

EXECUTIVE BUYER'S GUIDE

Can you hold your AI accountable?

20 questions and ATOM's answers for evaluating document intelligence and workflow agents



EXPLORE
INTELLIGENCE CREATION

The standard we hold ourselves to



Carol Valentic

FOUNDER AND CEO

In claims and care coordination, the hardest problem is often not finding more information. It is knowing what can be trusted, where it came from, and when human judgment must remain in control.

AI can help work move faster, but speed alone is not the standard. This guide is intended to help leaders evaluate whether the systems entering their workflows can be held accountable.



Avinash R. Boodoosingh

PRESIDENT AND COO

An impressive demonstration is not an operating model. The real test begins when the system meets existing workflows, exceptions, governance requirements, security controls, and production volume.

Buyers need clear answers about ownership, change control, integration, and cost before they commit. The questions in this guide are designed to make those answers visible and comparable.

What should support the answer?

A credible answer should be supported by evidence appropriate to the risk. Depending on the question, that may be:

01 A live demonstration

02 A measured pilot using agreed documents and definitions

03 Architecture, security, or governance documentation

04 A contractual commitment

Not every answer needs to be proven during the first meeting. Every material claim should be capable of verification before approval.

SCORECARD KEY



Demonstrated

Shown in a live workflow, a measured pilot, documentation, or the contract.



Claimed

Stated by the vendor and not yet supported by evidence.



Not available

The vendor cannot provide it.

01

Does the system return structured data that maps to our schema and flows directly into our existing systems, rather than only producing a summary document?

WHY IT MATTERS

A summary helps a person read. Structured, usable data lets the organization route, compare, validate, and use information in downstream workflows. Output that only exists as a PDF or inside the vendor's interface preserves manual handling and deepens vendor dependency.



ATOM'S ANSWER

Record Ranger OS returns structured, usable data mapped to the agreed schema. Delivery can use REST API, SFTP, Azure Blob Storage, AWS S3, or another agreed method. Existing claims, content, and workflow systems can remain in place.

02

Can every extracted value be traced to the exact page and location, with the value displayed beside the source for review?

WHY IT MATTERS

A reviewer should be able to resolve a disputed value without searching the entire record. Source-level traceability shortens review, supports defensibility, and makes the output easier to trust.



ATOM'S ANSWER

Record Ranger OS preserves source references and presents source evidence beside the extracted output for review. ATOM can demonstrate the experience against the document types included in the evaluation.

COMPARE EVERY ANSWER. ASK EVERY VENDOR FOR EVIDENCE.

03

Does every field include a confidence score, with configurable thresholds that hold low-confidence values for qualified human review before release?

WHY IT MATTERS

An overall confidence score can hide weak fields inside an otherwise strong document. Field-level confidence makes uncertainty visible. Configurable thresholds let the organization apply more control where a field carries greater operational, clinical, financial, or compliance consequence.



ATOM'S ANSWER

Record Ranger OS assigns confidence to extracted fields. ATOM configures thresholds around the client's workflow and risk requirements. Values below the agreed threshold go to qualified subject matter experts for review before delivery.

04

Will the vendor measure accuracy using our actual documents during a pilot and put the agreed results in the contract?

WHY IT MATTERS

Published accuracy figures often use a different document mix, field definition, or testing method. A measured pilot creates a relevant baseline. A contractual definition keeps the production standard aligned with what the buyer evaluated.



ATOM'S ANSWER

ATOM measures accuracy against the agreed document and workflow scope and guarantees **98.5% accuracy, with performance backed by an SLA**. Pilot measurement and contractual language use an agreed definition so the evaluation and production standard remain comparable.

COMPARE EVERY ANSWER. ASK EVERY VENDOR FOR EVIDENCE.

05

Does the evidence record link each delivered value to its source, confidence, routing and reviewer actions, with prior values still visible after correction?

WHY IT MATTERS

An evidence record should explain how a value moved from extraction to delivery. Without the source, confidence, routing, reviewer activity, and correction history, the organization can see the final answer but cannot reconstruct the process that produced it. The evidence record shows what happened. It does not independently prove that an extracted value was correct; it preserves the evidence needed to examine it.



ATOM'S ANSWER

Record Ranger OS records the evidence as the work happens, including source references, confidence, routing, reviewer actions and delivery. Corrections create new recorded events. **Nothing is overwritten.** The prior value remains visible alongside the corrected value, reviewer and timestamps.

06

Are the models hosted within the vendor's sealed processing environment, without routing client documents through third-party AI services?

WHY IT MATTERS

"Hosted securely" does not explain where documents travel. If a vendor sends documents to an external model service, the buyer inherits another dependency, another data-processing relationship, and another set of retention and change-control questions.



ATOM'S ANSWER

ATOM runs open-weight vision-language models inside a **sealed processing environment**. Client documents do not route through third-party AI services.

COMPARE EVERY ANSWER. ASK EVERY VENDOR FOR EVIDENCE.

07

Can the models evolve without workflow redesign, third-party model dependency, or unpredictable per-token AI charges?

WHY IT MATTERS

The underlying model will change over the life of the engagement. The workflow should not require redesign each time the model layer evolves. The buyer also needs to understand whether model use introduces variable charges that become difficult to forecast at production volume.



ATOM'S ANSWER

ATOM can evolve the model layer without requiring the client to redesign its workflow. Processing does not depend on a third-party model API. Pricing can be structured per page or per document for the engagement, and **no per-token AI charges are passed through**.

08

Are client documents processed and returned rather than retained, excluded from models serving other customers, and is the evidence retained afterward free of readable PHI or field values?

WHY IT MATTERS

Retention, training use, and the evidence record are separate questions. A vendor may return a document while keeping content in prompts, embeddings, logs, or shared training data. The buyer needs one answer for the client documents themselves and a separate answer for the evidence record that remains afterward.



ATOM'S ANSWER

Client documents are processed and returned, not retained, and are excluded from model training. What remains for seven years is a **non-readable evidence record** of what happened. It contains no documents, page images or readable field values.

09

Who performs human review, where are reviewers located, how is access controlled, and are they trained on our specific document types?

WHY IT MATTERS

Human review only strengthens quality when the reviewers have relevant subject matter knowledge and operate under defined access controls. Reviewer location can also affect contractual, privacy, and data-residency requirements.



ATOM'S ANSWER

ATOM routes low-confidence values to qualified subject matter experts before release. Reviewers train for the agreed document scope, receive role-based access to the items required for review, and operate within defined quality and escalation procedures. Reviewer locations are disclosed during security review and reflected in the agreement.

10

Does the contract state who owns our structured data, schemas, configurations, evidence history and improvements from our corrections, and what we can export if we leave?

WHY IT MATTERS

Portability determines whether the value created during the relationship stays with the organization. A data export alone may be insufficient if the schema, configurations, evidence history, and validated correction history remain trapped in the vendor's environment.



ATOM'S ANSWER

The client owns its data and its **Intelligence Layer**. Export rights cover client-owned structured data, configured schemas, workflow configurations, evidence history, and validated correction history. ATOM retains ownership of its proprietary models, domain training, and reusable intellectual property.

PART 2

Before an agent enters the workflow

An agent should inherit the same evidence, controls, and accountability expected of the people it supports.

The next 10 questions test whether an agent works from validated information, connects recommendations to evidence, respects human authority, manages uncertainty, and remains governable as models and rules change.

01

**VALIDATED
INPUTS**

02

**→ TRACEABLE
RECOMMENDATIONS**

03

**→ HUMAN
CONTROL**

04

**→ AUDITABLE
OUTCOMES**

01

Does the agent use structured, validated data before analyzing a claim or supporting workflow activity?

WHY IT MATTERS

An agent cannot repair unreliable inputs through confident language. If it works directly from unstructured, duplicated, or weakly validated documents, the workflow inherits uncertainty before analysis begins.



ATOM'S ANSWER

ATOM Agents work from structured, validated data created by Record Ranger OS and held within the client's **Intelligence Layer**, together with other agreed structured sources.

02

Can every recommendation be traced to its source evidence and the applicable policy, guideline, or rule?

WHY IT MATTERS

A recommendation becomes useful in a regulated workflow when the responsible professional can inspect the evidence and understand which rule or guideline shaped the result.



ATOM'S ANSWER

Within the configured workflow, ATOM Agents connect findings to source evidence and the applicable rule, policy, guideline, or protocol. ATOM Agents support decisions. **The responsible professional remains in control.**

COMPARE EVERY ANSWER. ASK EVERY VENDOR FOR EVIDENCE.

03

Can permissions, approval requirements, thresholds, and escalation paths be set differently for low-risk and high-risk decisions?

WHY IT MATTERS

One control model rarely fits every workflow. A routine administrative step and a high-consequence clinical or claims recommendation require different permissions, review thresholds, and escalation paths.



ATOM'S ANSWER

ATOM configures permissions, approval requirements, confidence thresholds, routing rules, and escalation paths around the workflow, risk level, and client requirements. **Human authority remains explicit.**

04

Can the agent decline to answer when evidence is insufficient and route the case to the appropriate reviewer with the evidence, rationale, and confidence assembled?

WHY IT MATTERS

An accountable agent needs a controlled way to stop. When evidence is incomplete, conflicting, or below threshold, escalation protects the workflow from a confident guess and reduces the reviewer's work to reconstruct context.



ATOM'S ANSWER

ATOM Agents can stop when evidence does not meet the configured requirement and route the exception to the appropriate reviewer with the relevant evidence, rationale, confidence, and reason for escalation assembled.

COMPARE EVERY ANSWER. ASK EVERY VENDOR FOR EVIDENCE.

05

Does the audit record capture the model or rule version, input references, recommendation, reviewer decision, and final action?

WHY IT MATTERS

The organization needs to reconstruct what the agent saw, which logic it applied, what it recommended, how a reviewer responded, and what happened next. A final status without that history cannot support meaningful governance.



ATOM'S ANSWER

ATOM's audit architecture records source references, confidence, model or rule version, routing, recommendation, human review, and final workflow action without retaining client document content in the audit record.

06

Are jurisdictional, clinical, and organizational rules versioned and tested, and does the contract identify who is responsible for keeping them current?

WHY IT MATTERS

Rules change. A correct recommendation under an old rule can become a governed failure. Buyers need to know how rules enter the system, who approves them, when they take effect, how the vendor tests them, and who owns ongoing maintenance.



ATOM'S ANSWER

ATOM separates jurisdictional, clinical, and organizational logic from the underlying model so rules can be versioned, tested, approved, updated, and retired. The implementation and agreement assign responsibility for source validation, maintenance, approval, and change notification.

07

When an underlying model is updated or retired, does the vendor run regression testing against our workflows before introducing the change and notify us beforehand?

WHY IT MATTERS

A model change can alter output even when the workflow configuration stays the same. Buyers need protection from untested changes and from a third-party provider's release schedule becoming their production schedule.



ATOM'S ANSWER

ATOM's self-hosted architecture avoids forced third-party model updates. ATOM validates model changes against approved test sets before they reach a client workflow and manages release notification and rollback under agreed change controls.

08

Can multiple agents work in sequence with monitored handoffs, conditional routing, exception management, and failure handling?

WHY IT MATTERS

Complex workflows depend on reliable handoffs. An orchestration claim should explain what happens when one agent lacks data, returns an exception, exceeds a threshold, or fails before the next agent begins.



ATOM'S ANSWER

ATOM's architecture supports ordered agents, monitored handoffs, conditional routing, exception management, error handling, and escalation. The production sequence and controls are defined for the client's workflow during implementation.

09

Does the contract confirm that our data, configurations, corrections, audit history, and workflow outcomes remain ours and can be exported?

WHY IT MATTERS

The organization's people, workflows, and validated corrections create value over time. The contract should prevent that value from becoming a vendor-controlled asset or disappearing when the relationship ends.



ATOM'S ANSWER

The client owns its data and its **Intelligence Layer**. ATOM defines export rights for client-owned workflow configurations, validated correction records, audit history, and outcome history while preserving ATOM's proprietary models, domain training, and reusable intellectual property.

10

Is our data excluded from models serving other customers, and are model usage, orchestration, integration, and exception-handling costs predictable at volume?

WHY IT MATTERS

Data use and pricing both shape long-term control. Cross-customer training can weaken ownership and privacy commitments. Unbounded model, token, orchestration, or exception charges can turn a successful pilot into an unpredictable production expense.



ATOM'S ANSWER

ATOM excludes client documents from shared model training. Its sealed, self-hosted processing model avoids third-party per-token billing. ATOM scopes pricing around defined implementation, integration, volume, orchestration, and exception-handling requirements so the client can evaluate the complete operating cost.

Four standards should survive procurement

A procurement process can generate hundreds of requirements and still miss the operating standard that matters. The 20 answers should leave the executive team with 4 clear conclusions.

01

Precision you can define

Accuracy uses an agreed method, actual document types, field-level confidence, and qualified review before release. ATOM guarantees **98.5% accuracy, with performance backed by an SLA**, within the agreed scope.

02

Proof you can inspect

Every important value, recommendation, exception, and change leaves evidence. Source references, confidence, routing, versions, review, and final workflow status remain visible.

03

Predictability you can plan around

The organization understands where documents travel, how model and rule changes reach production, what happens when evidence is insufficient, and how cost behaves at volume.

04

Progress your organization keeps

The client owns its **Intelligence Layer**. Validated correction patterns improve what comes next through quality assurance, threshold tuning, exception review, and human-in-the-loop feedback. Client documents are processed and returned, not retained.

COMPARE EVERY ANSWER. ASK EVERY VENDOR FOR EVIDENCE.

Record RangerOS[✓]

DOCUMENT INTELLIGENCE YOU CAN VERIFY

The buying decision should stand up in production.

Before approving a vendor, make each answer visible in a live workflow, measured in a pilot, documented in the architecture and security materials, or written into the contract.

The companion scorecard helps the team compare vendors. This guide explains what each check should mean.

ATOM included its answers because comparison should begin with what a vendor will put on the record. We can apply the same questions to the workflow, document mix, and operating requirements that matter most to your organization.

Compare every answer. Ask every vendor for evidence.



EXPLORE INTELLIGENCE CREATION

See how accountable document intelligence works in practice.