



A GUIDE FOR EXECUTIVES AND DECISION MAKERS

Who Owns the Process?

What it costs when nobody owns the process, and why four numbers change the answer

When nobody owns a process, the cost lands in places no report shows: the handoffs, the queues, the migration, the AI programme. This guide sizes that bill, explains why appointed owners fade, and shows how four numbers give a cross-functional executive a process they can see, measure and run.

Who this is for	The executive who will sponsor process ownership, usually the CEO or COO; the CFO and CIO who fund and staff it; and the executives who will be appointed as owners of processes they do not fully control
Reading time	About 15 minutes for Sections 1–3 and 8; 30 minutes in full
Author	David van den Berg, Managing Engineer, Impira Pty Ltd
Version	v6 · September 2026

If you read one page

Ask three questions at your next executive meeting. Who owns Procure-to-Pay, end to end? What does one purchase order cost us, fully loaded? How much of that process runs in the system we paid for? The pauses are the findings.

For an organisation processing 120,000 orders a year, a A\$45 gap to top quartile is A\$5.4 million a year, and it appears on nobody's scorecard.

Process ownership fails when it is a title. It works when it is four numbers, a signature and a meeting. Section 8 describes the test: two or three processes, four weeks, A\$80,000–150,000 ex GST, fixed, paid on accepted deliverables.

If you are being asked to own a process, read 6.2 first.

Contents

1. The bill for an unowned process

Every business process worth the name, Procure-to-Pay, Order-to-Cash, plant maintenance, crosses functions, and no single function owns any of them.

Functions are owned. Processes are not. Geary Rummler and Alan Brache called the gaps between functions the “white space” on the organisation chart: unowned, and therefore unmanaged. Work that leaves one department is nobody’s responsibility until it arrives at the next. Fewer than half of large organisations have an end-to-end owner for Procure-to-Pay (APQC). That is the diagnosis, and it has been known for thirty years. What follows is the bill.

It costs cash every day, in a place no report shows

Impira’s Australian bands put the fully loaded median cost of a purchase order at A\$60–100 and the top quartile under A\$40. APQC’s Open Standards data reports a median of US\$55 per purchase order (n=4,622). Much of the gap between the median and the top quartile is in the handoffs between procurement and accounts payable: the approval by email, the match done in a spreadsheet, the re-keying between two systems. The remainder is catalogue compliance and touchless matching, which are also easier to fix when one person owns the whole chain. Multiply the gap by annual volume and you have the bill. For an organisation processing 120,000 orders a year, a A\$45 gap is A\$5.4 million a year. It appears on nobody’s scorecard, because nobody owns the whole process. The extreme case is the debt that accumulates in the white space: an SAP-published case study of an Australian food manufacturer reported A\$53 million of goods-received-not-invoiced value identified in one process segment. It had accumulated unnoticed. No system flagged it; someone had to go looking.

It traps working capital

In the processes we measure, only 5–10% of elapsed time is work, in line with the flow-efficiency figures Modig and Åhlström report. The rest is waiting: in an approval queue, in a batch run, in a handoff between two teams that report to different people. Each wait sits in a different function’s queue, so no function sees the total. On Order-to-Cash, the waiting that costs cash is the lag between delivery and invoice and the time a disputed order or credit re-issue sits unresolved. Each day of that lag releases the affected revenue ÷ 365, once; on A\$1 billion, one day across the whole book is A\$2.7 million, so even a fraction of a day matters.

It turns improvement programmes into reports

Process mining and process intelligence tools can now show an organisation exactly how its processes run. That evidence is worthless if nobody owns the process it describes. One practitioner in a recent study of process mining put it plainly: without owners, process mining “can end up just being a report.” The findings arrive, it is nobody’s job to act on them, and the report ages on a shelf. The improvements that do get made fade the same way: a system upgrade pushes work back into spreadsheets, and nobody is measured on noticing until the next annual review.

It carries the workarounds into your new ERP

Horváth’s 2025 survey of 200 S/4HANA programmes finds fewer than one in ten finish on schedule and about two thirds report significant budget overruns. The most-cited cause is process complexity discovered late. It is discovered late because nobody could say what the process was before the programme started. Without an owner, every gap between the process and the vendor’s standard is found during the build, at the most expensive moment, and the programme defaults to funding the extension.

It is why the AI investment returns nothing

BCG puts 70% of the AI problem in people and process. Gartner expects more than 40% of agentic AI projects to be cancelled by 2027 for unclear value and inadequate risk controls. An agent acts on the process it is given. Give it an unowned process and it inherits the workarounds and optimises the fragment it can reach. Australia’s Voluntary AI Safety Standard, ISO/IEC 42001 and the NIST AI Risk Management

Framework each require named accountability for an AI system and its use. Where there is no owner, that accountability is unassigned, and the exposure is the organisation's.

It lands on IT as fire, and on the auditor as findings

Problems in the white space reach IT only when they break. An email approval chain adds four days to a process for two years, and the first anyone hears of it is a complaint about the system. The same off-system approvals and spreadsheet matching are control weaknesses, and auditors find them eventually. Meanwhile the executive table argues from three versions of the truth: Finance talks budget, IT talks technical debt, operations talks headcount, and investment follows whoever argues best.

Consequence	What it looks like	How to size it
Cost in the handoffs	Cost per transaction well above benchmark; nobody's budget shows the gap	Cost to Serve against APQC or Hackett, times annual volume
Trapped cash	Days of waiting in queues no single function sees	Cycle Time Efficiency; delivery-to-invoice and dispute days × affected revenue ÷ 365, one-off release
Reports, not change	Process mining findings with no owner; improvements that fade	Count the findings older than a quarter with no named owner
ERP migration risk	Process complexity found during the build; extensions funded by default	Digital Work Ratio per process before scoping locks
AI that returns nothing	Agents deployed on processes that run mostly in email; no accountable owner	Digital Work Ratio as the first gate; a baseline before the agent
Fires and findings	Escalations to IT; off-system approvals in the audit report	First Time Right and the share of approvals with no system record

An unowned process is a line of cost no budget shows and a migration no one can scope.

Size your own bill in ten minutes

Ask three questions in your next executive meeting. Who owns Procure-to-Pay, end to end? The pause before the answer is the first finding. What does one purchase order cost us, fully loaded? The guess is the second. How much of that process actually runs in the system we paid for? Nobody will know, and that is the third. Then do the arithmetic: annual purchase-order count × (your fully loaded cost per order – A\$40). If you cannot fill in your own cost, use A\$85, the APQC median converted, and label the result indicative. The rest of this guide is about turning those three findings into an owner with a scorecard.

2. Why the role fades

If the bill is that large, why has appointing an owner not settled it? Most organisations have tried. APQC has tracked the obstacles to process ownership for years. We group what they and we see into five design faults. Each has a design fix.

2.1 The role is part-time

The owner is a functional head with a day job. Process ownership is added on top. When the two conflict, the day job wins, because the day job is what the owner is measured and paid on.

2.2 The owner has no authority

Hammer's process maturity model (PEMM) rates the process owner on identity, activities and authority at four levels. At the first, the owner is informal and can only lobby. At the second, the role is official, held by "a senior manager who has clout and credibility" who can convene a redesign and has some say over the technology budget. At the third, the owner controls the systems and projects that change the process and

gives it priority over functional duties. At the fourth, the owner controls the budget and influences how people in the process are evaluated. Most organisations appoint at level one and expect level-three results, and a role created without authority fails to achieve its purpose.

2.3 The metrics belong to the functions

Procurement is measured on price variance. Accounts payable is measured on invoices per FTE. The plant is measured on wrench time. Each metric is rational inside its function and blind to the process. Hammer called these the sins of provincialism (letting functional boundaries set the metrics) and pettiness (measuring a small piece of what matters). A process owner handed functional metrics has nothing to own.

2.4 The scorecard is too big to carry

When organisations do build process metrics, they build too many. The SCOR supply-chain model has about 250 metrics. APQC's measure library has thousands. Process intelligence tools expose dozens of configurable KPIs per process. A senior executive who is offered forty numbers will own none of them.

Notice what every mature standard does at the top. SCOR collapses its 250 metrics into its classic five level-one attributes: reliability, responsiveness, agility, cost, and assets. APQC distils its accounts payable library into four core KPIs. Lean shop-floor boards run on four letters: safety, quality, delivery, cost. The pattern is consistent. Executives act on a vital few. Everything else is diagnostic detail for the people who work in the process.

2.5 There is no forum

An owner with a scorecard and no meeting has a hobby. Ownership needs a place where the numbers are reviewed, decisions are taken and the owner is seen to be accountable. Annual reviews are too slow. A steering committee that meets quarterly to look at last quarter's data is a history lesson.

Five faults, one root cause: the role is designed as a title, without the scorecard, authority and forum that make it a job.

3. What a working process owner needs

Strip the literature back and a process owner needs three things. None of them is a title.

Need	What it means in practice	What happens without it
A scorecard they can carry	A small set of end-to-end numbers that describe the whole process, that the CFO, COO and CIO each recognise, and that no single function can improve at another's expense.	The owner is offered functional KPIs or forty process KPIs, and owns neither.
Authority to change the process	An official appointment by the executive, authority over the process rules within the delegation of authority from day one, a named counterpart in each function, and a route to a say over the budget for changes to the process and its systems.	The owner can observe and recommend. Functional heads decide. The role fades within a year.
A forum with a cadence	A regular review where the numbers are read, interventions are approved and their effect is verified, chaired at a level that can hold functions to account.	Findings become reports. Improvements are not verified. Regression goes unnoticed until the next crisis.

Authority and forum are governance decisions. The executive sponsor grants them or does not. The scorecard is a design problem, and it is the one that has been solved badly for thirty years. The rest of this guide is about solving it, and about how a well-designed scorecard makes the other two easier to grant.

One more point before moving on. Process ownership is a part-time role with a real cost. Expect one to two days a month from the owner in the first quarter and about one day a month after that, plus a defined block of time from their team during the baseline (Section 8 gives the figures). A move to Hammer's third level, where the process takes priority over the owner's function, is a separate decision in quarter two, with its own role statement, and the owner can decline it. The design goal is a role a busy executive will accept with that commitment stated up front, not a role that needs a new department and not one sold as a few hours a month.

4. Portfolio Process Mining, from the beginning

This section introduces the method. Readers who know process mining can skip to 4.2.

4.1 Process mining in one paragraph

Every transaction in an enterprise system leaves a timestamped record: a purchase order created, approved, changed, received, invoiced, paid. Process mining reads those records and reconstructs how the process actually ran, case by case, rather than how the process manual says it runs. It shows every path, every loop, every wait and every rework step, with volumes and timings attached. Tools such as SAP Signavio and Celonis do this well and at scale. The result is a true picture of the process as the system saw it.

The limitation is in the last phrase. Process mining reads what the system records. Most process waste never reaches a system: the approval by email, the spreadsheet used to match invoices, the phone call to chase a delivery, the re-keying between two applications. A system-only view does not just miss this work. It reports the process as healthy. Impira's reference Procure-to-Pay data set, a three-region manufacturer with about 120,000 purchase orders a year, modelled on client engagements, shows the effect. All three regions read close to 100% digital on the system extract. After the off-system work was documented and added back, one region fell to 40% digital. The activity view of cost showed why: system touches account for A\$9–13 of the A\$140 fully loaded cost per order in that region. The system-only cost counts only the cost of on-system activities; the enriched cost adds the off-system effort documented in the sessions. The reconciled cost to serve, the process functions' cost pool divided by orders and reconciled to the general ledger, does not move with enrichment. The activity view shows where that cost is consumed, including in cost centres outside the process functions.

4.2 What Portfolio Process Mining adds

Portfolio Process Mining is Impira's method for turning process mining into a management system. It runs on the process intelligence platform an organisation already has. It adds four things.

- **Enrichment.** Your process analysts and ours run short, structured sessions with the people who do the work and capture each off-system step, with its evidence, as a separate, tagged event source in the data model (system, observed or reported). Nothing is written to the system log, so every metric can be read with and without the enriched events. The picture becomes complete, and it stays auditable. Appendix C states the lineage.
- **Four metrics.** Every process, every variant and every business unit is measured on the same four numbers, described below.
- **A diagnostic.** Three of the four numbers, read together, place each process in one of eight patterns; the fourth flags quality risk across all of them. Each pattern calls for a different intervention.
- **Roll-up and continuity.** The same four numbers are read at every level from variant to process to business to enterprise, aggregated as volume-weighted ratios. System-recorded events update as transactions flow; enriched off-system events are dated to the session and re-validated each quarter. Improvement and regression are visible in weeks, not at the next annual review.

4.3 The four metrics

Metric	The question it answers	Where it comes from	Whose question it answers
Digital Work Ratio (DWR)	What share of the work in this process happens on-system, rather than in email and spreadsheets?	Impira's measure. System-recorded activities divided by all activities in the case, after enrichment: a count. The bands (below 40% immature, above 80% leading) are Impira's working bands; no external benchmark exists, and we say so.	CIO
Cost to Serve (CTS)	What does one transaction cost, fully loaded, and how does that compare with benchmark?	Reconciled cost to serve: the process functions' cost pool divided by transaction count, reconciled to the general ledger within $\pm 5\%$ and unchanged by enrichment, mapped to APQC and Hackett Group benchmark scope. An activity view beside it shows where the cost is consumed.	CFO
Cycle Time Efficiency (CTE)	What share of the elapsed time is real work, rather than waiting in queues and approvals?	Lean flow efficiency: value-added time $\div$ elapsed time. Value-added classification is applied in the enrichment protocol; rework, re-keying and chasing are waste. Below 5% is severe; above 20% is lean.	COO
First Time Right (FTR)	How often does a case complete without correction, re-run or rework?	First-pass yield, the standard quality measure, with rework defined per process. Below 60% is critical; above 90% is leading.	Risk and Internal Audit

The process owner carries all four numbers for their process. The executives named here read one number across every process, sponsor it at the executive table, and arbitrate when it conflicts with a functional KPI.

Three of the four are established measures with decades of use behind them. Digital Work Ratio is Impira's addition, and it exists because the off-system blind spot is the largest source of hidden cost in most processes we see. It is not automation rate or touchless rate. Those measure how much on-system work runs without a human. DWR measures how much of the work reaches the system at all, and a process can be 90% touchless on the 40% the system sees. It is also the first gate for AI: an agent acts on the events it can reach, so DWR tells you how much of a process leaves a record an agent can be governed and measured against. Appendix C gives the formal definition of each metric, its known weaknesses and the control for each.

Read together, the four numbers describe a process the way a doctor's four vital signs describe a patient. High DWR with low CTS and high CTE is a healthy process: protect it and copy it. Low DWR with high CTS and low CTE is a process that has left the system, costs too much and waits too long: rebuild it before you automate it. High DWR with high CTS is a process that is used but over-engineered: simplify it. The eight patterns and their interventions are in Appendix A.

Four numbers for every process, at every level of the organisation. That is the whole scorecard, and it is small enough for an executive to carry.

5. How the four metrics make ownership work

Return to the three things an owner needs. Portfolio Process Mining solves the first directly and makes the other two much easier to grant. It does not replace them, and this section is honest about that.

5.1 A scorecard a cross-functional executive can own

The owner of Procure-to-Pay does not need to run the matching rules or the approval workflow; the people in the process do that. The owner needs to know that DWR is 58%, CTS is A\$85 per order against a benchmark of A\$40, CTE is 6% and FTR is 71%. And they need to know which way each is moving. That is a scorecard an executive will accept. It is four numbers. Each answers a question the CFO, CIO or COO gets asked about the process. None can be improved by one function at another's expense. The C-suite reads one number across all processes; the owner carries all four for one process.

5.2 Metrics that are hard to game, and where gaming shows up

Three of the four are end-to-end by construction. Procurement cannot improve cost to serve by pushing work onto accounts payable; the cost moves but the total does not. The plant cannot improve cycle time efficiency by closing work orders early; first time right falls. This is Hammer's provincialism and pettiness designed out, and it is what makes the owner's position tenable in a room of functional heads: the owner reads the process number that sits above everyone's functional numbers, and second-guesses none of them.

Hard to game is not impossible to game, and a CFO will ask. Cost to serve moves if you change the cost pool, split transactions, or move spend out of scope onto purchasing cards. Cycle time efficiency moves if you change where the clock starts. First time right cannot count rework that happens off-system, so a process with low DWR will over-report its quality until it is enriched. And Digital Work Ratio is the exception to end-to-end: its denominator is built from the enrichment sessions, so a thinner enrichment pass raises it without anything changing. The controls are simple and they are written into the charter: definitions and scope are frozen at baseline and changes are restated; every metric is published twice, system-only and enriched, so the size of the off-system layer is always visible; the enrichment protocol is fixed and repeated the same way each quarter; and DWR is a diagnostic, never an incentive. Appendix C sets out each metric's definition, failure mode and control.

5.3 A common language for the executive table

Picture the conversation. The CFO asks why cost to serve is A\$85 per order against a benchmark of A\$40. The CIO answers that DWR is 58% and a third of the work is in Excel. The COO adds that CTE is 6% and approvals are the bottleneck. That is one problem seen from three chairs, not three arguments. The same four numbers roll up to the view the CFO uses to allocate capital, and drill down to the variant a team leader uses to fix a step. They roll up as ratios, volume-weighted, never as averages of averages. Cost to serve rolls up within a process; across processes with different transaction units it becomes a dollar gap to benchmark, not a rate. Everyone argues from one page.

5.4 Targets and a forum come with the method

Cost to serve is benchmarked externally and the other three have maturity bands. So every owner starts with a target for each metric and, on cost to serve, a gap in dollars. Because the system-recorded events update continuously, the forum has something new to read every week; the enriched off-system events are dated and re-validated each quarter. A weekly trend review with process owners, a monthly intervention review, and a quarterly portfolio review, both chaired by the executive sponsor, is enough. Each intervention is registered before it starts with the metric it should move, the size of move expected, and the period needed to see it at that process's volume. The forum judges it against that, not against last week's number. Reacting to ordinary week-to-week variation is how measurement programmes make processes worse.

5.5 What the metrics do not do

They do not confer authority. In Hammer's model a scorecard fixes the Metrics enabler; the Owner enabler moves only when the sponsor grants identity and authority. An official appointment by the executive, with authority over the process rules within the delegations framework, takes the owner to the second level on identity and activities. A say over the systems and projects that change the process takes their authority there too, and toward the third. Those are decisions the sponsor makes, and Section 6 asks for them explicitly. Measurement without accountability produces what practitioners call watermelon metrics: green on the outside, red inside. The four numbers make ownership tractable. The sponsor makes it real.

6. How to stand up process ownership in 90 days

This is the sequence Impira uses. It is designed so that the organisation names owners and quantifies recoverable value in the same quarter, and so that ownership matures on evidence rather than on a restructure. It gives the owner authority over the process rules, within the delegation of authority, on day one, and asks them to earn authority over budget.

Step	When	What to do	Decision or output
1. Choose two or three processes	Week 0	Pick processes with high volume, visible pain and event data in the core system. Procure-to-Pay, Order-to-Cash, plant maintenance and payroll are the usual candidates. Resist picking six.	Named processes. One executive sponsor for the programme.
2. Appoint interim owners, with authority	Week 0, signed at kick-off	The sponsor appoints one senior executive per process, in writing, for 90 days. Owners are named and charters issued with the scope note, and signed at kick-off. The appointment carries three things from day one: authority to set and change the process rules within the delegation of authority (exception routing, step sequence, rework rules), and the standing right to propose changes to approval thresholds and release strategies to the DoA owner, who must answer within the monthly review; a named co-accountable counterpart in each function the process crosses, who owns the sub-process their function runs and carries cost to serve, cycle time efficiency and first time right for it, in the charter now and in their objectives at the next review cycle, with DWR reported beside them and never targeted; and a statement of which of the owner's existing functional KPIs conflict with the four, and who arbitrates.	A named owner per process, with counterparts: P-2 on Hammer's ladder for identity and activities; authority reaches P-2 at step 7.
3. Baseline the four numbers	Weeks 1-4	Extract event data, with user identifiers pseudonymised. Your process analysts and ours run short, structured enrichment sessions with the people who do the work, and calculate DWR, CTS, CTE and FTR per process and major variant. Publish each metric twice: system-only and enriched. Each process team gives about 5-8 person-days, in weeks two and three. Hand each owner a worklist of recoverable items in week 3, with amounts and evidence.	A one-page scorecard per process. Worklists in the owners' hands. A pattern and intervention for each process. Written metric definitions.

Step	When	What to do	Decision or output
4. Set targets against benchmark	Week 4	Compare cost to serve with external benchmark and each metric with the internal best performer. Propose a 90-day and a one-year target per metric. Put a dollar figure on the cost-to-serve gap. Size the worklist recovery effort from the week-3 list; Finance releases it. Finance reviews the value model before the readout.	Targets proposed at the readout and signed by the owner at the first monthly review. A quantified intervention roadmap. Recovery effort sized and released.
5. Start the cadence	Week 5	Weekly trend review with owners. Monthly intervention review chaired by the sponsor, where owners present and decide for their process. Quarterly portfolio review, also chaired by the sponsor. Keep each short; the analysis is prepared for the owner, not by them.	The forum exists and has a rhythm.
6. Run Wave 1 and verify	Weeks 5–12	Execute the quick interventions: exception routing, threshold proposals to the DoA owner, tracked work moved on-system, master-data fixes. The owner's teams work the worklists within the released effort. Register each intervention with the metric it should move, the expected size of the move and the period needed to see it. Judge it against that, not against the next week's number. Value is banked when the Finance sponsor posts it to the benefits register (cash recovered, working capital released, cost removed, capacity redeployed).	The first verified value moves, typically weeks 8–12. Evidence that the owner's decisions change the numbers.
7. Grow the owner's authority	Quarter 2	At the first quarterly review, the sponsor extends the interim appointment and gives owners who have moved their numbers a say over projects and budget that change the process and its systems. Ownership is earned on evidence, which makes the grant easy to defend to the functions. After the first quarter the owner's time is about one day a month. A move to Hammer's third level is a separate decision, with its own role statement, and the owner can decline it.	Owners at Hammer's second level on authority, with the third open to them. The next processes enter the framework.

6.1 The owner charter

Each appointment should fit on one page. Impira uses the template below. The form matters less than the content: the four numbers, the decision rights, the delegation of authority, the counterparts, the conflicts and the forum are written down and signed by the sponsor, the owner and each counterpart before the baseline starts.

Field	Content
Process	Name and boundaries (first event to last event). Business units in scope. Level in the process hierarchy at which the owner sits (usually the end-to-end value stream), and the level at which the metrics are diagnosed (usually the sub-process level, two levels below).
Owner	Name, role, appointment date, term (interim 90 days, then a two-way review). Time: one to two days a month in the first quarter, about one day a month after.
Sponsor	The executive who made the appointment. Chairs the monthly and quarterly reviews.

Field	Content
Co-accountable counterparts	One named person in each function the process crosses: the owner of the sub-process their function runs. Each carries cost to serve, cycle time efficiency and first time right for that segment, and the end-to-end numbers as a shared measure, in the charter now and in their objectives at the next performance cycle. DWR is reported beside them and never targeted.
Scorecard	DWR, CTS, CTE, FTR: baseline (system-only and enriched), 90-day target, one-year target, dollar gap on cost to serve. Targets are re-set, with the owner's signature, on any restatement of definitions, scope or cost pool, and on ERP cutover. Seasonal processes carry a stated comparison period.
Decision rights	Decides: exception routing, step sequence and rework rules within existing controls and the delegations framework. Proposes with priority: changes to approval thresholds (to the CFO as DoA owner) and to system configuration (to IT change control), each answered within the monthly review. Threshold and routing changes are notified to Finance and Internal Audit before implementation; Finance may veto. Escalates to the sponsor: any counterpart who has not implemented an agreed change within the stated period.
Delegation of authority	Which thresholds and release strategies the owner may change within the delegation of authority, and which need the CFO as DoA owner. A change to the DoA goes to the CFO with the owner's evidence.
Conflicting incentives	Which of the owner's and counterparts' existing functional KPIs pull against the four numbers, and who arbitrates. The four numbers are recorded as a shared objective at the next performance cycle; for the interim term, the sponsor confirms in writing that they take precedence where they conflict.
Forum	Weekly, monthly and quarterly reviews: who attends, what is decided at each, and the turnaround for proposals to the DoA owner and to IT change control.
Support and run cost	Who prepares the analysis for the owner, who maintains the data pipeline, and who re-validates the enriched events each quarter. Whose budget carries the analyst run cost (the sponsor's or the process CoE's). Who sizes and releases the worklist recovery effort.
Use of data	Metrics are used to assess the process and the owner's decisions. Enrichment data is not used to assess the people who do the work, and is not used to build a headcount case.
Signed by	Sponsor, owner and each counterpart; HR where objectives change.

6.2 For the person being appointed

If you are the executive being asked to own a process that runs through functions you do not control, ask for these before you accept.

1. Counterparts named and signed by the sponsor: one person in each function, carrying the numbers for the segment their function runs.
2. Authority over the process rules within the delegation of authority: exception routing, step sequence and rework rules from day one, and a standing right to propose threshold changes to the CFO, answered within the monthly review. Most of cycle time lives there.
3. You review the baseline, and how it was measured, before it goes to the executive.
4. The KPI-conflict statement, and the precedence of the four numbers for the term, agreed with your line manager before you accept.
5. Your time commitment in writing, including after the first quarter: one to two days a month in the first quarter, about one day a month after.
6. Worklist recovery effort sized from the week-3 list and released to your teams by Finance. It is additional to the baseline effort.
7. The analyst run cost in the sponsor's budget, or the process CoE's, before the sprint starts.

8. A two-way review at day 90. The sponsor can extend; you can step back.

An owner who accepts without these is accepting the title, and Section 2 describes how that ends.

6.3 If you are about to migrate your ERP

A core-system migration is the strongest forcing function for process ownership an organisation will ever have, and the most expensive place to discover it is missing. Modern programmes run fit-to-standard: each process is compared with the vendor's standard, and every gap is either closed by changing the process or funded as an extension. Without an owner, nobody is accountable for that choice, and the programme defaults to funding the gap. Horváth's 2025 survey of 200 S/4HANA programmes finds fewer than one in ten finish on schedule and about two thirds report significant budget overruns, with process complexity discovered late as the most-cited cause.

The four numbers do not replace the programme's own scoping method. They tell it where to look. A process with low DWR and high cost to serve is where fit-to-standard workshops should start, and where an owner should be in the room. A process at the top of the bands is a candidate to adopt standard with minimal debate; the others are re-engineered or retired, with the evidence on the table. When an owner exists, the programme has someone who can decline an extension request with evidence. During a migration, the monthly intervention review reports into the programme's design authority. It is not a second steering committee. Baseline before scoping locks. Re-baseline after cutover, because the events the new system records are not the same events, and only a fresh baseline makes before-and-after comparable.

If you are already in Realize, the baseline serves the next wave and the cutover: take it on the processes still to be scoped, and on the first wave's processes as the pre-cutover reference for hypercare.

6.4 If you are investing in AI

The larger forcing function is the AI programme most large organisations are now funding. The money is going in and the value is not coming out. MIT NANDA's 2025 GenAI Divide report, from about 150 executive surveys and 50 interviews, found 95% of organisations reporting no measurable return, and named the cause: tools that "integrate poorly, or match workflows" badly, rather than weak models. BCG puts 70% of the AI challenge in people and process and 10% in algorithms. McKinsey found that redesigning the workflow is the single attribute most associated with profit impact from generative AI, and that only 21% of organisations have done it. Gartner expects more than 40% of agentic AI projects to be cancelled by 2027 for unclear business value and inadequate risk controls.

An AI agent needs three things from the process it is put into, and they are the same three things a process owner needs. First, the process must be visible. An agent acts on the events it can reach. At 40% on-system, 40% of the work leaves a record you can govern, audit and measure. An agent in the mailbox may touch the rest, but nobody can show what it did there, and no governance framework will accept that. Work that lives in email and spreadsheets has no system record for the agent to act on or for you to audit. Digital Work Ratio is the first gate, and almost nobody measures it. Second, the process must have an owner. The frameworks Australian boards are asked about each require named accountability for an AI system and its use: Australia's Voluntary AI Safety Standard (guardrail 1), ISO/IEC 42001 (clause 5.3), the NIST AI RMF (GOVERN 2.1) and, where it applies, human oversight under EU AI Act Article 14. None of them says 'process owner'. In our view, for an agent inside a business process, all of them need one. Bain's advice on agentic transformation is blunter: charge general managers with the targets, not the CIO. Third, the process must have a baseline. Both the Australian and UK government Copilot trials measured self-reported time savings and stated plainly that they collected no data on organisational outcomes. Adoption is not value. The four numbers, taken before the agent arrives, are how an executive shows what it did to cost, cycle time and rework afterwards.

The practical sequence is the one in this guide, with one addition. When the owner's scorecard is set, mark each process by pattern. Processes in Protect & Scale or Streamline Flow have the on-system share an agent needs; the AI programme's own risk and integration checks still apply. Processes in Re-engineer or Adopt & Train do not; bring the work on-system first, or the agent inherits the workarounds. This is also the

argument for the platform investment already made: the process intelligence vendors position their platforms as the context layer for agents, Celonis with AgentC and SAP with Signavio as the process context for Joule. The platform is already there. What is missing is the owner and the baseline.

7. Objections we hear, and what we say

Objection	Response
"We tried process owners. It didn't work."	Ask which of the five faults in Section 2 applied. In our experience it is almost always the scorecard and the authority. The role was a title with functional KPIs. Fix those two and try again with an interim appointment.
"Our owners don't control the people in the process."	They never will, and they do not need to. Hammer's process enterprise kept people in their functions. What the owner needs is a process-level number no single function can improve alone, a named counterpart in each function who carries the number for their segment, and a sponsor who will hold the functions to it.
"My general managers run P&Ls. They won't give up budget authority."	They keep it. Step 2 gives the owner authority over the process rules within the delegation of authority: exception routing, the sequence of steps, rework rules. Thresholds move inside the DoA; a DoA change goes to the CFO with the owner's evidence. It does not touch functional headcount or budget. Step 7 asks the sponsor to give owners who have moved their numbers a say over projects that change the process and its systems. It is a say, granted on evidence the GMs have seen, and budget stays where it is.
"IT owns the system, so IT owns the process."	Digital Work Ratio is the CIO's number to read across every process, and the evidence IT needs to fund the fix. The owner still carries it for the process. IT owns the system-only/enriched gap. Paired with the cost-to-serve dollar gap, that is the evidence for a capital case to fix the systems that push work into spreadsheets. Cost to serve and cycle time belong to the business that runs the process.
"We have hundreds of KPIs already."	Keep them. They are diagnostic detail for the people in the process. The owner carries four. Every mature standard, from SCOR to APQC to Lean, does the same thing at the top.
"My ops review already runs on safety, DIFOT, OEE and cost."	Keep it. The four numbers are one line per cross-functional process, read once a month above the functional page: P2P, O2C, maintenance. They do not replace plant metrics; they show what the plant metrics cannot, which is the cost and waiting between the plant and the functions it depends on.
"We already have a process CoE, a repository and a CI programme."	Good. The sprint runs on your repository's hierarchy and identifiers. Your analysts pair with ours from week one and run the metrics from month two. The four numbers become the prioritisation layer for your improvement programme, and the eight patterns tell your belts where to start. Your CoE keeps the definition change log.
"Measurement will feel like surveillance."	The four metrics diagnose process design, not people. Low DWR is a system that has driven people to spreadsheets, not people who refuse to use the system. Say that in the appointment, write into the charter that enrichment data is not used to assess the people who do the work or to build a headcount case, and mean it.
"Process mining showed us the problems. Nothing happened."	That is the report problem, and it is why ownership comes before analysis in this guide. Findings without an owner are findings. Findings with an owner, a target and a monthly forum are a programme.

Objection	Response
"We are doing AI, not process."	Then you need this more. An agent acts on the process it is given. If the process runs 40% in the system and 60% in email, 40% of the work leaves a record you can govern and measure, and nobody can show what the agent did in the rest. If nobody owns the process, nobody owns what the agent does to it. Australia's Voluntary AI Safety Standard, ISO/IEC 42001 and the NIST AI RMF each require named accountability for an AI system and its use. In our view that person is the process owner. Baseline first, so you can show the agent's effect in cost, cycle time and rework rather than in licences activated.

Process ownership fails when it is a title. It works when it is four numbers, a signature and a meeting.

8. Where to start

The fastest way to test this is a Baseline Sprint: two or three processes, four weeks, on the process intelligence platform you already run. Owners are named and charters issued with the scope note, and signed at kick-off. The sprint baselines the four metrics, and recoverable worklists reach the owners in week three. It closes with targets proposed for the owner's signature at the first monthly review, a quantified roadmap, and the metric definitions written down so that your team can re-measure without us. Value is banked when the Finance sponsor posts it to the benefits register, typically weeks 8–12.

The sprint needs Signavio Process Intelligence (or Celonis) with custom event log entitlement. If you hold Process Insights only, we confirm the uplift with your SAP account team before the scope note, and it is your cost, not ours.

It is priced as a fixed fee, paid on accepted deliverables, and it asks the following of you. The figures are typical; the scope note fixes them.

Item	Typical commitment
Impira fee	A\$80,000–150,000 ex GST for two or three processes, fixed and capped, paid in three milestones (30/40/30) against accepted deliverables. The fee varies with process count and number of source systems; a process that straddles ECC and S/4, or SAP and Ariba, counts each. You receive a perpetual, non-exclusive licence to use the method, metric definitions, enrichment protocol and working artefacts internally. Everything sits on your tenant and stays there. Impira retains ownership of the method. Twelve months of updates to the benchmarks and method, and support for your analysts, are included from handover; renewal after that is optional.
Sponsor	Kick-off, two sprint reviews, the readout: about five hours over four weeks, then two hours a month.
Interim owners	One to two days a month in the first quarter: sprint reviews, the worklist, the weekly trend read, the monthly review. About one day a month after that.
Process teams	About 5–8 person-days per process, across six to ten people in weeks two and three (two 90-minute sessions, a ten-question survey, validation of a 20-item sample from each worklist).
Worklist recovery	The owner's teams work the lists after week 3; effort sized from the week-3 worklist, released by Finance, additional to the above.
IT	Three to five days of Basis or data-engineering time per source system for the extract, plus security review. User identifiers are pseudonymised at extraction; session records carry roles, not names; retention is agreed in the extract specification, which we issue in week 1. If the extract needs a transport, allow for your change-freeze and CAB calendar; that is why the specification comes first. Confirm early that your platform licence covers the event volumes and the enriched event source; a platform licence is not the same as a data entitlement.

Item	Typical commitment
Finance	About two days to agree the cost pool and rates for cost to serve. The value model is issued to the Finance sponsor five business days before the readout; reviewed means the Finance sponsor has initialled the assumptions register.
After handover	Half an FTE to one FTE of analyst time to run the metrics, prepare the monthly review and re-validate the enriched events each quarter. This is the run cost of ownership; put it in the sponsor's or the process CoE's budget before you start. Where you have a process CoE, its analysts pair with ours from week one and run the metrics from month two.

If you would like to talk through how this would apply to your processes, contact David van den Berg at david@impira.com.au.

Appendix A. The eight patterns

High and low states of DWR, CTS and CTE place a process in one of eight patterns within four quadrants. The cut points are 60% for DWR, the APQC or Hackett median for the process for CTS, and 10% for CTE; a process near a cut point is read on its trend, not its position. FTR is deliberately held outside the eight patterns: it flags quality risk across all of them rather than multiplying the patterns to sixteen. The patterns also read as an AI-readiness map: a process in Protect & Scale or Streamline Flow has the on-system share an agent needs, and the AI programme's own risk and integration checks still apply; a process in Re-engineer or Adopt & Train does not, and an agent deployed on it will optimise the fragment it can reach.

Quadrant	Pattern	CTE	What it looks like	Intervention
Optimise (high DWR, low CTS)	Protect & Scale	High	Best in class	Make it the internal benchmark. Copy it.
	Streamline Flow	Low	Low cost, but waiting dominates	Approval thresholds, parallel approvals, exception routing
Automate (high DWR, high CTS)	Simplify	High	Over-engineered, but flows	Variant analysis. Remove low-value steps.
	Redesign & Automate	Low	Used, but poorly designed, and waits	Redesign first, then workflow or RPA
Investigate (low DWR, low CTS)	Evaluate Shadow IT	High	Cheap, fast workaround	Understand it before forcing adoption
	Low-Priority Dormant	Low	Slow, off-system, cheap	Intervene only if volume grows
Redesign (low DWR, high CTS)	Adopt & Train	High	Fast, but manual and costly	Remove UX barriers. Move work on-system.
	Re-engineer	Low	Off-system, expensive and slow; the workarounds are the process.	Rebuild. Question the application.

Appendix B. Hammer's process ownership ladder

Adapted from Michael Hammer, "The Process Audit", Harvard Business Review, April 2007. Use it to be honest about where each owner sits, and to plan the next grant of authority.

Level	Identity	Activities	Authority
P-1	An individual or group informally charged with improving the process	Identifies and documents the process; communicates it	Lobbies for the process but can only encourage functional managers to change
P-2	An official process owner role, held by a senior manager with clout and credibility	Sets performance targets and sponsors redesign	Can convene a redesign team and implement the new design; some control over the technology budget
P-3	The owner gives priority to the process over functional responsibilities	Works with other owners to integrate processes toward enterprise goals	Controls the IT systems that support the process and projects that change it; some influence over personnel assignments and budget
P-4	The owner is a member of the senior-most decision-making body	Develops a rolling strategic plan for the process	Controls the process budget and exerts strong influence over personnel assignments and evaluation

Appendix C. Metric definitions and controls

This appendix is for the CFO, the CIO and the architecture review. It states each metric precisely, names the way it can be moved without the process improving, and the control that stops that. Definitions are frozen at baseline and any change is restated.

Metric	Definition and unit	Known weakness or gaming mode	Control
Digital Work Ratio	DWR-count, the published measure: activities that leave a timestamped record in a system of record ÷ all activities in the case, after enrichment. Unit: activity, aggregated to case, then volume-weighted to process. System of record: the list of in-scope systems agreed at baseline and frozen; an activity is on-system if it leaves a timestamped record in one of them. DWR-effort, weighted by duration where durations are known, is a secondary diagnostic and is never mixed with the count. DWR measures where work is recorded, not whether it is automated; touchless rate is a separate diagnostic.	The denominator depends on how completely off-system work is captured. A thinner enrichment pass raises DWR. The count equates a click with a long manual step; DWR-effort shows the difference where durations exist.	Fixed enrichment protocol repeated identically each quarter. Every DWR published beside its system-only value. Duration estimates carry a range. DWR is a diagnostic, never an incentive. No external benchmark exists; bands are Impira's working bands.

Metric	Definition and unit	Known weakness or gaming mode	Control
Cost to Serve	Two numbers, never mixed. (a) Reconciled cost to serve: the process functions' fully loaded cost pool ÷ transaction count. Cost pool: named GL cost elements, internal FTE at loaded rate, application run cost and amortised licence. Reconciled to the general ledger within ±5%. Unchanged by enrichment. This is the number in the value model. (b) Activity view: the cost of on-system activities alone against all activities, which shows where the reconciled cost is consumed, including cost centres outside the process functions. Denominator stated per process (purchase order, sales order, work order) and split by transaction type where mix hides the gap.	Moves with the allocation basis, with scope (spend shifted to purchasing cards, work shifted to requesters), and with transaction splitting. Benchmark scope may differ from the cost pool.	Cost pool and denominator agreed with Finance and frozen. Scope changes restated. Cost pool mapped to APQC scope before comparison; where scope differs, the benchmark is reported as indicative and the variance disclosed.
Cycle Time Efficiency	Value-added time ÷ elapsed time per case. Value-added time is a standard duration per activity, agreed at baseline for system activities (SAP timestamps are instants, so touch time is an assumption with a stated range), and taken from the enrichment range for off-system activities classified value-adding in the enrichment protocol; rework, re-keying and chasing are waste; waiting is everything else. The process figure is the ratio of summed times; the median-case CTE is published beside it. On P2P, faster flow can shorten DPO; payment terms are held constant when the working-capital effect is reported.	Depends on where the clock starts and stops and on which steps count as value-added. Enriched durations are estimates. Batching can flatter it. Long-open cases dominate the summed ratio.	Case start and end events fixed per process. Value-added classification agreed once and held. Reported with a sensitivity range on the enriched share and on the system-step assumption.
First Time Right	Cases completing with no correction, re-run or rework event ÷ all cases, with a stated look-back window. Rework events defined per process (for example price or quantity changes after release, blocked-invoice releases, credit re-issues) and include post-payment exceptions (duplicate payments, credit re-issues, retrospective approvals), so a loosened control surfaces as an FTR fall.	Cannot count rework that happens off-system, so low-DWR processes over-report quality until enriched. A late credit note can retrospectively fail a closed case.	Rework definitions written per process before baseline, post-payment exceptions included. Look-back window stated. Read together with DWR; an FTR rise with a DWR fall is a warning, not a win.

Aggregation. All four are ratios. They aggregate only as the ratio of summed numerators to summed denominators, weighted by volume, and an aggregate can move against every one of its components when the mix shifts. Cost to serve aggregates within a process family only; across processes with different transaction units, report total process cost and the dollar gap to benchmark.

Hierarchy. The owner sits at the end-to-end value stream as defined in the client's process architecture (Procure-to-Pay, Order-to-Cash). Each counterpart owns the sub-process their function runs. Metrics are diagnosed at the sub-process level (PCF level 3 or the client's equivalent) and interventions at variants. Each metric, its definition and its owner are attached to the repository element by its identifier; the enriched as-is variant is recorded against the same element; the CoE owns the definition change log.

Enrichment protocol. Nothing is written to SAP or to the system event log. Off-system steps are captured per variant as step definitions (actor, frequency, duration range, evidence) and loaded as a separate event source, tagged system / observed / reported, joined to cases by variant in the process data model and excluded from conformance checks. Enriched events are dated to the session, held as a custom event source on your tenant, written by named analysts, versioned by quarter; retention and re-validation are stated in the specification. Impira holds no copy after handover. Every metric is computed on the system source alone and on both together. Reported durations carry a range. A sample of reported steps is corroborated against aggregate, de-identified counts (shared-mailbox volumes, ticket counts, file timestamps) agreed with HR and Privacy before the sessions. Two analysts enriching the same variant should land inside an agreed band. The protocol is repeated the same way each quarter. Enrichment data is not used to assess the people who do the work or to build a headcount case, and the charter says so.

Verifying interventions. Each intervention is registered before it starts with the metric it should move, the expected size of the move and the observation period needed to detect it at that process's volume and variation. It is judged when the period elapses or the signal clearly exceeds the baseline range, whichever comes first. Week-to-week noise is not a verdict.

Sources

- Rummler, G. and Brache, A., Improving Performance: How to Manage the White Space on the Organization Chart (Jossey-Bass, 1990; 3rd ed. 2012).
- Hammer, M. and Stanton, S., "How Process Enterprises Really Work", Harvard Business Review, November 1999.
- Hammer, M., "The Process Audit", Harvard Business Review, April 2007 (the PEMM maturity model).
- Hammer, M., "The 7 Deadly Sins of Performance Measurement and How to Avoid Them", MIT Sloan Management Review, Spring 2007.
- APQC, "Why Procure-to-Pay Needs End-to-End Ownership"; "Answering the Challenges of Process Ownership"; "4 KPIs That Set Good Accounts Payable Organizations Apart"; Open Standards Benchmarking, cost per purchase order (median US\$55, n=4,622); Process Classification Framework.
- Supply Chain Council / ASCM, SCOR model: level-1 performance attributes.
- Mamudu, A. et al., "Process Mining Success Factors and Their Interrelationships", Business & Information Systems Engineering (2024).
- Horváth & Partners, "SAP S/4HANA transformations rarely go as planned" (April 2025, n=200).
- SAP, customer story: "GWF: Unlocking working capital through new levels of transparency, accountability, and insight into a critical financial process" (August 2025), on goods-received-not-invoiced value identified with SAP Signavio: www.sap.com/asset/dynamic/2025/08/d8f5cd33-1a7f-0010-bca6-c68f7e60039b.html
- Modig, N. and Åhlström, P., This is Lean (Rheologica, 2012), on flow efficiency. ASQ Quality Glossary, first-pass yield.
- MIT NANDA, The GenAI Divide: State of AI in Business 2025 (July 2025). BCG, Where's the Value in AI? (October 2024). McKinsey, The State of AI (March 2025). Bain & Company, State of the Art of Agentic AI Transformation (September 2025).

- Gartner, press release on agentic AI project cancellations (June 2025). Celonis, AgentC announcement (October 2024). SAP, “Welcome to the (Process) Conversation: Joule with SAP Signavio Solutions Now Generally Available”, SAP News, 10 February 2026, on SAP Signavio as the process context for Joule.
- Australian Government Digital Transformation Agency, Evaluation of the whole-of-government trial of Microsoft 365 Copilot (October 2024). UK Government, Microsoft 365 Copilot cross-government experiment (2025).
- Australian Government, Voluntary AI Safety Standard (August 2024). NIST AI Risk Management Framework 1.0, GOVERN 2.1. ISO/IEC 42001:2023, clause 5.3. EU AI Act, Article 14.
- Impira Pty Ltd, Portfolio Process Mining: a data-driven framework for application and portfolio management (2026).