



Loan Intelligence

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

Agent Operator Brief / July 1, 2026 / Loan Intelligence Team

Agent-Ready Loan Package Workflow

An MCP endpoint alone does not make a product agent-ready. A real agent operator needs discovery, account creation, credential setup, package submission, durable status, result retrieval, billing visibility, and cleanup proof.

Loan Intelligence is being built around that standard. The public product explains what the MCP server does. The authenticated workspace creates API keys. MCP tools let an agent quote, submit, poll, retrieve results, list packages, inspect credit balance, and manage credit policy.

What the agent path must prove

The minimum credible proof is a disposable end-to-end run:

- create an account through OTP;
- receive starter credits;
- create an API key;
- configure an MCP client with the public endpoint and bearer credential;
- submit an anonymized synthetic LDIS loan package;
- process through the worker and legacy engine boundary;
- capture credits only after outputs are durable;
- retrieve source-linked results and generated outputs;
- verify no queue keys, staged payloads, AppData/cache heads, or idempotency internals leak;
- clean disposable package, user, key, audit, credit, and queue rows.

Local proof already demonstrates this flow against the LDIS `synthetic/123-easy-st` package. Public production proof is blocked until production OTP email delivery is configured.

Agent-readiness checklist

MCP is necessary but not sufficient. The full workflow must prove account setup, credentials, package processing, billing, result retrieval, and cleanup.

Requirement	Weak MCP wrapper	Production-grade agent path
Discovery	Endpoint exists but the operator must infer what it does.	Public MCP page, metadata, status route, agents guide, and credential instructions.
Credential setup	Manual support ticket or hidden API key flow.	OTP account access and workspace API key creation from the key page.
Real package run	Tool list works, but no end-to-end fixture proof.	Synthetic LDIS package proof submits, processes, captures credits, retrieves outputs, and cleans up.
Billing visibility	Agent can start work but cannot explain credit state.	Credit balance, policy, history, quote, reserve, capture, and release are first-class surfaces.
Internal leakage	Queue/cache details leak into responses or logs used by clients.	Client responses expose durable product ids and source links, not coordination internals.

Why agents need source-linked outputs

Agents are good at condensing results. They are not a substitute for source evidence. When an agent reports a missing lease, a pending DSCR input, a policy issue, or a borrower/property fact, the human reviewer needs to know where that information came from.

Loan Intelligence is designed to return source references and generated output links so the agent can summarize without becoming the source of truth.

What should stay hidden

Agent-facing outputs should be stable product outputs, not infrastructure traces. Loan Intelligence intentionally hides:

- raw queue names and queue keys;
- staged source payload text;
- cache heads and AppData paths;
- internal idempotency keys;
- worker coordination locks;
- provider secrets and operator cleanup credentials.

That separation lets the product expose enough state for automation without turning internal coordination into a public contract.

Current production-readiness boundary

The public runtime exposes MCP discovery, status metadata, agents guidance, Paddle readiness diagnostics, and agent-readiness status diagnostics. The status diagnostic reports whether production is safe for a mutating agent smoke. Today it confirms the operator cleanup token is configured and the remaining blockers are production email provider setup and an approved disposable inbox domain for cleanup verification.

The next production proof is to configure Resend-backed OTP delivery, set the approved disposable inbox domain, verify `/api/ops/agent-readiness/status``, then run the public ``agent:smoke:readiness-http`` path against ``loanintelligence.ai``.

Executive takeaway

Agent readiness is operational readiness. A product is not ready for AI agents until it can prove real package processing, billing behavior, source-linked retrieval, and cleanup from the same public path an operator would give to an agent.