PRODUCTIVITY · AGS SOLUTIONS

Intelligent Productivity

Skilled people spend a surprising share of their week on work that does not need their judgement: rekeying, chasing, formatting, summarising. Intelligent productivity removes that load with governed AI assistance and cleaner workflow design, so the expensive hours go to the decisions only people can make.

Find the load before you automate it

Automating a badly designed task simply produces the wrong output faster. We start by measuring where hours are actually spent: which steps repeat, which wait, which get redone, and which exist only because a previous system could not do something.

That measurement usually finds three categories — work to delete, work to simplify, and work to assist. Only the third needs AI.

- Delete: steps that exist for a system that no longer exists.
- Simplify: hand-offs, duplicate approvals, and forms nobody reads.
- Assist: drafting, summarising, classification, and retrieval.

Governed assistance, not unattended guessing

Every AI-assisted step in an AGS design has a confidence score, a visible source, and a human decision point where it matters. Assistance is logged in the same event stream as the rest of the process, so its effect on cycle time and quality is measurable rather than anecdotal.

Where confidence is high and the consequence is low, the assistant acts. Where either changes, it drafts and a person approves.

- Confidence-rated output with the source visible.
- Human approval on anything consequential or contested.
- Every assisted action recorded and measurable.

What teams get back

The typical first wave targets admin time and approval latency because both are easy to measure and easy to feel. Reduced rekeying and drafting frees hours per person per week; cleaner approval routing removes days of waiting from end-to-end cycle time.

The second wave is harder to quantify but matters more: people spend their attention on the work that actually needs it.

- Hours per person per week returned from admin.

- Approval wait time cut by removing serial hand-offs.
- Higher-value work measured by share of time on judgement tasks.

Video walkthrough — key points

1. Measure the load first

- Which steps repeat
- Which steps wait
- Which steps get redone

Automating a badly designed task just produces the wrong output faster. So we start by measuring where the hours actually go: which steps repeat, which wait, and which get redone. That measurement is what tells you where assistance is worth having.

2. Delete, simplify, then assist

- Delete legacy steps
- Simplify hand-offs
- Assist with AI where it pays

Most of what we find falls into three buckets. Work to delete, because it exists for a system that no longer does. Work to simplify, like duplicate approvals and forms nobody reads. And only then, work to assist with AI.

3. Governed assistance

- Confidence scores
- Visible sources
- Human approval where it counts

Every assisted step carries a confidence score and a visible source. Where confidence is high and the consequence is low, the assistant acts. Where either of those changes, it drafts and a person approves. Nothing is unattended by accident.

4. The measurable return

- Admin hours returned
- Shorter approval waits
- More judgement work

The first wave targets admin time and approval latency, because both are easy to measure and easy to feel. The second wave is harder to count but matters more: skilled people spend their attention on the work that genuinely needs it.