and audit boundaries.
How are calculations controlled?	Prompted reasoning that must be rechecked every run.	Deterministic DSCR, LTV, payment, and policy checks when inputs exist.
How does cost stay predictable?	Token usage varies with document volume and retry behavior.	Credits are quoted before work starts and captured only after durable billable outputs exist.
Can agents operate it directly?	Usually through custom prompt wrappers or ad hoc API glue.	Native MCP/API tools for quote, process, status, results, packages, balance, and policy state.
What is hidden from users and agents?	Implementation details vary by assistant and tenant setup.	Queue keys, staged payloads, cache heads, AppData paths, and idempotency internals stay private.

What Loan Intelligence does

- Accepts package sources through workspace, API, or MCP.
- Quotes the package before processing so cost is visible.
- Processes source documents behind a worker and legacy-engine adapter boundary.
- Calculates DSCR, LTV, payment, and related metrics deterministically where inputs are present.
- Flags missing items, warnings, policy alignment issues, and tasks.
- Returns workspace, PDF/report, spreadsheet, JSON, and agent-ready outputs.
- Exposes status, results, source references, credit balance, and credit policy through MCP/API.

What changes for the executive buyer

Loan Intelligence turns loan package review into an operational system. That matters for executives because the work becomes easier to measure, repeat, delegate, and integrate.

What changes operationally

Loan Intelligence is strongest where lending teams need the same review method to survive incomplete files, revised documents, and agent handoffs.

Operating need	Manual / generic AI workflow	Loan Intelligence workflow
Incomplete package intake	Reviewer rebuilds context every time another document arrives.	Package state persists; missing items and supported calculations update from the same record.
Source-linked answers	Notes and summaries drift away from the source file.	Facts, tasks, and outputs are returned with source references for human verification.
Policy alignment	Reviewer or prompt has to restate the policy each time.	Uploaded credit policy can become a reusable rule set for package checks.
Reprocessing	A new prompt or spreadsheet pass often starts over.	A package can be rerun against the same workflow as new documents arrive.
Agent handoff	Agent needs bespoke instructions and fragile context management.	Agent gets stable MCP tools, status, result, and credit-state contracts.

Current commercial readiness boundary

The public runtime is live on `loanintelligence.ai`, and production-safe diagnostics prove public web health, MCP status, Paddle readiness diagnostics, and agent-readiness status routes. Local agent-readiness proof has processed the LDIS synthetic `123-easy-st` fixture end-to-end through OTP account creation, API key creation, MCP submission, worker/legacy-engine processing, credit capture, result retrieval, staged-payload purge, and cleanup.

The remaining commercial launch gates are explicit: production email delivery for OTP, public production agent-processing proof, real Paddle checkout/webhook revenue proof, Drive production proof, and final executive collateral verification.

Decision-support boundary

Loan Intelligence supports lending teams. It does not approve or decline loans, replace human credit judgment, or guarantee that every source document is correct. Its job is to make package facts, missing items, calculations, and policy context easier to verify.

> The strategic position is model plus engine: enterprise AI assistants remain useful at the interface, while Loan Intelligence provides the durable package-processing system behind them.