

PREVIEW SAMPLE - FORMAT DEMONSTRATION ONLY

Not a client engagement, testimonial, case study, or reported result. No client or confidential data was used. Every organization, fact, metric, risk, owner, and recommendation is illustrative and was created solely to demonstrate the method.

1. DECISION AND SOURCE BOUNDARY

Decision under review
Should an illustrative professional-services firm approve a 30-day pilot of an AI-assisted proposal-and-delivery intake workflow? Leadership wants faster qualification and handoff, while delivery, security, and legal have not agreed on the data boundary, review standard, or ownership of AI-assisted output.

- Illustrative inputs**
- Leadership request for a rapid pilot
 - Current intake workflow sketch
 - Security and data-handling excerpts
 - Standard confidentiality terms
 - Delivery-team handoff notes

Evidence represented in the preview inputs

- Intake crosses three functional handoffs and two systems.
- Some source material has confidentiality restrictions.
- No shared acceptance test exists for an AI-assisted intake summary.
- Final proposal and delivery decisions remain human responsibilities.

- Assumptions requiring confirmation**
- Staff will review every generated summary.
 - Prior client material may be reused for context.
 - Existing tool logs are sufficient for audit.
 - Faster intake will not merely move ambiguity downstream.

2. CONSTRAINT AND CONTRADICTION MAP

C1 - Speed vs. human review
Time savings depend on review work that is not measured or assigned. Gate: name the reviewer and standard.

C2 - Learning vs. data minimization
Central reuse may improve learning while increasing access and confidentiality risk. Gate: approve permitted sources and exclusions.

C3 - Reuse vs. confidentiality
Operational convenience does not establish contractual permission. Gate: confirm source rights before ingestion.

D1 - Pilot depends on an output test
Without a baseline, plausible-sounding error can look like acceleration. Gate: define pass, fail, and escalation conditions.

Decision integrity rule
No unresolved assumption is promoted to evidence, and no local workflow result is promoted to legal, security, compliance, or enterprise validation.

3. RISK, OWNERSHIP, AND NEXT DECISION

Restricted material enters workflow
Owner: Security + Legal | Control: Synthetic or approved internal inputs only.

AI text is treated as an approved decision
Owner: Delivery lead | Control: Visible human approval on every output.

Corrections have no owner
Owner: Operations | Control: Named correction path and response time.

Success is declared without a baseline
Owner: Executive sponsor | Control: Predeclared measures and manual comparison.

Recommended next decision
Approve a **shadow-mode pilot only**, contingent on an approved data boundary, a named human decision owner, and a written output acceptance test. No client-facing output or delivery commitment is permitted during shadow mode.

Illustrative 30-day path
Days 1-5: Scope, inputs, owners, failure conditions.
Days 6-12: Configure shadow workflow and comparison set.
Days 13-22: Run bounded cases with review and exception logging.
Days 23-27: Compare usefulness, error modes, and review effort.
Days 28-30: Stop, revise, or advance under separate governance.

What this demonstrates
SynthianAI separates evidence, assumptions, contradictions, dependencies, risks, and decision ownership before more staffing or implementation is committed. It supports a decision; it does not guarantee an outcome or replace licensed review.