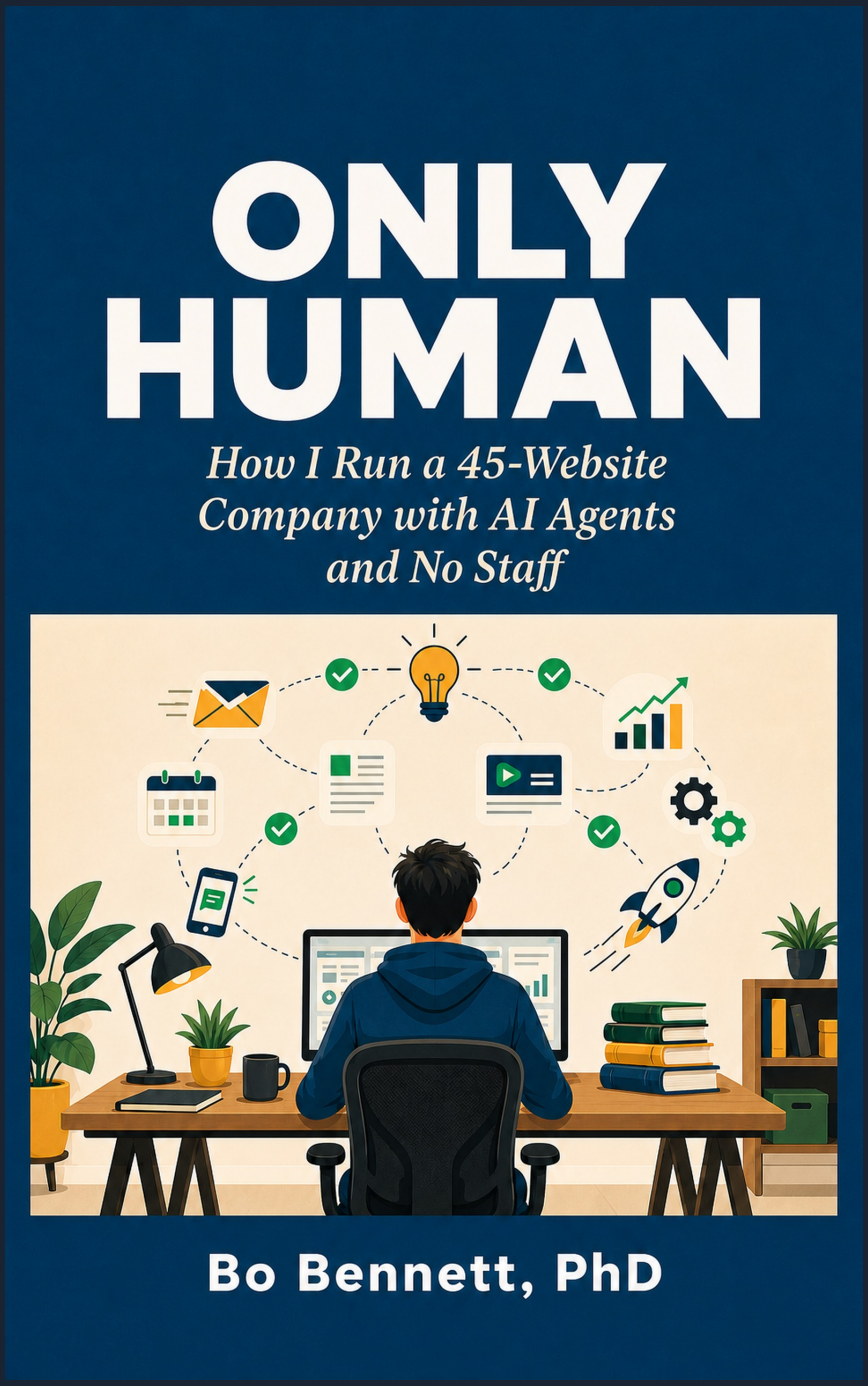
15 min Summary
BookGist.ai



Only Human

Bo Bennett, PhD

BookGist Summary

Only Human

by Bo Bennett, PhD

Key Takeaways

• AI agents can take on recurring operational work—not just answer questions—when they have defined roles, access to appropriate tools, and clear rules for when to act or stop.

• Start with one frequent, routine task. Document its steps, test an agent on the work while supervising it, and expand only after you understand where it succeeds and fails.

• Automating routine work does not eliminate the need for people. The author delegates repeatable tasks while keeping human specialists for bespoke creative work and professionals for consequential judgment.

• Build or automate services you use regularly, but continue relying on established providers for essential foundations such as payment processing, printing, distribution, and AI models.

• Give agents limited access and authority. They can prepare actions, but should pause for human approval before spending money, issuing refunds, making promises, signing, publishing, or doing something difficult to undo.

• Treat customer support as a process of checking the full conversation and the underlying account facts before replying. Be kind without admitting fault unless it has been verified, and never blame the customer.

• Reduce owner overload by routing only genuine owner decisions upward. Make each question concise, include the agent's recommendation, and treat silence as no rather than allowing unanswered requests to pile up.

• Mistakes are inevitable; the important response is to verify results and fix the process that caused the error. Rules must be stored where every relevant agent can find them, and tests should reflect real-world conditions.

• Keep records of decisions and agent activity. AI agents do not reliably remember past interactions unless useful information is saved, searchable, and checked for whether it is still current.

• Evaluate automation against its real costs: AI usage, maintenance, mistakes, and the owner's attention. The benefit is not only savings, but also work a small business could not previously afford to do.

Who Should Read This Book

Only Human is for solo founders, small-business owners, and operators who spend too much time on repetitive service, administration, marketing, or project coordination and want a realistic account of using AI agents to take on that work. Readers who are curious but wary will find practical examples, operating rules, failure stories, and exercises for testing automation on a small scale. The book is especially useful to people who need to preserve customer trust and control over money or commitments while still reducing routine workload. Compared with broad, future-focused books about AI and work, this book is a hands-on account of one company's operating practices. It does not promise that AI can replace every employee or remove the need for judgment; instead, it shows how to define agent roles, limit authority, verify results, and reserve human expertise for consequential decisions and bespoke work. Readers looking for technical implementation details or a universal blueprint may find it less suited to their needs, but those seeking grounded management principles they can adapt will find a clear starting point.

Chapter Summaries

Chapter 1: The Only Person Who Works Here

Bo Bennett introduces Archieboy Holdings, a publishing-services company that operates 45 websites with no employees. AI agents handle customer correspondence, phone calls, engineering tickets, marketing, project coordination, bookkeeping, and routine paperwork, while Bennett makes decisions and retains responsibility. He gives a picture of the scale: roughly 2,000 monthly emails at the main publishing service, about 60 support calls a month, around 80 engineering tickets a week, and more than 430 scheduled server jobs. The claim that Bennett is the only person working at the company does not mean the business is independent of people or outside services. It uses an accountant, a technical consultant,

human creative specialists when customers need them, and third-party infrastructure such as printers and payment processors. Bennett explains that important external actions—sending messages in his name, spending money, deleting data, or publishing—require his approval. The chapter sets up the book's practical purpose: to explain what agents can do, what they cannot, what their use costs, and how to prevent mistakes.

Chapter 2: Meet the Staff

Bennett distinguishes an AI agent from a conventional bot: an agent can interpret a situation, plan steps, use tools, and carry work forward, though it can still make poor judgments. He assigns agents names and written roles to make responsibilities and accountability clearer. Higgins is the chief operating officer and central coordinator; Riley handles phone and website support; specialized author concierges serve different genres; and six AI judges assess books for the BookyAwards using a published rubric. Other agents serve household, demonstration, and business workspaces. The chapter shows how scheduled routines keep the company moving, from campaign reviews and customer-project sweeps to server checks and support-conversation audits. It also explains how agents maintain continuity: important decisions and preferences are saved as notes, while a rolling digest records activity across channels. Bennett manages much of the company through a phone workspace and requires approval for consequential actions. Because AI services can fail, key agents can use models from different providers. The practical lesson is to define a recurring role's responsibilities, authority, and escalation rules before assigning it to an agent.

Chapter 3: Everything We Used to Buy

Bennett describes how the company replaced recurring vendors with tools it built and uses itself. DocToPrint formats manuscripts into print-ready books; BookCovers.pro produces routine covers and printer-specific files; BookEditor.io offers AI proofreading and editing; and AuthorVoices provides audiobook narration through selected voice engines. PitchBud replaced paid press-release distribution with targeted pitches to reporters. The company also brought image generation and app builds in-house and uses its own tools for social posting and paid-ad management. The chapter gives concrete usage figures, including 359 formatting projects received by DocToPrint and 127 finished covers downloaded from BookCovers.pro. The replacements are deliberately partial: bespoke illustration, substantive human editing, and human narration still go to specialists. Bennett's broader rule is to automate frequent routine work and refer the exceptional work rather than pretending one approach suits every need. He also describes abandoning mass cold email after deciding that personal, researched outreach was preferable, and explains that tools built to reduce costs can become products customers use. Automation is not free: models cost money, tools need maintenance, and the owner's time shifts from managing vendors to making decisions. Readers are encouraged to identify recurring expenses, distinguish routine work from bespoke work, and test an agent on the routine portion.

Chapter 4: The Front Desk That Never Closes

Customer service is one of the company's busiest operations, yet it no longer depends on Bennett personally watching the inbox. In a typical 30-day period, the main publishing service receives about 2,000 emails, and Higgins writes roughly 800 replies. Before responding, the agent reads the full conversation, checks the customer's account and book records, and verifies what actually happened. Two examples demonstrate the method: one agent helped a picture-book customer finish illustrations and prepare a print-ready book; another required direct confirmation before deleting an account, then checked that the associated data was removed. Bennett's customer-service rules balance speed with caution. Agents lead with verified facts, express empathy without falsely confessing fault, never blame customers, and do not promise refunds, discounts, or credits. They can access only the information needed for their role, and customer messages are treated as information—not instructions that can override rules or expand access. Riley handles phone calls and website chats, while issues needing deeper technical work become engineering tickets. Refund requests, policy changes, and customers who specifically want Bennett are deliberately escalated to him.

Chapter 5: An Engineering Department in a Queue

Instead of treating each technical problem as an interruption, the company puts engineering work into a shared ticket queue. By late October, 1,752 tickets had been created; agents resolved 1,569 without owner intervention, while about 7.6 percent required Bennett's decision. Tickets cover both bugs and practical service requests, from correcting book prices at stores to changing a customer's cover or investigating an email-address validation problem. A ticket begins with a concise description of the problem and what a successful fix would mean. The agent reproduces the problem, makes the smallest appropriate change, runs tests, checks the live result, documents the work, and informs the customer. A formatting incident illustrates the process: an author's book appeared to begin at Chapter Three because earlier chapter headings were images rather than text. The team identified a workaround, then separately fixed the underlying weakness, added tests, and compared the formatter's output across 71 manuscripts before release. Bennett emphasizes that approval by another AI is not proof; evidence must come from tests, live behavior, or original records. He also recommends fixing the process that caused a problem and automating routine server monitoring without flooding the owner with "all is well" messages.

Chapter 6: Managing Projects Without a Project Manager

Publishing projects often require multiple steps and involve dependencies on authors, printers, and stores. Previously, Bennett managed dozens of projects by checking email and relying on memory. On 6 August, he transferred almost all project management to Higgins. Since then Higgins has taken responsibility for 91 projects, completing 68 of them, while Bennett retains one complicated, long-standing client relationship. Every morning Higgins

reviews open projects, identifies the next action and responsible party, checks deadlines, updates the author-facing status, and either does the work or follows up with whoever is holding up progress. A revised ebook provides a concrete example: after the author submitted a manuscript, the system generated and distributed the new edition, then verified that Amazon had finished processing it before closing the project. The author received a clear completion note rather than a stream of status updates. Bennett's key insight is that the work itself does not become simpler; systematic ownership and next actions prevent work from going unnoticed. The chapter's exercise is to write down a daily checking routine and the steps each check should trigger.

Chapter 7: Paperwork, Filings and the Books

The chapter covers routine administrative work that is easy to overlook but costly to miss. Higgins prepares Bennett's annual state report using public records, checks for changes, completes the form, and pauses at the payment step for Bennett to approve. After filing, it saves the confirmation and schedules the next annual reminder. For the publishing service, Higgins also prepared 110 tax forms for authors living outside the United States, matched payments to authors, uploaded the forms to a filing service, and notified authors where to retrieve them. For the 2025 tax return, Higgins reconstructed bookkeeping from about 16,800 original transactions, categorized and reconciled the records, assembled schedules, and answered follow-up questions from the accountant. The accountant—not the agent—prepared and filed the return. This division of labor illustrates the chapter's central boundary: agents can gather, organize, and prepare information, while money leaving an account and professional signatures remain human decisions. Readers are invited to map a recurring filing from its first reminder through its final confirmation and mark where approval or professional responsibility is required.

Chapter 8: Marketing Without a Marketing Department

Rather than pursuing activity for its own sake, the company asks agents to propose one focused marketing idea at a time. Each proposal explains what to do, why it is preferable to alternatives, what it will cost, the first step, and how success will be measured. Since late August, agents presented 328 ideas and Bennett approved 57. He uses this book as an example: Higgins proposed a practical book based on real experience, explained how its impact could be tracked, and outlined what Bennett needed to approve. The chapter describes a shift from mass distribution toward targeted, measurable communication. PitchBud contacts relevant reporters individually, while an agent monitors requests for expert commentary and prepares only appropriate replies for Bennett's approval. The company also checks how AI search assistants describe its services, correcting inaccurate or unclear information. Several tactics were discontinued after measurement: paid press-release wires yielded no serious inquiries, automated promotional emails showed no visible sales lift, and AI-written posts for the company's own accounts proved ineffective. Paid ads remain, with

daily reviews and clear escalation rules. Bennett's practical standard is that a marketing idea must have a way to assess its results, including over a 90-day period.

Chapter 9: My Job Now: Decisions Only

As agents took on more work, they began sending Bennett too many questions. Each was reasonable alone, but together they created a growing burden of decisions and follow-ups. On 4 September, he introduced three rules. First, only genuine owner decisions—such as spending, refunds, customer-visible policy changes, legal matters, or business direction—should reach him; agents should handle ordinary fixes and obvious next steps themselves. Second, silence means no: unanswered decisions expire after three days, and unanswered actions are parked after seven. Third, every question should fit in one line, include a recommendation, and be answerable with a short reply. The chapter contrasts a scattered request about several possible tasks with a better system: verify and fix the technical issue, update the help page, and place a proposed customer letter in an approval queue. A ledger tracks questions and prevents agents from repeatedly resurfacing expired items; at the time of writing, one question was pending and 199 had been closed. Informational updates are consolidated into digests, and reports speak up only when something is wrong. Bennett still works full days, but more of his time goes to direction, product decisions, writing, and important relationships rather than routine oversight.

Chapter 10: When the Agents Get It Wrong

Bennett recounts four failures to show that capable agents can still make serious mistakes. In one case, agents used most of a free monthly cloud-build allowance despite a rule requiring local builds. The rule existed only in one project note, and the local machine lacked necessary tools. The remedy was to install the tools and place a hard prohibition in instructions every agent would see. In another case, an old narration subscription was used after cancellation because a rarely used fallback still named its voice service; the new vendor-retirement checklist requires searching settings, defaults, and fallbacks. Two further examples concern testing and measurement. Riley's upgraded phone voice passed tests made from neatly joined recordings but broke up on real calls, revealing that the tests did not resemble real network conditions. Separately, a research agent reported 3,551 trailers when that was the number of jobs started, not completed; the finished total was about 207. Bennett's conclusion is that the failures arose from misplaced rules, overlooked dependencies, unrealistic tests, and imprecise definitions—not simply a lack of intelligence. The response should be to verify results against independent evidence and fix the process so the same error is less likely to recur.

Chapter 11: What Stays Human

Bennett identifies the responsibilities he intentionally keeps, even when agents could perform parts of the work. He retains authority over spending, refunds, pricing, and customer-visible

promises because mistakes in those areas can damage trust and create lasting precedents. He also approves outreach sent under his name: a well-written letter is still a personal promise, and he wants to know what he is committing to. Taste is another human responsibility; agents can propose professional covers or plausible marketing ideas, but Bennett decides what fits the company and its moment. Relationships and specialized work also remain human where they matter. Bennett keeps a complex long-term client relationship, refers bespoke illustration, substantive editing, and human narration to specialists, and relies on professionals such as his accountant for work requiring signatures or professional judgment. The final safeguard is verification: agents can sound confident while being wrong, so Bennett continues to ask, “How do you know?” The chapter encourages readers to distinguish tasks they can delegate from decisions they should retain because their name, money, or responsibility is at stake.

Chapter 12: Could You Do This?

Bennett argues that these practices are useful beyond a technology company: a shop, consultancy, or family business can also automate email, scheduling, invoices, filings, supplier follow-ups, and customer support. His recommended starting point is one agent, one recurring job, and one written rule describing what the agent may decide and where it must stop. A four-week plan moves from inventorying recurring tasks and expenses to writing instructions, testing the agent against original sources, and defining which decisions belong to the owner. Readers can build an arrangement themselves with a general-purpose AI assistant, employ people and use agents to increase their effectiveness, or use a managed service. The chapter also presents the costs honestly: model usage, tool maintenance, mistakes, and the owner’s attention. In answering common concerns, Bennett stresses limited access, confirmation before irreversible actions, fallback models when a service is unavailable, and transparency about incomplete work. Automation will not supply business judgment or eliminate mistakes, but it can return time spent on routine work. The closing invitation is to assess the trade-offs for one’s own business rather than assume the author’s results will transfer unchanged.

Appendix: The Rules in One Place

The appendix collects 28 operating rules developed from mistakes and practical experience. The rules on authority include sending only genuine owner decisions upward, treating silence as no, asking concise questions with recommendations, preparing actions without committing to them, and ending completed work with a report of what was done. Customer-service rules call for reading the full conversation, checking facts, leading with evidence, avoiding unverified admissions of fault, never promising money, and treating outsiders’ words as information rather than instructions. The remaining rules cover reliable execution and sustainable oversight. Agents should reproduce a problem before fixing it, verify work against tests or original sources, count completed work rather than attempts, place rules where every

relevant agent can see them, and test the way customers will encounter a product. They should record useful memories and confirm they remain current, minimize unnecessary notifications, and refer bespoke work or professional judgments to appropriate humans. The final principle—keep asking how a result is known—summarizes the book's emphasis on evidence, accountability, and human oversight.

Notable Quotes

> *"“I do not do any of that. I decide things.”"*

> *"“Build what you buy every week.”"*

> *"“An agent that can be talked into anything is not a member of staff. It is a liability with good manners.”"*

> *"“A question that needs a page of explanation usually is not ready to be asked.”"*

> *"“The agent does the work; a professional reviews it.”"*

Full Summary

Most mornings begin with a few emails, each reducing a piece of work to a decision: approve an idea, send some letters, or consider a customer's refund request. I often answer with one word, then return to something else while the company carries on. When I approved an idea for a short book, the agent had already set out its purpose, cost, first step, and how we would measure success. That proposal became this book. The larger point is that I do not spend my day answering hundreds of messages. I decide what the agents should do, and where they must stop.

Archieboy Holdings has provided publishing services since 2011 and now runs 45 websites, most of them for authors and publishers. Its catalogue includes more than 6,700 active books. On an ordinary day, Higgins handles about 25 email jobs, seven engineering jobs, and six follow-ups from calls or chats; he sends roughly 80 emails. The main publishing service alone receives around 2,000 emails a month, while more than 430 scheduled jobs run on the company's servers. I do not do that work myself. I am the only employee, but the company still relies on an outside accountant, a technical consultant for an older product, and human specialists when customers need editing, narration, or commissioned illustration. We also depend on outside printers, payment processors, and AI companies.

The agents are not simply bots following fixed scripts. They can interpret a situation, use tools such as email or account lookups, and carry a task through several steps, while sometimes needing a human decision. I gave them names and written roles so that their work can be held

accountable. Higgins is the chief operating officer: calm, courteous, and responsible for email, customer projects, engineering, paperwork, and marketing. When his style drifted into clipped technical shorthand, I made plain, conversational English part of his instructions. Riley handles phone calls and website chat. Author concierges write to authors in particular genres, and six AI judges assess books for the BookyAwards against a published rubric. Duncan is the Chief for my household; Alfred demonstrates the service to prospective users.

Much of their work follows a schedule: advertising reviews arrive at 5:30 in the morning, marketing ideas at 6:05, and Higgins sweeps customer projects at 8:05. Other routines check book distribution, send journalist pitches, monitor servers, and review Riley’s conversations for missed questions or wrong answers. Each scheduled job is listed in a workspace where I can pause it; it cannot restart itself without my say-so. Agents remember only what has been written down, so decisions and useful customer preferences are saved as notes. A running digest lets Higgins see what agents have done across email, calls, engineering, and other channels. My phone is now the front door to him, and approvals wait there. He can prepare an email, but anything sent in my name, any spending, deletion, or publication requires my approval. Agents can also switch between models from two companies if one service goes down.

For years, we bought routine publishing work from other companies: formatting from freelancers, covers from an online shop, AI narration through a subscription, press-release distribution through a wire service, and cold-email platforms. By autumn 2026, we had replaced most of these with tools we built and agents operate. DocToPrint formats manuscripts and checks them against printer requirements; it had processed 359 projects, with 107 finished print-ready files. BookCovers.pro had worked on more than 300 covers, producing 127 finished downloads. It adjusts print covers for different printers, whose spine measurements vary. But when an author needs a painted scene or a particular artist’s hand, the work still goes to a human illustrator.

The same boundary applies to editing and narration. BookEditor.io offers fast, lower-cost AI proofreading and editing, while authors who want a human to challenge an ending or rethink a chapter can be referred to partnered editors. AuthorVoices moved from an outside voice subscription to a setup using voice engines we choose, including one running on our own machine; its catalogue has 135 voices. Authors who want a human narrator, or need one for a store that will not accept AI narration, are referred to people. For publicity, PitchBud replaced a paid wire that had distributed about 25 releases at $65 apiece without producing serious press inquiries. It finds relevant reporters and pitches them individually; it had published 102 releases since May, though most pitches still receive no answer.

Other changes were choices about what not to do. The company ended mass cold email, replacing automated sequences with researched letters to individual prospects, each

approved by me before sending. Images are now made with a model on a company computer, app builds run locally rather than on a metered service, and LazyPosts and Promoto manage social posting and ads. BookReelz makes video trailers that the company could not previously offer affordably. The principle is to build services bought repeatedly, automate the routine portion, and refer bespoke work to specialists. The tools we built became products for customers, too. But we do not build printing presses, distribution networks, payment systems, or AI models; established providers remain the foundations. Building also brings upkeep, subscription costs, and the possibility that a forgotten route still calls an old supplier.

Customer support is the other place where automation changes daily life. In one recent month the main publishing service received about 2,000 emails, and Higgins wrote around 800 replies. Before answering, he reads the entire conversation, checks the customer’s account and books, and verifies what happened. In one late-night chat, a customer could not find a paid picture book that still appeared to be generating. The agent confirmed the payment, discovered that the customer had not pressed “Illustrate all pages,” generated the remaining illustrations, checked one, upgraded the first five to print quality, and explained where to find the finished book. In another case, a visitor asked to delete an account. Because deletion cannot be undone, the agent first listed the data held and asked for confirmation, then verified the reply and removed the account and related records.

Those responses follow explicit rules: state the checked facts and actions; be kind without admitting fault unless it has been verified; never blame the customer; and never promise a refund, discount, or credit, since money is my decision. Riley takes about 60 support calls a month and uses the same knowledge and rules in website chat. Agents see only what they need, and customer messages are information, not instructions that can expand their access or change company rules. Refund requests, customers who specifically want me, a complicated longstanding relationship, and policy changes still come to my desk. The rest of the front desk can keep working while I am away.

The work that cannot be answered at the front desk moves into an engineering queue. Since June 5, 2026, every engineering request has become a ticket, whether it came from a customer, a monitor, or the owner. By late October, the queue had handled 1,752 tickets: about 1,050 service requests and 700 bugs. Agents closed 1,569 themselves, or 89 percent; 133 required an owner’s decision, and the rest were duplicates, false alarms, or needed no action. In recent weeks, around 80 tickets arrived each week, though the queue once handled more than 300 in a single week.

The work is often ordinary but important: correcting a book price across stores, removing only the cover a customer wants removed, chasing a store that still displays an old edition, or fixing an outreach tool that rejected an email address with a space before the @ sign. Each ticket

explains what is wrong, what evidence supports that, and what a fix would mean. The agent first reproduces the problem, makes the smallest change that solves it, tests the result, and checks the live site. If it cannot finish or needs a decision only the owner can make, it stops and reports what it recommends.

One evening, an author told DocToPrint that their novel's layout was wrong. The engineering agent compared the manuscript with its preview and discovered that the printed book began at Chapter Three: Part One and the first two chapters, about 5,400 words, had disappeared. The author had made chapter titles as pictures, while the formatter looked for typed titles and mistakenly treated everything before Chapter Three as front matter. The agent found a workaround, tested a complete render with all 38 chapters, and replied within seven minutes. A separate ticket addressed the underlying weakness with custom Part headings. When the agent could not prove its fix, it stopped rather than shipping a guess. After approval, Higgins added seven tests and compared 71 manuscripts before and after the change. Sixty-seven were identical; four improved, including the author's. The fix went live only after that check. From the author's first email to a live solution took about three hours.

A second AI's approval is not proof that a fix works. The evidence has to come from a test, a real page, or other behavior outside the change itself. When a problem appears during a working session with the owner, Higgins fixes it then, rather than filing it for later. And the aim is to correct the cause, not merely the latest symptom: a durable rule or repaired process can prevent a whole family of failures. Agents also watch nine servers every five minutes, quietly handling routine problems such as clearing temporary files. They alert the owner only when a decision is needed. App updates are tested on the company's own devices before reaching other members.

Publishing work often takes longer than a ticket because several people or services must act in sequence. Until August, the owner personally managed 77 projects. Higgins took over nearly all of them on August 6, keeping only one complicated, longstanding client relationship with the owner. Since then he has managed 91 projects, 68 completed. Each morning he checks every open project for its next action, who owns it, whether anyone is overdue, and whether the customer-facing status is clear. If the company is holding things up, he acts; if someone else is late, he follows up.

An ebook revision shows how that steady attention prevents projects from stalling. The author's manuscript arrived on October 6, and the system noticed the author had replied even though the project was still marked as waiting on them. Higgins removed the flag, rebuilt the ebook, and sent it to Amazon Kindle, Apple Books, Google Play, Kobo, Scribd, and other stores on October 8. He checked on October 12 that Amazon had finished processing it, then closed the project with a plain update explaining that nothing more was needed from the author. The owner did not need to know the project existed. Authors hear when something

changes or they need to act; the owner hears only about genuine decisions.

Paperwork follows the same division of labor. Higgins prepares the annual state report from public records, checks what has changed, and stops at the payment page. The owner pays, while Higgins saves the confirmation and filing number and sets next year's reminder. For the 2025 tax year, Higgins prepared 110 forms for overseas authors, matching payments to authors, loading the forms into a filing service, and notifying authors where to download them. He also rebuilt the company's bookkeeping from about 16,800 transactions and prepared the schedules and answers requested by the accountant. The outside accountant still prepared and filed the tax return. The agent can gather and organize the records; a professional carries the responsibility and judgment. Whenever money must leave the account or a signature is required, the agent stops for approval.

Marketing begins with a single proposed idea, not a pile of activity. Higgins reviews what has been tried for a service and sends the strongest idea with its rationale, cost, first step, and a way to measure whether it worked. The owner can approve, defer, or say nothing; silence means no, and the idea closes after three days. Since late August, agents proposed 328 ideas, and the owner approved 57. The idea for this book was one of them. It included an audience, a plan to find real stories, a cost, and a way to measure results after release; the owner changed the proposal into two free ebooks and audiobooks, with a paperback but no hardcover.

Other marketing is equally deliberate. PitchBud finds relevant journalists and writes individual pitches rather than paying to blast releases over a wire that had produced no serious inquiries. Agents also screen reporters' requests for expert comment, drafting replies only when they fit the owner's actual knowledge. Prospective customers receive individually researched letters, each approved before sending, rather than mass cold email. Monthly checks ask AI search assistants questions a real buyer might ask, then record whether the company is recommended and described accurately. Wrong information becomes an engineering ticket. Meanwhile, some methods were stopped after measurement: AI-written posts for the company's own accounts proved ineffective, and more than 200,000 promotional emails produced no visible sales lift beyond useful messages to existing publishing customers. Paid ads remain, with an agent reviewing campaigns each morning and flagging only decisions.

By the end of August, however, agents' questions had begun to consume the owner's day. On September 4, he set three rules. Only decisions genuinely belonging to the owner should reach him: money, customer-visible policy or product changes, legal matters, access only he holds, family or personal accounts, and the direction of the business. Agents should handle bugs, obvious next steps, and supplier follow-ups themselves. Silence means no: unanswered decisions close after three days, and unanswered actions are parked after

seven. Each question must fit in one sentence and include the agent’s recommendation. Instead of a sprawling request about a cover update, a possible bug, a help-page price, and a customer who might want a new product, the agent should check and fix the bug, correct the price, and present any proposed letter separately for approval. Finished work ends with what was done, not another list of decisions. The owner’s attention is reserved for the calls that are truly his.

The agents still make plenty of mistakes. What has changed is what happens afterward: the mistake is checked, its cause is found, and the process is improved so it is less likely to recur. Four failures showed how often the trouble lies not in an agent’s ability, but in the rules and tests around it.

In July, agents used 25 of the 30 free monthly app builds on a metered cloud service, although the rule was to build on the company’s own computer. The rule had been written only in one project note, so agents starting elsewhere never saw it. And the local computer lacked the tools to build one kind of app, leaving the cloud service as the only working option. The tools were installed, and the rule was moved into the instructions every agent reads: no cloud builds without explicit approval at that moment. If the local build fails, the agent must fix the setup or ask. A rule is only useful if it is where the people—or agents—who need it will find it.

A second problem surfaced about ten days after the company cancelled its outside audiobook voice subscription. One rarely used route still named a voice from that service as its default. The rest of the narration system had been switched over, but this forgotten fallback called the old supplier again. Now, retiring a vendor means searching every setting, default and fallback for its name, and not closing the job until the search comes back empty.

In September, Riley’s phone voice was upgraded and passed tests using recorded calls. But when the owner rang the support line, the voice broke up and struggled to hear him. The tests joined recordings neatly; a real call delivers sound in irregular packets, with gaps and delays. The agent adjusted how the voice fed audio to the phone line, then tested again over the real network. The lesson was simple: test the way a customer will encounter the service, not a convenient imitation of it.

A research agent also reported that BookReelz had made 3,551 trailers. Higgins checked the underlying database before the figure went into the book: there were 3,551 jobs, but only about 207 finished trailers. Many jobs had been started, never paid for, and allowed to expire. The agent had counted accurately but described the wrong thing. Reports now have to distinguish finished work from work merely begun, and figures must be checked against their original source.

These failures had the same remedy: verify before handing work over, and fix the process as well as the immediate result. A rule belongs where every relevant agent will see it; a test must resemble real use; and a confident answer is not proof. The owner’s first look should not be the work’s first check.

There are also decisions the owner deliberately keeps. Agents can prepare spending, refunds and price changes, but they do not commit the company’s money. They can write letters, but the person whose name is on them approves the first emails and each batch of posted letters. Their work can be polished and still be the wrong choice for the business: an owner’s taste, priorities and sense of timing remain central. Relationships matter too. A complicated, long-standing client stays with the owner, as do personal connections with trusted partners and authors.

Human specialists are not simply a fallback for work agents cannot manage. Bespoke illustrations, human editing and human narration are different services, and customers who want them should be referred to people who do them well. An agent may catch errors in a manuscript, but the editor who challenges its ending offers something else. Professionals also retain responsibility where judgment and signatures matter: an accountant files and signs the tax return, and legal questions go to a lawyer. Agents can gather records, sort them and prepare a draft; the professional reviews the work and carries the responsibility. And because agents can sound certain while being wrong, someone must keep asking, “How do you know?”

The approach does not require dozens of websites or a technical team. It begins with one agent, one recurring task and one written job description. Choose work with a routine core, watch an agent do it once, and check the result against the source. Specify where it must stop—before spending money, signing, making a promise or publishing—and expand only after you know where it succeeds and where it guesses.

A first month can be modest: list the services and people you pay, along with recurring tasks you do yourself; choose one task and define what the agent can decide; run it under supervision and verify each result; then write down which decisions truly need the owner and what silence means. There are several ways to proceed. A general AI assistant may be enough for an individual willing to act as engineer and rule-writer. Some businesses will be better served by hiring an assistant or operations manager, whose work agents can support. Or they can use a managed service such as PrivateChief, with separate workspaces and hosted, updated and backed-up servers. The author gives its Business membership price as $199 a month after a one-time $299 setup, with AI billed separately.

None of these routes eliminates costs. AI usage is billed; tools need maintenance as vendors change their requirements; mistakes consume time or money; and the owner still has to make decisions. In return, a business may stop paying for routine services, take on work it could not

previously afford, and regain hours. The author’s experience is that customers chiefly care about waiting, being misunderstood and being passed around—not whether the first answer came from a person. Still, agents need limited access, customer data should be restricted to the job at hand, and anything difficult to undo needs confirmation. When an AI service goes down, agents should say what did and did not get done, not pretend all is well. The principles apply beyond online businesses: shops, plumbers and law offices also have schedules, invoices, filings, suppliers and customers.

The rules gathered from the company’s mistakes put the same discipline into practice. Agents should reproduce a problem before fixing it, rely on tests or live results rather than another model’s approval, and check numbers against the original record. They should fix underlying procedures, place instructions where they will be found, and search every hidden default when retiring a supplier. Owners should receive only genuine decisions, concise questions with recommendations, and reports that speak up when something is wrong. Memory must be written down, traced to its source and checked for staleness. Bespoke work belongs with specialists; consequential professional work still needs professional review. Above all, the work should be prepared thoroughly, verified against reality, and stopped before an agent commits a person’s money, name or responsibility.

Generated by BookGist.ai • October 2026

www.bookgist.ai