

# Process Intelligence

Process intelligence reconstructs how work actually flows from the records your systems already keep. Not the diagram, not the interview — the measured path, with every hand-off, wait, exception, and loop in it, and a cost attached to each.

## What the event log knows

Three columns are enough to begin: a case identifier, an activity name, and a timestamp. From those, process intelligence reconstructs every path taken, how often, by whom, and how long each step waited before it started.

Most organisations discover they have four or five times as many real variants as their documentation admits.

- Case, activity, timestamp — the minimum viable event log.
- Resource and cost columns make attribution sharper but are optional.
- Variant counts are usually a shock, and always useful.

## Scoring the bottleneck, not guessing it

A step is a bottleneck when the wait in front of it is large, frequent, and expensive. Scoring combines those three so the ranking reflects business impact rather than the loudest complaint.

The same scoring surfaces the quiet costs: the approval that is fast but wrong a fifth of the time, the exception path that carries a third of the volume.

- $\text{Wait time} \times \text{frequency} \times \text{unit cost}$  gives a comparable score.
- Rework loops are scored on total cost, not on how annoying they feel.
- Exception paths are sized before anyone argues about them.

## Decision support that survives contact

Insight only counts if somebody acts on it. Each finding is presented with the population it applies to, the change it implies, the owner who can make it, and the metric that would prove it worked.

That structure is what turns an analytics project into an operating habit.

- Finding, population, implied change, owner, proof metric.
- Comparable dashboards for operations and for assurance.
- Monthly cadence beats a single large study.

## Video walkthrough — key points

### 1. Three columns are enough

- Case identifier
- Activity name
- Timestamp

You need surprisingly little to start. A case identifier, an activity name, and a timestamp. From those three columns we can reconstruct every path your work actually takes, how often, and how long each step waited before it began.

### 2. The variants nobody drew

- Real paths far outnumber documented ones
- Volume concentrates in a few
- Exceptions get sized

Most organisations find four or five times as many real variants as the documentation admits. A few carry most of the volume, and the rest are exceptions. Once those exceptions are sized, the argument about them stops being an argument.

### 3. Score the bottleneck

- Wait time
- Frequency
- Unit cost

A step is a bottleneck when the wait in front of it is large, frequent and expensive. Combining those three gives a comparable score, so the ranking reflects business impact rather than whoever complained most recently.

### 4. Make it actionable

- Finding and population
- Implied change and owner
- Proof metric

Insight only counts if somebody acts on it. So every finding comes with the population it applies to, the change it implies, the owner who can make it, and the metric that would prove it worked.

© Automated Governance Systems. Shared for use inside your organisation. For a walkthrough against your own processes, book a session at [ags-apip.lovable.app/book-demo](https://ags-apip.lovable.app/book-demo).