

# VGB INSIGHTS

## SCALEUP VIEWS

## Deep Diving on AI Commercialization

### Table of Contents

#### [I. Key Takeaways](#)

1. Current market leaders have cost-effective, diversified monetization schemes
2. Large-scale partnerships, often backed by equity investments, provide concept validation
3. A new trend among SaaS leaders to buy and package has emerged
4. Growth equity investors stand to benefit from the barbell-shaped market
5. Horizontal AI tools favor user-based pricing (for now)
6. AI data and compute "picks and shovels" critical to broader industry commercialization
7. Current sector-oriented investments less focused on monetization timelines
8. GTM approach critical when commercializing AI-powered hardware/physical assets
9. With generative AI, value determined by metrics, not revenue
10. Open-source could expedite AI commercialization

#### [II. Methodology / Sample Details](#)

#### [III. Select AI Company GTM Profiles](#)

---

**Tyler Holly**

[tyler.holly@lazard.com](mailto:tyler.holly@lazard.com)

**Nick James**

[nick.james@lazard.com](mailto:nick.james@lazard.com)

**Ali Birkby**

[Alastair.birkby@lazard.com](mailto:Alastair.birkby@lazard.com)

**Matthew Sykes**

[matthew.sykes@lazard.com](mailto:matthew.sykes@lazard.com)

**FinTech | Enterprise Software | Consumer | InfraTech | Healthcare | DeepTech**

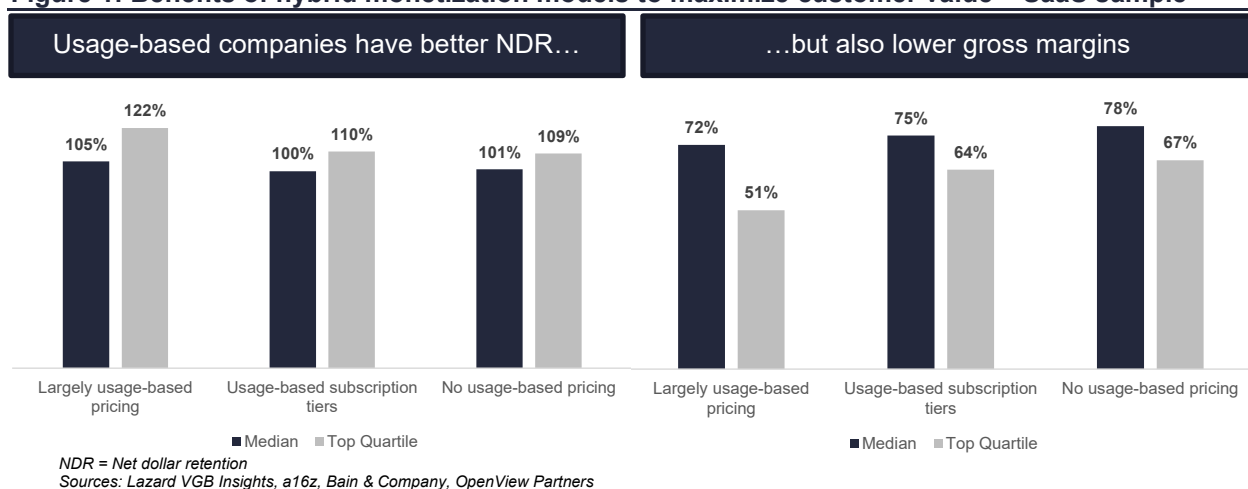
## Section I: Key Takeaways

We evaluated pricing and go-to-market (GTM) models for 150 of the most promising VC-backed AI companies, ranging from pre-revenue to unicorn-stage innovators. Here are our top 10 learnings:

1. **Current market leaders have cost-effective, diversified monetization schemes** – based on the subset of late-stage growth (\$40-100M estimated revenue) and scaled (\$100M+ estimated revenue) AI companies we examined, today's market leaders have maximized their commercialization capabilities through layered pricing mechanisms and GTM strategies. These companies averaged almost 3x the number of combined pricing models and third-party sales channels leveraged compared to their early- and mid-stage (<\$40M estimated revenue) peers. We also found that hybrid GTM models encompassing both user- and usage-based/pay-as-you-go pricing, often coupled with feature add-ons or additional products offered at flat or tiered rates, were most correlated with higher revenue scale and market value (most recent post-money valuation). In traditional SaaS, usage-based models have been associated with some of the highest-value market leaders, particularly in the infrastructure category (i.e. Snowflake, Datadog, Zscaler).

Because gross margins of high-performing AI software companies (~50-60%) often trail traditional SaaS benchmarks (~75-85%) – primarily due to higher input costs and the need for customer-specific services – hybrid subscription/consumption-based models can improve overall margin profiles by encouraging upsells while also providing flexibility in onboarding new customers who are still in an AI discovery/experimentation phase. AI software can generate substantial upside value through monetization models based on data consumption – especially when layered onto flat-fee access subscriptions – as data volume and quality are direct determinants of AI system success and ROI. We think AI companies that adopt hybrid pricing strategies as part of their initial GTM strategies will have a competitive advantage in capturing maximum value from early flagship customers.


**Figure 1: Benefits of hybrid monetization models to maximize customer value – SaaS sample**



2. **Large-scale partnerships, often backed by equity investments, provide concept validation** – the natural extensibility of AI's core value propositions – automating manual processes, personalizing

customer experiences, and making predictive interpretations of the data flowing through a hardware/software platform – enable AI providers to reach a much broader audience by working with channel partners and third-party marketplaces. This trend held true in our analysis, as the monetized partnerships category was the greatest differentiator between the “growth-scaled” and “early-mid” staged companies we examined, as shown in Figure 2. Additionally, we found that growth-stage companies more frequently offered add-on features sold independently, activating expansions into their sales motions. Early- and mid-stage counterparts, in contrast, relied heavily on freemium, tiered monthly rates – mostly based on team size and use case – with comparably simpler product/feature packaging.

**Figure 2: Greatest monetization category disparities – early-mid vs. growth-scaled AI companies**




















A success story leveraging this model in its early growth phase was **Databricks**, which used partnerships and integrations, along with new feature development, to scale from an open-source project to an at-scale industry leader with over \$100M in ARR in just three years. The company partnered early with Microsoft—who also became a strategic investor—to collaboratively develop the Azure Databricks offering and enable 700M+ Azure customers to consume their products without friction through the Azure Marketplace. Databricks also deftly took a cloud-agnostic approach in its GTM strategy by enabling customers to pay compute and data storage costs directly to the cloud providers rather than collecting this as direct revenue, which established friendly dynamics with the hyperscalers. Additionally, in its quest to further differentiate from Snowflake as it scaled to over \$1B in ARR, Databricks continuously layered new features into its enterprise platform to become the leading “data lakehouse”—converging many of the capabilities of a data lake and warehouse into one.

We’ve also seen this approach validated at the infrastructure layer through equity commitments and strategic partnerships executed by leading cloud/SaaS vendors. These tech giants are operating under the thesis that they can create an “economy of scale effect” to commercialize AI by leveraging their existing GTM engines, deep balance sheets, and computing resources. Rather than build in-house,

current market trends suggest partnerships and integrations are widely viewed by legacy vendors as the best entry point for upscaling AI research and development, creating purpose-built models, and incorporating AI capabilities into existing commercial-ready consumer and enterprise products.

**Figure 3: Select cloud/SaaS leaders' AI/ML infrastructure investment and partnership activity**

Cloud/SaaS Leader	Primary Partner(s)	Details
	NA	Launched \$500M fund dedicated to AI strategic investments and led \$100M round for Typeface
		Equity investment and exclusive partnership to embed Cohere in existing services
		Multi-year, \$10B+ partnership; MSFT deploys OpenAI models across consumer and enterprise products while providing supercomputing support for OpenAI's research
		Joint GTM strategy enabling product integrations across AI, low code/no code app development, and data governance
		Equity investment and strategic partnership to build the largest GPU cluster in existence to develop large-scale AI models
		Partnered to provide businesses a platform to create customized generative AI apps within the Snowflake Data Cloud using a business's proprietary data
		AWS agreement to become Stability.AI's preferred cloud provider
	 <b>Hugging Face</b>	Enables Hugging Face and AWS to accelerate ML adoption using the latest models hosted on Hugging Face with the capabilities of Amazon SageMaker
	<b>ANTHROPIC</b>	Equity investment and strategic partnership to integrate Anthropic's model into Zoom's Contact Center portfolio
	<b>ANTHROPIC</b>	\$750M+ investment and partnership to integrate Anthropic's advanced AI assistant into software products and provide further R&D support
		Replit developers get access to Google Cloud infrastructure, services, and foundation models via Ghostwriter, the start-up's software development AI
	NA	Added \$250M to existing fund to invest specifically into AI/ML
	NA	New \$50M fund dedicated to AI and new integrated AI features











Sources: VentureBeat, Financial Times, TechCrunch, Company Websites, Press Releases, Tomasz Tunguz, Contrary Capital

- A new trend among SaaS leaders to buy and package has emerged** – despite the proliferation of partnerships dominating the market leaders' growth strategies thus far, we have recently seen enterprise software companies pay upfront – often at a premium or flat value to a recent highly-priced equity raise – to acquire and integrate AI/ML solutions into their existing platforms with the stated intent of quickly commercializing a bundled offering (see Figure 4). The dynamic nature of today's AI foundational models and data infrastructure – which serve use cases across nearly every industry and attract both tech and non-tech users – presents a challenge for B2B software companies seeking defensible AI/ML strategies. This fight to establish a competitive moat has accelerated M&A timelines and set a foundation for valuation multiples that is highly inconsistent with current market precedents. Buyers increasingly seek to own powerful model/infrastructure assets and industry-oriented applications before they have gained proven commercial traction, believing their existing breadth of GTM capabilities will create an "economies of scale" effect and enable them to be first-to-market with pre-packaged solutions specific to their customer base.

Sources: Richard Waters (FT), Lazard VGB Insights




**Figure 4: Select AI/ML infrastructure M&A activity – 1H 2023**


Acquiror	Target	Amount (\$M)	Step-Up from Last Round	Rationale
 databricks	 mosaic <sup>ML</sup>	\$1,300	~6x	MosaicML will become part of the Databricks Lakehouse Platform, providing generative AI tooling alongside Databricks' existing multi-cloud offerings
 THOMSON REUTERS	 casetext	\$650	~5x	Thomson Reuters' thesis was that Casetext will accelerate and expand the company's market potential for generative AI offerings
 ThoughtSpot.	 MODE	\$200	~1.2x	The combination enables data teams to transform their business intelligence to be AI-first, while reducing bottlenecks and increasing access to insights that drive tangible business results
 snowflake	 neeva	\$150	NA	Snowflake plans to infuse and leverage Neeva's generative AI-enabled search experience across its Data Cloud platform
 NVIDIA.	 OMNI <sup>ML</sup>	Undisclosed	NA	By integrating OmniML's tech into its edge offerings, NVIDIA can optimize models for efficient deployment on lower-end hardware. Additionally, NVIDIA can create custom profiles, maximizing the utilization of its edge hardware suite

Sources: Transaction Press Releases, Pitchbook Data, Inc.

- Growth equity investors stand to benefit from the barbell-shaped market** – disproportionately significant volumes of early-stage AI companies are flooding the market, largely enabled by the model and infrastructure providers who have borne the high costs of model-building and data-tuning to facilitate rapid, low-cost shipment of new AI/ML applications. This nascent, yet fast-growing product development activity – coupled with a highly concentrated funding environment – has created a barbell dynamic in the AI/ML market (see Figure 5). We believe this is likely to create a wave of investment opportunities at the Series B and C stages over the next 6 – 12+ months as more sector-specific AI enterprise tools will attract more generalist B2B SaaS growth investors into the AI funding race. The most active early-stage AI investors to date – including brand names like Sequoia, Index Ventures, a16z, Tiger Global, and Khosla – will provide a validated pipeline for top-tier growth investors to mine, while the bulk of mid-quartile growth investors will need to follow Softbank's lead in crafting separate AI mandates enabling them to invest earlier than in other verticals.

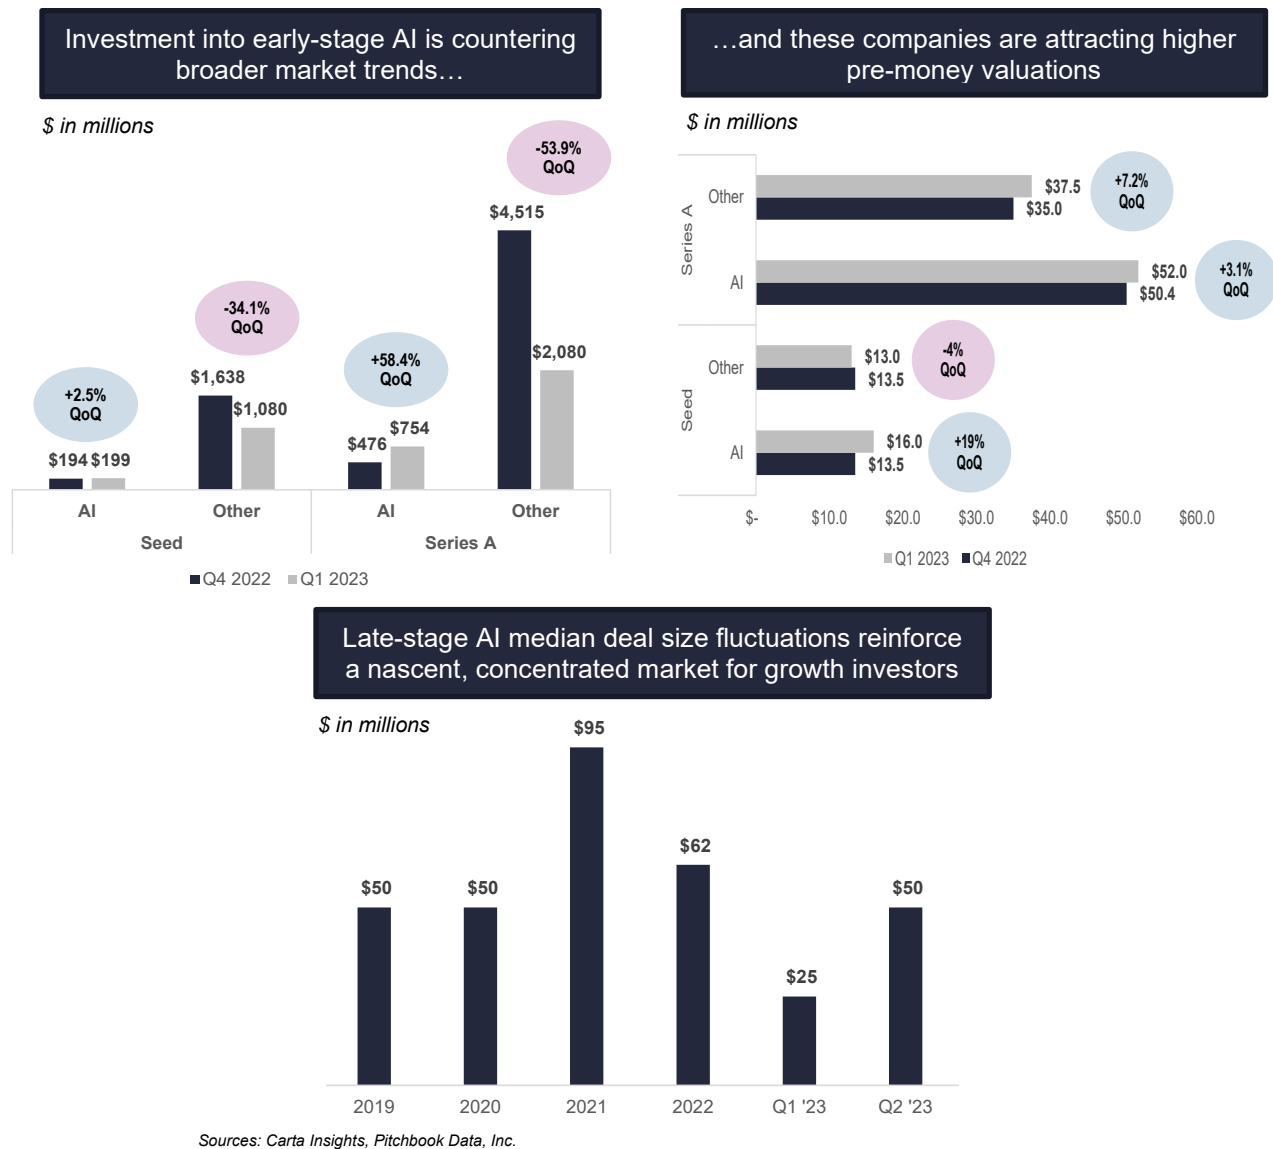
**Figure 5: Barbell dynamics of the AI market – historical deal counts and capital invested**

>50% of total AI funding through 1H '23  
sourced from 7 late-stage infrastructure deals



Source: Crunchbase

Funding data displayed in Figure 6 suggests that the tier 1 VCs investing at the early stages are also looking beyond current market conditions when underwriting new AI investments, adopting a long-term view that backing potential category leaders early amid a new platform shift justifies premium pricing, even at the risk of incurring higher-than-average loss ratios. In contrast, median late-stage deal sizes have followed a lumpy trajectory as investors have recently focused on the 20 – 40 top-funded AI/ML infrastructure providers. Median deal sizes contracted significantly in Q1 '23 to sub-pandemic levels before doubling back to pre-COVID norm levels in Q2.


**Figure 6: Early-stage AI funding trends suggest investors are taking a long-term view**

5. **Horizontal AI tools favor user-based pricing (for now)** – our analysis found that industry-agnostic AI software – encompassing general enterprise workflow tools and generative capabilities to deliver and enhance form content (text, image, and video) – skewed heavily towards seat-based pricing methodologies by a factor of 3x over others included in our sample. Nearly 80% of all GTM strategies centered around user-based pricing that we studied came from this horizontal subset. There are several potential explanations for this trend: **1)** usage-based models are more challenging to implement as companies need to reward high-volume consumers while also finding ways to drive higher engagement among low-average users; **2)** generative tools have myriad use cases for content development with unclear ROI distribution to inform which usage metrics to incentivize through pricing schemes; and **3)** general back-office/ERP solutions often have standardized user profiles with limited upsell potential on a per-user basis. This is consistent with Bain & Company's recent findings on horizontal applications

being a relatively poor fit for usage-based models relative to infrastructure platforms that leverage data as the core asset.

Sources: VentureBeat, Lazard VGB Insights

**Figure 7: Share of customers using consumption models vs. those who want to use them (2022)**



Source: Bain & Company Technology Report 2022

6. **AI data and compute “picks and shovels” critical to broader industry commercialization** – enterprises are rushing to identify differentiated ways to incorporate LLMs and vector databases into their tech stacks, and as recently noted by NVIDIA CEO Jensen Huang, are increasingly focused on cloud-first AI strategies that enable fast development and scalable deployment. The quickest scalers that we studied in the AI/ML infrastructure category all had a common trait of being the early facilitators—“the picks and shovels”—enabling enterprises to leverage models with their own proprietary, unstructured data and access necessary compute power to build scalable applications. While much of the industry’s focus has been centered around the model providers themselves, the data and resource optimizers are the ones most impacting commercialization across the broader market.

**Scale AI** is one example of a company that has maximized its market value (>\$7B) by being the go-to platform that sits between the raw data and the AI models themselves, acting as an enabler for companies seeking to leverage smarter AI capabilities but without the technical resources to implement them. The Scale AI platform automates the manually-intensive process of annotating and labeling enterprise data before it can be fed into AI models, and through its own rigorous back-end ML model training, is able to do so in smarter ways than if humans controlled the process. Having grown from an image and video-tagging business in its early days, the company has expanded its GTM strategy over time by focusing on volume-based pricing that scales with the data labelled for customers (sticky expansion opportunity) and adding offerings such as data debugging tools and synthetic data generation to fill in gaps from customers’ existing datasets. Being the foundational access vector and democratizer to AI models has enabled the company to amass significant market share and fend off competition from earlier-stage players such as **Labelbox** and **DataLoop**.

On the compute side, **Coreweave** has differentiated itself by being the first at-scale access provider to NVIDIA GPUs (highest quality for AI models), and claims to do so at 80% less cost than existing cloud providers. This has led to large-scale, monetized strategic partnerships, including a recent deal with Microsoft that is reported to be worth billions of dollars over multiple years. The company has grown its 1,000+ customer base across four verticals: generative and open-source AI/ML, batch processing, pixel streaming and visual effects, and rendering. Even while competing head-to-head with the major cloud providers – AWS, Google Cloud and Azure – Coreweave has successfully marketed itself as the leading hardware provider specifically for inference. Ultimately, Coreweave's early GTM strategy of offering access to best-in-class GPUs at customer friendly, usage-based rates has been the differentiating factor enabling its scale. Whether the company can maintain that pricing advantage, or diversify its generative AI-led customer base, will ultimately determine its growth potential in a crowded market.

Sources: Wall Street Journal, Contrary Capital, Company Websites, Lazard VGB Insights, CNBC, TechCrunch

**Figure 8: Select AI data and infrastructure providers**

Company	Total Capital Raised (\$M)	Description
Coreweave	\$ 482	Specialized cloud provider powering GPU-accelerated workloads (AI, VFX, and HPC) at scale.
Lightmatter	\$ 266	Changes chip architecture, powering faster, energy-efficient computing with photonic processors for sustainable AI advancement.
Anyscale	\$ 260	Accelerates the development and productionization of any AI app, on any cloud, at any scale.
Weights & Biases	\$ 200	Provides a developer-first MLOps platform that offers performance visualization tools for machine learning.
LangChain	\$ 10	LLM application development library.
OctoML	\$ 133	Offers an acceleration platform that helps engineering teams deploy machine learning models on any hardware.
Weaviate	\$ 68	Builds, maintains, and commercializes the open-source vector database Weaviate
InstaDeep	\$ 109	Delivers AI-powered decision-making systems for the Enterprise, to solve complex industrial problems.
Celestial AI	\$ 164	ML accelerator company that develops data center and edge AI computing solutions.
Comet	\$ 69	Allows data scientists to automatically track their datasets, code changes, experimentation history, and production models.
RelationalAI	\$ 122	Creator of a breakthrough relational knowledge graph system.
Resistant AI	\$ 43	Helps to protect AI systems from targeted manipulation, adversarial machine learning attacks and advanced fraud.
Scale AI	\$ 603	The data platform for AI, providing training data for leading machine learning teams.
Pinecone	\$ 138	Develops a vector database that makes it easy to connect company data with generative AI models.
Replit	\$ 208	Browser-based integrated development environment for cross-platform collaborative coding.
LightOn	\$ 5	Develops "extreme-scale" AI (LLMs, Foundation Models) for the enterprise
Synthesis	\$ 25	On-demand synthetic data for computer vision
MostlyAI	\$ 32	Building an AI-powered global B2B healthcare marketplace
BeeKeeperAI	\$ 24	Zero trust collaboration platform protecting both algorithm IP and regulated data.


Sources: Crunchbase, Pitchbook Data, Inc.

**7. Current sector-oriented investments less focused on monetization timelines** – our sample found that vertical-focused AI deals to date have looked more like DeepTech investments rather than traditional vertical B2B software plays, judging from the risk profiles, long lead-times, and uncertainty around customer adoption inherent to these companies. We believe this dynamic will evolve as the application layer continues to be built out and as industry-specific models enable more widespread integration of AI functionality onto enterprises datasets. Healthcare, mobility, InfraTech (industrial + logistics/supply chain), and clean energy technologies were the predominant vertical-focused solutions covered in our analysis. The common traits of these businesses include:

- High barriers for market entry with required trials/proof-of-concepts
- Proprietary, self-generating data sets – often with a hardware component
- High CAPEX requirements for product development and to reach operational scale

- Industry-specific regulatory and customer adoption hurdles
- Longer relative sales cycles, though often bringing long-term, high-upside contracts

**Figure 9: Industry distribution of vertical-focused AI companies in our sample**



Source: Lazard VGB Insights

Despite this, we think a new wave of verticalized AI applications that look more “SaaS-like” – automating more traditional B2B workflows specific to an industry – is likely to flood the market in the next 12 months. Rather than being developed to solve complex technical problems or enable novel product creation (i.e. new medical therapeutics, innovative infrastructure projects), entrepreneurs are likely to focus on delivering solutions trained on highly-specific datasets that provide instant ROI to industries with content-based workflows at the center of their core operations – this is why the legal industry has been a proving ground for verticalized AI to date. However, in the near-term, more economic value will likely be created by incumbent SaaS leaders integrating AI functionality into their existing offerings. Index Ventures’ Paris Heymann’s recent [vertical AI analysis](#) outlined several factors that could determine which of these two strategies commercializes faster:

- New AI-native vertical applications have a natural wedge against traditional SaaS in automating manual processes and delivering quick time-to-value, particularly in knowledge work industries
- Existing vertical SaaS companies with extensive customer data sets and robust distribution channels can more easily layer AI functionality into existing products, and risk falling behind the competition if this is not made a top strategic priority
- Ultimately, pricing new native vertical applications will be a challenge until customer value can be accurately forecasted, whereas legacy vertical SaaS companies can quickly package new AI features into existing product packages or offer them as tiered add-ons









Sources: Lazard VGB Insights, Index Ventures

8. **GTM approach critical when commercializing AI-powered hardware/physical assets** – the AI-powered hardware and physical asset providers we analyzed were among the most creative in designing GTM strategies. Despite facing more complex sales cycles landing large customers such as government agencies or major health systems with AI-powered hardware leases or sales, these

companies benefit from being able to collect and analyze new types of data through the deployment of physical assets. Coupled with AI-powered software that integrates existing customer data, these companies have an advantage in bringing personalized AI solutions to market. While often bearing significant upfront CAPEX for product development, asset providers often sign multi-year, custom master services agreements (MSAs) with individual customers that offer higher average values than are seen with typical annual software licensing subscriptions. However, given the sensitivities around automated physical asset deployment in industries like defense and healthcare – where lives are potentially at stake – there are often extended lead times to commercialization.

That said, creative GTM approaches have enabled some industry leaders to circumvent these hurdles, including autonomous aircraft and weapons manufacturer **Anduril**. The company scaled to \$10M in revenue in its first 22 months, a milestone which took SpaceX five years to reach, and one Palantir struggled to surpass while selling to the same defense industry customer base. Anduril learned from its industry peers' experiences to accelerate the company's sales motions into the federal government, averaging a much quicker 9 – 12 months to build and ship minimal viable products, and centering its sales around the LatticeOS software platform ("the brain") that powers and integrates the various sentry towers and autonomous aircraft systems it sells. The business model of building "software-native defense systems"—hardware that is operating almost entirely on intelligent software—and selling its software as the initial "land" to later "expand" with its physical assets, has differentiated the company from other field asset providers and enabled the company to scale above \$150M in revenue as of 2022.


**Figure 10: Select AI hardware/software provider GTM models**

	Use Case / Industry	Trials / Development	Unit-Based Sales / Leases	MSAs / Custom Contracts	Partnerships / Integrations	Subscription	Feature Add-Ons / Services
 <b>ANDURIL</b>	Aerospace & Defense		✓	✓	✓	Custom	
<b>CANVAS</b>	Construction	✓	✓		✓		
 <b>FarmWise</b>	AgTech		✓		✓	Usage-Based	Non-Recurring Services
 <b>VERDANT ROBOTICS</b>	AgTech		✓			Usage-Based	
 <b>XtalPi</b>	Healthcare	✓	✓	✓	✓		Feature Add-Ons
<b>refiberd</b>	Recycling	✓	✓				
 <b>Shield AI</b>	Aerospace & Defense		✓	✓	✓		Feature Add-Ons
 <b>AUGMENTA</b>	Building Design		✓	✓		Freemium Usage-Based	
 <b>AMP ROBOTICS</b>	Recycling		✓		✓		
<b>FIGURE</b>	Robotics	✓	✓				
 <b>BeeHero</b>	AgTech		✓			Freemium Usage-Based	Feature Add-Ons

Sources: Company Websites, The Generalist, Lazard VGB insights

9. **With generative AI, value determined by metrics, not revenue** – lofty generative AI valuations relative to the broader market is not a new trend, however the widening delta since pre-2021 has illustrated investors' flight to identify the most promising early-movers (see Figure 11). Through 1H '23, AI startups comprised over half of all VC-backed unicorns newly minted this year.

**Figure 11: Median early-stage pre-money valuations (\$M) – generative AI vs. all startups**



Source: Pitchbook Data, Inc. "In the world of startup valuations, there's generative AI—and everything else"

Our sample of generative startups included >90% of companies at the early-mid stages in terms of revenue traction. The nascency of the market, coupled with the consumer-like nature of their business models (upfront focus on user acquisition), has made revenue an almost obsolete metric in terms of assessing a company's current and future value. So how are investors evaluating "traction" in the generative AI space? Our analysis illustrated the following insights:

**Figure 12: Sample generative AI companies' commercial traction metrics**

Company	Offering	Total Capital Raised (\$M)	Valuation (\$M)	Key Traction Metrics
<b>Character AI</b>	Chatbot application	\$ 193	\$ 1,000	14 million chatbots in usage in first 18 months
<b>Copy.AI</b>	Marketing copy	\$ 14	\$ 41	575% YoY unique search growth
<b>Descript</b>	Video and podcast editing	\$ 100	\$ 575	Significant investment from OpenAI, simple user acquisition via social media affiliate program
<b>Imagen</b>	Personalized photo editing	\$ 34	--	Ingests 150M+ photos annually for tens of thousands of customers
<b>Jasper AI</b>	Content creation platform	\$ 147	\$ 1,500	100,000+ paying customers in 2 years
<b>Midjourney</b>	AI image generator	--	--	15M+ active users, 1.2M average unique searches per month
<b>PolyAI</b>	Voice chatbots	\$ 112	\$ 300	Early contracts with large enterprise customers and key SaaS integrations
<b>Replit</b>	Software code creation	\$ 209	\$ 1,160	Over 22M active developer users; commercial partnership with Google
<b>Synthesia</b>	Synthetic video creation	\$ 155	\$ 1,000	Over 50,000 customers and 450%+ YoY user growth
<b>Tavus</b>	Personalized sales videos	\$ 11	--	Landed 3 early enterprise customers before enterprise pricing was even formalized
<b>Tome</b>	Presentation creation software	\$ 81	\$ 300	Reached 1M active users in first 6 months
<b>Triple Whale</b>	AI for Shopify sellers	\$ 52	\$ 350	Over 5,000 brands managed generating \$14B+ in cumulative sales
<b>Writer</b>	Copywriting software	\$ 26	--	Tripled ARR growth and added 150 enterprise customers in 12 months
<b>You</b>	Personalized search engine	\$ 45	--	Hundreds of thousands of active users, 70% growth in sign-ups, 30% growth in unique searches MoM

Sources: Pitchbook Data, Inc., Funding Press Releases, TechCrunch, Company Websites




10. **Open-source could expedite AI commercialization** – based on the data we could gather, at least 20% of the AI companies we examined had inputs and/or development roots from open-source projects; what's notable about these startups is that the majority have quickly adopted enterprise monetization playbooks by introducing “freemium” tiering that incorporates flat-fee or user/usage-based subscriptions to sell premium features. As noted by CRV's [recent analysis on open-source in AI](#), historical precedents show the majority of thriving open-source companies take 2-3+ years before introducing monetization strategies (i.e. Confluent, MongoDB, Elastic). With AI, free open-source LLMs like Stanford University's Alpaca can leverage the outputs of closed-source models (ChatGPT) to offer enterprises a quicker on-ramp to AI model-integration into existing products.

A recently leaked internal Google document validated that the company's senior engineers are concerned about the rapid progress made from their open-source competition. The letter noted that open-source models are faster to develop, more customizable, more private, and as – if not more – capable. Open-source models are also free, unrestricted, and equally friendly to non-tech users, creating the potential for them to play a central role in democratizing AI's commercialization within the enterprise. Regulatory challenges – coupled with bottlenecks created by resource scarcity (compute, engineers) and cloud provider relationships among closed-source competitors – will ultimately determine whether open-source models and infrastructure can maintain a steady pace of growth and adoption.

Sources: The Guardian, CRV, Lazard VGB Insights

**Figure 13: Historical % open-source vs. closed-source by technology type**



Sources: Kelvin Mu (Translink Capital), Lazard VGB Insights

## Section II: Methodology / Sample Details

Our sample of 150 AI-centered companies was generated by names primarily included in the [Forbes AI 50](#), [CB Insights AI 100 2023](#), and [NFX Hot AI 75 lists](#)—all of which were released in Q2 '23 and were compiled by credible AI investors and technical experts. Before jumping into the actual sample data, we used the following criteria and definitions to segment the market:

Estimated Revenue Stage		
Early – Mid – <\$40M estimated current revenue	Growth – \$40 – 100M estimated current revenue	Scaled – \$100M+ estimated current revenue
AI Category		
AI/ML Infrastructure	Application: Horizontal	Application: Vertical
Enterprise Infrastructure	Hardware / Applications: Horizontal	Hardware / Applications: Vertical

**AI/ML Infrastructure** – Includes foundational models (i.e. Anthropic, Cohere) and other “picks and shovels” like databases and model-tuners/enhancers



**Application: Horizontal** – AI application software (including generative) offering an industry-agnostic functional capability




**Application: Vertical** – AI application software (including generative) catering to a specific industry use case



**Enterprise Infrastructure** – AI-powered enterprise solutions for security, data, compute, and other infrastructure needs



**Hardware/Application: Horizontal** – AI-powered hardware/physical assets, often coupled with software, delivering value to customers across industries



**Hardware/Application: Vertical** – AI-powered hardware/physical assets, often coupled with software, delivering value to customers in specific industries



**Sample Breakdown by Category, Pricing Models, and Estimated Revenue Stage**

While heavily skewed towards companies in early stages of development and commercialization, the sample was designed to reflect the nascency of the broader AI market, which includes ~125 unicorns in total. The following Figures 14 - 16 provide further details on the distribution of companies sampled by AI category, pricing model, and estimated revenue stage:

Figure 14: Breakdown of sample by business model and estimated revenue stage




Figure 15: Sample breakdown by AI category





Figure 16: Monetization strategies used by sample participants by AI category



Source: Lazard VGB Insights

## Section III: Select AI Company GTM Profiles

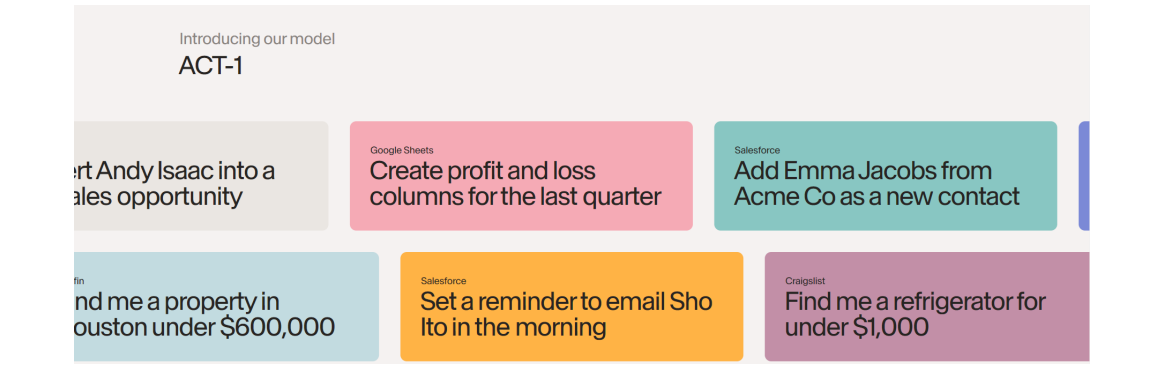
AI/ML Infrastructure			
<a href="#">Adept</a>	<a href="#">Cohere</a>	<a href="#">OpenAI</a>	<a href="#">Granica</a>
<a href="#">Anthropic</a>	<a href="#">Hugging Face</a>	<a href="#">Snorkel AI</a>	<a href="#">Datarobot</a>
<a href="#">Arize AI</a>	<a href="#">MosaicML</a>	<a href="#">Weights &amp; Biases</a>	
Application: Horizontal			
<a href="#">Abnormal Security</a>	<a href="#">Descript</a>	<a href="#">Jasper</a>	<a href="#">RevComm</a>
<a href="#">AlphaSense</a>	<a href="#">Glean</a>	<a href="#">Midjourney</a>	<a href="#">Runway</a>
<a href="#">Character.AI</a>	<a href="#">Gong</a>	<a href="#">Moveworks</a>	<a href="#">Synthesia</a>
<a href="#">Clari</a>	<a href="#">ImagenAI</a>	<a href="#">Neeva</a>	<a href="#">Tome</a>
<a href="#">Coactive AI</a>	<a href="#">Inflection</a>	<a href="#">PolyAI</a>	<a href="#">Writer</a>
<a href="#">Ironclad</a>			
Application: Vertical			
<a href="#">Bayesian Health</a>	<a href="#">Trigo</a>	<a href="#">Unlearn.AI</a>	<a href="#">Waabi</a>
<a href="#">Eightfold AI</a>	<a href="#">Pachama</a>	<a href="#">Vannevar Labs</a>	<a href="#">EvenUp</a>
<a href="#">Harvey</a>	<a href="#">PathAI</a>	<a href="#">VIZ.AI</a>	
<a href="#">Insitro</a>	<a href="#">Slingshot Aerospace</a>		
Enterprise Infrastructure			
<a href="#">Databricks</a>	<a href="#">Surge AI</a>	<a href="#">Builder AI</a>	
<a href="#">Scale AI</a>	<a href="#">Vectra AI</a>		
Hardware / Applications: Vertical			
<a href="#">Anduril Industries</a>	<a href="#">FarmWise</a>		
<a href="#">Canvas</a>	<a href="#">Shield AI</a>		

AI/ML Infrastructure

AI/ML Infrastructure: Models

ADEPT

**Estimated Revenue Stage:** Early-Mid  
**Current Commercialization Model:** Trial/Early Commercialization



Adept's models are designed and trained specifically for taking actions on computers in response to natural language commands. Adept's MVP, ACT-1, is the company's first step towards a foundation model that can use every software tool, API, and website that exists. Unlike giant models that generate language or make decisions on their own, Adept's are much narrower in scope, acting as an interface to existing software tools and making it easier to mitigate bias issues

**Pre-Revenue Development**

- Adept's models remain in development, pre-commercialization phase

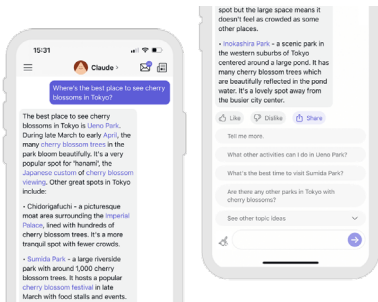
Sources: Company Website, Forbes

## AI/ML Infrastructure: Models

### ANTHROPIC

**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Usage-Based Pricing, Partnerships/Integrations, Non-Recurring Services



Claude is an AI assistant based on Anthropic's research that is capable of a wide variety of conversational and text processing tasks while maintaining a high degree of reliability and predictability with state-of-the-art high-performance model.

Claude Instant is a lighter, less expensive, and much faster option compared to Claude-v1 and is optimized for low latency, high throughput use cases at one-sixth the cost of the Claude family of models.

In 2023, Anthropic began deploying its Claude technology publicly. Anthropic is working with a broad range of early strategic partners/investors to deploy Claude and will look to monetize these relationships through product integrations.

#### Usage-Based

- **Context Window** – 9,000 tokens
- **Pay-As-You-Go Pricing**
  - Prompt – \$11.02/M tokens
  - Completion – \$32.68/M tokens

#### Low Requirement, Usage-Based

- **Context Window** – 9,000 tokens
- **Pay-As-You-Go Pricing**
  - Prompt – \$1.63/M tokens
  - Completion – \$5.51/M tokens

#### Partnerships

- **Google:** Anthropic has selected Google Cloud as its preferred cloud provider, giving it the compute power necessary to build reliable and trustworthy AI systems. Additionally, Google Cloud intends to build large-scale, next-generation TPU and GPU clusters that Anthropic plans to use to train and deploy its cutting-edge AI systems

Sources: Company Website, Ars Technica, Contrary Capital



## AI/ML Infrastructure: Models



**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Freemium, Usage-Based Subscription, Partnerships/Integrations



Sources: Company Website, Google, SaaSwothy

## AI/ML Infrastructure: Models



**Reported Revenue:** Early-Mid

**Current Commercialization Model:** Flat-Fee Subscription Tiers, Usage-Based Subscription

### Surface. Resolve. Improve.

Analytics and workflows to catch model issues, troubleshoot the root cause, and continuously improve performance



ML observability platform designed to help ML engineers and data scientists better understand and resolve model issues and improve model performance.

#### Flat-Fee + Usage-Based

- **Free:** free for personal projects, 2 models, \$500K/month production predictions
- **Pro:** \$100/month for small teams, includes 2 models, 10M annual production predictions, volumetric pricing for incremental predictions, and Slack and email support
- **Enterprise:** custom pricing plan, custom number of models, committed volume-based pricing, API access, on-prem deployment option; Slack, email, and phone support

Sources: Company Website, ZoomInfo

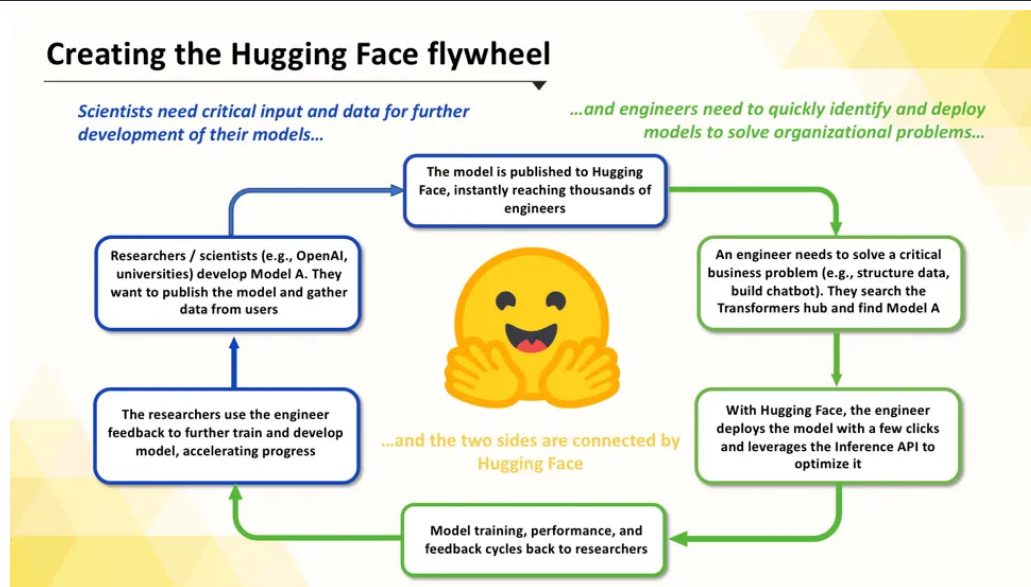
## AI/ML Infrastructure: NLP



### Hugging Face

**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Flat-Fee Subscription Tiers, Usage-Based Subscription



Hugging Face develops AI algorithms and products aimed at software developers. It acts as a community and platform for anyone interested in data science. The majority of the models it makes available on an open-source basis are situated in the field of NLP. Furthermore, Hugging Face offers thousands of data sets, such as IMDb movie scores, that developers can use to train their models on. Meanwhile, developers can exchange tips and tricks with each other or share their ML apps via Hugging Face's dedicated spaces. All the work that users work on is stored within a dedicated GitHub repository, which can be freely accessed by anyone in the community (if the developer chooses to open it up). Lastly, Hugging Face also directly sells enterprise-level products to business customers, including models that don't require code (AutoTrain) or Private Hub, which enables organizations to collaborate in private.

### Flat-Fee + Usage-Based

- Hugging Face generates a majority of its revenues through licensing fees for its models and software, which accompany a range of free open-source tool offerings
  - **Enterprise:** dependent on use case, includes access to ML expert guidance and dedicated, secure deployment hubs
  - **Pro Account:** \$9/month, subscription
  - **Inference Endpoints:** \$.06/hour, enables deployment of endpoints in seconds
  - **Autotrain:** up to \$3/model, codeless AI model creation and automatic search and training
  - **Spaces Hardware:** up to \$4/hr by CPUs, GPUs, Accelerators used
  - **HF Hub:** free model hosting and open-source tool access

Sources: Company Website, Microsoft, Contrary Capital

## AI/ML Infrastructure: Models



**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Usage-Based Subscription, Non-Recurring Services, Feature Add-Ons, Partnerships/Integrations

### MosaicML Training

Allows customers to train their own LLMs and other generative AI models in a secure environment with data privacy and model ownership. Customers can train multi-B parameter models in hours rather than days. Multi-cloud orchestration allows users to deploy on MosaicML infrastructure, public clouds including AWS and Azure, or on-premise

#### Usage-Based

- Pricing based on per GPU-minute of usage
  - Indicative pricing for a case study of a 24,000 A100 hour project is \$2/A100 hour for a total cost of \$48,000

### MosaicML Inference

Mosaic also allows customers to add AI to apps for up to 15x less than competing technologies such as OpenAI. The Starter Edition offers the choice between six open source ML models, while the Enterprise edition can deploy any open source model or the clients' own LLM

#### Usage-Based

- **Starter Edition** – charge based on API usage and varies by model; priced per 1k of tokens
- **Enterprise**
  - Charge is per GPU-minute of usage (pricing undisclosed)
  - 24/7 support and 1 hour response time for SEV-1, 4 hours for SEV-2

Sources: Company Website, Reuters

## AI/ML Infrastructure: Models



**Estimated Revenue Stage:** Growth

**Current Commercialization Model:** Freemium Tiering, Usage-Based Pricing, and Partnership/Channel Profit-Sharing Agreement



Chat GPT is an NLP tool driven by AI technology that enables users to have human-like conversations with the chatbot. The language model can answer questions and assist with tasks such as composing emails, essays, and code.

### Flat-Fee Tiering

- **Chat GPT** – open access
- **Chat GPT Plus** – subscription
  - \$20/month
  - Faster response times
  - Access to ChatGPT during peak times
  - Priority access to new features and improvements



Foundry is a platform for running OpenAI models on a dedicated capacity. It allows users to process larger workloads than the previous GPT models, with full control over the model configuration and performance profile.

### Usage-Based Compute

- **Compute Unit Rentals** – running individual model instances requires specific compute unit minimums
  - Price / Unit: \$260/month for 3 months or \$220/month for annual
  - Depending on the model instance (GPT-3.5 Turbo or DV), annual usage costs range from \$265k - \$1.5m



Multi-billion dollar partnership enables Microsoft to leverage OpenAI to help build out its Azure supercomputing infrastructure, and to deploy OpenAI's models into its consumer and enterprise products.

### Partnership Profit-Share

- **Azure OpenAI Service** – fully-managed, B2B repackaging of OpenAI tech enabling AI model-building and deployment
- **Supercomputing** – ongoing effort to build infrastructure capable of training the largest, most complex AI models
- **MSFT Products** – OpenAI models integrated within MSFT consumer and enterprise products (GitHub, Bing, Office)

Sources: Company Website, Contrary Capital, Sopmac Labs

## AI/ML Infrastructure: Models

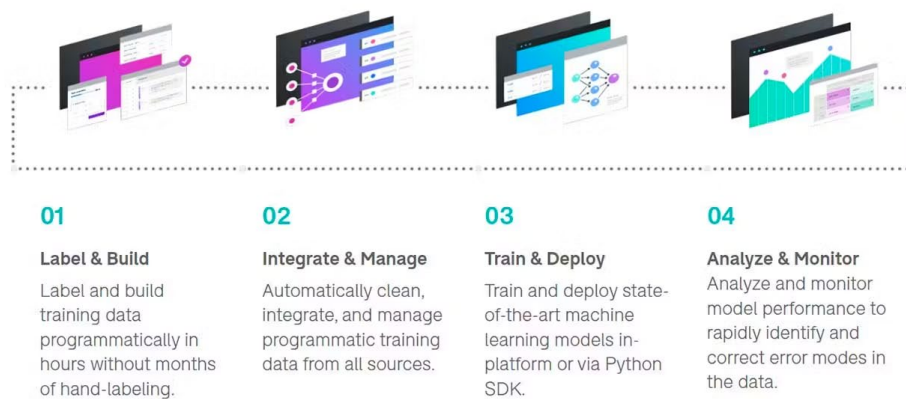


**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Partnerships/Integrations, MSA / Custom Contracts, Feature Add-Ons

### Snorkel Flow

Snorkel Flow is the first data-centric AI platform powered by the programmatic labeling pioneered by Snorkel AI.



Snorkel AI equips enterprises to build or adapt foundation models and large language models to perform with high accuracy on domain-specific datasets and use cases, by supporting the most critical component of developing models: the datasets they learn from.

Custom Contracts	Partnerships/Integrations	Professional Services
<ul style="list-style-type: none"> <li>Snorkel AI has confirmed that customized enterprise contracts are central to the company's pricing strategy, including a "seven- to eight-figure contract with a large BioTech customer"</li> </ul>	<ul style="list-style-type: none"> <li>Snorkel AI partners with cloud providers, technology innovators, and world-class consulting and service providers to build an ecosystem around data-centric AI solutions. Snorkel AI recently teamed with Snowflake to help shared customers transform raw, unstructured data into actionable, AI-powered insights</li> </ul>	<ul style="list-style-type: none"> <li>As part of its customized contracts, Snorkel AI offers managed services to assist with specific data labeling requirements</li> </ul>

Sources: Company Website, Contrary Capital, TrustRadius

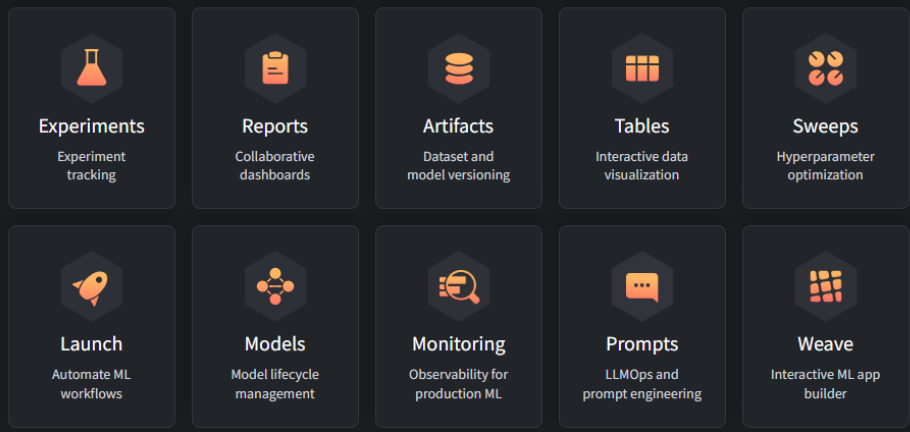
## AI/ML Infrastructure: MLOps



**Estimated Revenue Stage:** Growth

**Current Commercialization Model:** Freemium, User-Based Subscription

The Weights & Biases platform helps you **streamline** your ML workflow from end to end



Weights & Biases provides a developer-first MLOps platform that offers performance visualization tools for ML.

Customer managed deployment options with updates released twice monthly. and command line interface tools for development teams.

### User-Based (Cloud-Hosted)

- **Personal** – Free for personal projects only, unlimited experiments, unlimited tracked hours
- **Starter** – \$50/month per user with tiered pricing
  - Tier 1: \$50/user/month from 250 to 5,000 cumulative hours tracked
  - Tier 2: \$100/user/month from 5,000 to 10,000 cumulative hours tracked
  - Tier 3: \$150/user/month from 10,000 to 15,000 cumulative hours tracked
- **Enterprise** – Custom pricing plan, dedicated support team, flexible deployment options

### Freemium/Enterprise Hosting

- **Personal** – Free for personal projects only, unlimited experiments, unlimited tracked hours
- **Enterprise** – Custom pricing plan, dedicated support team, flexible deployment options

Sources: Company Website, Contrary Capital

## AI/ML Infrastructure: Data

**granica**

**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Outcome-Based, Partnerships/Integrations

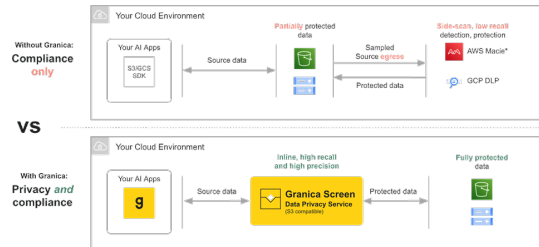


Crunch is a data reduction service for enterprise AI and pre-process training data before it is used in downstream stages to provides deep, and fast, inline data reduction of large-scale image and textual data which reduces the cost to store and access the data.

### Outcome-Based

- Free to deploy with no upfront costs
- Cost is based on a % of savings which is calculated relative to S3/GCS baseline
  - Savings = S3/GCS without Crunch – S3/GCS with Crunch – Crunch infrastructure costs
- Key features:
  - Byte-granular
  - Adaptive and cloud-native
  - Scalable, fast, and resilient with secure and lossless data protection with S-3 compatible API

Sources: Company Website



Future planned API service will provide data privacy via high recall and high precision identification and protection of sensitive data with the goal to automatically detect, classify, encrypt, and control access to all sensitive information to reduce costs.

### Integrations

- Pricing undisclosed, but the API service will provide:
  - Seamless and secure point-of-entry protection
  - Scalable, customizable, and cost-efficient
  - Data-agnostic



Application: Horizontal

AI/ML Infrastructure: Models

DataRobot

**Estimated Revenue Stage:** Scaled  
**Current Commercialization Model:** Flat-Fee Subscription Tiers, User-Based Subscription, Feature Add-Ons, Partnerships/Integrations



DataRobot’s enterprise AI platform democratizes data science with end-to-end automation for building, deploying, and managing ML models. DataRobot the “intelligence layer,” of the data stack, a developing category between the data layer (such as Snowflake, DataBricks, hyperscalers) and application layer (including SAP, Salesforce, and ServiceNow).

DataRobot engages a range of cloud, technology, and services partners to deliver strategy, data science, and IT expertise to help customers build in, deploy, and migrate to the DataRobot AI Platform. Partners create complementary platforms, applications, integrations, or connectors that help extend and enhance the DataRobot AI Platform

Flat-Fee + User-Based

- Up to \$100k annually for 5 users

Partnerships

- DataRobot bundles its ML model development engine with cloud and other data software providers to enable customers to leverage their existing tech stack and data science teams for rapid delivery of new models

Sources: Company Website, AWS Marketplace

## Application: Horizontal – Security



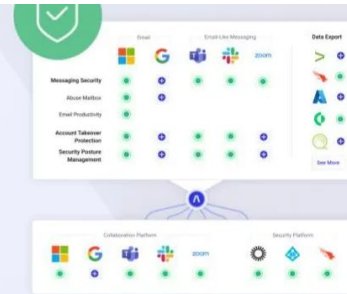
**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** User-Based Subscriptions, Partnerships/Integrations, Feature Add-Ons

### Abnormal Security Expands Behavioral AI Platform to Protect Against Multi-Channel Attacks

New platform enhancements enable detection of email-like attacks across Slack, Microsoft Teams and Zoom

[Learn More](#) →



Unlike secure email gateways, Abnormal's API architecture analyzes unique signals and empowers security teams with behavioral AI detection that stops more attacks. It combines three key features: an identity model, a relationship graph and content analysis. These methods help create a high-confidence decision engine to reduce the number of false positives.

Abnormal leverages various third-party marketplaces to accelerate their GTM strategy and broaden sales channels.

#### User-Based

- **Abnormal Cloud Email Security Platform:** starts at \$35/user/year – pricing includes SaaS service with 24/7 support and a technical account manager
- Additional, optional features added per-user:
  - Email Account Takeover Protection
  - Email Security Posture Management
  - Abuse Mailbox Automation
  - Email Productivity
  - Posture Management
  - Email-Like Messaging Security

#### Partner Sales

- Abnormal Security partners with a range of tech and security providers to extend sales reach via integrations and marketplaces:
  - **Microsoft Azure:** the company has seen success with the Azure Marketplace sponsorship benefit, as Abnormal used more than \$100k in Azure sponsorships to close over \$1M in deals
  - **Crowdstrike:** received strategic funding, announced new platform integrations, and joined the CrowdXDR Alliance to bring a joint XDR solution to market

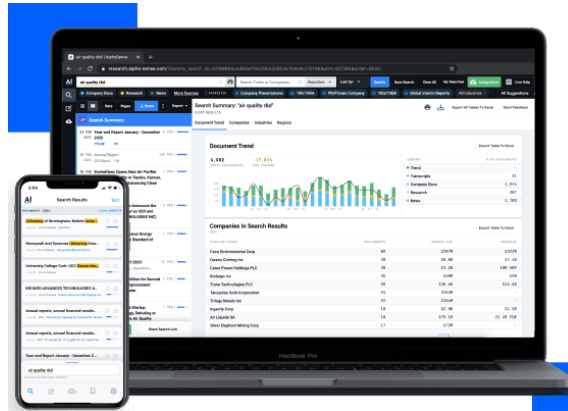
Sources: Company Website, G2 Software, Microsoft, SourceForge

## Application: Horizontal – Market Intelligence

# AlphaSense

**Estimated Revenue Stage:** Scaled

**Current Commercialization Model:** Flat-Fee Subscription Tiers, User-Based Subscriptions, Partnerships/Integrations



Proprietary AI and NLP technology enables customers to easily surface and track insights from Ms of documents across earnings, broker research, company documents, expert calls, etc.

### Flat-Fee and User-Based

- **Per-User:** ~\$5,000 - \$7,000 on average
- **Enterprise Tier:** custom rates for larger teams

### Partnerships

- AlphaSense partners with leading industry data providers (i.e. largest U.S. investment banks) to gain access to proprietary data sets. These arrangements likely enable partners to share in the upside from sales generated through these products

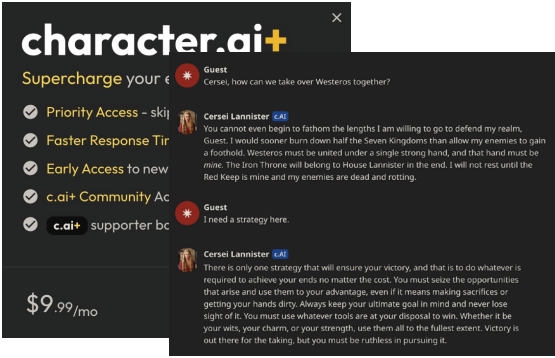
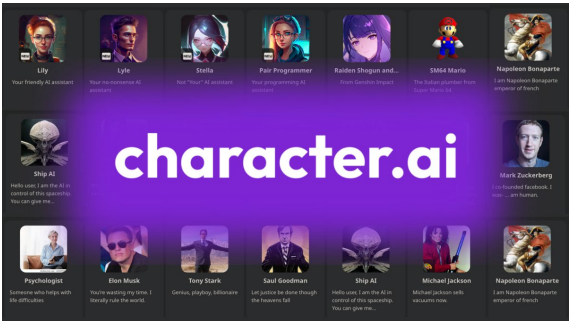
Source: Company Website

Application: Horizontal – Text



**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Freemium, User-Based Subscription, Partnerships/Integrations



Character.ai allows you to interact and chat with various AI-generated characters. Users can create their own characters or interact with pre-built characters.

Character AI Plus, also called character.ai+ or c.ai+, is a paid subscription tier that offers priority access (no more waiting room), faster character response times, early access to new features, and exclusive access to the Plus member community.

**Free Tier**

- Free and unlimited messaging for all users
- Users placed in a waiting queue when the site experiences high-load
- Content remains the same compared to the c.ai+ model

**User-Based**

- Character AI Plus – subscription**
  - \$9.99/month
  - Faster response times
  - Priority access
  - c.ai+ Community Access

Sources: Company Website, Reuters

## Applications: Horizontal – RevOps



**Estimated Revenue Stage:** Scaled

**Current Commercialization Model:** User-Based Subscription, Feature Add-Ons, Partnerships/Integrations

### ACHIEVE REVENUE PRECISION. PREDICTABLY. REPEATEDLY.

Run Revenue like a process to win more deals, protect your customer base, and achieve revenue precision—even in a downturn.

[See Clari in action](#)

**\$52M**

QUOTA

**\$54-58M**


AI PROJECTION RANGE

**\$49M**

PAID COMMIT

**1.6X**

COVERAGE



Basic package including Clari Inspect, Inspect+, and Analyze. Provides the basic functionality to sales team to track and monitor accounts with easy-to-use analytics, dashboards and reporting.

Access to all products in Execution bundle with additional revenue forecasting and pipeline management functionality.

Access to all products in Execution and Forecasting packages with additional CRM data autocapture, collaborative workspaces, sales conversation intelligence, and training environment.

#### Execution

- **Per user pricing** – Similar service providers usually charge \$12 to \$50 /month per user

#### Forecasting

- **Per user pricing** – Similar service providers usually charge \$80 to \$120 /month per user

#### Platform

- **Per user pricing** – Similar service providers usually charge \$120+ /month per user

Sources: Company Website, G2 Software, TrustRadius

## Applications: Horizontal – Data Labeling



**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Usage-Based Subscription

PRODUCTION-READY API AND SDK

```

1 demo.py
2 def classify_assets(image_path: str) -> None:
3     ...
4     Evaluates the visual safety of an image
5     ...
6     with coactive.ApiClient(CONFIG) as api_client:
7         api = ClassificationApi(api_client)
8         image = open(image_path, 'rb')
9         api.classify_assets_for_visual_safety([image])
10

```

### Transform your insights into action

Your data is constantly evolving, and Coactive is here to help. Use our API and Python SDKs to understand and monitor your data as it's coming in.

Visual content labeling platform that leverages pre-trained models to give an initial understanding of visual content without labeling thousands of images. Enterprises can upload raw images or videos to Coactive through an API or data lake connection and the platform will then embed and index the data without need for supervision. Data is then available to query using Coactive's fully hosted image search API and SQL interface with only a little bit of labeling by the end customer.

### Usage-Based

- API access with per-image labeling fees. Comparable companies charge between \$.02 and \$.06 cents per image

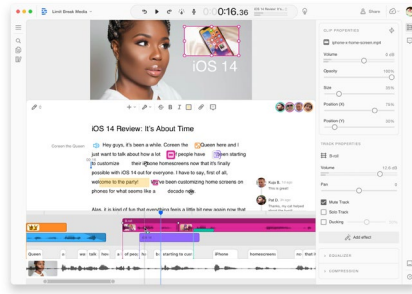
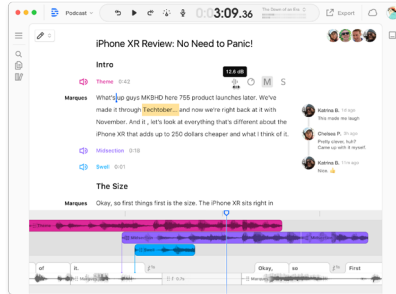
Sources: Company Website, Bessemer Venture Partners

## Application: Horizontal – Audio & Voice



**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Freemium, User-Based Subscription, Partnerships/Integrations



Descript Free Plan offers a cost-effective solution for beginners and casual users to experience the platform's basic features.

The Creator, Pro, and Enterprise Plans each offer an additional level of features for more advanced and extensive multimedia editing needs.

### Freemium

- Unlimited number of projects
- Automatic transcriptions of up to 3 hours of content
- Provides essential editing tools for audio and video files

### User-Based

- **Creator Plan:** Designed for more advanced features and resources which includes 10 hours of automatic transcription /month, additional collaboration tools, and other features
  - \$12 per user/month; billed annually
- **Pro Plan:** Designed for users and businesses with extensive multimedia editing needs with 30 hours of automatic transcriptions/month and other benefits
  - \$24 per user/month; billed annually
- **Enterprise Plan:** Tailor-made solution specifically for large organizations and businesses with unique requirements that include all the features of the Pro Plan

Sources: Company Website, Contrary Capital

## Application: Horizontal – Enterprise Search



**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Freemium, User-Based Subscription

Search across all your company's apps.

### Semantic understanding.

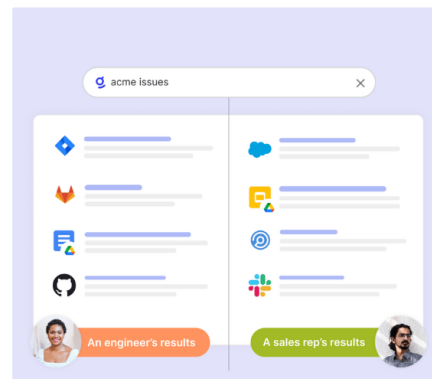
Vector search powered by deep learning-based LLMs enables semantic understanding for natural language queries.

### Personalized results.

Results are based on who you are, what you're working on, and who you're working with.

### Permissions-aware.

AI grounded in your company's own trusted knowledge model delivers the answers you need, aligned to your organization's data privileges.



Glean is a work assistant platform that searches across all of its users' company's apps with the aim of finding exactly what they need. Its enterprise search platform understands context, language, behavior, and relationships with others and helps users find personalized answers to their questions.

## Freemium Subscriptions

- **Glean Chat:** Glean Chat will be priced on a per-seat basis as premium add-on to Glean's core search product. This generative AI-based assistant is in early access for all Glean customers and will soon be made generally available

Sources: Company Website, ComputerWorld

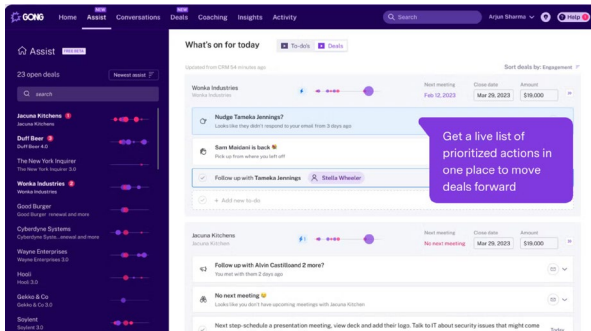


## Application: Horizontal – Sales



**Estimated Revenue Stage:** Scaled

**Current Commercialization Model:** Partnerships/Integrations, User-Based Subscription, Non-Recurring Services, Feature Add-Ons



The Gong Reality Platform uncovers what's happening in customer conversations so revenue teams can do more of what's working and set themselves apart from the competition. Gong's AI-powered platform captures the most comprehensive set of customer interactions in the industry – from video calls, to phone calls, emails, and more – and goes beyond summarization to extract critical insights from customer interactions

Gong's platform offers add-on features to unlock greater levels of visibility across customer interactions. The company recently launched new proprietary generative AI models built specifically for revenue teams and customizable by individual companies. Unlike off-the-shelf systems, Gong's proprietary generative AI models are able to accurately categorize data specific to sales deals, such as customer objections, and understand context and intent in sales conversations

### User-Based (Licenses)

- Gong operates on a subscription model. All licenses offered by Gong are priced per user, and a platform fee is assessed based on the number of users supported. Total costs per user depend on the license type. They do not charge extra for any integrations.
  - **Platform Subscription:** a fixed annual subscription—on top of the seat price—based on number of users supported (\$5k-50k depending on size)
  - **Per-User:** \$1.3k-1.6k, scales down as volume increases
  - **100 Users:** ~\$150-165K annually

Sources: Company Website, G2 Software, Contrary Capital

### Optional Features: Per Add-On

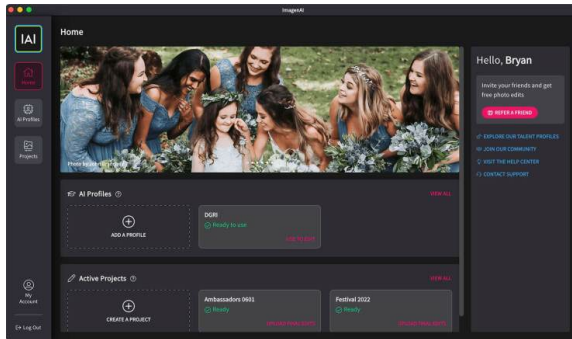
- Gong's platform offers add-on features to unlock greater levels of visibility across customer interactions.
- **Per-Feature Pricing:** varies by offering
- **Training and Services:** additional costs for reinforcement training can range up to \$90k per year, based on Forrester customer interviews

## Application: Horizontal – Images

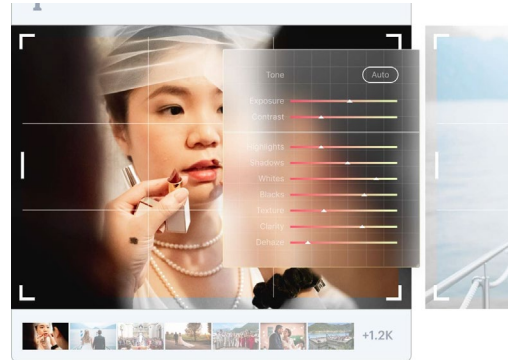


**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Freemium, Usage-Based Subscription, Flat-Fee Subscription



AI-powered platform that aims to simplify and enhance the photo editing process for photographers and content creators by learning personal editing styles over time, resulting in more accurate editing



The Enterprise Plan is a tailor-made solution for larger teams that has all the core functionality of Imagen AI but also has additional features and benefits around customization, scalability, compliance, and integration

### Freemium and Usage-Based

- Free plan with 1,000 free edits to allow for plan discovery
- **Pay-as-You-Go** – \$0.05/photo
  - Includes additional AI tools, culling, and back up
  - Billed monthly per usage
- **Annual Commitment** – 18,000+ photos/year
  - Keeps monthly costs stable
  - Same features as the pay-as-you-go plan

### Flat-Fee Subscriptions

- **Enterprise** – 500,000+ photos/year
  - \$29+/month fixed charge
  - Tailored, high-volume plan
  - Also includes custom adaptation, dedicated account managers, and personalization capabilities

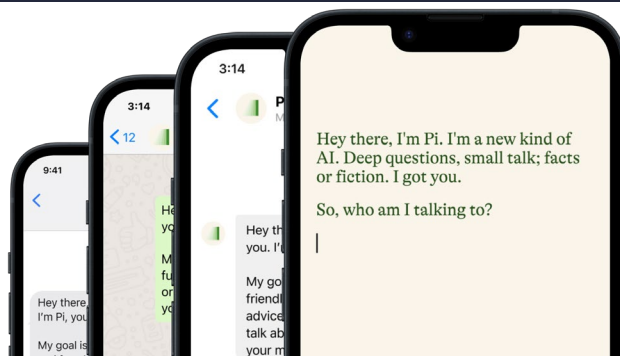
Sources: Company Website, TechCrunch

## Application: Horizontal – Speech

# Inflection

**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Pre-Commercialization



Talk to Pi,  
your personal AI

heypi.com

Inflection is an AI-first company that redefines human-computer interaction by enabling users to speak to a computer in a conversational language. Pi is a personal chatbot designed to assist users in various tasks and provide a seamless, human-like interaction experience and can understand user queries, offer personalized recommendations, and engage in meaningful conversations.

### Pre-Monetization

- Pi is free for all users today. Inflection envisions moving to a subscription model in the future

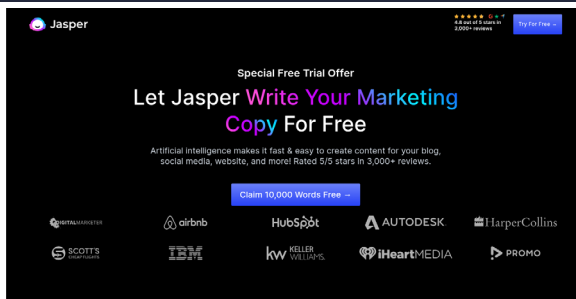
Source: Company Website

## Application: Horizontal – Text



**Estimated Revenue Stage:** Growth

**Current Commercialization Model:** Tiered Monthly/Annual User-Based Pricing + Usage-Based, Partnerships/Integrations



Jasper.ai enables businesses to create different forms of high-quality text and image content with the help of advanced AI algorithms. The software uses NLP to generate content that is optimized for search engines, engaging, and easy to understand. Jasper operates on a mixed seat- and usage-based pricing model with multiple tiers. Pricing scales up with the number of words generated.

Jasper's Partner Program is designed for third-party software companies and consultants that wish to integrate and/or represent Jasper to their customers. Jasper's AI Engine is offered via API and is also accessible as a web browser extension.

### User + Usage-Based

- **Creator Package:** Freelancers, marketers, and entrepreneurs who are ready to apply Generative AI to automate their daily work in their personalized Brand Voice
  - \$39/month billed annually
  - Only single user
- **Teams Package:** Small marketing teams who need to generate, create, and repurpose content for multiple campaigns for business outcomes
  - \$99/month billed annually
  - 3 – 10 users total
- **Business:** Departments and organizations who need a secure, scalable, and customizable AI platform that meets technical needs and aligns your team
  - Custom pricing or ~10k+/year
  - >10 users
  - Includes custom AI templates and onboarding and support from Jasper

### API/Partner Program – Usage-Based

- Jasper offers discounts and commissions to affiliate partners who refer and/or market Jasper alongside their products and services

Sources: Company Website, Contrary Capital

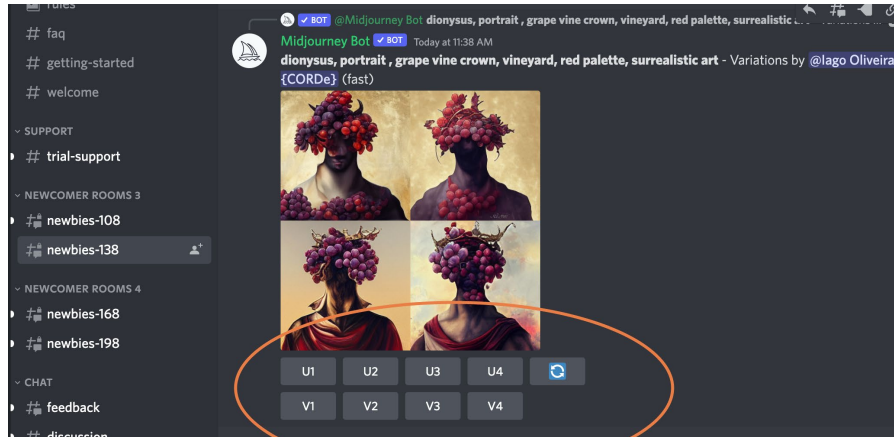
## Application: Horizontal – Images



**Midjourney**

**Estimated Revenue Stage:** Scaled

**Current Commercialization Model:** Freemium, Flat-Fee Subscription Tiers, Feature Add-Ons, Partnerships/Integrations



Midjourney is a self-funded, independent AI art generator that turns text-based prompts into images run on Discord.

### Freemium Tiering

- Midjourney has three subscription tiers and can be paid month-to-month or for the entire year for a 20% discount. Each subscription plan includes access to the Midjourney member gallery, the official Discord, and general commercial usage terms. Additional GPU time can be purchased for \$4/hr
  1. **Free Trial**
    - Fast GPU Time – 0.4 hr/lifetime
  2. **Basic Plan** – \$10/month
    - Fast GPU Time – 3.3 hr/lifetime
  3. **Standard Plan** – \$30/month
    - Fast GPU Time – 15 hr/lifetime
    - Unlimited Relax GPU Time
  4. **Pro Plan** – \$60/month (companies >\$1M revenue)
    - Fast GPU Time – 30 hr/lifetime
    - Unlimited Relax GPU Time


Source: Company Website

## Application: Horizontal – Customer Experience



**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** User-Based Subscription, Partnerships/Integrations



Enterprise co-pilot platform that integrates with all business systems to facilitate enterprise search, automate HR requests and resolve IT issues. Moveworks has hundreds of out-of-the-box use cases, as well as developer tools to help create completely new use cases if necessary.

Moveworks is used by >100 of the Fortune500, and has helped IT teams eliminate 75% of their help desk's workload.

### User-Based

- Pricing information taken from the AWS Marketplace is tiered by number of users:
  - 1,000-2,500 users: \$150 per user/year
  - 2,500-5,000 users: \$108 per user/year
  - 5,000-10,000 users: \$73 per user/year
  - 10,000-15,000 users: \$47 per user/year
  - 15,000-20,000 users: \$40 per user/year
  - 20,000-25,000 users: \$37 per user/year

Sources: Company Website, AWS Marketplace

## Application: Horizontal – Search



**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Freemium, User-Based Subscription, Feature Add-Ons



Subscription-based, ad-free, and privacy-focused search engine that enables individuals to search the internet without ads. Individual privacy is protected by personal data not being disclosed to advertisers. Snowflake recently acquired Neeva in May 2023.

### Pre-Acquisition User-Based

- **Basic Tier:** offered free, search questions are replied to with an answer generated from relevant source websites that are displayed and linked to below the answer
- **Premium Tier:** \$4.95/month, comes with additional features including a VPN and password manager

### Post-Acquisition Integration

- Snowflake announced it plans to “infuse and leverage” Neeva’s generative AI solution across its Data Cloud product, bringing search and conversation in Snowflake to a new level for customers

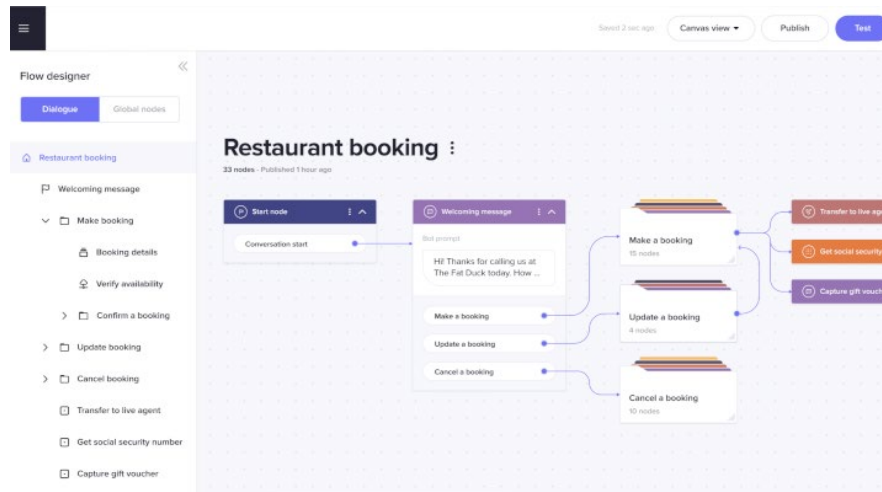
Sources: Company Website, The Verge, TechCrunch

## Application: Horizontal – Speech



**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Usage-Based Subscription, Flat-Fee Subscription Tiers, Non-Recurring Services



Developer of customer-led voice assistants enabling customer service teams to serve customers with real-time voice assistants. Customers can talk to the assistants within the flow of conversation and the assistant will serve customer requests and questions naturally.

### Flat-Fee and Usage-Based

- **Upfront Fee:** based on call complexity and requirements agreed to during an initial discovery period
- **Per-Minute Pricing:** tiered pricing available based on additional volume and committed use
- **Maintenance and Support:** no additional fees for tuning, improvements, maintenance, and 24/7 enterprise-grade support

Sources: Company Website, TechCrunch, Gartner



## Application: Horizontal – Voice Transcription



**Reported Revenue:** Growth

**Current Commercialization Model:** Freemium, User-Based Subscription, Partnerships/Integrations



AI-powered voice transcription and analytics platform that helps businesses streamline their voice communication. Software can transcribe and record meetings enabling teams to perform communication analysis down to the specific keyword or phrase.

### User-Based

- **RevComm Phone (MiiTel) – \$69 /month**
  - Business telephone system that integrates with your CRM and automatically transcribes conversations that can then be later analyzed and used to provide real-time coaching
- **RevComm Meetings – \$29 per user /month**
  - Link meetings to RevComm, meetings will automatically be recorded and stored in RevComm Analytics. Notes and analysis can be reflected and imported to Salesforce

Sources: Company Website, Forbes

Application: Horizontal – Video



**Estimated Revenue Stage:** Scaled

**Current Commercialization Model:** Freemium, User-Based Subscription, Non-Recurring Services, Partnerships/Integrations



The free Starter tier limits the resolution, cloud storage, and the number of projects a user can work with and limited access to the AI Magic Tools.

Runway is a full-fledged web-based editor that allows for real-time collaboration and offers a growing suite of AI tools developed through its research arm, Runway Research.


Free	User-Based
<ul style="list-style-type: none"><li>125 credits without option to buy more</li><li>Watermarked product</li><li>Generated videos up to 4 seconds</li><li>5GB assets</li><li>3 video projects</li></ul>	<ul style="list-style-type: none"><li><b>Standard</b> – \$12 per user/month<ul style="list-style-type: none"><li>Max 5 users per workspace</li><li>625 credits/month</li><li>Generate videos up to 15 seconds and 100GB assets</li><li>Additional AI Magic Tools and custom AI generators</li></ul></li><li><b>Pro</b> – \$28 per user/month<ul style="list-style-type: none"><li>Max 10 users per workspace</li><li>2,250 credits/month</li><li>Generate videos up to 15 seconds and 500GB assets</li><li>All AI Magic Tools and custom AI generators</li></ul></li><li><b>Enterprise:</b> Available for larger teams and organizations with custom pricing</li></ul>

Sources: Company Website, Forbes, Contrary Capital

Application: Horizontal – Video



**Estimated Revenue Stage:** Early-Mid  
**Current Commercialization Model:** User-Based Subscription, Non-Recurring Services, Partnerships/Integrations



Synthesia offers an AI video creation platform, with AI avatars, voice cloning, and text-to-speech capabilities. Typical use cases include training or product demo videos.

Synthesia maintains partnerships across a range of tech, service provider, and affiliate groups.

User-Based
<ul style="list-style-type: none"><li>▪ <b>Personal</b><ul style="list-style-type: none"><li>– \$17.25/month for 1 seat and 10x video credits /month</li></ul></li><li>▪ <b>Enterprise</b><ul style="list-style-type: none"><li>– Custom number of seats and unlimited video credits each month</li><li>– Includes collaboration features and priority support</li></ul></li></ul>

Partnership-Related Fees
<ul style="list-style-type: none"><li>▪ <b>Technology Partners</b><ul style="list-style-type: none"><li>– Integrate Synthesia videos into partners' applications (i.e. add Synthesia videos to a Shopify page)</li></ul></li><li>▪ <b>Service Providers</b><ul style="list-style-type: none"><li>– Offering features for Synthesia customers such as video course creation, shooting customer avatars, and translation</li></ul></li><li>▪ <b>Affiliate Partners</b><ul style="list-style-type: none"><li>– 20% commission on personal plan payments in first 12 months via an affiliate link</li></ul></li></ul>

Sources: Company Website, SaaSwothy


## Application: Horizontal – Presentations



**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Freemium, User-Based Subscription

**Future Monetization Opportunities:** Feature Expansion, API Integrations/Marketplaces



Built for storytellers, Tome is the world's first generative storytelling format to truly harness the power of AI — enabling anyone to tell a compelling story. Tome combines a fluid, interactive canvas with an AI storytelling partner so users can create and share immersive narratives that look great on any device.

### Freemium + User-Based

- **Free:** starter plan that allows unlimited creation activity for individual users, 500 AI credits upon sign-up. Limited creation for 2+ users
- **Pro:** \$8/user/month billed annually (or \$10/user if billed monthly); mid-tier pricing for heavy creative users and small to mid-sized teams.
- **Enterprise:** custom pricing for large, enterprise teams that need advanced controls. Includes unlimited creation and compute credits for users

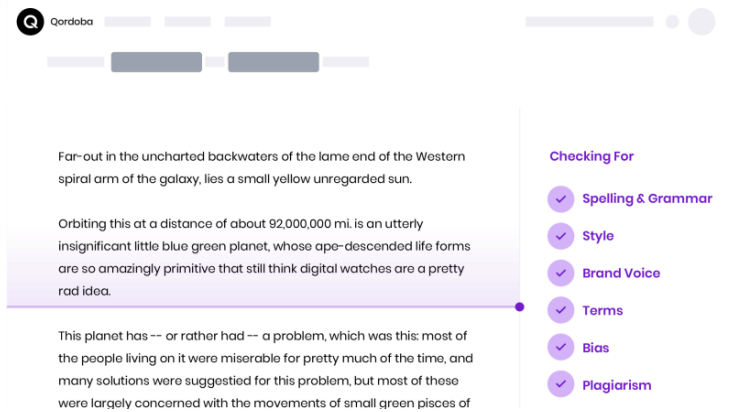
Sources: Company Website, SaaSworthy

## Applications: Horizontal – Text Generation

### WRITER

**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** User-Based Subscription, Feature Add-Ons



Writer Platform is a content generation platform that leverages their Palmyra family of LLMs to generate content in the same look and feel of your enterprise. Platform also enables users to create custom templates and style guides to maintain consistency across the company.

#### User-Based

- **Team (1 – 5 people)** – \$18 per user /month with 14-day free trial including:
  - Automated content generation
  - API access
  - Writing revision (Rewrite)
  - Draft writing (CoWrite)
  - Term management and Style guides
- **Enterprise** – Flexible pricing plan including:
  - Everything in Team, plus:
    - Dedicated support team
    - Boosted API access
    - Custom language model training
    - Multiple teams and style guides under a single organization

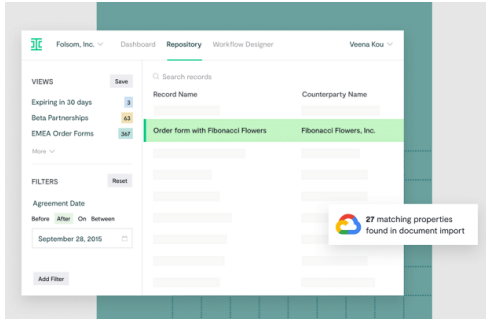
Sources: Company Website, SaaSwothy

## Application: Horizontal – Contract Management



**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** User-Based Subscription, Partnerships/Integrations, Feature Add-Ons



Ironclad covers the entire contract lifecycle management (CLM) process and their core CLM product includes several modules: (1) workflow designer, (2) editor, (3) repository, and (4) reporting. Additional products include ClickWrap and an AI layer that automatically extracts contract data and is capable of reviewing and negotiating contracts with clicks and replacement features.

### User-Based

- 30-day free trial
- Pricing varies based on the type of products/features, specific use cases, and the number of users with a starting price of \$500 per user/month

Sources: Company Website, Contrary Capital



Zip has a strategic integration with Ironclad which enables employees to initiate a spending request through Zip that is then automatically routed, along with the accompanying contract, to Ironclad for legal review and approval.

### Strategic Partnership with Zip

- Likely receive a take-rate from inbound sales generated through Zip

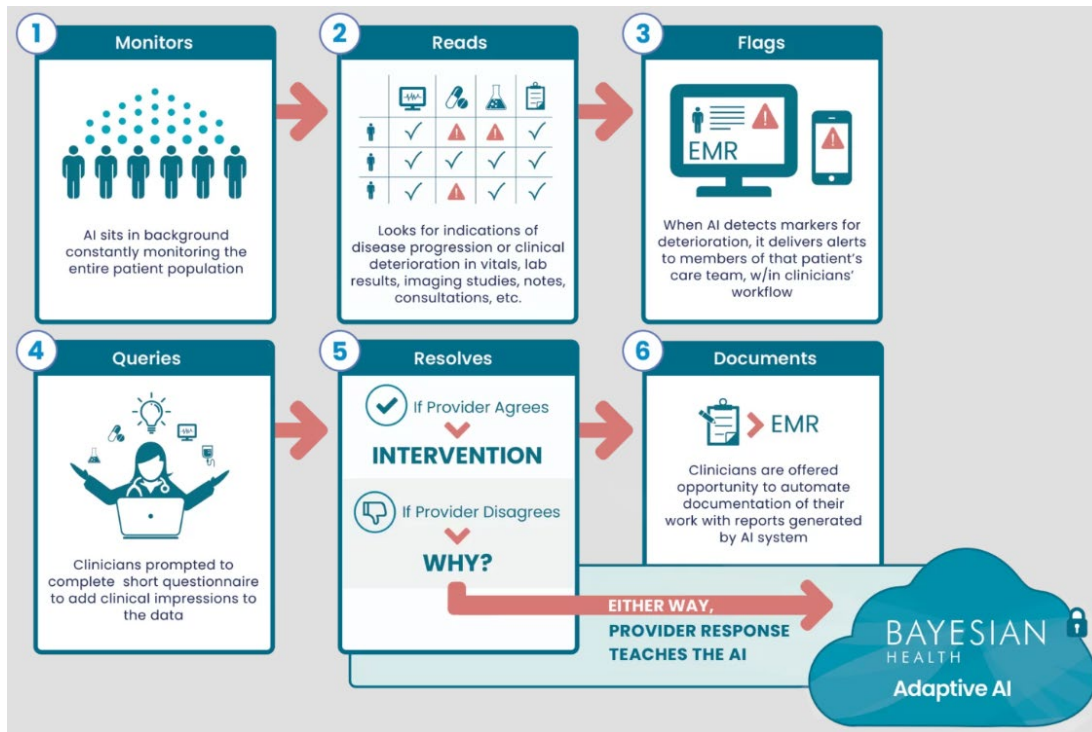
## Application: Vertical

## Application: Vertical – Healthcare

BAYESIAN  
HEALTH

Estimated Revenue Stage: Early-Mid

Current Commercialization Model: Trial/Early Commercialization



Bayesian Health's Targeted AI system integrates within a hospital's EMR system where it analyzes patient data, sending accurate and actionable clinical signals within existing workflows to empower physicians and care team members to catch life-threatening complications early. Early use cases include building models for sepsis detection, patient deterioration, and bedsores.

## Model Commercialization

- Bayesian's algorithms are in early stages of commercialization. The company will first go-to-market with its sepsis detection tool, and is also developing models for clinical deterioration, transitions of care, and pressure injuries

Sources: Company Website

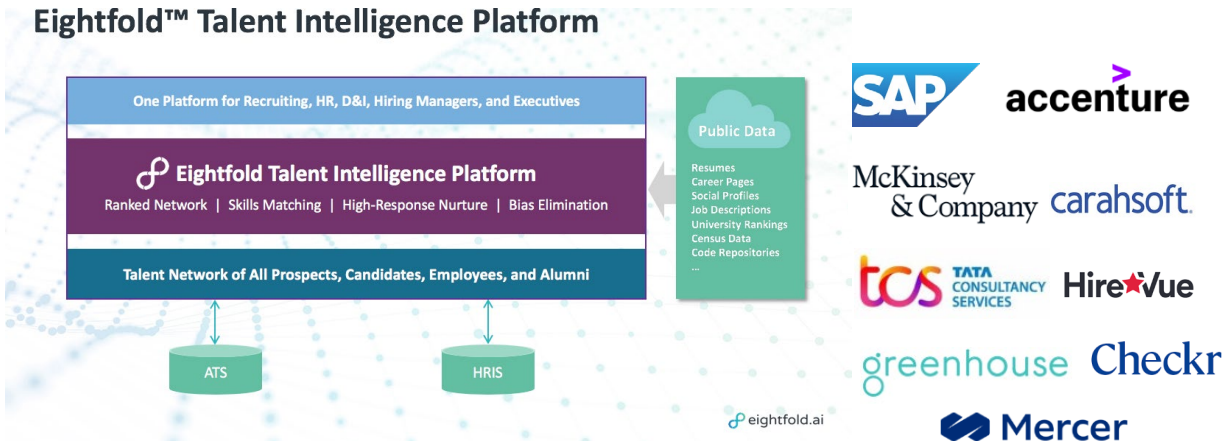
## Application: Vertical – HR



**Estimated Revenue Stage:** Scaled

**Current Commercialization Model:** User-Based Subscription, Non-Recurring Services, Partnerships/Integrations

### Eightfold™ Talent Intelligence Platform



Powered by AI, the Eightfold Talent Intelligence Platform empowers organizations to recruit and retain a diverse global workforce and provides applications for enterprises to build on their talent's existing skills.

Eightfold's professional services team helps organizations launch impactful implementations, drive specific business outcomes, and train and educate the HR workforce on how to leverage the platform.

Eightfold's platform is also made available to customers through a variety of professional service and HR software providers who act as channel partners.

#### User-Based Licenses

- **Talent Intelligence Platform:**
  - **Starter Edition:** fixed fee of \$25k per quarter, pricing based on number of users

#### Implementation Services

- Optional add-on at installation, custom pricing

#### Partner Sales

- Platform licenses sold through partners directly or through marketplaces

Sources: Company Website, Capterra, G2 Software




## Application: Vertical – Legal

Harvey.

**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Trial/Early Commercialization, Partnerships/Integrations



Harvey helps generate insights, recommendations, and predictions based on large volumes of data enabling lawyers to deliver more cost-effective solutions to their clients. Harvey was the first company to be built on top of Chat GPT-4. More than 15,000 law firms are on the waiting list currently.

### Partnerships

- Harvey currently works side-by-side with legal and professional service providers to customize AI solutions for each clients' needs while training its models on a vast volume of proprietary data. Select partner agreements include:
  - **PwC:** Harvey will give PwC's professionals across 100+ countries access to leading generative AI technology. This will enhance the ability of PwC's network of more than 4,000 legal professionals to deliver human led and technology enabled legal solutions in a range of areas, including contract analysis, regulatory compliance, claims management, due diligence, and broader legal advisory and legal consulting services
  - **Allen & Overy (A&O):** Harvey will empower more than 3,500 of A&O's lawyers across 43 offices operating in multiple languages with the ability to generate and access legal content with unmatched efficiency, quality and intelligence

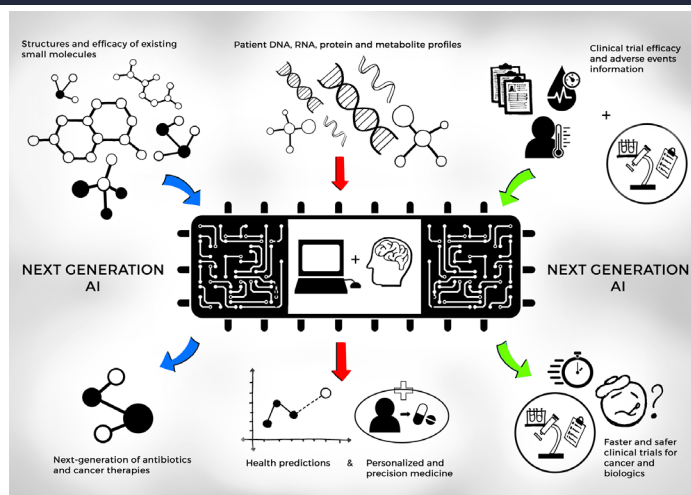
Sources: Company Website, Clio.com, TechCrunch, MartinDale

## Application: Vertical – Healthcare

insitro

**Estimated Revenue Stage:** Scaled

**Current Commercialization Model:** Trial/Early Commercialization, Partnerships/Integrations



Insitro brings an engineering approach to biological and chemical data generation, embracing advanced technologies such as robotics, CRISPR, stem cells, and cellular engineering. This yields cell-based disease models that are accurate reflections of the disease phenotype, allowing for rapid yet predictive assays of potential therapeutics, mutations, and other aspects of the disease biology.

### Discovery and Development Contracts

- The company has entered partnerships with major Pharma players including Gilead Sciences and Bristol-Myers Squibb to discover new biological targets for drug treatments, and the drugs themselves
  - **Gilead:** Insitro will receive an upfront payment of \$15M with additional near-term payments up to \$35M based on operational milestones. Insitro will be eligible to receive up to \$200M for the achievement of preclinical, development, regulatory and commercial milestones for each of the five Gilead targets; and up to low double-digit tiered royalties on net sales
  - **Bristol-Myers Squibb:** Insitro will receive \$50M in an upfront payment and will be eligible to receive an additional \$20M in near term operational milestones and up to an aggregate of more than \$2B in discovery, development, regulatory and commercial milestones in addition to royalty payments on net product sales

Sources: Company Website, Press Releases, FierceBiotech

## Application: Vertical – Energy







Pachama

**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Usage-Based Subscription, Non-Recurring Services, Partnerships/Integrations

### Solutions

- 
**Our evaluation process**  
 How we measure forests
- 
**Explore projects**  
 View our global tech-verified projects
- 
**The Pachama® Marketplace**  
 Offset on behalf of your organization
- 
**Make a purchase**  
 Buy from our portfolio of high-quality forest projects





Pachama trains ML models using satellite imagery, Lidar, radar, and field plots to estimate the carbon stored in forests today and into the future. Using algorithms combined with satellite imaging, Pachama monitors a project, figuring out the amount of carbon concentration in the project area in comparison to the region, providing a quality of assurance that claims to offset carbon are true or not.

### Project + Usage-Based

- Upon purchase of a carbon offset project, 80% of the transaction goes to the project. Pachama takes a 20% fee to verify and monitor the project, run the platform, and provide customer support
- The prices in the marketplace range anywhere from between \$5 – 40 per metric ton, depending on location and the type of project. In addition, the cost of Pachama's portfolio is constant, no matter what project a buyer purchases.

Sources: Company Website, Contrary Capital

### The Exchange API will empower you to:

- 
**Easily incorporate sustainability**  
 Seamlessly link with your product or service. Manually tracking emissions and calculating carbon credits is a thing of the past.
- 
**Purchase with trust and transparency**  
 Select from carbon sequestration and conservation projects accredited by top registration bodies and verified by Pachama's technology.
- 
**Access data analytics on your impact**  
 Access the complete history of purchases through Pachama's API, including information about the quantity of credits retired over time.
- 
**Explore and select projects**  
 Sort, analyze and select forest carbon projects based on attributes including location, biomass, and community involvement.

Pachama is building an API so other companies can easily integrate their marketplace into their product or service. It enables companies to offer their customers and partners vetted carbon credits from forest projects verified by Pachama's remote sensing and ML technology.

### API Integrations

- **Access Charges:** Pachama's API will enable companies to offer their customers and partners quality credits from forest projects verified by Pachama's remote sensing and ML technology
  - Could be offered as fixed-fees (monthly) or as a consumption-based model

Application: Vertical – Healthcare



**Estimated Revenue Stage:** Growth  
**Current Commercialization Model:** Trial/Early Commercialization, Partnerships/Integrations, Non-Recurring Services



AI-powered technology to improve pathology and patient outcomes. The software can be used at every stage of the drug development process, from biomarker and drug discovery to clinical development and commercialization. PathAI's technology is used by 90% of the world's top 15 biopharmaceutical companies.

PathAI also sells into clinical and commercial labs, where its AISight Dx pathology platform increases diagnostic accuracy, improves efficiency, and therefore improves patient outcomes. The software is used by clinical labs across the US. The company also maintains a lab services business previously acquired.

**BioPharma Partners**

- PathAI works with a range of tech partners to develop solutions tailored to BioPharma clients
  - Diagnostic and Research Services
  - Clinical Development Services
  - Diagnostic Development and Commercialization
  - BioPharma Lab Services

**Provider/Lab Partners**


- PathAI Diagnostics
- Lab Diagnostic Research and Clinical Studies
- Diagnostic Assays

Sources: Company Website, SourceForge

Application: Vertical – Space



**Estimated Revenue Stage:** Early-Mid  
**Current Commercialization Model:** MSA / Custom Contracts, Freemium, Feature Add-Ons

Vantage	Beacon	Laboratory	Seradata 
The world's only day/night LEO-to-GEO optical satellite tracking and monitoring service.	The world's first collision avoidance coordination and communications platform for space.	Next-generation space simulation and training environment to learn the fundamentals of astrodynamics and space domain awareness.	Web-based launch and satellite analytical system to provide research and information or leading players in the space industry.
<b>MSAs / Custom</b>	<b>Free</b>	<b>MSAs / Custom</b>	<b>MSAs / Custom</b>
<ul style="list-style-type: none"><li>Data-as-a-service offering powered by Slingshot's Global Sensor Network</li></ul>	<ul style="list-style-type: none"><li>Free version of its space traffic control software to satellite operators worldwide</li></ul>	<ul style="list-style-type: none"><li>Awarded a \$25M, 39-month contract by the US Space Force to develop Slingshot Laboratory</li><li>Custom pricing per deal</li></ul>	<ul style="list-style-type: none"><li>Custom pricing</li></ul>

Sources: Company Website, SpaceNews, Datarade

## Application: Vertical – Retail



**Estimated Revenue Stage:** Growth

**Current Commercialization Model:** Flat-Fee Subscription Tiers, Partnerships/Integrations

**EasyOut®**

**EasyStock**

Trigo retrofits existing grocery stores with hardware including sensors, transforming them into “smart” stores. EasyOut tracks shoppers as they walk through the store, sees the items that they pick up and then allows shoppers to pay by app without stopping at a checkout. EasyStock manages inventory with real-time on-shelf availability, applying its proprietary algorithm to data from its in-store sensors.



Trigo joined the Google Cloud Partner ecosystem in June 2021 to leverage Google Cloud's relationships with leading retailers. Trigo works alongside the Google team (and others) to support retailers with digital transformations.

### Flat-Fee Subscriptions

- Supermarket operators pay a monthly fee which covers both software and hardware required for the store

### Partnerships / Co-Sales

- Work with global system integrators and large consulting firms to expand GTM reach
  - **Google:** Trigo joined the Google Cloud Partner ecosystem in June 2021 to leverage Google Cloud's relationships with leading retailers. Trigo works alongside the Google team (and others) to support retailers with digital transformations


Sources: Company Website, TechCrunch, Google

## Application: Vertical – Healthcare



**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Trial/Early Commercialization, Partnerships/Integrations



Unlearn develops digital twins for randomized control trials, helping to accelerate drug development by prediction outcomes for each patient if they were to receive the placebo or standard of care. Participants' digital twin outcomes are compared to actual outcomes to calculate an unbiased estimate of the treatment effect. Gives higher power of testing with a smaller control group, helping to encourage patients to take part as there is a smaller chance that they receive the placebo treatment.

In September '22, the EMA released its final favorable clinical qualification opinion providing a regulatory application for the TwinRCT solution developed by Unlearn AI.

### Multi-Year Pharmaceutical Partnerships

- Unlearn has signed a multi-year partnership with Merck. Unlearn will receive milestone payments tied to the progress of Merck drugs developed using its digital twin technology

Sources: Company Website, TechCrunch

## Application: Vertical – Defense



**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** MSA / Custom Contracts

**Future Monetization Opportunities:** New Product Development, Commercial Market Entry

### VANNEVAR Decrypt

A foreign text workflow platform  
built for national security.

Vannevar's flagship product, Decrypt, focuses on foreign language NLP, organizing overseas data and resources that are collected by the intelligence community, and then immediately translating and interpreting those documents for foreign policy decisionmakers.

#### MSAs / Federal Contracts

- **Agency/Military Sales:** Decrypt is sold into the DoD and IC agencies via direct R&D contracts (usually multi-year with extensions). Vannevar begins new federal relationships by launching pilot programs, deploying into federal customer tech stacks to solve specific use cases. The product is currently used at more than 15 military bases for tasks such as processing battlefield information and countering malign activity by adversaries

Sources: Company Website, TechCrunch, HigherGov

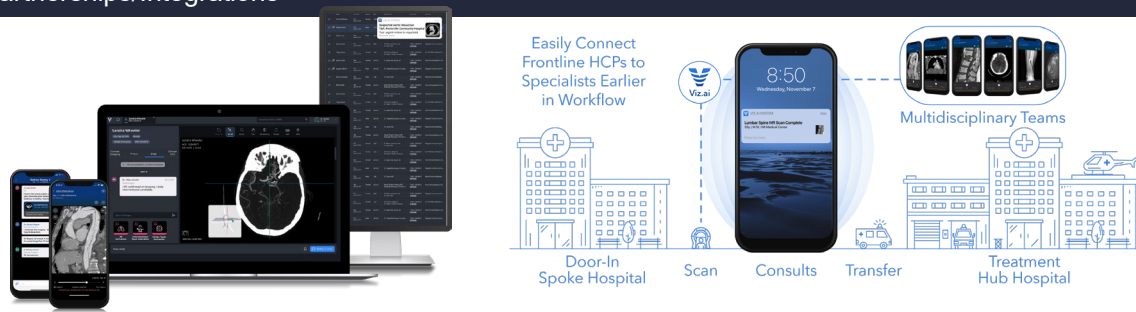


## Application: Vertical – Healthcare



**Estimated Revenue Stage:** Scaled

**Current Commercialization Model:** Flat-Fee Subscription Tiers, Feature Add-Ons, Partnerships/Integrations



The Viz Platform uses AI to connect care teams earlier, increasing the speed of diagnosis and care, and optimizing clinical care pathways across hospitals and health systems to ensure the right patient gets to the right specialist at the right time. Current solutions are focused on neurovascular, vascular, and cardiology diseases.

### Flat-Fee Subscriptions

- Platform subscription pricing dependent on features and use cases. Average annual costs reported to be ~\$25k/year/hospital and higher
  - AI-Powered Care Coordination
  - AI-Powered Workflows
  - Life Sciences
- Currently Viz LVO is eligible for reimbursement up to ~\$1.0k per patient

### Partnerships

- Viz.AI also partners with MedTech manufacturers to expand their market reach. Examples include:
  - **Medtronic:** Viz is priced into Medtronic's Solitaire offering as a subscription. The world's largest device maker sells Viz AI's stroke identification software, which uses AI to spot suspected large vessel occlusion strokes on CT scans, and automatically notify neurologists and surgeons.

Sources: Company Website, Out of Pocket Health

## Application: Vertical – Auto



**Estimated Revenue Stage:** Growth

**Current Commercialization Model:** Trial/Early Commercialization, Partnerships/Integrations



Developing the Waabi Driver, an AI-first autonomy stack combined with sensors and compute as hardware, which serves as an operating system for self-driving trucks. The Waabi Driver has been developed by Waabi World -> instead of testing on roads with real-world data, the company uses AI to improve the model by exposing it to common situations and safety-critical edge cases. On-road driving is only used for validation and verification.

Waabi Driver will be integrated directly with OEMs on their existing assembly lines without disrupting the manufacturing process.

### OEM Partnerships

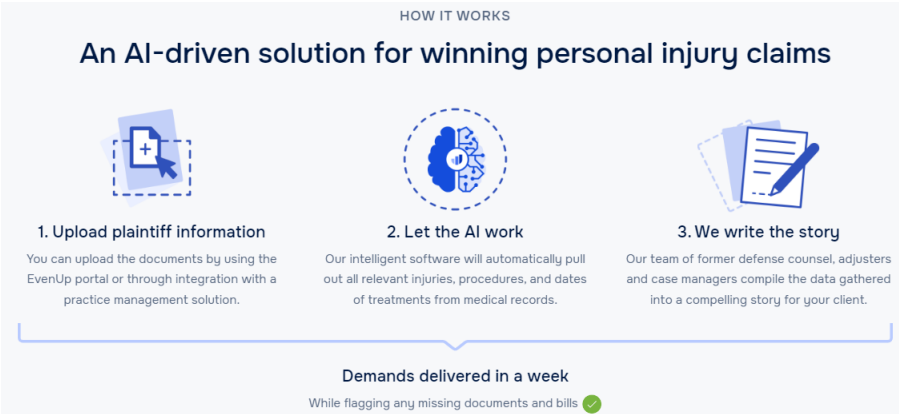
- First partnership thought to be with Peterbilt (although undisclosed). The trucks will be used initially for data collection, testing and commercial pilots
- Ultimately the Waabi Driver will be integrated into vehicles at the factory level

Sources: Company Website, TechCrunch

Applications: Vertical – Legal



**Estimated Revenue Stage:** Early-Mid  
**Current Commercialization Model:** Flat-Fee Subscription Tier, Partnerships/Integrations



Powered by Ms of records, EvenUp turns medical docs and case files into AI-driven demand packages for injury lawyers.

**Flat-Fee Subscriptions**

- **Flat Fee Subscription:** starts at \$499

**Partner Sales and Integrations**

- **Partner Marketplaces/Stores:** EvenUp subscriptions are offered through several adjacent, industry-focused partners such as Clio
- **Integrations:** EvenUp integrates with third-party legal solutions such as Litify, the leading legal practice management software, expanding access points to its platform

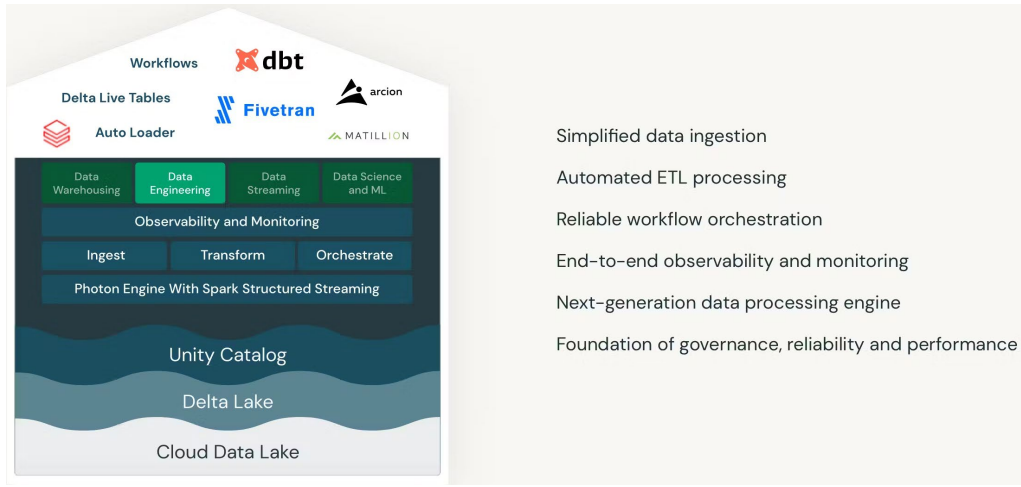
Sources: Company Website, Reuters

Enterprise Infrastructure

Enterprise Infrastructure: Data



**Estimated Revenue Stage:** Scaled  
**Current Commercialization Model:** Freemium, Usage-Based Pricing, Partnerships/Integrations, Services, Feature Add-Ons, Non-Recurring Services



Databricks Lakehouse Platform is a unified set of tools for building, deploying, sharing, and maintaining enterprise-grade data solutions at scale on all the major cloud providers and built on Apache Spark. Pricing is primarily based on compute usage. Storage, networking, and related costs will vary depending on the services you choose and your cloud service provider, but typically has a standard, premium, and enterprise tier.

Freemium	Usage-Based	Partnerships/Integrations	Services
<ul style="list-style-type: none"><li>▪ <b>Community Edition:</b> ability to share and host notebooks for free without incurring AWS costs</li></ul>	<ul style="list-style-type: none"><li>▪ <b>Consumption-Based:</b> pricing based off a Databrick unit (DBU), which is a normalized unit of processing power, billed monthly</li><li>▪ Workflows and Streaming, Data Warehousing, Data Science and ML</li></ul>	<ul style="list-style-type: none"><li>▪ <b>MSFT Azure:</b> fully managed first-party service that enables an open data lakehouse in Azure</li></ul>	<ul style="list-style-type: none"><li>▪ <b>Professional Services:</b> Includes expert support for: migrating / connecting to a data lake, building a lakehouse, or driving analytics</li></ul>

Sources: Company Website, TechCrunch, Contrary Capital

Enterprise Infrastructure: Data



**Estimated Revenue Stage:** Scaled  
**Current Commercialization Model:** Flat-Fee Subscription Tiers, Usage-Based Subscription, Feature Add-Ons, Partnerships/Integrations

**Fine-Tuning and RLHF**

Adapt best-in-class foundation models to your business and your specific data to build sustainable, successful AI programs and data from your enterprise.

**Foundation Models**

Scale partners or integrates with all of the leading AI models, from open-source to closed-source, including OpenAI, Anthropic, Google PaLM, Cohere, and more.

**Enterprise Data**

Scale's Data Engine enables you to integrate your enterprise data into the fold of these models, providing the base for long-term strategic differentiation.

Data annotation service enabling companies to collect, curate, and annotate data for their ML models. Also provides dataset management tools, testing, model evaluation, and model comparison tools.

Data fine-tuning service for enterprises looking to deploy generative AI models with enterprise data.

Platform for deployment of LLM applications enabling developers to build custom LLM models with enterprise data, deploy, and test them all in a single environment.

**Scale Data Engine**

- **Rapid**
  - \$.02-.13 cents per Image, annotation, video, text

**Scale Generative AI**

- **Rapid**
  - \$.02-.13 cents per Image, annotation, video, text

**Scale Spellbook**

- **Pay as you go** – Base model pricing per token with unlimited seats, fine tuning capabilities, and private apps
- **Enterprise** – Same as pay as you go with dedicated support and flexible deployment options

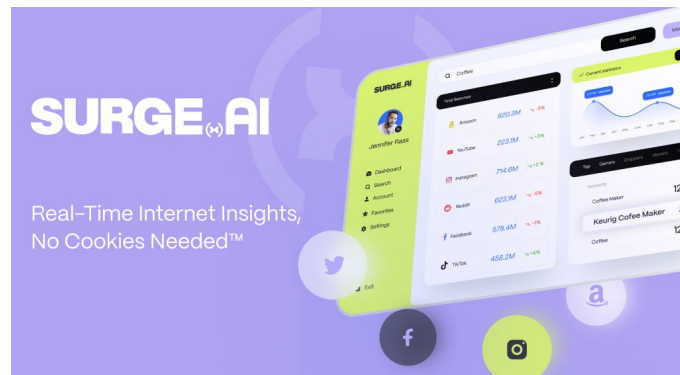
Sources: Company Website, CNBC

## Enterprise Infrastructure: Data Labeling



**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Usage-Based Subscription, Non-Recurring Services, Feature Add-Ons



Data-labeling platform and service to enable customers to build datasets. Use cases include Instruct and RLHF Training, Search Evaluation, Code Generation, Hate Speech Detection and auto-responses to Customer Queries. Data labeling is offered in 30+ languages, and uses custom data-labelling teams to ensure higher quality of data and better outcomes from ML models. Surge estimates that it offers 10x faster iteration, 2x data quality which in turn improves ML models by 50%.

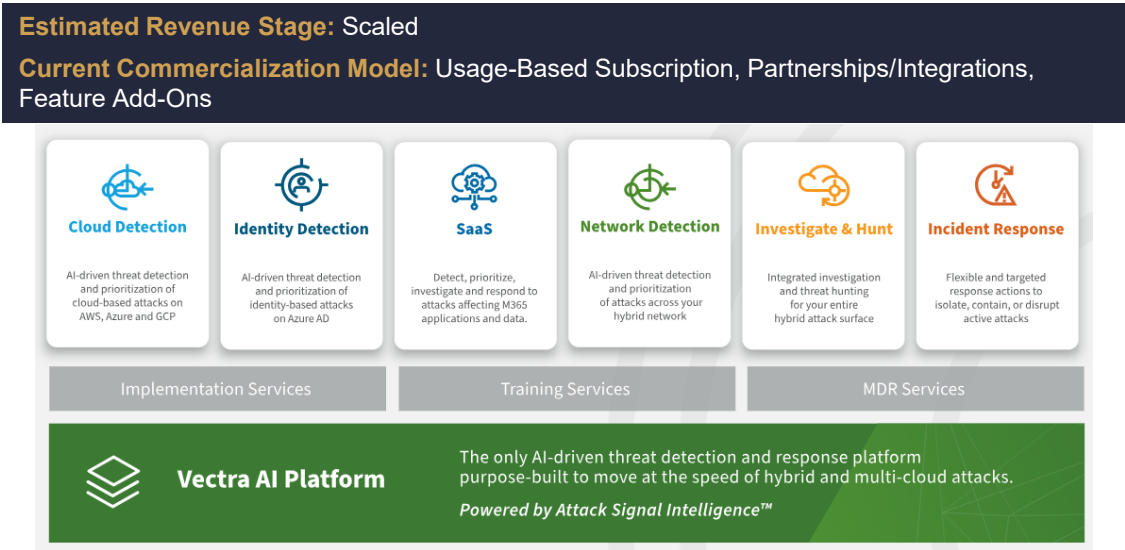
### Usage-Based

- Pricing is based on the size of projects, with two tiers, Individual and Enterprise:
- **Individual**
  - Free to sign up, and pay-per-response once a project is live
  - The individual sets the payment-per-response and how many workers should label each row of data to give an overall cost for the project
  - Surge recommends setting a price per task that equates to 30-40 cents per minute
- **Enterprise:** for teams that need high-volume, fully managed data labeling services with guaranteed SLAs
  - Projects are priced by Surge on a per-task basis
  - Two components of cost: requirements for labeling tools and the Surge professional services team

Sources: Company Website

Enterprise Infrastructure: Security

VECTRA®



Threat detection and response platform helping IT teams identify threats with AI and respond quickly to isolate the threat. Additionally, as the enterprise grows and changes the platform is modularized to enable additional product integrations as needed. Security analysts are able to track lateral threat movement and design a unique response to prevent the threat from spreading. Pricing is based on an IP License model with a fixed dollar amount for each active IP monitored.

Partnerships	Usage-Based Licensing
<ul style="list-style-type: none"><li>▪ <b>Resell Partners:</b> VARs, resellers, and distributors delivering detection and response security solutions based on the Vectra platform</li><li>▪ <b>Managed Service Partners:</b> managed security services partners leverage Vectra to provide their customers 24/7 advanced threat-detection and response services</li><li>▪ <b>Service Partners:</b> leverage Vectra platform via APIs that can be integrated into existing subscription security solutions</li><li>▪ <b>Technology Partners:</b> API tools that integrate the Vectra platform into clients' existing technology stacks; partners include AWS, Fortinet, Vmware, and Juniper Networks</li></ul>	<ul style="list-style-type: none"><li>▪ <b>IP Licensing:</b> fixed dollar amount for each active IP monitored</li></ul>

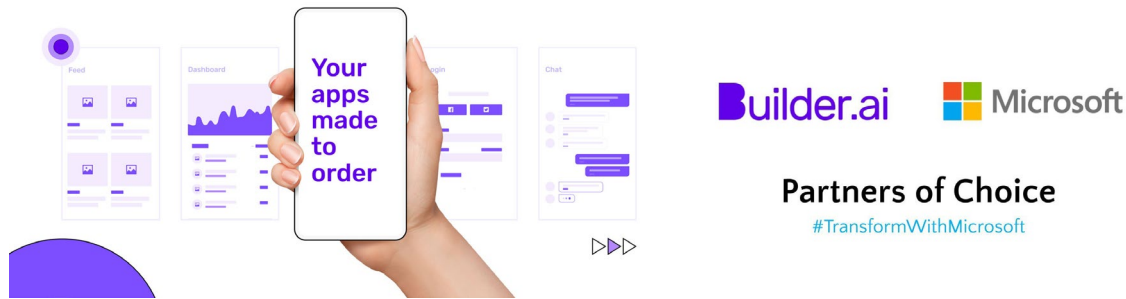
Sources: Company Website, AWS Marketplace

## Enterprise Infrastructure: App Development

# Builder.ai

**Estimated Revenue Stage:** Scaled

**Current Commercialization Model:** Flat-Fee Subscription Tiers, Usage-Based Subscription, Partnerships/Integrations, Non-Recurring Services



Builder.ai offers both small and larger business the ability to build a web or native mobile app in a modular fashion.

Web-Only	Web + App	Hosting and Support
<ul style="list-style-type: none"> <li>▪ <b>Starter plan</b> <ul style="list-style-type: none"> <li>– \$100/24 months</li> <li>– Take up to 100 orders a month on a personal responsive web store</li> </ul> </li> <li>▪ <b>Premium Plan</b> <ul style="list-style-type: none"> <li>– \$225/24 months</li> <li>– Make unlimited sales on a personal responsive web store</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>▪ <b>Pre-Packaged Apps</b> <ul style="list-style-type: none"> <li>– From \$200/24 months</li> <li>– Online starter kits for small businesses, delivered in two weeks</li> </ul> </li> <li>▪ <b>Custom Apps</b> <ul style="list-style-type: none"> <li>– Pay-per-feature, built with bespoke features built for specific business needs</li> </ul> </li> <li>▪ <b>Enterprise Apps</b> <ul style="list-style-type: none"> <li>– Pay-per-feature, Builder.ai handles all development</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>▪ <b>Builder Cloud</b> <ul style="list-style-type: none"> <li>– Usage-based pricing based on hosting</li> </ul> </li> <li>▪ <b>Studio One</b> <ul style="list-style-type: none"> <li>– Subscription support and maintenance services</li> </ul> </li> </ul>
Partnerships		
<ul style="list-style-type: none"> <li>▪ Builder.ai maintains a broad set of strategic partner relationships, expanding their GTM strategy through integrations and third-party marketplaces. Select partners include:</li> <li>▪ <b>Microsoft:</b> <ul style="list-style-type: none"> <li>– Builder's "Natasha" (AI product manager) is made available to users of Teams</li> <li>– Integrations across Azure OpenAI Service and other Azure Cognitive Services with Builder.ai's software assembly line and adoption of the Microsoft Cloud and AI</li> </ul> </li> <li>▪ <b>Etisalat UAE:</b> <ul style="list-style-type: none"> <li>– Gives customers access to Builder.ai's Studio Store product, including a range of pre-packaged apps designed to help SMBs take a digital first approach</li> </ul> </li> </ul>		

Sources: Company Website, TechCrunch



## Hardware / Application Vertical

## Hardware/Application: Vertical – Defense



**Estimated Revenue Stage:** Scaled

**Current Commercialization Model:** MSA / Custom Contracts, Unit-Based Sales/Lease, Partnerships/Integrations




Anduril sells a range of autonomous warfare hardware assets to the defense industry as the “land” part of its land-and-expand sales strategy (to get in the door with customers). While the LatticeOS AI software is the brain of Anduril’s product, the intelligence of the solution is only as effective as the quality of the hardware assets and sensor technologies collecting the data.

## Hardware Sales – Unit-Based

- Sentry towers
- Sensors
- Unmanned Aerial Systems
- Autonomous Underwater Vehicles

## Lattice: A combat OS

Anduril’s software system improves understanding, streamlines decision-making, and acts fast



Lattice autonomously parses data from thousands of sensors and data sources into an intelligent common operating picture in a single pane of glass. Lattice uses sensor fusion, computer vision, edge computing, and AI/ML to detect, track, and classify every object of interest in an operator’s vicinity.

## Software Integrations

- **MSAs:** often couple hardware sales with integration to Lattice platform. Lattice is also offered independently (hardware-agnostic)
  - **Example:** \$1B SOCOM contract to lead its counter-drone systems integration work, which includes developing prototypes for new counter-UAS tech


Sources: Company Website, The Generalist, Contrary Capital

## Hardware/Application: Vertical – Construction

# CANVAS

**Estimated Revenue Stage:** Early-Mid

**Current Commercialization Model:** Unit-Based Sales/Lease, Partnerships/Integrations



Canvas partners with contractors to re-envision the drywall trade by deploying semi-autonomous machines capable of completing drywall finishing (mudding and sanding) projects via remote-worker control. The robots use ML to install drywall.

### Hardware Leases – Unit-Based

- **Multi-Year Lease Agreements:** Canvas partners with drywall service providers through multi-year, multi-machine enterprise lease agreements that grant customers priority access to new machines as they are produced. Canvas management believes leasing is a strategic move in the current market environment in which owners are increasingly looking to contractors to reign in costs and boost productivity

### Strategic Partnership Sales

- Canvas has partnered with a range of robotics and materials producers (i.e. Hilti, USG) to outsource Canvas machine production and co-sell their machines alongside partners' related construction robotics technologies

Sources: Company Website, Construction Junkie

Hardware/Application: Vertical – AgTech



**Estimated Revenue Stage:** Early-Mid  
**Current Commercialization Model:** Unit-Based Sales/Lease, Non-Recurring Services



Vulcan is a new weeding implement that packs cutting-edge computer vision into a lightweight, open cultivator frame. It is powered by the Intelligent Plant Scanner, a high-performance, tightly integrated combination of camera, lighting, and computation elements. The implements are backed by sophisticated ML and actuation control software for reliable and consistent sub-inch weeding accuracy.

Titan was FarmWise’s initial product in market that resembles a tractor with an implement on the front, but is operated remotely via an iPad. As the Titan traces up to six rows at a time, it uses sensors and machine vision to distinguish weeds from crops. Its cultivation blades quickly respond by creating and executing a plan to snip the weeds within less than an inch of the crops without impacting their integrity.

Unit-Based Sales

- FarmWise’s two Vulcan implement models are currently offered for pre-order, and field demos began in March 2023. Machines will be sold on a per-unit basis beginning in Q3 2023 – pricing has yet to be publicly announced.

“Robot-as-a-service”


- Rather than selling the Titan to farms, FarmWise has operated them under a robot-as-a-service (RaaS) model, bringing them to farms for weed-whacking on an as-needed basis. FarmWise maintains its Titan fleet and moves machines around the country to match the growing cycles in various regions.

Sources: Company Website, TechCrunch

Hardware/Application: Vertical – Defense



**Estimated Revenue Stage:** Scaled  
**Current Commercialization Model:** MSA / Custom Contracts, Feature Add Ons, Unit-Based Sales/Lease



Shield AI's Hivemind software is an AI pilot for military and commercial aircraft that enables intelligent teams of aircraft to perform missions ranging from room clearance, to penetrating air defense systems, and dogfighting F-16s.



Shield AI also offers several autonomous flight vehicles powered by Hivemind software.

**Flat-Fee (Software)**

- **Hivemind:** Shield AI offers software-only contracts that enable customers to integrate Hivemind's AI pilot features into their own aircraft (F-16 and next-gen jets)

**Hardware + Software (MSAs)**

- Shield AI also offers its own specialized, autonomous vehicles powered by Hivemind
- **V-BAT:** vertical take-off and landing unmanned aircraft system capable of operating like a helicopter while undertaking forward travel as a fixed-wing aircraft
- **Nova 2:** mission-capable indoor drone to autonomously navigate complex subterranean and multi-story buildings

Sources: Company Website, Contrary Capital

## Disclaimer

This document has been prepared by Lazard Frères & Co. LLC ("Lazard") solely for general information purposes and is based on publicly available information which has not been independently verified by Lazard. The information contained herein is preliminary and should not be relied upon for any purpose. No liability whatsoever is accepted, and neither Lazard nor any member of the Lazard Group (being Lazard Ltd and its direct and indirect subsidiary and associated undertakings) nor any of their respective directors, partners, officers, employees, representatives or other agents is, or will be, making any warranty, representation or undertaking (expressed or implied) concerning the accuracy or truthfulness of any of the information, ideas, forecasts, projections or of any of the views or opinions contained in this document or any other written or oral statement provided in connection herewith or for any errors, omissions or misstatements contained herein or for any reliance that any party may seek to place upon any such information. Lazard undertakes no obligation to provide the recipient with access to any additional information or to update or correct any information contained herein. Interested parties should conduct their own investigation and analysis of the companies and information referenced herein. Lazard only acts for those entities whom it has identified as its client in a signed engagement letter and no-one else and will not be responsible to anyone other than such client for providing the protections afforded to clients of Lazard nor for providing advice. Nothing contained in this document constitutes, or should be relied upon as, (i) the giving of financial, investment or other advice or recommendations by, or the issuance of research by, Lazard, or (ii) a promise or representation as to any matter whether as to the past or the future. Recipients are recommended to seek their own financial and other advice and should entirely rely solely on their own judgment, review and analysis of this document. Lazard or other members of the Lazard Group (A) may have acted in the past, or act currently or in the future as adviser to some of the companies referenced herein, (B) may receive fees in connection with any such advisory engagements, (C) may at any time be in contact with such companies in order to solicit them to enter into advisory engagements and (D) may from time to time have made, and may in the future make, investments in such companies. Nothing contained in this document constitutes, or should be deemed to constitute, an offer or solicitation for the purchase or sale of any security. You undertake to keep this document confidential and to not distribute it to any third party, or excerpt from or reproduce this document (in whole or in part), without the prior written consent of Lazard.