AI · AGS SOLUTIONS

AI-Enabled Process Improvement

Your systems already record how work happens. AI can read that record far faster than a team of analysts, surface the friction that costs the most, and propose specific changes — which your people then review, accept, or reject. Nothing changes without a human saying yes.

Reading the process, not the slide deck

Process documentation describes the intended path. Event data describes the real one, including the variants nobody drew. AI is good at exactly this kind of reading: clustering variants, finding the rework loops, and attaching cost to delay.

The output is not a picture of your process. It is a ranked list of specific frictions with the evidence attached.

- Variant clustering shows the handful of paths that carry most of the volume.
- Rework loops are named with their frequency and their cost.
- Delay is attributed to a step and an owner, not to the process in general.

Every suggestion is rated and reversible

An AI suggestion in APIP arrives with a confidence rating, the evidence it was drawn from, and an estimate of the effect. It stays in a pending state until a named person accepts it, and accepted changes are logged so they can be measured and reversed.

This is what makes the approach usable in a regulated environment: the model proposes, the organisation decides, and the decision trail survives the audit.

- Confidence rating and supporting evidence on every suggestion.
- Pending until accepted by a named person.
- Accepted changes measured against the baseline they claimed to improve.

Proving the gain

Improvement work fails politically when nobody can show what it bought. Each accepted change is tied to the metric it targeted — cycle time, rework rate, first-pass yield, breach count — and the before-and-after is reported automatically.

Changes that do not deliver are rolled back with as little ceremony as they were made.

- One metric per change, agreed before the change is made.
- Automatic before-and-after against the same population.

- Rollback is a normal outcome, not a failure.

Video walkthrough — key points

1. Read the real process

- Event data, not documentation
- Variants nobody drew
- Cost attached to delay

Documentation describes the intended path. Event data describes the real one, including the variants nobody ever drew. AI reads that record quickly, clusters the variants, finds the rework loops, and attaches cost to the delay.

2. Ranked friction, with evidence

- A ranked list, not a picture
- Frequency and cost named
- Owner attributed

The output is not another process diagram. It is a ranked list of specific frictions, each with the evidence behind it, the frequency, the cost, and the step and owner it belongs to.

3. Human in the loop

- Confidence-rated suggestions
- Pending until accepted
- Full decision trail

Every suggestion arrives with a confidence rating and the evidence it came from, and stays pending until a named person accepts it. The model proposes, the organisation decides, and the decision trail survives the audit.

4. Prove the gain

- One metric per change
- Automatic before and after
- Rollback is normal

Every accepted change is tied to the metric it targeted, and the before and after is reported automatically against the same population. Changes that do not deliver get rolled back with as little ceremony as they were made.