



INDUSTRY SOLUTIONS

Operational AI for Accounting Firms

A guide to the recurring work inside an accounting firm, and the AI roles that can take it on.



A practical reference for firm leaders deciding where AI is worth building and where it is not.

6

representative AI
roles mapped

45 min

AI Workforce Review
to find the first role

50

open-source workflow
patterns published

RECURRING WORK · AI ROLES · SYSTEMS · OVERSIGHT · VALUE CASE · FIRST ROLE



Start with **the recurring work**, not the technology.


A surprising amount of firm capacity is lost in the handoffs between clients, preparers, reviewers, and the systems that hold the engagement. When those handoffs depend on someone rebuilding a list or re-tracing an entry, the cost shows up as reviewer time, missed deadlines, and busy-season overtime.

Four recurring patterns appear across accounting firms:


Client records arrive late or wrong Files come in mislabeled, duplicated, or for the wrong period. Staff rebuild the missing-items list again and again before work can start.	Close status that looks finished but is not Tasks show complete while support, dependencies, or reviews remain open. The gap appears in the close meeting, when it is hardest to fix.
Investigations repeated at every level Staff trace source entries before they can explain an unmatched balance, and reviewers repeat the same search to check it.	Review knowledge lost between engagements Comments are routed one engagement at a time, so the same issue comes back next month for a different client.

How an AI role responds

The work keeps moving while accountable people remain in control.



Every Metacto AI role follows the same operating loop.





What an AI role in an accounting firm **actually** does.

Illustrative example: a bank reconciliation prepared before review

ANATOMY OF AN AI ROLE

reconciliation prep ai role · client 0412

✓ SYSTEM OF RECORD

✓ EVERY LINE CITED

1

TRIGGER

Bank feed · Oct close · Client 0412

Three October deposits do not match the ledger.

No explanation yet. Review is scheduled for Thursday.

2

Brief R-118 · October bank reconciliation

CLOSE DAY 3

RECONCILIATION PREP

Deposit 10/03 · matched to invoice 2207	cited: bank statement line 14, AR invoice 2207	\$12,480
Deposit 10/11 · no invoice match	cited: bank statement line 31, client question drafted	\$4,150
Wire 10/19 · duplicate posting suspected	cited: bank statement line 52, GL 1010	\$27,300
Items in the brief		\$43,930

3

SENIOR DECISION

Approve

awaiting senior

edit · override · reject, all logged

4

WRITE-BACK

written back

Workpaper R-118 updated · client question sent · approved by your senior · timestamped

!

EXCEPTION

No support in the client portal? The item is held, and the request goes to the client as a question.

✓ HUMAN IN THE LOOP

✓ ASKS, NEVER GUESSES

✓ MEASURED TO A BASELINE

Success metric: Reconciliation items explained before review, not during it.



The same standard, across the work that runs an engagement.

Each AI role is deliberately narrow: it owns one repeated piece of work, uses approved records, cites its sources, and routes decisions to an accountable owner.

Engagement Setup AI Role	Assembles new client and engagement records for staff to confirm.	4 setup items · ready to confirm	Engagement start time
Document Desk AI Role	Matches uploads to the PBC list and drafts only the follow-ups still needed.	11 requests outstanding	Collection delay
Reconciliation Prep AI Role	Preps bank reconciliations and builds the exception brief with source support.	R-118 · 3 items flagged	Review time per rec
Close Watch AI Role	Separates missing support, dependencies, and reviewer waits before the close meeting.	7 blockers by owner	Close days
Review Notes AI Role	Packages each review note with its evidence and proposes the preparer response.	19 comments with preparers	Repeat review notes
Delivery Pack AI Role	Compiles only approved outputs, with completion evidence, into the client package.	Client 0412 · ready	Corrections after delivery

Any work that recurs across engagements and carries real cost in reviewer time, deadlines, or client experience may be a candidate. Before funding a build, evaluate it against the Five Conditions of a Fundable AI Workflow:

THE FIVE CONDITIONS OF A FUNDABLE AI WORKFLOW

Business impact	Workflow repeatability	Source readiness	Decision ownership	Implementation path
-----------------	------------------------	------------------	--------------------	---------------------

Do not automate everything at once. Start with one AI role that has a credible baseline, a clear owner, and a measurable outcome.



AI roles work **inside the systems** you already run.

An AI role should not create another destination for your team. It reads and writes at the record level inside approved systems, using role-based access, explicit approval gates, and controlled write-back.


Common system categories

PRACTICE & CLIENT MANAGEMENT	ACCOUNTING, TAX & AUDIT	DOCUMENTS & SIGNATURES
Practice management engagements, deadlines, staffing	General ledger client books, bank feeds, entries	Document management binders, workpapers, archives
Client portal PBC requests, uploads, messages	Tax software returns, organizers, e-file status	E-signature engagement and representation letters
CRM clients, entities, services	Audit software workpapers, tickmarks, sign-offs	File storage client folders, shared drives
Workflow tracker close checklists, task status	Payroll registers, filings, reconciliations	

Categories only. The specific systems, access, and write-back path are confirmed during the AI Workforce Review.

Governance & risk controls

Production AI requires explicit boundaries: what the system may read, what it may draft, who must approve, where the result is written, and how every action is recorded.





Establish **value** before you build.

The business case should begin with the workflow baseline, not a generic estimate of AI ROI.

Value shows up at the workflow level

Each workflow targets a specific operational lever: speed, quality, margin, or risk control.

DOCUMENT COLLECTION

Requests closed without staff rebuilding the list.

RECONCILIATIONS

Exceptions explained before review.

CLOSE

Blockers visible before the close meeting.

REVIEW NOTES

Repeat notes stop coming back.

CLIENT DELIVERY

Packages leave complete the first time.

The stakes are concrete, per event

In the illustrative example, brief R-118 puts three unexplained items worth \$43,930 in front of the reviewer with every source attached. The question is not whether AI can draft a workpaper. It is whether your senior reviews a finished brief or rebuilds the trail from scratch.

How the value case is built

In the AI Workforce Review and Launch & Prove, Metacto baselines the selected work using measures such as collection delay, prep and review time, repeat notes, corrections, and manager touches. That baseline supports a build-or-no-build decision before implementation budget is committed, and it stays in place after launch.

A credible production workflow has to work for each stakeholder.



Owner / CFO: Measurable economics

An explicit baseline, value model, and measurement plan before investment.



CEO / COO: Operational leverage

More engagements per reviewer, earlier visibility into blockers, and fewer late nights in busy season.



Partners / Managers: Process control

Clear ownership, review gates, and exception handling without removing people from accountable decisions.



IT / Security: Access & governance

Role-based access, source citations, audit logs, and controlled write-back into existing systems.



Proof gets built on your work, not borrowed.

There is no borrowed case study here. Every AI role Metacto builds is proven on your own work: a baseline is set first, the role runs behind your review queue, and it scales only when the numbers hold. This is how that looks for a first role.

The test is simple: the first role beats its baseline inside your review queue, or it does not scale.

FIRST ROLE	FUNCTION	MEASURED AGAINST	PRIMARY USERS
Document Desk AI Role	PBC requests and client follow-up	Collection delay and hours spent chasing	Seniors, staff, and client service teams


The starting queue

Most firms already know where the queue is: the PBC list. Requests go out, uploads come back incomplete or mislabeled, and staff rebuild the outstanding list by hand before every follow-up. That is the work the first role takes on.

The roles mapped around it

Around that first role, six AI roles are mapped across engagement setup, collection, reconciliation, close, review, and delivery. Each one is representative; yours are designed around your firm's methodology.

Representative roles. Yours are designed around your work.



What gets measured

Baseline set before the build and kept after launch, measured across three tracks	01 Speed Collection delay and close days, measured from request to complete	02 Quality Repeat review notes and corrections after delivery	03 Capacity Staff and reviewer hours recovered from chasing and re-tracing
---	---	---	--

No results are claimed here. Each track is measured on your own baseline during Launch & Prove.



What you can check before you buy.

Open by default

Metactio publishes 50 workflow patterns through Vociion, its open-source platform, so your team can read how an AI role is built before committing to one. Roles run in your own environment. If you stop the service, you keep the workforce: deployed roles stay yours on their last deployed version.

What every AI role ships with

Eleven standards, built into every role from day one:

- | | |
|---|---|
|  Named human owner |  Defined scope of work |
|  Source citations |  Approval gates |
|  Exception routing |  Role-based access |
|  Write-back to your systems |  Audit trail |
|  Baseline and success metric |  Cost and usage tracking |
|  Monthly value review | |

Oversight by design

01 Citations

Every claim and number ties back to a source your reviewers can check.

02 Review Gates

External actions are held for human approval; the system stages drafts and recommends next steps.

03 Audit Trail

Activity is actor-attributed and traceable: what AI suggested, what a human approved, and what moved.

04 Utilization & Cost

Usage is tracked by role, workflow, and surface, providing visibility into adoption, value, and operating cost.



Start with the first AI role **worth funding**.

AI in an accounting firm should begin with one repeated piece of work where reviewer time or deadlines still depend on rebuilding a list or re-tracing an entry. Pick the work with the clearest economics, build the role around your real process, and expand from measured results.

What a first build has to answer

01

Where is value leaking?
Identify the requests, reconciliations, and handoffs where staff and reviewer time is repeatedly lost.

02

Which work is most fundable?
Select the candidate with the strongest combination of business impact, repeatability, source readiness, decision ownership, and implementation path.

03

What context and controls does it require?
Define the records, systems, permissions, rules, approvals, and exception paths required for reliable operation.

04

How will value be measured?
Establish the baseline, success metric, and value model that will determine whether the build is worth funding.

Estimate the opportunity with your own numbers

Use a small set of operating inputs to create a preliminary annual estimate. Metacto validates the assumptions against your actual baseline during the AI Workforce Review.

Annual opportunity = prep and chase hours × loaded cost × recoverable share + reviewer value

PREP & CHASE • YOUR INPUTS	
Prep and chase hours per year (across the team)	—
Loaded cost per hour (\$ per staff hour)	\$ —
Recoverable share (vs. your baseline)	— %
Reviewer capacity value (recovered hours × loaded cost)	—
Illustrative gross annual opportunity (\$)	—

WORKED EXAMPLE	ILLUSTRATIVE
Prep and chase hours per year mid-size firm, full team	7,200
Loaded cost per hour \$ per staff hour	\$75
Recoverable share handled by the role	35%
Reviewer capacity value 240 hrs × \$120 loaded	\$28,800
Illustrative gross annual opportunity approximate opportunity	≈ \$218K / yr

Illustrative gross opportunity before implementation and operating costs. Validate that benefit categories do not overlap before adding them together. The same approach can be applied to close tracking, review notes, and client delivery.



Is this work ready for an AI role?

Select one piece of recurring work and rate it from 1 to 5 against each condition. A total of 21 or above, with no factor below 3, signals work ready for an AI Workforce Review.

THE FIVE CONDITIONS OF A FUNDABLE AI WORKFLOW

Business impact	materially affects reviewer time, deadlines, or client experience	___ / 5
Workflow repeatability	the same core steps recur across engagements in a repeatable sequence	___ / 5
Source readiness	the records exist in systems you can integrate	___ / 5
Decision ownership	one accountable role that signs off today	___ / 5
Implementation path	a clean destination for the approved record	___ / 5
Your total		___ / 25

21 to 25 · Ready to map

With no factor below 3, bring it to an AI Workforce Review. The numbers are likely there.

15 to 20 · Needs discovery

Worth mapping. Expect some source or ownership work first.

Below 15 · Prepare first

Line up the records and the approval owner, then score again.

Is this the right fit?

GOOD FIT

You run recurring work like PBC collection, reconciliations, close, and review

Reviewer and staff time is lost to chasing and re-tracing

Your engagement records and client files sit in systems you can integrate

A partner, manager, or senior can stay in the approval loop

You want a production system that runs on real work

NOT A FIT

You want a generic chatbot or a one-off prototype

There is no source of truth for engagement records and client files

No one owns the workflow or the outcome it affects

You are not ready to measure impact against a baseline

Put AI roles to work across your firm.

45 minutes. We look at where recurring work is costing your team the most time and map the first one or two AI roles around it. Nothing to prepare.

[Book an AI Workforce Review](#)