CONFIDENTIAL · CLIENT PROPOSAL

# Booksquire

Outsourced AI Accountants · It prepares. You post.

## Architecture, Trust & Commercial Proposal

How a governed AI accountant works inside your NetSuite, email, Slack and document stores — without violating your AI policy.

| | |
|---|---|
| PREPARED FOR | Prospective NetSuite finance team |
| DOCUMENT TYPE | Architecture & Trust Pack · Version 1.0 |
| DATE | September 2026 |
| COMMERCIAL SEAT | USD 2,000 per month · 1 entity · human-approved posts |
| AUDIENCE | CFO / Controller · CISO / DPO · Legal · IT |
| CONTACT | booksquire.co · info@booksquire.co |

> THE SENTENCE THIS DOCUMENT EXISTS TO DEFEND
>
> The agent prepares. A named human accountant approves. NetSuite posts only after that approval. Consumer chat products are never the system of record. Customer data does not train any model.

This pack is written so an AI-policy committee can say yes. Work and price follow the controls — not the other way around.

## 1. Why this exists

Most mid-market NetSuite teams already know what good looks like. Journals should be supported. Invoices should be coded before they hit the AP queue. Bank recs should start from exceptions, not a blank spreadsheet. Master data should be reviewed before it becomes a fraud or audit finding. Insights should arrive before the Controller has to ask.

The work still lives in inboxes, Slack threads, SharePoint folders and heroic accountants. Hiring a competent NetSuite accountant costs six to ten thousand dollars a month fully loaded. Fractional specialists cost three to eight. Software close tools still leave the labour and the AI-policy risk with you.

The reason most finance teams have not put an AI accountant on that work is not capability. It is policy. Consumer chat tools train on inputs, leave no audit trail, and can post into an ERP if someone connects them carelessly. That is a rational reason to say no.

Booksquire is built for that no. It is not a chatbot pasted onto NetSuite. It is a controlled accounting operator: enterprise model APIs behind a policy gateway, official NetSuite Model Context Protocol (MCP) tools under a least-privilege role, connectors that use delegated OAuth, and a named human who is accountable for every posting.

### What it is, and what it is not

| IT IS | IT IS NOT |
|---|---|
| An AI-augmented accountant seat on your NetSuite | A consumer Grok / Claude / ChatGPT window pointed at the GL |

| IT IS | IT IS NOT |
|---|---|
| A preparer that drafts journals, bills, rec packs and insights | An autonomous closer that posts without a human |
| A service with a named reviewer included in the fee | Software you still have to operate and govern alone |
| Aligned with Oracle's own bring-your-own-AI Connector | A shadow integration that bypasses NetSuite roles |

## 2. Eight commitments Legal can hold us to

Give this section to Legal and the CISO first. If any commitment cannot be evidenced in the sandbox, write-access does not switch on.

1. Your data does not train models. Customer Content is excluded from training, fine-tuning and model improvement at xAI, Anthropic and OpenAI under their commercial API terms. That obligation is flowed into our DPA.
2. Production traffic uses enterprise APIs, not consumer chat. Nothing goes through grok.com, Claude.ai, Claude for Work UI, or ChatGPT consumer accounts. Those products are out of scope.
3. Zero Data Retention at the model layer. Eligible Grok, Claude and OpenAI API endpoints run with ZDR so prompts and completions are not kept by the vendor after inference. The system of record for evidence is your audit store.
4. The agent prepares. A named human approves. NetSuite posts only after that approval. Journals, invoices, vendor bills and master-data changes are drafted by the agent. A named human clicks Approve, and NetSuite posts only that approved entry. The same identity cannot draft and approve the same transaction.
5. NetSuite permissions stay yours. The connector uses a custom MCP role. Oracle blocks Administrator and full-access roles from the AI Connector. Tools cannot run elevated scripts or call the public internet.
6. Least privilege on every other system. Email, Slack, SharePoint and Drive use user-delegated OAuth with mailbox, site and channel allow-lists. Tenant-wide application permissions are not the default.
7. Every material action is reconstructable. Source document IDs, model and version, tool calls, proposed payload, approver, timestamp and NetSuite internal ID are written to an immutable log you can export to Internal Audit or your external auditor.
8. You choose the models and you can stop them. Grok only, Claude only, both, or OpenAI as a third option. A kill-switch disables a vendor or all write actions without waiting on us.

> **POPIA, GDPR AND SOX IN ONE LINE**
> Human approval on postings is not a product preference. POPIA section 71 restricts solely-automated decisions with legal or substantial effect. SOX / ICFR treat this as a semi-automated control (AI prepares, human disposes) — a pattern auditors already know how to test. Cross-border transfers to US model vendors are covered by DPA + SCCs + data minimisation + ZDR, with residency pinned where the vendor offers it.

## 3. Architecture

The product is the control plane. Models never speak to NetSuite, email or Slack directly. Every request is classified, minimised, routed, logged, and — if it would change a record — queued for a human.

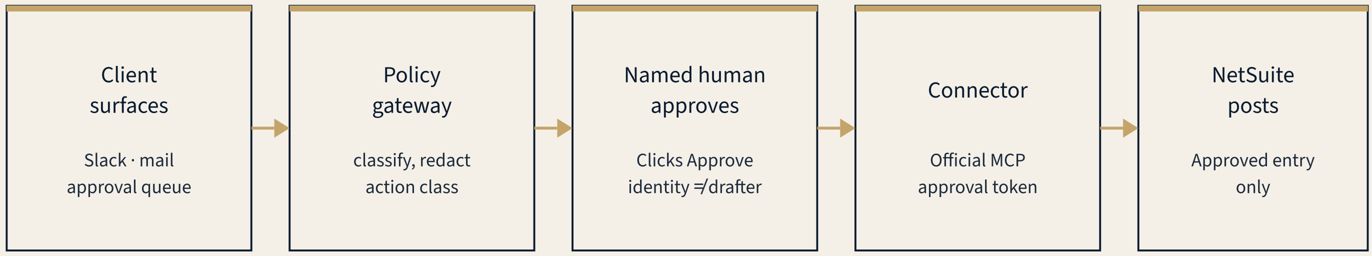


**FIGURE 1.** Writes require a named human approver. Consumer chat UIs are excluded. Customer data is not used for training.

## 3.1 Layer by layer

**Client surfaces.** Users work where they already work: Slack, Outlook / Gmail, and a web approval queue. The agent can be invoked by a message, by an inbound invoice, or on a schedule (nightly master-data scan, weekly controller pack). Users authenticate with the client's SSO (Microsoft Entra or Google) and MFA.

**Policy gateway — the product.**

- Tenant isolation: one customer, one logical tenant, no commingled stores.
- Data classification and PII / bank-detail / national-ID redaction before any model call.
- Purpose limitation: only accounting tasks on the approved list.
- Model allow-list and region pin per client. Fail-closed if residency cannot be verified.
- Action classes L0–L3 (observe, draft, post, prohibited).
- Immutable audit log, exportable to the client SIEM.
- Human approval queue with dual-control and approver routing.

**Model router — API only, ZDR.** Grok, Claude and optionally OpenAI sit behind the gateway as interchangeable inference engines. The service is ours. The models are approved tools. Task routing is conservative: a cheap model classifies and extracts; a mid model analyses and drafts; a flagship model is reserved for material journals and master-data changes.

**MCP fabric — delegated OAuth.** NetSuite is reached through Oracle's official AI Connector Service (Model Context Protocol). Microsoft 365 (Outlook, SharePoint, OneDrive) and Google Workspace (Gmail, Drive) are reached through their APIs with user-delegated OAuth. Slack is limited to allow-listed channels. Scopes are the minimum required for the subscribed workflows.

**Systems of record.** NetSuite remains the book of record. SharePoint and Google Drive remain the document stores. Bank feeds remain the bank's. Booksquire does not become a second ledger.

## 3.2 Action classes

This is the answer to "will the bot post journals by itself?" No. It prepares. A named human approves each entry, and NetSuite posts only that approved entry. There is no threshold below which the agent posts on its own.

| CLASS | EXAMPLES | MODEL MAY | HUMAN REQUIRED |
|---|---|---|---|
| **L0 Observe** | Aging, trial balance Q&A, saved searches, controller pack | Read and summarise, with citations | No — review at leisure |
| **L1 Draft** | JE pack, invoice / bill draft, rec worksheet, vendor exception | Create a draft in the pending queue only | Yes — before anything posts |
| **L2 Post** | Approved JE, invoice, vendor update | Nothing on its own. The tool runs only after Approve. | Named human clicks Approve; identity ≠ preparer |

| CLASS | EXAMPLES | MODEL MAY | HUMAN REQUIRED |
|---|---|---|---|
| L3 Prohibited | Payments, bank-file release, tax filing, period close, delete master data | Blocked at the gateway | Not applicable |

### 3.3 Why this survives an AI-policy review

| TYPICAL OBJECTION | HOW THE ARCHITECTURE ANSWERS IT |
|---|---|
| **The model will train on our books** | Commercial API no-training clauses + ZDR + our DPA. Consumer products are banned. |
| **Staff will paste the P&L into a public chatbot** | Approved path only. Connectors are governed. Raw mailboxes are not dumped into a model. |
| **No audit trail of what the model saw or did** | Every prompt hash, tool call, draft, approval and NetSuite ID is stored in your log. |
| **It will post into the ERP by itself** | L1/L2 split. Write tools run only after a named human's Approve, carry that approval token, and never run on the model's own initiative. NetSuite posts only that approved entry. |
| **Over-privileged service account** | Custom MCP role. Admin blocked by Oracle. Recommended role is narrower than the user's day job. |
| **Cross-border processing without a basis** | DPA, SCCs, POPIA s72 assessment, minimisation, ZDR, residency pin where offered. |
| **Vendor lock-in / model-risk** | Client picks the allow-list. Router fail-closes. Model upgrades that touch posting require re-validation. |
| **Shadow-IT connectors that bypass DLP** | SSO, allow-lists, delegated OAuth, kill-switch, quarterly access review. |

## 4. NetSuite: what is official, what we build

Oracle shipped the NetSuite AI Connector Service as a standards-based MCP server. That is production infrastructure, not a prototype. We use it as the only write path into the ERP.

### 4.1 What Oracle already provides

- **AI Connector Service.** Account-scoped MCP endpoint under SuiteTalk, OAuth 2.0, user-delegated.
- **MCP Standard Tools SuiteApp.** ns_createRecord / ns_getRecord / ns_updateRecord, ns_runCustomSuiteQL (read-only), ns_runReport, ns_runSavedSearch, plus metadata and filter helpers.
- **Role-based access.** The model sees and does only what the connected role can see and do. Administrator cannot use the connector. MCP permission is off by default.
- **MCP-ready roles.** CFO, Controller, Accounts Receivable Analyst, Accounts Payable Analyst, Treasury Analyst — used as templates, then stripped further.
- **Companion prompts and Skills.** Finance-aligned prompt templates and NetSuite-specific agent instructions so the model talks chart-of-accounts, not generic ERP.
- **Full MCP usage logging inside NetSuite.** Every tool call is already an ERP event, in addition to our control-plane log.

### 4.2 What we build on top — the Accountant SuiteApp

Standard tools are a USB port. The accountant product is a SuiteApp plus the control plane. Write tools require an approval_id issued by us; without it they refuse.

| CUSTOM TOOL | PURPOSE |
|---|---|
| **accountant.proposeJournalEntry** | Returns a balanced draft, source evidence, period-open check, debit=credit proof, impact preview. |
| **accountant.postApprovedJournal** | Runs only after a named human's Approve, carries that approval token, and never runs on the model's own initiative. NetSuite then posts the approved journal entry through the connector, under the client role, in an open period. |
| **accountant.proposeInvoice / proposeVendorBill** | Draft customer invoice or vendor bill after master-data match and duplicate-bill check. |
| **accountant.postApprovedInvoice / Bill** | Runs only after a named human's Approve, carries that approval token, and never runs on the model's own initiative. NetSuite then posts the approved invoice or bill. |
| **accountant.findDuplicateMasterData** | Vendors, customers, items: duplicates, missing fields, inactive-on-open-txn, bank-detail changes. |
| **accountant.proposeMasterDataFix** | Draft change only. Bank-detail changes always dual-control. |
| **accountant.listUnmatchedBankLines** | Unmatched bank lines vs unmatched NetSuite transactions, with proposed matches and evidence. |
| **accountant.getControllerPack** | Runs approved reports and searches. Returns cited figures only — no reconstructed statutory P&L from raw SQL. |

> **AN HONEST LIMITATION THAT BUILDS TRUST**
> SuiteQL is excellent for operational lists (overdue invoices, material JEs, dormant vendors). It is the wrong tool for statutory financial statements. Eliminations, consolidation, accounting book and period rules live in NetSuite reports. Official numbers always go through ns_runReport. We will not ship a "P&L" that is technically valid SQL and financially wrong.

### 4.3 REST and period constraints we disclose up front

- Tax-related fields are not supported on journalentry via REST. GST/VAT coding that REST cannot carry is handled in a custom tool or left for the reviewer.
- GL Impact and System Notes sublists are read-only / unavailable via REST — we do not pretend otherwise.
- Closed or locked periods still block posting. The AI role never receives period-override permission.
- MCP writes skip client-side form scripts. Server-side workflows and our custom tools carry the validations that matter.
- SuiteTalk governance units still apply. High-volume jobs are scheduled, not fired ad hoc.

## 5. What the accountant actually does

| WORK | AGENT PREPARES | HUMAN DOES | MECHANISM | CLASS |
|---|---|---|---|---|
| **Journals** | Balanced draft from source docs or prior pattern. Period, subsidiary, debit=credit checked. | Approve, edit or reject. | proposeJournalEntry, then postApprovedJournal only after Approve | L1 / L2 |
| **Invoices & bills** | Ingest PDF/email. Extract header and lines. Match vendor/item. Duplicate check. Code to COA. | Confirm vendor, tax, coding. Click Approve. | createRecord invoice / vendorbill only after Approve. NetSuite posts. | L1 / L2 |
| **Reconciliations** | Unmatched lines vs NS. Proposed matches with evidence. Exception list. Draft clearing JEs. | Accept matches. Finish the rec in NetSuite. | SuiteQL + saved search + proposed JE. Not a native "complete rec" click. | L1 |

| WORK | AGENT PREPARES | HUMAN DOES | MECHANISM | CLASS |
|---|---|---|---|---|
| **Master data** | Nightly scan: duplicates, missing fields, bank-detail changes, inactive on open txns. | Approve any change. Bank details always dual-control. | SuiteQL + getRecord + approved updateRecord | L1 |
| **Insights** | Weekly controller brief and ad-hoc Q&A. Every figure cited to a report, search or record ID. | Use the pack. Ask follow-ups. | ns_runReport for official numbers. SuiteQL for lists. | L0 |
| **Inbox / Slack** | Watches allow-listed mailboxes and channels. Classifies. Opens a ticket with a source link. | Only sees items that need a decision. | Graph / Gmail / Slack → gateway. Raw HTML never enters a write loop. | L0 / L1 |

### 5.1 Three workflows a Controller will recognise

**Invoice arrives at finance@.** The message lands in the allow-listed mailbox. The gateway strips signature banners and extracts the PDF. A classifier labels it “vendor bill.” Extraction returns vendor identifiers, dates, lines and tax. SuiteQL matches the vendor and flags a duplicate bill number if one exists. A draft vendor bill sits in the approval queue with the source PDF and proposed GL coding. Slack sends the named accountant a card: Approve / Edit / Reject. A named human clicks Approve. NetSuite posts the bill under the AP-draft role. The audit record stores the source hash, model version, tool calls, approver and NetSuite internal ID.

**Controller in Slack: “Draft the September rent journal.”** The agent reads the allow-listed lease folder in SharePoint and last month’s journal via ns_getRecord. It proposes a balanced journal with the same accounts, the current open period and the rent-schedule amount. It shows the period-open check and debit=credit proof. A named human clicks Approve. postApprovedJournal runs only after that Approve, carries the approval token, and never runs on the model’s own initiative. NetSuite posts the approved journal and returns an internal ID. The agent replies in-thread with the ID and a link.

**Month-end bank rec exception.** The agent pulls unmatched bank lines and unmatched NetSuite payments and deposits. It proposes matches (date within two days, amount exact, reference fuzzy). Leftovers become an exception list — for example a Wise fee with no NetSuite entry, historically coded to bank fees. It drafts that journal. The human approves the draft and finishes the reconciliation in NetSuite. The exception pack is filed in the month-end folder. We do not claim the bot “closed the bank rec.” We claim the rec started from exceptions instead of a blank sheet.

### 5.2 Hard exclusions

These are blocked at the gateway, not left as a setting someone might forget.

- Payment initiation or bank-file release
- Period close
- Tax opinion or statutory audit sign-off
- Payroll processing
- Deleting master data
- Administrator-role access to NetSuite
- Any solely-automated decision with legal or substantial effect on a person (credit-like vendor onboarding included)

## 6. Model governance

Multi-model is a policy feature, not a marketing feature. The client’s AI policy names the approved vendors. The router obeys that list.

| CONTROL | XAI GROK API | ANTHROPIC CLAUDE API | OPENAI ENTERPRISE API |
|---|---|---|---|
| **Product we use** | API only — not grok.com, not Grok Bot as system of record | Messages API only — not Claude.ai / Claude for Work UI | API / Enterprise — not ChatGPT consumer |
| **Training on customer content** | No, by default | No, commercial terms | No, by default |
| **Zero Data Retention** | Enterprise feature, team-wide toggle | Commercial org; Messages + token-count APIs | Available, approval-gated |
| **Default retention if no ZDR** | ~30 days abuse monitoring | ~30 days | ~30 days |
| **DPA / SCCs** | Available | Available | Available |
| **SOC 2** | Type II | Type II | Type II |
| **Residency** | US default; enterprise residency is sales-gated — confirm in writing | US default; EU via Bedrock / Foundry if required | US + EU option for eligible projects |
| **Fit in this service** | Fast analysis, tool use, controller packs | Long documents, policy reasoning, high-stakes drafts | Optional third vote if client policy allows |

Default routing inside an approved set: a fast Grok or Claude Haiku-class model classifies and extracts; Grok 4.6 or Claude Sonnet analyses and drafts; Claude Opus is reserved for material journals and master-data changes. OpenAI is off until the client's policy names it.

Grok Bot and Claude Cowork may later be offered as an analyst workstation. They are not the posting path. Grok Bot runs on US-hosted Cursor VMs, does not support Legacy Privacy Mode, and has Enterprise-only network and audit controls. That is a separate CISO review.

### 6.1 Threat model we design against

- **Prompt injection from invoices and email HTML.** Retrieved content is sanitised. Raw HTML never enters a write-capable agent loop. Write tools ignore model-suggested approval tokens.
- **Hallucinated journals.** Debit=credit and period-open checks run in our layer and again in NetSuite. Dual control is the product, not a setting.
- **Confused deputy / over-broad role.** Custom MCP role is narrower than the reviewer's day-job role. Separate Read / Draft / Post roles.
- **Tool drift.** Custom tools are versioned. Model upgrades that touch posting workflows require re-validation before production use.
- **Sensitive disclosure.** Minimisation: the model receives extracted fields and IDs, not the whole general ledger or the whole mailbox.

## 7. Commercial offer

### 7.1 The seat — USD 2,000 per month

Positioned against a loaded NetSuite staff accountant at roughly $6,500–$10,000 per month, and against a competent fractional NetSuite specialist at $3,000–$8,000. This is an AI-augmented accountant seat, not a Controller or CFO replacement.

**Included in the $2,000 Accountant tier**

- Always-on Accountant Agent on Slack and email.
- NetSuite MCP connection with custom least-privilege roles.
- Inbox, Drive and SharePoint intake for invoices, statements and supporting documents.

- Draft journals, invoice / bill coding and reconciliation exception queues.
- Master-data review flags (vendors, items, customers).
- Weekly insights brief and month-end exception pack.
- Named human reviewer with dual-control on all posts.
- Immutable audit log, exportable to the client.
- Quarterly access and role review.
- Client-selectable model policy (Grok / Claude / OpenAI / mix).

### 7.2 Scope that protects quality — and the economics

The seat is scoped so the human reviewer can actually review. Default envelope:

- One legal entity.
- Up to approximately 500–1,000 operational tickets per month.
- Standard write workflows: journals, AP bills, AR invoices, rec exceptions, master-data flags.
- Month-end pack targeted for business-day 5 to 7, depending on client close calendar.

| ITEM | COMMERCIAL |
|---|---|
| **Monthly — Accountant tier** | USD 2,000 · 1 entity · draft + human-approved posts |
| **Monthly — Monitor tier (optional start)** | USD 1,250 · read-only insights and exception flags, no posts |
| **Monthly — Close Partner (upsell)** | USD 3,500 · full month-end pack, flux narratives, auditor PBC support |
| **Additional legal entity** | + USD 600 per entity per month |
| **Onboarding (one-time)** | USD 3,500 · 1 entity, standard connectors, core custom tools, policy workshop |
| **Two-week sandbox** | Credited against onboarding if the client converts |
| **Customer VPC / customer-managed keys** | Scoped separately for stricter buyers |

> **HOW TO START IF AI POLICY IS STILL A DEBATE**
> Begin on Monitor ($1,250): read-only insights, no write path, same control plane. The CISO sees logs, residency, ZDR and the kill-switch on real work. Upgrade to Accountant when Legal has signed the eight commitments.

### 7.3 Compared with the alternatives

| OPTION | TYPICAL COST | WHAT YOU GET | GAP |
|---|---|---|---|
| **In-house NetSuite accountant** | $6.5k–10k / mo loaded | One person, 40 hours, leave risk | Slow to scale; no system memory |
| **Fractional NetSuite accountant** | $3k–8k / mo | Human hours only | No 24/7 prep; no governed AI |
| **Close-automation software** | Mid-market software quote | Tools the client still operates | Labour and AI-policy risk stay with you |
| **QBO / Xero bookkeeping bots** | $400–1,500 | Bookkeeping on the wrong ERP | Not NetSuite |
| **Booksquire** | $2,000 | AI prep + named human on your NetSuite | Scoped seat — not a Controller |

## 8. Implementation — the path that gets past “policy says no”

We do not jump from a slide to posting rights. Phase 0 exists specifically so the AI-policy committee has a decision packet before any financial data moves.

| PHASE | WHEN | WHAT HAPPENS | CLIENT EFFORT |
|---|---|---|---|
| **0 Trust review** | Week 0 | Architecture pack, DPA, subprocessor list, AI-policy mapping, named approvers, approver routing, ZDR confirmation. | CISO + DPO + Controller, 2–3 hours |
| **1 Read-only sandbox** | Week 1–2 | Custom MCP role = view only. One mailbox or Slack channel. Controller pack and rec exception list. No writes. | Controller + IT, ~4 hours |
| **2 Draft mode** | Week 3–4 | L1 on. Invoice and JE drafts land in the approval queue. Posts blocked. Accuracy measured. Vendor aliases tuned. | Reviewer practises the queue |
| **3 Supervised post** | Week 5–6 | L2 on: every post requires a named human Approve. Dual control on. | Sign-off on approvers and routing |
| **4 Steady state** | Month 2+ | $2,000 retainer live. Month-end pack, rec copilot, master-data scan. Quarterly role recertification. Model-change revalidation. | ~1 hour / month governance |

Sandbox exit criteria: CISO sign-off on the eight commitments, and Controller sign-off that drafts are usable. Only then does the paid seat go live.

# 9. RACI and operating model

| ACTIVITY | BOOKSQUIRE AGENT | BOOKSQUIRE REVIEWER | CLIENT CONTROLLER | CLIENT CISO / IT |
|---|---|---|---|---|
| **Classify inbound docs** | R | C | I | C (allow-lists) |
| **Draft JE / bill / rec match** | R | A | C | I |
| **Post to NetSuite** | R (runs the tool after Approve) | A (named approver) | A (named approver) | I |
| **Model allow-list / kill-switch** | R (operates) | C | C | A |
| **Role & access review** | C | R | A | A |
| **Accounting policy / COA** | C | C | A / R | I |

R = Responsible, A = Accountable, C = Consulted, I = Informed. The Booksquire reviewer is a named human on our side. Approver routing names who clicks Approve. Segregation of duties: the identity that prepared a draft cannot be the identity that approved it.

# 10. Deployment options

- **SaaS isolated tenant — default.** Fastest path. Dedicated logical tenant, encryption at rest and in transit, optional customer-managed keys as an add-on.
- **Customer VPC.** Control plane runs in the client's AWS or GCP project for buyers who will not send orchestration logs off-network. Models still infer in the vendor region the client approved.
- **Hybrid.** Audit store and approval queue in the client VPC; model calls egress through the client's existing SSE / DLP.

## 10.1 Offboarding

Keys revoked. Connectors dropped. Logs exported or deleted on instruction. NetSuite remains the client's. Nothing proprietary is trapped in a model. That is a contractual commitment, not a slogan.

# 11. Suggested next step

A two-week sandbox is the cheapest way for a CISO and a Controller to disagree in public and then agree.

1. **Week 1 — read-only.** Custom MCP role, SuiteQL + reports + saved searches only. One mailbox or one SharePoint library. A controller pack and a rec exception list. No writes.

2. **Week 2 — draft-only.** Invoice and journal drafts land in the approval queue. The client reviewer accepts or rejects. Nothing posts unless they click Approve.

Exit: CISO sign-off on the eight commitments, Controller sign-off that drafts are usable. Then the $2,000 seat, or remain on Monitor if write-access is still a debate.

> **HOW TO TALK ABOUT GROK IN THE ROOM**
> Booksquire is our service. Grok is one approved inference engine behind the control plane, alongside Claude and optionally OpenAI. We do not rent a consumer Grok bot and point it at the general ledger. That distinction is what makes the architecture acceptable to a conservative AI policy — and it is the distinction we will keep in every statement of work.

## Document control

| | |
|---|---|
| **Classification** | Confidential — prospective client |
| **Version** | 1.0 · September 2026 |
| **Rebrand** | Brand: Booksquire — Outsourced AI Accountants. booksquire.co. Architecture is independent of the name. |
| **Sources of platform fact** | Oracle NetSuite AI Connector Service and MCP Standard Tools documentation; xAI, Anthropic and OpenAI commercial API privacy terms as at September 2026. |