From Prompt to Practice

Inside the Pezigns AI Operating Model

White paper · Revised editorial draft · October 16, 2026

How we use AI across our business

How we use AI across our business

At Pezigns, we use AI across our business: research, planning, writing, design, implementation, project coordination, and review. Some of that work supports our own operations. Some supports client projects. A website redesign and a promotional-products project give us one way to explain how the pieces fit together, but they represent only a small part of that broader use.

The practical problem is keeping client requests moving while I'm also responsible for decisions inside our own business. I still make the business and creative decisions, maintain relationships, and take responsibility for the work. We use AI to help carry the steps around those decisions: gathering information, preparing drafts and project materials, carrying out approved tasks, and checking what comes back.

Our setup includes more than 60 AI agents across several platforms. The agents have different responsibilities, and an assignment uses the roles it needs.

This paper explains how we organize that work, what we pass from one assignment to the next, and where I remain involved.

Why we organize AI around the work

Much of running a business happens between the obvious tasks. A customer's email needs to become a clear brief. Approved copy needs to reach the person or tool implementing it. A review finding needs a correction, and someone needs to check that correction.

If I have to reconstruct the background and relay every decision at each step, I am still doing much of the coordination myself. A useful draft helps, but the surrounding work remains.

Our approach gives that surrounding work a place in the assignment. Along with the requested result, we identify the current materials, the person or role responsible, the decisions already made, and anything that could prevent progress. When the work comes back, I need to see what was produced and what still needs attention.

That applies to a short writing task as much as a larger project. The level of detail changes. The need to know which version is current and whether the task is actually finished does not.

How we divide responsibilities

We organize the AI team into management, project, specialist, and review roles. Those roles help us decide who should handle an assignment and where questions or problems should return.

Responsibility	What it covers
Business and creative direction	I set priorities, make the decisions reserved for me, and remain accountable for the work.
Coordination	Organize assignments, identify dependencies, and bring back questions that need a decision.
Project management	Keep requirements, current materials, responsibilities, and next steps together.
Specialist work	Research, write, design, implement, or prepare the material requested.
Review	Check the result against the brief and identify corrections or unresolved issues.

A small assignment may need one specialist and a check. A larger project may involve several roles, with work moving back and forth as it is refined. Adding more agents is useful only when they have a clear contribution to make.

A role also needs appropriate access. Giving an agent a title does not grant permission to use an account, share information, send a message, or publish work. Those permissions have to be established separately.

When work passes between platforms, the handoff needs the current materials, a clear status, and a way to return the result.

One example: our website and a client project

This is just one example of how we've put AI to work.

Our website and a client promotional-products project illustrate two different kinds of responsibility within the business. The website needs my attention on positioning, content, and design. The client project needs careful handling of a request, project information, and follow-up.

Working on our own website

For our website, I need to decide what we are saying, who we are speaking to, and how we want the work presented. AI can help develop and carry out that direction, but those choices still need my judgment.

Once a direction is approved, the next assignment needs more than a general instruction to build it. The implementation work needs the current copy and assets, the agreed structure, and any requirements for how the page should behave. The reviewer needs the specific version to

inspect and the conditions it should meet.

If a review finds a problem, the finding needs to return with enough detail to correct it. After the correction, the relevant part needs another check. That is different from starting the whole project over or assuming a change worked because it was made.

Accessibility and privacy belong in those requirements and checks. Neither is settled by a polished design or a single automated test.

My attention can stay on a decision while other approved, independent tasks proceed. Where work depends on that decision, the dependency should remain clear.

Preparing and following up on a client request

For one client project, we used AI to summarize the promotional-products request, prepare a recap for my review, and organize the request and follow-up.

The customer’s request needed to become a clear set of next steps. We recovered the original request, saved the recap as a draft, created a shared project folder, and linked it to a project record with tracking and dispatch follow-ups.

That gave me the request, a draft to review, and a record of the next actions. At that preparation stage, receipt and quality confirmation were still outstanding. The customer later confirmed receipt, and I closed the project.

For the customer, our responsibility extended through the project, not just the reply. Keeping those details together helped us identify what still needed attention before we could consider the work complete.

What the example shows

Across both projects, the work needs to come back with a usable account of its progress: the current result, checks performed, unresolved questions, and next action.

That is the approach we apply more broadly across the business. Different assignments need different tools and expertise, but I should not have to piece together their status from scattered conversations.

An independent task can proceed within its approved scope while I focus elsewhere. If work needs my decision, that question should come back with enough context for me to answer it.

Keeping the right information with an assignment

A useful handoff includes the information needed to continue the work without making the next contributor guess.

For a writing assignment, that might be the audience, approved message, current draft, and claims that need support. For implementation, it might be the agreed design, exact files, requirements, and known issues. For a client follow-up, it includes the original request and what has or has not been confirmed.

We distinguish between a few kinds of records:

- **The work record:** what is owed, who is responsible, and the current next step.
- **Source material:** the original request, relevant correspondence, and supporting information.
- **Working files:** the specific draft, design, or implementation being reviewed.
- **Decisions and checks:** what was approved, what was tested, and what remains unresolved.

These records serve different purposes. A conversation can explain a decision without being the best place to track an open commitment. An earlier draft can provide context without being the version someone should edit.

Sharing context also requires restraint. An agent needs the information relevant to its assignment, not every conversation or every client's material. Internal and client work retain their own information boundaries.

If work is interrupted, the next step is to check what was saved before repeating an action. Otherwise, recovery can create duplicate drafts, conflicting updates, or a claim that something is complete when its result is uncertain.

Where human decisions stay

I want delegated work to reach a useful review point without repeated requests to confirm decisions I have already made. That requires a clear brief and a clear limit on what the assignment allows.

Research, drafting, organizing materials, implementing approved changes, and checking results can all be part of delegated work. Sending a customer message, making a purchase, publishing, or releasing a change needs the authority appropriate to that action.

For example, asking for an email draft does not authorize sending it. Asking for a website change does not automatically authorize a public launch. Passing an assignment to another role does not expand its permissions.

I remain responsible for business priorities, creative judgment, relationships, and the decisions reserved for me. The workflow needs to bring those decisions back with enough context for me to make them.

This takes ongoing work. Instructions need maintenance, connections can fail, and a result can look convincing while missing a requirement. Review remains part of using the system.

How we judge a result

When work comes back, I need to know whether it meets the brief and what still needs to happen. For a website, design approval, implementation, review, and launch are separate steps. For a client project, preparing the materials is one part of following through on the commitment.

The check should fit the task. A draft needs a review of meaning, accuracy, and voice. An implemented change needs checks against its requirements. A delivery commitment needs confirmation appropriate to what was promised.

Having another agent review work can help, but it is not a guarantee. Reviewers can repeat the same assumptions or overlook the same issue. A useful review explains what was checked, what evidence supports the finding, and what could not be verified.

What we need to measure

The reason to build this system is to make business work more manageable while keeping its quality and responsibilities clear. Whether it does that well needs to be assessed.

Useful measures include:

- How long an assignment takes to reach a satisfactory review.
- How much preparation, checking, and correction I still need to do.
- Whether handoffs lose information or create duplicate work.
- Whether commitments reach a confirmed outcome.
- What the tools, integrations, maintenance, and review cost.
- Whether I can protect time for focused work and personal commitments without losing track of obligations.

Time spent waiting for a decision or customer response should be separated from time spent doing the work. A high volume of generated material is not enough to establish a benefit.

We need to look at representative assignments over time, including the ones that require correction. That is how we can judge whether the help is worth its coordination and maintenance.

What another business can take from this

There is no need to start with a large agent team. A recurring piece of work is a more useful starting point.

Choose a task with a clear result. Identify the information it needs, the part AI can help with, the decisions a person keeps, and how the result will be checked. Try it with real work, then

examine what still required explanation or correction.

If that process becomes dependable, consider how it connects to the next step. More roles or tools should address a specific need. They also add coordination and maintenance, which belong in the assessment.

For us, the broader aim is to use AI throughout the business without losing sight of the work itself. Research should inform a decision. A draft should be usable. A handoff should let someone continue. An open commitment should stay visible until its outcome is confirmed.

Our website and the promotional-products project offer one view of that approach. The same questions apply wherever we use AI: what needs to happen, what help is appropriate, and how will we know the work is ready?

A practical assignment checklist

Before delegating, we need enough clarity to answer:

1. What result are we asking for?
2. Which source materials and current files should be used?
3. What decisions have already been made?
4. What is allowed, and what requires a separate decision?
5. What information or access is still missing?
6. How will the result be checked?

The return should include the work itself, the checks performed, any limitations, and the next action or decision. That gives me something concrete to review and keeps unfinished work from disappearing behind a confident summary.

About this account: The description of our use and the 60+ agent scale is my firsthand account, not an independently audited roster. The website and client examples draw on separate historical records, not a verified timeline of concurrent activity. The earlier client record documents preparation and tracking, including an unsent recap. I later reported the customer's confirmation of receipt and authorized project closure. That owner-reported receipt is not independent provider verification or a separate quality confirmation. No measured savings, automatic cross-platform synchronization, or unattended-operation result is claimed. Customer and implementation details have been omitted.