

# Runtime Control Before AI Commitment

LimFlex Wrapper is a reversible runtime control layer for AI workflows. It helps evaluate what should proceed, hold, re-check, escalate, fallback, or learn before AI outputs become business commitments.

Licensing / evaluation-scope discussion only.  
No model replacement. No full migration.  
No implementation outsourcing.

## Reduce premature commitment risk

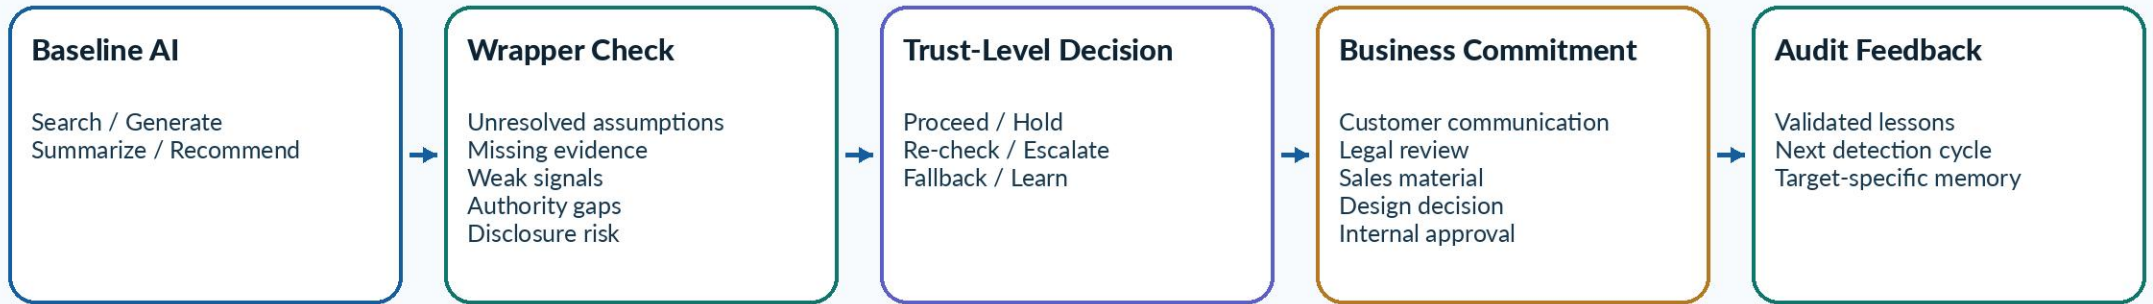
Check unresolved assumptions, missing evidence, weak signals, authority gaps, and disclosure risks before outputs move forward.

## Keep the baseline

Existing AI remains in place. The wrapper can be evaluated around one bounded evaluation target.

## Turn audit into next-cycle learning

Audit is not only after-the-fact explanation. It feeds the next detection cycle.



### PLAN

#### Bounded evaluation plan

1. Select one candidate workflow
2. Keep the baseline
3. Evaluate the wrapper
4. Measure what the baseline would miss
5. Remove it if value is not measurable
6. Scale validated memory if value is measurable

### QUESTIONS

#### Evaluation questions

- What did the baseline miss?
- What should have been held?
- What should have been re-checked?
- Where should human judgment return?
- Can audit feedback improve the next detection cycle?
- Can validated lessons be reused?

### WHAT WE DO NOT CLAIM

#### LimFlex Wrapper is not:

- a replacement for existing AI
- a guarantee of correctness or safety
- automatic approval
- legal / medical / investment / driving judgment replacement
- implementation outsourcing
- operational consulting
- managed service
- human accountability replacement