

# OIP Insurtech - Site Audit

---

One Roll Studios

2026-05-19

# CONTENTS

Site Audit .....

oipinsurtech.com and boundai.ai .....

Introduction .....

oipinsurtech.com .....

Finding 1. The H1 describes the company, not the outcome .....

Finding 2. Three service pillars compete for the same click .....

Finding 3. Statistics section shows literal zeros .....

Finding 4. "Let's Talk" vs. "Book a call". Two CTAs, no clear priority .....

boundai.ai .....

Finding 1. The hero nails the problem but buries the product .....

Finding 2. "Production Ready. System Integrated." is the best claim on either site, and it's buried

Finding 3. No demo, no trial, no screenshot, no video .....

Finding 4. The blog is a single-author monolith. No customer voices, no case study narrative

Close .....

# Site Audit

## oipinsurtech.com and boundai.ai

Prepared by One Roll Studios. 2026-05-19.

---

### Introduction

This audit reviews the public-facing homepages of oipinsurtech.com and boundai.ai with two questions in mind:

1. Does each site convert the prospect it is built for?
2. If not, where specifically does the conversion break down, and what is one concrete direction that would fix it?

We approached the audit as a first-time visitor would. Landing cold, scanning for relevance, looking for proof, deciding whether to commit to a next step. We did not run heatmaps or session recordings; those require access we do not yet have. What follows is structural critique grounded in the actual copy, layout, and CTAs as they render on a desktop browser on the audit date.

Eight findings follow. Four for oipinsurtech.com, four for boundai.ai. Each is structured the same way:

- **What we see.** The observable behavior on the live site.
- **What it costs.** The specific conversion impact we'd expect.
- **One direction.** A concrete first move, not a complete redesign.

This is not the redesign. This is the read.

---

# oipinsurtech.com

## Finding 1. The H1 describes the company, not the outcome

**What we see.** The homepage H1 is “All-in-One Partner for Comprehensive Insurance Operations Solutions”. Nine words describing *what OIP is*, not what it enables. The actual value proposition (“eliminate your limiting factors”) doesn’t appear until two scrolls down, after visitors have already encountered generic copy about “streamlining workflows” and “future-proofing processes.”

**What it costs.** E&S operators scanning for help don’t know whether OIP solves their capacity problem, their tech gap, or their process drag until paragraph three. The hero subheading uses four abstract nouns in one sentence (“workflows,” “stability,” “processes,” “expertise”) without naming a single concrete pain point. A VP of Operations in a hurry will bounce before discovering the “lack of talent or tech limiting growth” framing that actually resonates.

**One direction.** Lead with the operational reality, not the category. *“Underwriting and claims teams stretched thin. Trained operators and tech-enabled execution under one partner. Founded in 2012, over 100 clients across US, Canada, and the UK.”* That names what OIP actually delivers: insurance-trained teams, ops infrastructure, and scaling capacity without adding headcount. The hero earns the right to detail by leading with the buyer’s pain and OIP’s posture, not by listing solutions.

---

## Finding 2. Three service pillars compete for the same click

**What we see.** The homepage presents three service cards in a row. “Underwriting Support,” “InsurTech Services,” and “BoundAI”. Each with identical “Learn more” CTAs. No visual hierarchy distinguishes them. BoundAI gets a small subtitle (“Agentic AI Solutions”), but a first-time visitor has no signal about whether they need one service, two, or all three. The cards treat a fractional underwriting contract and a six-figure AI deployment as equivalent browse actions.

**What it costs.** Visitors with a specific problem (like *“I need underwriting capacity tomorrow”*) don’t know which door to open. The lack of differentiation creates decision paralysis instead of a clear path. And because BoundAI is presented as one option among three rather than the flagship product, it doesn’t receive the attention or conversion priority it likely deserves as the higher-ACV, higher-margin offering.

**One direction.** Stratify the CTAs by urgency and ticket size. BoundAI gets “Request a demo” as a product-page entry. Underwriting Support gets “See capacity options”. InsurTech Services gets “Explore solutions”. Different verbs, different commitment levels, different conversion paths. Same hero, three distinct journeys.

---

### Finding 3. Statistics section shows literal zeros

**What we see.** Midway down the homepage, four statistics appear under the “14 Years of Excellence” section: “Insurance Industry Clients,” “Process Improvement Projects,” “Market Platforms Proficiency,” and “Employees.” Every single number displays as “0” (or “> 0 %” for one metric). This suggests a JavaScript animation that loads the numbers dynamically. But if the script fails, is blocked, or simply hasn’t fired yet, visitors see a company claiming zero clients and zero employees.

**What it costs.** Social proof becomes social doubt. A prospect evaluating OIP’s credibility sees “Trusted by 120+ insurance companies” in one section, then scrolls down to “0 Insurance Industry Clients” two sections later. Even if the animation works for 95% of visitors, the remaining 5% see a company that can’t count its own customers. And for decision-makers using screen readers or browsing with JavaScript disabled, the zeros are permanent.

**One direction.** Static text beats animated text when trust is on the line. If the animation is non-negotiable, set default values in the HTML so the real numbers are visible before JavaScript loads, not after. The site’s other proof points (Insurity SI partnership, ISO 27001, SOC II) should sit in the same module, not buried lower on the page.

---

### Finding 4. “Let’s Talk” vs. “Book a call”. Two CTAs, no clear priority

**What we see.** The top-right navigation features a “Let’s Talk” button, while the hero section below it displays “Book a call.” Both CTAs presumably route to the same contact form, but the copy divergence creates micro-confusion: is “Let’s Talk” for general inquiries and “Book a call” for qualified prospects? Are they different commitment levels? The site never clarifies. Further down the page, “Learn More” appears 8+ times across service cards, testimonials, and content blocks, becoming the default action for everything.

**What it costs.** CTA proliferation dilutes conversion intent. When every section offers a “Learn More” link, no section feels like the conversion moment. And when the primary CTA changes copy between the nav and the hero (“Let’s Talk” to “Book a call”), visitors intuitively assume a hi-

erarchy that doesn't exist, which slows decision-making. High-intent visitors who are ready to book may hesitate because they're not sure which button is "correct."

**One direction.** One primary CTA, one verb, repeated everywhere. For the OIP umbrella site, "Book a Call" is a confident, low-friction primary that fits the operator brand. Different sub-pages can have product-specific secondary CTAs (e.g., BoundAI's "Request a demo"), but the home-page navigation and hero should be identical twins.

---

# boundai.ai

## Finding 1. The hero nails the problem but buries the product

**What we see.** BoundAI's homepage opens with "AI Agents Built for Specialty Lines". A clear, specific H1 that immediately orients the visitor. The subheadline, "Insurance Workflow Infrastructure: Built for Production," positions BoundAI as enterprise-grade, not experimental. The problem statement two lines down ("Underwriters shouldn't spend half their day moving data between platforms") is concrete and emotionally resonant. But the actual product, what BoundAI *is*, doesn't crystallize until after three scroll sections. The visitor knows the pain point before they know the platform's name or category.

**What it costs.** High-intent buyers scanning for "policy validation AI" or "submission intake automation" don't get a crisp product definition above the fold. The site optimizes for engagement (pull them in with the problem) rather than qualification (tell them exactly what we sell). For a marketing-qualified lead arriving from a conference or referral, this works beautifully. For a cold visitor Googling "AI for specialty insurance," they may bail before understanding that BoundAI is a document intelligence and workflow orchestration platform, not a chatbot or underwriting copilot.

**One direction.** One sentence, above the fold, that completes the frame. *"BoundAI is a structured execution layer that connects submission intake, clearance, policy validation, and compliance into your existing PAS/AMS environment with over 99% core field accuracy."* The problem hooks them; the product qualifies them; the accuracy claim closes the trust gap before they have to look for it.

## Finding 2. "Production Ready. System Integrated." is the best claim on either site, and it's buried

**What we see.** Two-thirds down the page, BoundAI stakes a sharp positioning claim: *"BoundAI is not a demo environment. Not a siloed extraction engine. Not a proof of concept waiting to scale."* This section directly confronts the buyer's most likely objection (that AI tools are science projects, not production infrastructure) and dismantles it with specific evidence: *"No rip-and-replace. No parallel shadow tools. No duplicate data stores."* The section title, *"Production Ready. System Integrated."*, is the clearest value claim on the entire site.

**What it costs.** Nothing. This is the win. But it's positioned in the middle of a long scroll, after "Document Intelligence" and "Broker Workspace" deep-dives, when it should be a hero-section argument. The site assumes visitors will read 2,000+ words of feature exposition before they need reassurance that BoundAI won't blow up their tech stack. In reality, *"Is this another tool my team won't use?"* is the first question, not the fifth.

**One direction.** Move "Production Ready. System Integrated." into the top three sections. It is the strongest deflection of the default "no" and it directly addresses the CFO/COO who controls budget. Feature details come after the buyer believes deployment is possible. Pair it with the Insurity SI partner badge and SOC II / ISO 27001 marks for an instant credibility row.

---

### Finding 3. No demo, no trial, no screenshot, no video

**What we see.** BoundAI's homepage is 100% text and static graphics (workflow diagrams, capability lists). There is no product screenshot, no screen recording, no embedded demo video, no interactive element that shows the platform in action. The site asks prospects to "Book a Demo" as the primary CTA, but offers zero preview of what that demo will show. The closest visual evidence is a three-box execution layer diagram (Submission to Cleared Risk to Structured Record), which is conceptual, not interface-level.

**What it costs.** Trust friction. B2B software buyers, especially insurance operators who have been burned by vendor promises, want proof before they book 45 minutes with sales. A 90-second screen recording showing a submission being parsed, validated, and written to a PAS would do more credibility work than three paragraphs about "clause-level deviation detection." Without it, BoundAI is asking prospects to take the entire value proposition on faith, which is a high bar for a category (agentic AI for insurance) that didn't exist two years ago.

**One direction.** Embed one 60 to 90 second screen recording or animated workflow asset in the "Document Intelligence" or "Execution Layers" section. Show a real submission flowing through the system. Make it silent, annotated, fast. Let the product sell itself before the demo call. The asset doubles as sales-deck material.

---

### Finding 4. The blog is a single-author monolith. No customer voices, no case study narrative

**What we see.** BoundAI's blog features six recent posts (as of May 2026), five of which are authored by Milos Paskas, one by Andrijana N. Ilic. Every post is editorial or thought leadership: *"The Myth of End-to-End Automation," "Decision Automation vs. Decision Support," "The P&C*

*Playbook Is Broken.*” There are zero customer case studies, zero *”How [Client X] reduced policy review time by 80%”* stories, zero guest posts from carrier CTOs or MGA operators. The two “Webinar Recap” posts suggest BoundAI is hosting events, but those recaps don’t feature customer speakers or testimonials.

**What it costs.** The blog reads like a manifesto, not a proof engine. Thought leadership establishes category authority, but case studies close deals. A VP of Underwriting evaluating BoundAI wants to see a peer at a comparable carrier saying *”Here’s what we deployed, here’s what broke, here’s what worked, here’s the ROI.”* Right now, the content strategy is 100% outbound perspective (BoundAI teaching the market) and 0% inbound evidence (the market validating BoundAI). The “Trusted by the Most Ambitious E&S Partners” logo section tries to fill this gap, but logos without narratives are weak signal.

**One direction.** Publish one customer case study per quarter with a named operator, a specific workflow (e.g., *”How [Syndicate X] automated ACORD 125 intake for 400 submissions per week”*), and quantified outcomes. If NDAs prevent named attribution, publish anonymized case studies with vertical and volume context (*”Mid-market MGA, 15 underwriters, \$200M GWP”*). Let customers do the selling.

---

# Close

These eight findings are not a complete redesign brief. They are the eight things we noticed first, on the first read. The friction a cold prospect encounters before the actual product gets a chance to land.

A redesign would do more. It would tighten the copy systematically, pull positioning research forward into the page, build a component system that makes future iteration cheap, and integrate evidence (case studies, demo recordings, customer quotes) as first-class content rather than afterthought.

But the findings above are where we would start. Each one moves a measurable thing. Bounce rate on the hero, click-through on the service cards, time-to-CTA on the BoundAI page. None of them require new content; they require restructuring what is already there.

If you'd like to talk through any of them in more depth, the call link is in the deal-room.

---

**One Roll Studios** *An AI-native product studio. Kali Nigam and Prerit Khanna.*

This audit was produced as part of Sprint 0. Work shipped before the engagement began.