How to Build an AI-Native Services Company

11 MIN · YOUTUBE · HTTPS://WWW.YOUTUBE.COM/WATCH?V=GSNFJBGOAHI
https://www.youtube.com/watch?v=gSNFJbgoaHI

SUMMARY

The video discusses the emergence of AI-native service companies, which are poised to dominate industries like insurance, law, and healthcare by leveraging AI to deliver outcomes rather than just tools. Founders are encouraged to consider specific market traits, assemble the right teams, and adopt operational rigor to build successful AI service businesses.

- *AI-native service companies are expected to thrive in markets like tax, audit, insurance, and healthcare, leveraging AI to automate tasks.*
- *Ideal markets for these companies have traits such as low trust, low judgment at the task level, high intelligence thresholds, and favorable regulations.*
- *Founders should prioritize building teams with domain and model fluency, as well as operational rigor.*
- *The product design should focus on scaling human work nonlinearly, with an emphasis on throughput and consistency to maintain customer trust.*
- *Pricing strategies should align with value rather than traditional cost-plus models, avoiding commoditization.*
- *Founders should be cautious of the early demand trap, limiting pilot customers to ensure product scalability.*
- *The profit and loss structure is critical, with a focus on managing costs of goods sold and operating expenses to achieve higher margins.*
- *Buying existing service businesses is generally discouraged; building from scratch is often more effective for achieving product-market fit.*

Some of the biggest companies of the next decade won't be software businesses at all. They'll be services companies like insurance carriers and law firms rebuilt from scratch with AI doing most of the work. These are what we call AI native service companies. And the markets are trillions of dollars in size. tax, audit, insurance, law, parts of healthcare, and so forth. This opportunity didn't exist even a couple years ago. But advances in the models have unlocked this new type of business where companies provide the outcome to the customer versus build a co-pilot that the customer uses internally. These

companies also look and feel different than most startups today. In this video, I'll walk through a playbook for founders starting AI services businesses from scratch. It's aimed at people thinking about starting a company, not if you're already running one. I'll share some obvious and non-obvious elements of building these businesses that we've observed here at YC. Topics include picking a market, forming a team, building the actual product, serving the customers, the P&L, and whether or not you should even buy a business. One general comment before we get started. We're still early here. Like most things in AI, the market is moving fast. We're learning as we go, but the early successes here should get you really excited. First, picking the right market. The same general advice for all startups applies here with some important caveats. You should pick a market you're excited to work in for a long time. These companies still take a decade or more. If you don't love some combination of the customers or the market or the technical problem, you're not going to make it. And that part isn't really new. But the best markets for AI services have four new pretty unique traits. The first is low trust, meaning the work is already outsourced and the customer cares about the final product, not how they got there. You're displacing a vendor, not asking the customer to do something fundamentally different. That's a huge deal because you're not changing behavior. You're showing up where the budget already lives and doing the work. Second, low judgment at the task level. If you can

break the work into pieces and every piece needs a human exercising actual judgment, you can't really scale. You need most of the steps to be automatable with judgment focused in a few places where humans stay in the loop. The third is a high intelligence threshold. This sounds contradictory, but actually it isn't. The overall work has to be hard. Hard enough that models plus humans are need to actually deliver an outcome the customer accepts. The fourth is regulation could actually be good. Regulated industries have higher expectations and legal accountability that raises the bar and the moat for founders. For instance, Panacea is a current YC company that provides FDA regulatory services for biotechs and medteex. They actually hire experienced FDA consultants, pair them with an AI platform to deliver faster, higher quality FDA approvals. So, what are some of the specific markets we have in mind here at YC? The known good fit markets include tax, audit, insurance, mortgages, parts of healthcare, and parts of logistics. But there are plenty more markets nobody has touched yet. Don't hold yourself to the obvious ones or what people talk about on X. And here are a few more things to keep an eye on when it comes to the markets. The first is on the models. Will the models disrupt these businesses? It depends on what I call the Sam Alman test. You should ask yourself as the models get better, does your service get stronger or does the model itself commoditize you? You want to be in the first camp where to be careful. Anything involving equipment and on-site labor. The

software margin math doesn't apply when you own and operate physical things. It's very hard to create real leverage, though these can be really good businesses. Let's leave this area to the robotics founders. One more honesty check. Ask yourself sincerely, are you using humans because the work genuinely needs judgment, or are you compensating for product gaps? Be honest here so you're not papering over product shortcomings with actual humans. There are still great massive technology businesses to be built with humans in the loop. Second, and maybe most importantly, the right founding team. The same advice applies for all startups here, but again, with some important caveats. You should build companies with people you already know and you've worked with. If you're so low, think about the best people you've ever worked with and ask them to join you. You'd be surprised who says yes. For AI services specifically, there's three attributes that all the best founders share. The first is domain fluency. Direct experience is best, but learned is actually okay. You're selling to skeptical buyers and often regulated spaces. You have to bleed credibility. How you acquire it matters somewhat less. The second is model fluency. You need to know what frontier models can do today and design the product to ride the curve as they get better. There is no substitute for great tech here. People underestimate this. Next is operational rigor. Topics like variance, throughput, cycle times, SOPs. This is not an exciting set of words for most founders, but you are fundamentally running an

operation. You have to learn that skill set and you have to enjoy it or at least you have to respect it. The product is an operation. A great example here is the general legal team which is an AI native law firm that YC recently backed. The founders have a unique mix of actual law firm experience at Cooling Fenwick as well as years of technical leadership at Caseex. But most importantly, they think deeply about throughput and how they staff their firm. They've integrated shift work into how they serve clients to reduce cycle times and attract the best lawyers on the team. This is a win-win for scale. Now, let's talk about building the actual product. With AI services, the setup is the opposite of most software. The human is the interface of the customer, not the product. The product helps the human scale their work nonlinearly. That changes pretty much everything around building the actual product. First, you need to apply an operations mindset. Find the bottlenecks and build for the bottlenecks. Throughput and cycle time are now product metrics. Track them like you would daily active users. Variance is the existential problem here. By variance, I mean non-uniform outputs from your actual service. Customers will fire you for variance faster than they will fire you for being a bit slower or a bit more expensive than the incumbents. They need to trust the output. Inconsistency destroys trust, which causes churn. Thirdly, humans in the loop should scale nonlinearly. If revenue scales just in line with the number of humans you add, you'll have major problems. The humans in the loop

also need to enjoy the software. They are your users. A general point, it's okay to do things that don't scale at the very beginning, but eventually you really do need to scale. Automating the process is the product. Okay. Sales and customer success. The biggest challenge facing founders here is what I'll call the early demand trap. It's easy to sign up a lot of pilot customers when you're just starting out and have nothing. But it can quickly overwhelm your ability to serve them and you won't be able to build the product to scale. You'll be stuck using humans. It is a literal trap. Our advice here is to cap your first pilot customers to a small handful. Resist the temptation to sign too many too quickly. Assuming you avoid this early demand trap, pre and post sales looks pretty different, too. You have to sell outcomes, not seats or tokens. The pilot is the product. For the first handful of customers, don't try and standardize too early. Use those pilots to learn. Find the spots where AI gives you unique leverage versus spots where you're just automating something obvious. Build the product accordingly and do it fast. Pricing is harder than traditional software because you're not competing with other software providers. You're competing directly with the cost of labor internal or outsourced. A few options here on pricing. There's perunit pricing. So per return, per claim, per loan. This is the cleanest. It's the easiest to explain. There's also outcomebased pricing. This aligns incentives beautifully, but it can be harder for you to forecast in your business. Panace prices on the completed

consultant study versus hourly, which is the norm in the industry. There's definitely two pricing strategies to avoid. Cost plus pricing caps your upside permanently. Don't do it. Straight line undercutting makes your work seem cheap and potentially low quality. Price on value. Next, the P&L or the profit and loss statement. This is where these companies live or die. So, let's do a quick walk through. If you haven't stared at one of these financial statements before, you're not alone among founders, and that's okay. The general structure is revenue minus cost of goods sold gets you to gross profit. Gross profit minus operating expenses gets you to operating income. Let's dive into each of these in the context of AI services. Revenue. Ah, the easy part relatively speaking. You will be able to sign contracts. Can you deliver on them repeatedly? I don't know. That depends on your product and your process. Eventually, you want smooth and predictable growth. A great product process is going to smooth out that lumpiness, but in the early days, it'll be spiky on a monthly basis. That's okay. Cost of goods sold or COGS. Obsess over this from day one. There's three main components here. There's model costs, hosting costs, and those humans in the loop. All three of them need a number, a trend line, and someone who owns them. Be deeply suspicious of zero margin or negative margin pilots. They're fine to learn from, but it's really dangerous to get hooked on those. The core bet here is the more the product is built, the lower the cogs, the better the gross margin. I call this

AI operating leverage. Okay. Opex. This includes research and development costs. So, building the product, sales and general administrative, which is like finance, legal, exact salaries. Pretty standard stuff here. Operating income. This is the revenue minus the COGS minus the opex. You will be judged on your operating income in these businesses faster than you might expect. Net income. This is the operating income less taxes and interest. That's less important in the medium term. So, here's the actual P&L opportunity for these businesses. Traditional services firms top out around 30% margins. Pure software and agent companies have more margin, but often smaller TAMs. The bet on these services companies is that that AI operating leverage gets you closer to software margin, say 50% plus on a market that's two to three times bigger than software. You don't need to be there right away. That's okay. But the trajectory has to be believable. Okay. Last but not least, don't try to buy your way in. There's a temptation we've seen, especially among founders with some operating background, to try and buy an existing services business, add some AI on top, short circuit the revenue. This is generally a trap. There's one decent reason to do it. You need a regulatory mode fast, insurance licensing for example, but otherwise this almost never works. So why is that? You just can't acquire a product market fit. Legacy service businesses are, you know, legacy. They have different expectations on metrics, hiring, and performance. Adding AI on top of that doesn't immediately change any of those

realities. Building is almost always better than buying. Okay, to recap, AI services companies present an incredible opportunity for today's founders, but these are fundamentally different startups to build. If you avoid the traps and focus on the process as the product and the product as the process, you have a chance to create a generational company. We're excited to see what you build and hope you apply to YC.