


Moving from linear pipelines to autonomous agentic loops



Eleven domain experts that explicitly prove their work

11 Independent SME Agents

1 Evaluate: Structured analysis with evidence-linked strengths and concerns



2 Reflect: Check evidence backing ratio (Target: 80%+, Confidence: 0.7+)




3 Tool Verification: Execute query_schema, web_search, and calculate



4 Final Gate: Output strict domain-scoped decisions (PASS, FAIL, NEEDS_CLARIFICATION)



Zero Hallucinations: Every finding traces to reality.



Source Document

Exact Page Location



Extracted Snippet

Final Decision



Quality Formula

$$(\text{Concerns Backed} \times 40) + (\text{Strengths Backed} \times 30) + (\text{Confidence} \times 20) + (\text{Summary Quality} \times 10)$$

The parliament of synthesis resolves conflict through debate

V2 vs. V3 Conflict Handling

- **V2:** Concatenates opinions blindly (Legal PASS + Risk FAIL = Confusion)
- **V3:** Triggers conditional debate to generate reasoned resolutions

5-Step Parliament Flowchart

1. Aggregate:

Run deterministic compiler for baseline packet

2. Synthesize:

Write cohesive executive summary from all 11 SMEs

3. Debate (Conditional):

Triggered on conflict. Generates structured arguments & reasoned resolution

4. Challenge (Conditional):

Validates CRITICAL concerns; downgrades if evidence is weak

5. Question Synthesis:


Semantic deduplication across all SMEs

Highest Bar: 90/100 Quality Threshold
(Human-ready output)



Reports drafted, critiqued, and revised by the same mind

V2: Single-Shot Generation



Flawed, unchecked final report

V3: 3-Phase Loop

1. Draft 
(from UnderwritingPacket)

2. Self-Review 
(critique missing sources and tone)

3. Revise 
(fix unsourced numbers and apply improvements)

The end of first-draft reports.  .
Every claim links to a source document with a severity grading (CRITICAL, HIGH, MEDIUM, LOW) and suggested mitigations. )

[EQUITY]
Focuses on IRR and upside scenarios.

[DEBT_SENIOR]
Focuses on credit rationale and covenant analysis.



Intelligent gap analysis ensures we only ask what matters

The Question Funnel

Gap Analysis: Identifies missing fields, conflicts, and drafts initial questions.

Web Filter: Checks if answers are publicly available online. If yes, deletes questions.

Prioritize:
Ranks the remaining questions by deal (CRITICAL first).



The Numbers (Arkane Cloud test case)

SME (gpu_engineer)	75K tokens ~5 min
IC Panel	20K tokens ~30 sec
Report (equity)	95K tokens ~3 min
Question Agent	21K tokens ~1 min

Total pipeline (single SME):
~211K tokens in ~10 minutes

Validated by **63 tests** (58 integration, 5 E2E). The pipeline didn't get smarter by using a better model—it **got smarter by giving agents** the ability to question themselves and each other.