

ARTIFICIAL INTELLIGENCE LANDSCAPE

PART III – Global Listed Companies

This report is part of a 4 part series on AI companies in the listed and unlisted space

January 2025



Table of Contents

S.No.	Company Name
1	Amazon
2	Google
3	SoundHound
4	IBM
5	Nvidia

Valuation

Company Name	Market Data			Financial Data - TTM Basis								Valuation				
	Price \$	Market Cap (\$ Mn)	EV (\$ Mn)	Sales (\$ Mn)	EBITDA (\$ Mn)	Margin %	Net Profit (\$ Mn)	Margin %	D/E (x)	ROCE %	ROE %	EV/Sales (x)	EV/EBITDA (x)	P/E (x)	PE/ROE (x)	
Amazon	222.1	23,72,748.23	24,43,232.23	6,20,128.00	1,11,583.00	18.0%	49,868.00	8.0%	0.6	26.7%	19.2%	3.9	21.9	47.6	2.5	
Google	195.4	24,44,582.91	23,80,641.91	3,39,859.00	1,23,470.00	36.3%	94,269.00	27.7%	0.1	36.0%	30.0%	7.0	19.3	25.9	0.9	
Oracle	163.1	4,63,807.02	5,52,530.02	54,933.00	22,336.00	40.7%	11,624.00	21.2%	7.2	19.6%	84.6%	10.1	24.7	39.9	0.5	
Nvidia	140.1	34,83,526.91	34,55,264.91	1,13,269.00	72,741.00	64.2%	63,074.00	55.7%	0.2	95.6%	95.7%	30.5	47.5	55.2	0.6	
Altair Engineering	110.2	9,512.85	9,258.95	639.50	71.70	11.2%	32.90	5.1%	0.3	6.5%	3.9%	14.5	129.1	289.1	73.6	
SoundHound AI	15.0	4,608.14	4,516.34	67.30	-88.40	-131.4%	-110.10	-	0.1	-26.0%	-37.2%	67.1	-51.1	-41.9	1.1	
IBM	223.2	2,08,316.21	2,54,822.21	61,867.00	11,951.00	19.3%	6,397.00	10.3%	2.5	14.1%	26.2%	4.1	21.3	32.6	1.2	
C3.ai	33.3	4,129.20	3,404.30	346.50	-301.40	-87.0%	-274.40	-	0.0	-34.9%	-31.9%	9.8	-11.3	-15.0	0.5	
Microsoft	424.6	31,71,887.76	31,90,296.76	2,54,190.00	1,36,552.00	53.7%	90,512.00	35.6%	0.3	35.5%	31.5%	12.6	23.4	35.0	1.1	
Meta	610.7	16,01,918.56	15,80,065.56	1,56,227.00	79,209.00	50.7%	55,539.00	35.6%	0.3	37.1%	33.8%	10.1	19.9	28.8	0.9	
Alibaba	83.7	2,07,182.96	1,99,885.66	1,37,078.30	24,894.80	18.2%	12,303.80	9.0%	0.2	13.5%	9.1%	1.5	8.0	16.8	1.9	
												Median	10.1	21.3	32.6	1.1
												Average	15.6	23.0	42.2	1.1



Amazon

Market Cap
\$ 23,72,748 Mn

Revenue
\$ 6,20,128 Mn

EBITDA
\$ 1,11,583 Mn

PAT Margin
\$ 49,868 Mn

As of Sept 2024

Product Offerings

[1] Amazon Q: Intelligent virtual assistant across all business functions and by employees of all levels and capabilities.

For Business- AI-powered assistant for employees to get answers to questions across business data—such as company policies, product information, business results, code base, employees, and many other topics—by connecting to enterprise data repositories and summarizing the data logically.

For Developers- Developers can use it for all their tasks like coding, testing, upgrading, troubleshooting, and security scanning. It helps in writing code faster by increasing productivity by 40%. With a 37% acceptance rate, which is the highest code acceptance rate among assistants who perform multiline code suggestions.

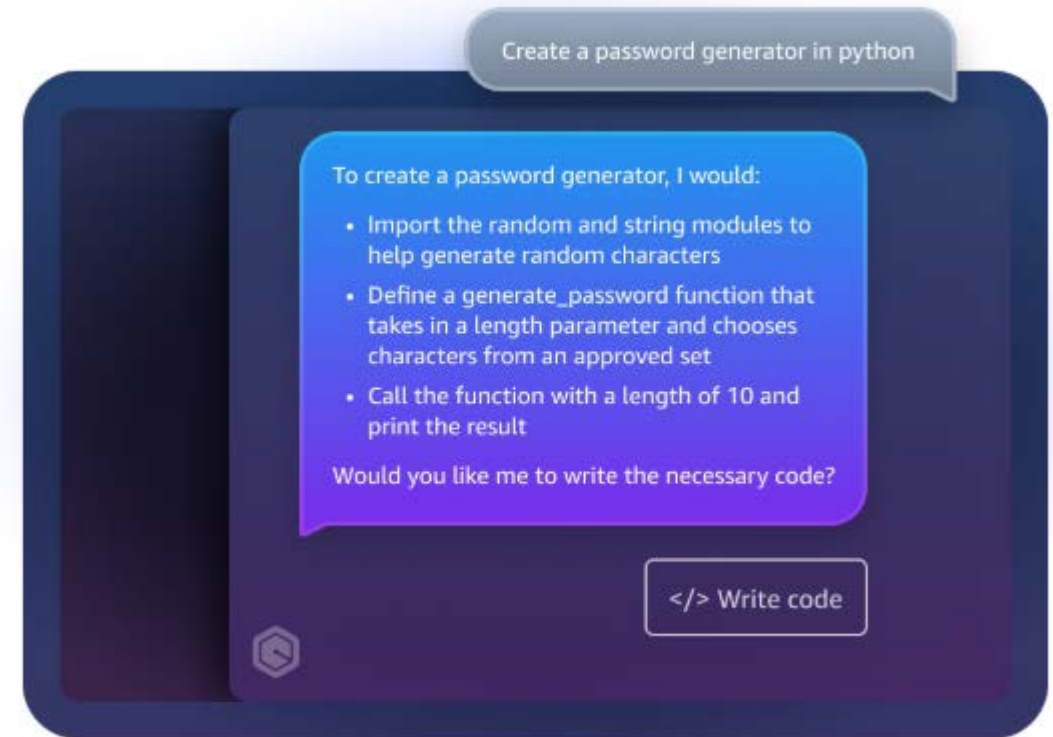
Amazon Q is built on **Amazon Bedrock**, a fully managed service for building generative AI applications that offers a choice of high-performing foundation models (FMs) from Amazon and leading AI companies. Amazon Q uses multiple Foundational Models to complete its tasks and uses logic to route tasks to the Foundational Model that is the best fit for the job.

Pricing:

- Amazon Q Business Pro charges \$3 per user/month.

Clients

- Bayer Crop Science scales agriculture up to 70% reduction in employee onboarding time, Over 2PB of imagery collected and transformed into insights for growers worldwide.
- Accenture utilizing Amazon Q Developer to create Infrastructure as Code (IaC) templates for one of their US based customer





Amazon

Market Cap
\$ 23,72,748 Mn

Revenue
\$ 6,20,128 Mn

EBITDA
\$ 1,11,583 Mn

PAT Margin
\$ 49,868 Mn

As of Sept 2024

Product Offerings

[2] Amazon SageMaker

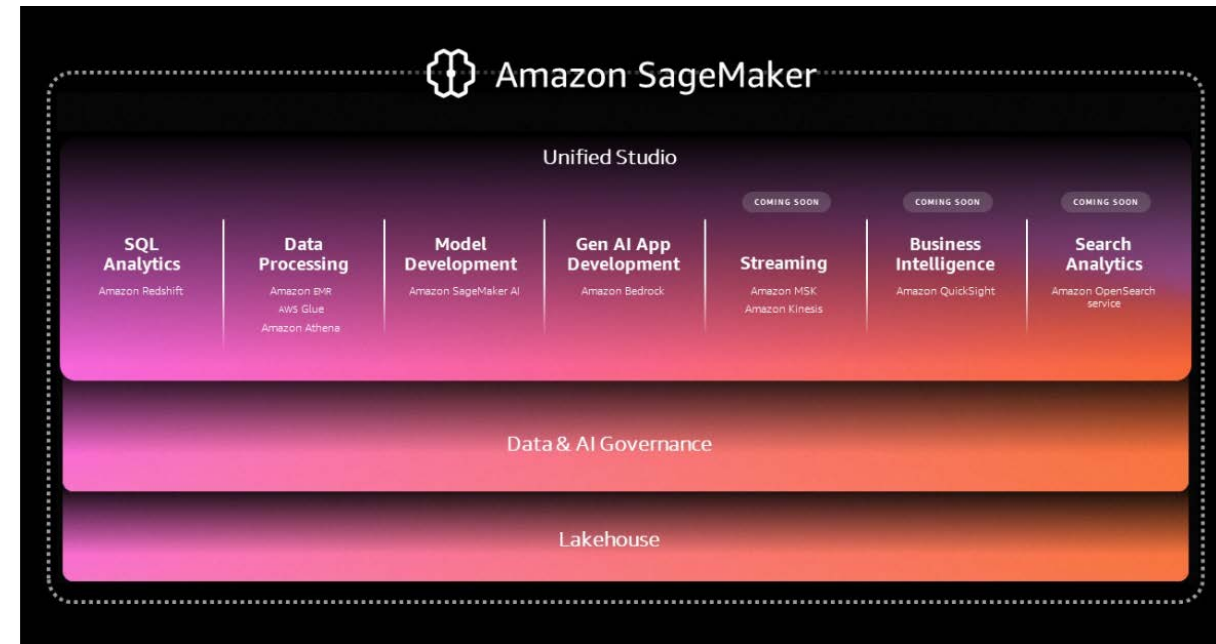
- It enables *data scientists and developers* to build, train, and deploy ML models efficiently
- Offers a *comprehensive suite of tools* and workflows to *pretrain FMs* from scratch so they can be used internally or offered to other businesses for generative AI applications
- Provides *access to hundreds of publicly available FMs* and tools to evaluate and fully customize models for user specific use case and data
- Harness the power of human feedback across the ML lifecycle to improve the accuracy and relevancy of FMs with human-in-the-loop capabilities

Pricing:

- Offering an on-demand pricing model starts at \$0.05 per hour.

Clients:

- LG AI Research - scales foundation model development using Amazon SageMaker AI
- Perplexity AI - uses SageMaker to power the Perplexity conversational answer engine that answers questions along with references provided in the form of citations
- Canva – started with generative AI, test and deploy various foundation models





Amazon

Market Cap
\$ 23,72,748 Mn

Revenue
\$ 6,20,128 Mn

EBITDA
\$ 1,11,583 Mn

PAT Margin
\$ 49,868 Mn

As of Sept 2024

Product Offerings

[3] Amazon Bedrock

- Fully managed service that offers a choice of industry leading foundation models (FMs) along with a broad set of capabilities that customers need to build generative AI applications, simplifying development with security and privacy
- It currently offers 47 models from leading AI companies like AI21 Labs, Anthropic, Cohere, Meta, Mistral AI, Stability AI, and Amazon
- It is serverless, customers don't have to manage any infrastructure, and they can securely integrate and deploy generative AI capabilities into their applications through a single API

Pricing:

- On-demand – charged for every input token processed and every output token generated
- Provisioned throughput - charged by the hour, user have the flexibility to choose between 1-month or 6-month commitment terms

Use Cases:

- Customer Support
- Data Analysis
- Text, image generation and search
- Virtual assistants

Clients:

- Lonely Planet cut itinerary generation costs by 80% using Amazon Bedrock.
- HappyFox automates support agent responses with Claude on Amazon Bedrock, increasing ticket resolution by 40%

[4] Amazon App Studio

- Generative AI-powered service that uses natural language to build business applications, empowering a new set of builders to create applications in minutes
- With App Studio, technical professionals such as IT Project Managers, Data Engineers, Enterprise Architects, and Solution Architects can quickly develop applications tailored to their organization's needs—without requiring deep software development skills
- App Studio provides conversation assistant to accelerate building new apps and modifying existing ones. Simply describe the application users want to build, such as, “Build an app to review and approve invoices,” and App Studio will generate the new application, including a user interface for listing and approving invoices, a data model for the approval records, and business logic for routing approval records and notifying owners.

Use Cases:

- Business applications for a wide range of activities, such as inventory management, multi-step approval processes, claims management, inspection and audits, and metrics and reporting

Pricing:

- Users will be charged \$0.25/user hour



Amazon

Market Cap
\$ 23,72,748 Mn

Revenue
\$ 6,20,128 Mn

EBITDA
\$ 1,11,583 Mn

PAT Margin
\$ 49,868 Mn

As of Sept 2024

Product Offerings

[5] Amazon Augmented AI

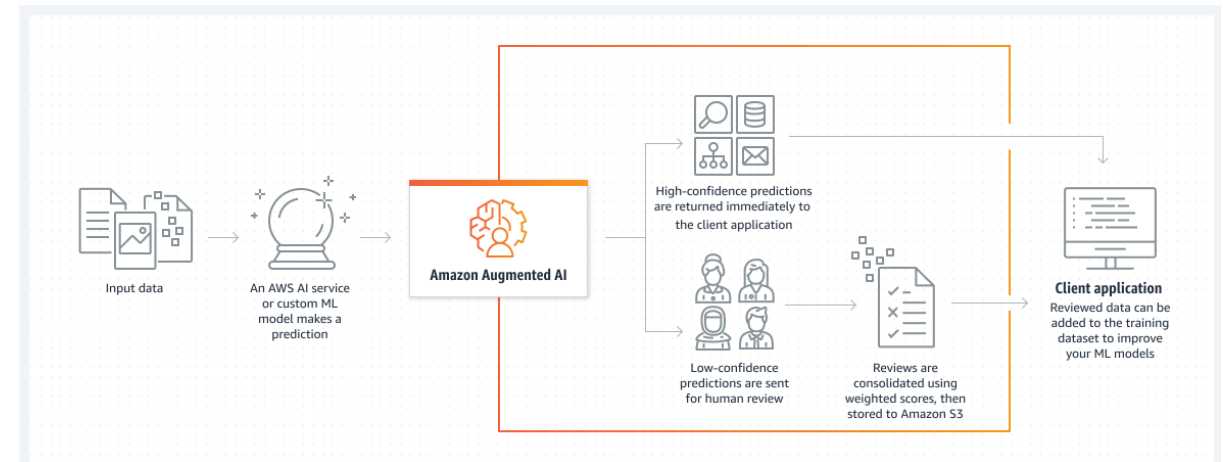
- Amazon A2I is a fully managed service that provides in *build workflows required for human review of ML predictions*.
- It brings human review to all developers, removing the undifferentiated heavy lifting associated with building human review systems or managing large numbers of human reviewers.
- It supports multiple choices for human reviewers, user's private team of reviewers for in-house review jobs, especially when handling sensitive data, access to an on-demand 24x7 workforce of over 500,000 independent contractors worldwide, a third-party workforce vendor through the AWS Marketplace.

Features:

- **Content moderation:** to identify potentially unsafe or inappropriate content on images
- **Form extraction:** review the extract key-value pairs from document images or online form
- **Image classification:** identify the object in the images
- **Data labelling**

Clients:

- Amazon AI is helping the NHS by combining AI-powered data extraction with human judgment through tools like Textract and Augmented AI to enhance decision-making in healthcare
- Amazon AI is helping Deloitte enhance client productivity and confidence in machine learning by integrating human-in-the-loop systems like A2I into ML workflows for accurate decision-making across industries.





Amazon

Market Cap
\$ 23,72,748 Mn

Revenue
\$ 6,20,128 Mn

EBITDA
\$ 1,11,583 Mn

PAT Margin
\$ 49,868 Mn

As of Sept 2024

Product Offerings

[6] Amazon Comprehend

- Amazon Comprehend is NLP service provided by AWS that leverages ML to extract insights from unstructured text.
- It identifies the language of the text, extract key phrases, places, people, brands, or events, understand sentiment about products or services, and identify the main topics from a library of documents.
- The source of this text could be web pages, social media feeds, emails, or articles. User can also feed Amazon Comprehend a set of text documents, and it will identify topics (or group of words) that best represent the information in the collection.

Use Cases:

- **Call centre analytics:** Detect customer sentiment, analyze customer interactions, and automatically categorize inbound support requests
- **Index and search product reviews:** Focus on context by equipping the search engine to index key phrases, entities, and sentiment, not just keywords.
- **Legal document management:** Automate the extraction of insights from legal briefs, such as contracts and court records
- **Process financial documents:** Classify and extract details from financial services documents such as insurance claims or mortgage packages

Pricing:

- Starting from \$0.0001 per unit (1 unit= 100 characters)

[7] Amazon Comprehend Medical

- NLP processing service that makes it easy to use machine learning to **extract relevant medical information** from unstructured text.
- Users can extract and accurately gather information, such as medical condition, medication, dosage, strength, and frequency from a variety of sources, like doctor's notes, clinical trial reports, and patient health records.

Use Cases:

- **Accelerate Insurance Claim Processing:** automate claim capture, validation, and approval workflows for healthcare insurance companies to accelerate and simplify claim processing.
- **Improve Population Health:** analyze unstructured population health data to improve health outcomes, like identifying gaps in care and improving hospital efficiency.
- **Scale and Accelerate Pharmacovigilance:** Identify adverse effects of pharmaceutical products to improve drug safety and meet regulatory requirements.
- **Perform Medical Cohort Analysis:** select the right group of patients for clinical trials to enable faster, more accurate patient trial selection and lower costs.

Pricing:

- Starting from \$0.00025 per unit (1 unit= 100 characters)



Amazon

Market Cap
\$ 23,72,748 Mn

Revenue
\$ 6,20,128 Mn

EBITDA
\$ 1,11,583 Mn

PAT Margin
\$ 49,868 Mn

As of Sept 2024

Product Offerings

[8] Amazon Kendra:

- It is an enterprise search service that enables organizations to index and search unstructured data across various content repositories.
- It allows employees and customers to find the information they need using natural language queries, enhancing productivity and satisfaction.
- Amazon Kendra is able to increase employee productivity by 25%. Reduce development costs by 80%.

Use Cases:

- Enhancing Employee Productivity by unified search experience
- Improving Customer Interactions
- Integrating Search into SaaS Applications

Pricing:

- Ranging from \$0.32 to \$1.4 per hour with a flat fee of \$30 per month for connectors.

Clients

- Amazon AI helps The Wall Street Journal Talk 2020 provide voters with seamless access to candidate information using Amazon Kendra for natural language search and real-time data updates.
- Amazon AI helps 3M accelerate innovation by using Amazon Kendra to enable scientists to quickly and accurately retrieve relevant research from their extensive knowledge base.

[9] Amazon Fraud Detector:

- It is a fully managed service by AWS that leverages ML to identify potentially fraudulent online activities, such as payment fraud and the creation of fake accounts without requiring prior ML experience.

Use Cases:

- Identifying Suspicious Online Payments
- Detecting New Account Fraud
- Preventing Trial and Loyalty Program Abuse
- Improving Account Takeover Detection

Pricing:

- Data processing \$0.10 per GB.

Clients:

- Truevo confidently uncovers fraud in real-time
- GoDaddy identifies deceptive sign-ups immediately



Amazon

Market Cap
\$ 23,72,748 Mn

Revenue
\$ 6,20,128 Mn

EBITDA
\$ 1,11,583 Mn

PAT Margin
\$ 49,868 Mn

As of Sept 2024

Product Offerings

[10] Amazon CodeGuru

- CodeGuru Security: It is a developer tool from AWS that uses ML to *automate code reviews* and *finds security vulnerabilities* in the application code. It also provides recommendations for improvement. It currently supports scanning Java, Python, JavaScript, TypeScript, C#, CloudFormation, Terraform, Go, and Ruby.
- CodeGuru Profiler: continually *searching for application performance optimizations*. It identifies user most expensive lines of code and recommending ways to fix them to reduce CPU utilization, cut compute costs, and improve application performance. For example, CodeGuru Profiler can identify when the application is consuming excessive CPU capacity on a logging routine instead of executing on core business logic.

Pricing:

- Starting from \$10 for first 100k line of code and then \$30 per 100k line of code.

Clients

- Cognizant enhance digital transformation by using CodeGuru to automate code reviews, identify issues, and ensure high-performance, secure, and compliant software deployments
- Wheel Pros optimizes application performance by using CodeGuru Reviewer and Profiler to streamline code reviews and provide ML-driven performance improvement recommendations.

[11] Amazon Personalize

- It enables developers to create *real-time, personalized user experiences* by generating product and content recommendations by using end-users' data (e.g., age, location, device type), items in the catalog (e.g., genre, price) and interactions (e.g., clicks, purchases).
- The algorithms analyze customer behavior and recommend products, content, and services that are likely to be of interest to them. This enhanced customer experience, customer engagement, loyalty, and sales

Use Cases:

- Personalize and rank movie, TV, and music recommendations
- Highlight trending retail products in real time.
- Deliver popular and seasonally relevant travel content across channels.

Pricing:

- Varies based on service starting from \$0.002 to \$0.05 per real time recommendations

Clients

- Discovery helps customers find curated content with the help of Amazon Personalize.
- Calm increased daily app use by 3.4%.



Amazon

Market Cap
\$ 23,72,748 Mn

Revenue
\$ 6,20,128 Mn

EBITDA
\$ 1,11,583 Mn

PAT Margin
\$ 49,868 Mn

As of Sept 2024

Product Offerings

[12] Amazon Transcribe:

- It is a speech-to-text capability for their applications. It accurately converts audio input into text, facilitating various use cases such as transcription of customer service calls, subtitling media content, and analyzing clinical conversations.

Use Cases:

- Contact Centers
- Media and Entertainment
- Healthcare
- Legal and Compliance

Pricing

- Ranging from \$0.00780 to \$0.02400 per minute based on usage.

Clients

- Amazon AI helps Intuit by using Amazon Transcribe to automatically convert speech to text, enabling accurate call transcriptions and enhancing customer support analytics
- Amazon AI helps T-Mobile by using Amazon Transcribe to convert voicemail messages into text, enabling faster and more accessible voicemail management for customers.

[13] Amazon HealthScribe

- It enables healthcare software vendors to develop clinical applications capable of automatically generating preliminary clinical documentation from patient-clinician conversations.
- By integrating speech recognition and generative AI., HealthScribe transcribes consultations, identifies speaker roles, extracts medical terms, and produces clinical notes, thereby reducing the administrative burden on healthcare professionals.

Pricing:

- HealthScribe operates on a pay-as-you-go model, charging \$0.001667 per second of audio processed, with a minimum per-request charge of 15 seconds.

Clients

- 3M enhance clinical workflows by using HealthScribe to streamline documentation, billing.
- ScribeEMR enhance healthcare documentation by using AWS HealthScribe to streamline patient visit interpretation, optimize EMR workflows, and improve efficiency and patient care.



Amazon

Market Cap
\$ 23,72,748 Mn

Revenue
\$ 6,20,128 Mn

EBITDA
\$ 1,11,583 Mn

PAT Margin
\$ 49,868 Mn

As of Sept 2024

Product Offerings

[14] Amazon Lex

- It enables developers to create conversational interfaces using voice and text for any application.
- It leverages the same deep learning technologies that power Amazon Alexa, allowing for the development of natural language chatbots and voice assistants.

Use Cases:

- **Customer services:** enable end-customer self-service capabilities across interactive voice response (IVR), chat, and SMS to solve customer queries quickly and efficiently.
- **Internal Business Applications:** complete end-to-end tasks like making a booking, raising a ticket, generating a report, submitting an application, and more using natural language or voice AI chat.

Pricing:

- \$0.004 per speech requests and \$0.00075 per text requests

Clients

- Vanguard optimized call routing for over 40,000 calls a day – Eliminated unnecessary transfers, reduced the long hold, first call resolution times.
- Morrisons delivered a self-serve contact center in 8 weeks.

[15] Amazon Rekognition:

- It automates image and video analysis, enabling developers to add visual search and discovery capabilities.
- It detects objects, scenes, activities, landmarks, faces, dominant colors, and image quality. It also extracts text, recognizes celebrities, and identifies inappropriate content in images/video

Use cases:

- **Content Moderation:** Process millions of images and videos efficiently while detecting inappropriate or unwanted content
- **Identity verification:** pretrained facial recognition and analysis capabilities that can be quickly added to user onboarding and authentication workflows to verify opted-in users' identity online
- **Media analysis:** streamline quality control, ad insertion, and content production
- **Workplace Safety:** Automate PPE detection to improve workplace safety practices

Pricing:

- Starting from \$0.001 per image

Clients

- Pinterest uses text detection to moderate user-uploaded images.
- PBS uses object detection to streamline media content operations.



Amazon

Market Cap
\$ 23,72,748 Mn

Revenue
\$ 6,20,128 Mn

EBITDA
\$ 1,11,583 Mn

PAT Margin
\$ 49,868 Mn

As of Sept 2024

Product Offerings

[16] Amazon Polly– It converts written text into speech it enables developers to create applications that can speak in a variety of languages and voices, enhancing user engagement through auditory experiences. With Lifelike voices.

Benefits:

- Enhanced User Engagement
- Scalability and Flexibility
- Cost-Effective Pricing

Use Cases:

- Content Creation
- Interactive Voice Response (IVR) Systems
- E-Learning Platforms
- Media and Entertainment

Pricing:

- Ranging from \$0.01 to \$100.

Clients

- Amazon AI helps GE Appliances enhance customer service by using Amazon Polly to streamline call handling and save time for consumers.
- Amazon AI helps WaFd Bank streamline customer interactions by using Polly, to automate account balance check inquiries and reduce it from 4:30 minutes to 25 seconds.

[17] Amazon Translate: It is a fast, high-quality, and affordable language translation. Also offers text-based real-time translation within chat, email, helpdesk etc.

Use Cases:

- Translating User-Generated Content
- Analyzing Online Conversations
- Cross-Lingual Communication

Pricing:

- Ranging from \$15 to \$60 per million characters.

Clients

- BMW Group improved translation time by over 75%
- CaptionHub uses Amazon Translate for real-time, automated translation of subtitles, enabling seamless localization of video content across multiple languages by increasing the production by 800%.



Amazon

Market Cap
\$ 23,72,748 Mn

Revenue
\$ 6,20,128 Mn

EBITDA
\$ 1,11,583 Mn

PAT Margin
\$ 49,868 Mn

As of Sept 2024

Product Offerings

[18] Amazon Monitron

- Provides customers an end-to-end hardware and software system comprising of wireless sensors to capture vibration and temperature data from equipment, a gateway device to securely transfer data to AWS
- It analyzes the data for abnormal machine patterns in industrial machinery, so users can implement predictive maintenance and reduce unplanned downtime.
- It has companion mobile app to set up the devices and receive reports on operating behavior and alerts to potential failures in machinery

Pricing:

- A one-time device purchase cost for the Sensors and Gateways, and an ongoing pay-as-you-go service fee for each Amazon Monitron Sensor in use.



[19] Amazon Panorama: It enables organizations with computer vision capabilities on their on-premises cameras, allowing for real-time video analysis at the edge. By deploying ML models directly on local devices, AWS Panorama reduces latency, enhances data privacy, and minimizes the need for extensive internet bandwidth.

Use Cases:

- Supply Chain Logistics
- Traffic Management
- Manufacturing Quality Control

Pricing

- Device costs \$4000 per device to \$8.33 per month per active camera system.

Clients

- Amazon AI helps the Port of Vancouver by using AWS Panorama for computer vision-based container tracking, improving efficiency, reducing costs, and enhancing data privacy through edge computing
- Amazon AI helps Tyson Foods automate inventory counting and production flow optimization at its plants using AWS Panorama for real-time computer vision processing at the edge.



Amazon

Market Cap
\$ 23,72,748 Mn

Revenue
\$ 6,20,128 Mn

EBITDA
\$ 1,11,583 Mn

PAT Margin
\$ 49,868 Mn

As of Sept 2024

Financials

Particulars	FY19	FY20	FY21	FY22	FY23	TTM
Revenue (In \$ million)	2,80,522	3,86,064	4,69,822	5,13,983	5,74,785	6,20,128
EPS Growth Rate	-	37.62%	21.70%	9.40%	11.83%	7.89%
EBITDA (In \$ million)	36,330	48,079	59,312	55,269	85,515	1,11,583
<i>EBITDA Margin</i>	12.95%	12.45%	12.62%	10.75%	14.88%	17.99%
PAT (In \$ million)	11,588	21,331	33,364	-2,722	30,425	49,868
<i>PAT Margin</i>	4.13%	5.53%	7.10%	-0.53%	5.29%	8.04%
EPS (In \$)	1.10	2.03	3.18	-0.26	2.90	4.67
EPS Growth Rate	-	84.08%	56.41%	-108.16%	-1217.74%	60.99%
Price (31st March)	94.16	164.29	168.64	83.04	153.38	222.13
Historic P/E	85.25	80.81	53.03	-320.08	52.89	47.58
Current P/E (TTM)			47.58			
Current PE/ROE			3.16			
Shareholder funds (In \$ million)	62,060.00	93,404.00	1,38,245.00	1,46,043.00	2,01,875.00	2,59,151.00
Minority Interest(In \$ million)	0.00	0.00	0.00	0.00	0.00	0.00
Debt (In \$ million)	77,533.00	1,04,740.00	1,39,757.00	1,69,938.00	1,61,574.00	1,58,535.00
Cash (In \$ million)	55,021.00	84,396.00	96,049.00	70,026.00	86,780.00	88,051.00
No. of shares (In million)	10,080.00	10,198.00	10,296.00	10,189.00	10,492.00	10,681.80
Closing Price	94.16	164.29	168.64	83.04	149.10	222.13
Market Cap (In \$ million)	9,49,132.80	16,75,429.42	17,36,317.44	8,46,094.56	15,64,357.20	23,72,748.23
Enterprise Value (In \$ million)	9,71,644.80	16,95,773.42	17,80,025.44	9,46,006.56	16,39,151.20	24,43,232.23

	Ratios					
D/E	1.25	1.12	1.01	1.16	0.80	0.61
PE/ROE	4.57	3.54	2.20	171.73	3.51	2.47
ROCE	26.03%	24.26%	21.34%	17.49%	23.53%	26.71%
ROE	18.67%	22.84%	24.13%	-1.86%	15.07%	19.24%
EV/Sales	3.46	4.39	3.79	1.84	2.85	3.94

	FY19	FY20	FY21	FY22	FY23	TTM
R&D	35,931	42,740	56,052	73,213	85,622	85,750
<i>% of revenue</i>	13%	11%	12%	14%	15%	14%



Google

Market Cap
\$ 24,44,583 Mn

Revenue
\$ 3,39,859 Mn

EBITDA
\$ 1,23,470 Mn

PAT Margin
\$ 94,269 Mn

As of Sept 2024

Product Offerings

Gemini Family

- Gemini is part of DeepMind's LLMs and multimodal AI. Designed for broad adaptability, Gemini combines LLMs with multimodal processing to power conversational agents, automation, and creativity tools. From natural image, audio and video understanding to mathematical reasoning, widely-used academic benchmarks used in LLM research and development.
- Gemini surpasses GPT 4 performance on a range of benchmarks including text and coding.

Model Variants:

- **Gemini 1.0 Ultra:** The largest model in the Gemini family, tailored for highly complex tasks requiring substantial computational resources. Gemini 1.0 trained at scale on AI-optimized infrastructure using Google's in-house designed Tensor Processing Units (TPUs) v4 and v5e. They have launched an efficient and scalable TPU system to date, Cloud TPU v5p, designed for training AI models.
- **Gemini 1.5 Pro:** Optimized for reasoning tasks, capable of processing and analyzing large volumes of information across multiple modalities, including text, images, audio, and video.
- **Gemini 1.0 Nano:** The most efficient model in the family, designed for on-device tasks where computational resources are limited. Such as the search for using on mobile.
- **Gemini 2.0 Flash Experimental:** An experimental model emphasizing low latency and enhanced performance, built to power agentic experiences. It introduces capabilities such as native tool use, image creation, and speech generation

Capabilities:

- Image generation
- Text-to-speech
- It also enables agentic workflows for example giving instructions to order something online.
- Setting the workflow conditions for real-time. Responding to live audio and video input.
- Enables the virtual world of video games faster, cheaper, and more options with the help of automation.
- **Spatial Understanding-** Gemini can give the location of the objects, and text (language detection).
- **Video Understanding-** Gemini can help in outlining key moments and summarizing the call.
- **GenExplainer-** Gemini can explain a topic in the given style. (Explain to me what is blackhole in magician style)
- **GenWeather-** Gemini can also explain the weather in the given style. (Explain today's weather in a wise poet style)
- **For Developers-** Gemini Multimodal API allows developers to build applications with better natural language interactions and video understanding.

Use Cases:

- Google Search's conversational features.
- Creative assistance in Google Workspace (e.g., Docs, Slides).
- Agentic tasks like summarizing meetings, automating workflows, or designing visual media.



Google

Market Cap
\$ 24,44,583 Mn

Revenue
\$ 3,39,859 Mn

EBITDA
\$ 1,23,470 Mn

PAT Margin
\$ 94,269 Mn

As of Sept 2024

Product Offerings

[1] **Gemini Conversational AI** –Google’s flagship conversational AI interface, leveraging advanced LLMs like PaLM(multilingual LLM). Designed for interactive, contextually aware exchanges, Gemini excels in facilitating human-like dialogues and generating creative content.

Capabilities:

- Addresses complex, open-ended queries in context-sensitive responses.
- Enhances user interaction by incorporating visual outputs alongside textual explanations.
- Assists in tasks ranging from brainstorming and planning to creative writing.

Features:

- **Adaptive Dialogues:** Continuously adjusts responses to align with conversational context and user intent.
- **Visual Augmentation:** Enriches explanations with images, charts, and other graphical elements when relevant.
- **Creative Generation:** Produces high-quality written content, such as essays, emails, and imaginative works, with minimal input.

Applications:

- **Enhanced Search:** Powers enriched conversational experiences in Google Search.
- **Educational Assistance:** Supports students by summarizing topics, answering questions, and aiding in research.
- **Customer Support:** Delivers real-time solutions to common user inquiries in service settings.

Clients

- Google Workspace helps **Uber** by automating repetitive tasks, enabling developers to focus on high-value work, reducing agency costs, and improving employee retention.
- Google Workspace helps **Finquery** streamline tasks like creating summaries by generating concise outputs, such as one-page summaries of observability and monitoring tools, thereby improving productivity and allowing the team to prioritize high-impact tasks.



Google

Market Cap
\$ 24,44,583 Mn

Revenue
\$ 3,39,859 Mn

EBITDA
\$ 1,23,470 Mn

PAT Margin
\$ 94,269 Mn

As of Sept 2024

Product Offerings

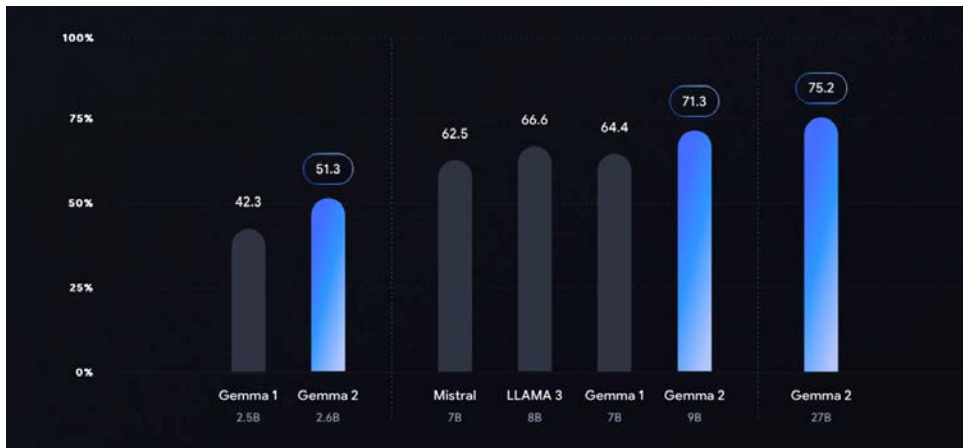
[2] **Gemma** – Gemma consists of lightweight AI models derived in comparison to Gemini, optimized for devices with limited computing power.

Capabilities:

- Performs tasks like summarization and translation efficiently on devices with limited resources.
- Processes different types of data, such as text and images, on local devices.
- Provides quick responses for real-time applications.

Applications:

- **Mobile Translation:** Offers real-time language translation on Android and Pixel devices.
- **Accessibility Innovations:** Powers features like live captioning and text-to-speech for assistive technologies.
- **On-Device AI Tools:** Enables voice assistants and context-aware summarization in mobile app



[3] **MetNet 3 and GraphCast** - This is an open-source product built by Google. It improves the weather forecast through a deep-learning model that analyses past weather data to understand trends and patterns. It is faster, more accurate, and provides more detail. It is useful for everyday weather and extreme weather forecast conditions such as cyclones. It is currently used by the European Centre for Medium-Range Weather Forecasting.

[4] **Vertex AI** – A Vertex AI tool for rapidly prototyping and testing generative AI models. Test sample prompts, design own prompts, and customize foundation models and LLMs to handle tasks that meet the application's needs.

Use Case:

- Prompt design and tuning with an easy-to-use interface
- Code completion and generation with Codey
- Generating and customizing images with imagen
- Universal speech models

[5] **SynthID** – Identifies AI-generated content by embedding digital watermarks directly into AI-generated images, audio, text or video. It identifies either through digital watermark directly into AI-generated content without compromising the original content.

-SynthID can scan images, audio, text, or video for digital watermarks to determine if the content or a part of it was generated by Google's AI tools. It will also work after modifications such as cropping, adding filters, changing colours, frame rates, etc.



Google

Market Cap
\$ 24,44,583 Mn

Revenue
\$ 3,39,859 Mn

EBITDA
\$ 1,23,470 Mn

PAT Margin
\$ 94,269 Mn

As of Sept 2024

Product Offerings

[6] **Veo**- Veo creates videos with realistic motion and high-quality output, up to 4K. Claims to significantly improve over other AI video models in terms of detail, realism reduction. Veo represents motion with accuracy as veo understands physics and can follow detailed instructions.

[7] **Imagen 3** – It creates the highest quality text-to-image, capable of generating images with better detail, lighting, and fewer distracting artefacts than previous models.

- Brightness and vibrancy
- Diverse art styles
- High-fidelity detail
- Better text rendering

[8] **AlphaFold** – AlphaFold is an AI system developed by Google DeepMind that predicts a protein's 3D structure from its amino acid sequence. AlphaFold allows scientists to make predictions quickly. AlphaFold has many potential applications, including stopping malaria, developing vaccines, understanding proteins linked to cancer and autism and Creating plastic-eating enzymes.

[9] **MusicFX** – It can create songs with little or no upload of audio. Helps in generating new genres and faster production and remixing.

[10] **AlphaCode 2** – It helps in advanced code generation systems useful for solving competitive programming problems that are beyond coding to involve complex math and theoretical computer science. Its ability to work across languages and reason about complex information makes it one of the leading foundation models for coding in the world.

[11] **PaLM 2** – PaLM 2 is a large language model by Google that supports over 100 languages and is fine-tuned for specific fields like medicine and law.

Capabilities:

- **Multilingual Proficiency:** Understands and generates text in multiple languages.
- **Domain-Specific Adaptability:** Specialized versions like MedPaLM 2 for healthcare and LegalPaLM for legal analysis provide accurate information in these areas.
- **Advanced Logical Reasoning:** Handles complex questions and problems, making it useful for technical documentation and problem-solving.

Applications:

- **Gemini/Bard:** Powers real-time conversational AI for interactive dialogues.
- **Google Translate:** Improves translation accuracy and fluency for global communication.
- **Technical Domains:** Assists in writing, summarizing, and solving problems in healthcare, legal, and technical fields.



Google

Market Cap
\$ 24,44,583 Mn

Revenue
\$ 3,39,859 Mn

EBITDA
\$ 1,23,470 Mn

PAT Margin
\$ 94,269 Mn

As of Sept 2024

Financials

Particulars	FY19	FY20	FY21	FY22	FY23	TTM
Revenue (In \$ million)	1,61,857	1,82,527	2,57,637	2,82,836	3,07,394	3,39,859
EPS Growth Rate	-	12.77%	41.15%	9.78%	8.68%	10.56%
EBITDA (In \$ million)	47,579	54,903	88,987	88,317	1,00,172	1,23,470
<i>EBITDA Margin</i>	29.40%	30.08%	34.54%	31.23%	32.59%	36.33%
PAT (In \$ million)	34,343	40,269	76,033	59,972	73,795	94,269
<i>PAT Margin</i>	21.22%	22.06%	29.51%	21.20%	24.01%	27.74%
EPS (In \$)	2.70	3.17	5.98	4.71	5.80	7.53
EPS Growth Rate	-	17.26%	88.81%	-21.12%	23.05%	29.90%
Price (31st March)	71.54	86.98	144.68	87.93	142.72	195.39
Historic P/E	26.50	27.48	24.21	18.65	24.60	25.93
Current P/E (TTM)			25.93			
Current PE/ROE			1.00			
Shareholder funds (In \$ million)	2,01,442.00	2,22,544.00	2,51,635.00	2,56,144.00	2,83,379.00	3,14,119.00
Minority Interest(In \$ million)	0.00	0.00	0.00	0.00	0.00	0.00
Debt (In \$ million)	16,082.00	27,872.00	28,508.00	29,977.00	29,867.00	29,289.00
Cash (In \$ million)	1,19,675.00	1,36,694.00	1,39,649.00	1,13,762.00	1,10,916.00	93,230.00
No. of shares (In million)	13,971.10	13,741.00	13,553.00	13,159.00	12,722.00	12,511.30
Closing Price	71.54	86.98	144.68	87.93	140.53	195.39
Market Cap (In \$ million)	9,99,492.49	11,95,192.18	19,60,848.04	11,57,070.87	17,87,822.66	24,44,582.91
Enterprise Value (In \$ million)	8,95,899.49	10,86,370.18	18,49,707.04	10,73,285.87	17,06,773.66	23,80,641.91

	Ratios					
D/E	0.08	0.13	0.11	0.12	0.11	0.09
PE/ROE	1.55	1.52	0.80	0.80	0.94	0.86
ROCE	21.87%	21.92%	31.76%	30.87%	31.98%	35.95%
ROE	17.05%	18.09%	30.22%	23.41%	26.04%	30.01%
EV/Sales	5.54	5.95	7.18	3.79	5.55	7.00

	FY19	FY20	FY21	FY22	FY23	TTM
R&D	26,018	27,573	31,562	39,500	45,427	48,323
<i>% of revenue</i>	16%	15%	12%	14%	15%	14%



SoundHound AI

Market Cap
\$ 4,608 Mn

Revenue
\$ 67 Mn

EBITDA
\$ -88 Mn

PAT Margin
\$ -110 Mn

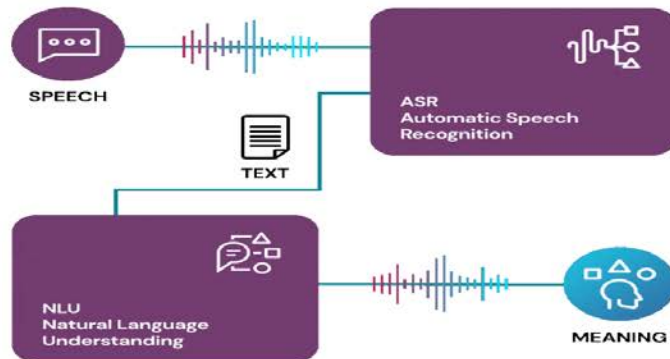
As of Sept 2024

Overview

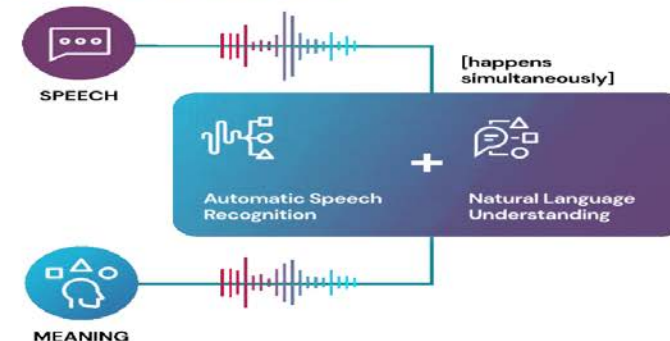
Overview: SoundHound a voice AI company that develops voice-enabled products and services. Their products are used in many industries, including automotive, TV, IoT, restaurants, customer service, healthcare, finance, and retail

- **Smart Answering:** - AI-Powered **Voice Assistant Solutions** tailored to help businesses **manage phone inquiries** efficiently. It ensures that customers always receive a response, even outside business hours, reducing missed calls and maximizing business opportunities. The system is designed to **automate answering customer questions, collecting leads, routing calls, and even providing links for bookings or information**
 - **Technology:** **Automatic Speech Recognition (ASR), Natural Language Understanding (NLU), Text-to-Speech (TTS), Generative AI + Proprietary Algos**
 - **Use Cases:**
 - **Restaurants:** Answers queries about **menu options, dietary preferences, operating hours, and reservations**
 - **Real Estate:** Provides details on **property listings, booking appointments, and agent availability**
 - **Personal Care Services (Salons, Spas):** Schedules appointments and shares service details
 - **Retail and e-Commerce:** Answers product-related questions, pricing, and store locations.
 - **Law Firms and Consulting Services:** Routes queries to the appropriate expert and handles appointment scheduling
 - **Pricing:** Single Location: **\$99/month (limited-time offer)**. For multiple locations, customized plans are available through direct contact with the sales team.

Typical voice AI technologies
involve two steps,



SoundHound
gets it done in one.



Faster, more efficient
and with greater accuracy



SoundHound AI

Market Cap
\$ 4,608 Mn

Revenue
\$ 67 Mn

EBITDA
\$ -88 Mn

PAT Margin
\$ -110 Mn

As of Sept 2024

Technology Stake

- **Core Technology:**
 - **Speech-to-Meaning:** Unlike most platforms that follow a two-step process—converting speech to text and then analyzing the text for meaning—SoundHound processes both simultaneously
 - This enables: **Faster response times**, and **Improved accuracy in understanding spoken commands**
 - **Example:** Alexa or Google Assistant rely on a two-step process (ASR + NLU) - Leads to slow down response time
 - **SoundHound's Speech-to-Meaning®** eliminates this inefficiency
 - **Deep Meaning Understanding:** Handles complex queries by understanding context and intent. The platform can:
 - Address multiple questions in a single query
 - Provide filtered results tailored to user needs
 - Manage nuanced, layered interactions
 - **Dynamic Interaction:** This enables real-time **understanding of user inputs**. This means users can interact with a device or system using both **voice and touchscreen inputs** simultaneously
 - **Audiovisual Feedback:** Audiovisual feedback refers to the system's ability to respond in real-time using both:
 - **Audio feedback:** Spoken responses or sound cues
 - **Visual feedback:** On-screen elements like animations, updated text, or highlighted options that reflect the user's query or actions
 - **Example:**
 - **Initial Interaction:** A user says, "Find Italian restaurants nearby. The voice assistant audibly responds, "**Here are some options**," while displaying a list of Italian restaurants on the car's touchscreen.
 - **Refining the Query with Touchscreen:** The user notices a specific restaurant on the screen and taps on it to learn more
 - **Updating the Query Using Voice:** The user decides to refine their search by saying, "Show me ones with outdoor seating instead."



SoundHound AI

Market Cap
\$ 4,608 Mn

Revenue
\$ 67 Mn

EBITDA
\$ -88 Mn

PAT Margin
\$ -110 Mn

As of Sept 2024

Product Offerings

Restaurant Phone Answering and Ordering:

- **Restaurant Smart Answering:** An advanced AI-powered voice assistant tailored for restaurants. It handles 100% of incoming phone calls with natural, conversational responses.
 - It provides essential information, **supports reservation and ordering tasks**, and seamlessly **transitions to human staff when necessary**
 - **Comprehensive Call Handling:** Automates responses for common inquiries such as **hours, menu, location, wait times, and delivery options**
 - **Intelligent Task Management:** Facilitates **reservations, ordering, and delivery scheduling via SMS or voice**. Integrates with **existing systems** to confirm availability and streamline the process
 - **Personalized Customer Engagement:** Customizes responses based on the **restaurant's branding**, ensuring a friendly and professional interaction.
 - **Efficiency and Scalability:** Operates around the clock, ensuring no call is missed, even during peak hours.
- **Restaurant Smart Ordering:** An AI-powered voice assistant designed to handle phone orders for restaurants and food businesses. It simplifies the ordering process by interacting naturally with customers, understanding their preferences, and processing orders automatically
 - **Voice-Driven Order Management:** Captures customer orders via voice and converts them into actionable items in the restaurant's POS system. Handles complex orders, including customization requests, dietary preferences, and add-ons.
 - **POS Integration:** Directly integrates with Square, Toast, Oracle, and Olo POS systems for automatic processing and ticket generation.
 - **Operational Efficiency:** Works 24/7, reducing dependency on human staff and ensuring no order is missed during peak hours.
- **Clients** – White Castle, Chipotle, Jersey Mike's, Panda Express, Krispy Kreme, Noodles & Company, Firehouse Subs, Applebee's, Church's Texas Chicken, Five Guys, Blake's Lotaburger, Chipotle Mexican Grill, Casey's, MOD, Papa John's, Luby's, Blaze Pizza, Hungry Howie's, Beef O' Brady's, Detroit's Original Chicken Shack, Toast, Square, Olo, Oracle Food and Beverage
- **Pricing:** Customized (Starts from \$99 per months)



SoundHound AI

Market Cap
\$ 4,608 Mn

Revenue
\$ 67 Mn

EBITDA
\$ -88 Mn

PAT Margin
\$ -110 Mn

As of Sept 2024

Product Offerings

- **SoundHound's Automatic Speech Recognition (ASR):** SoundHound's ASR technology converts **spoken words into text accurately and efficiently**, making it easier to use voice commands in various devices and applications. Here's what makes it special
 - **Unmatched Understanding and Accuracy:** Uses Speech-to-Meaning® technology, which understands the meaning of what's being said, not just the words, making it more accurate for complex speech
 - **Comprehensive Call Handling:** Automates responses for common inquiries such as **hours, menu, location, wait times, and delivery options**
 - **Large Vocabulary and Precision:** Supports millions of words, allowing it to handle a wide range of speech, from casual to professional contexts.
 - **Customizable for Your Needs:** The technology can be adapted for specific environments (e.g., factories, offices, or homes) to ensure it works perfectly with your product
 - **Use Cases:**
 - **Video Calls:** Imagine saying, **"Summarize this meeting,"** and the system provides a written summary, tagging each speaker
 - **Kiosks in Airports or Cities:** Speak to a kiosk in your native language, and it responds or gives instructions in your preferred language.
 - **Smart Appliances:** Say, **"Start the oven at 350 degrees,"** even in a **noisy kitchen**, and it works perfectly.
 - **Pricing:** Customized
- **Natural Language Understanding (NLU):** NLU is enabling real-time comprehension and natural, intuitive interactions between humans and voice assistants
 - **Tech:** Speech-to-Meaning® Technology + Deep Meaning Understanding + Support for **25+ Languages**
 - **What makes it different from others:**
 - **Speed of Comprehension** - Uses Speech-to-Meaning for faster & accurate response
 - **Simultaneous Question Handling:** Uses Deep Meaning Understanding to handles compound and complex queries effortlessly
 - **Negation and Context Awareness:** The ability to process double negations (e.g., **"I don't want anything that isn't vegetarian"**) sets it apart in terms of precision
 - **Customization for Brands:** Full customization for businesses, allowing them to integrate voice AI solutions tailored to their brand, unlike generic solutions provided by others
 - **Use Cases**
 - **Smart Cars:** User: "Turn on the air conditioning, lower the back windows, and navigate to the nearest coffee shop that's open now"
 - **Response:** The system simultaneously processes all commands and responds in real-time, while also filtering coffee shops based on operating hours
 - **Smart Home Devices:** User: "Dim the living room lights and play some relaxing music from Spotify."
 - **Smart Home Devices:** User: "Dim the living room lights and play some relaxing music from Spotify."



SoundHound AI

Market Cap
\$ 4,608 Mn

Revenue
\$ 67 Mn

EBITDA
\$ -88 Mn

PAT Margin
\$ -110 Mn

As of Sept 2024

Product Offerings

- **Automatic Content Recognition (ACR):** ACR technology is designed to automatically identify copyrighted material, particularly music, in UGC. This system helps platforms comply with copyright regulations by detecting, reporting, and addressing copyrighted material efficiently and accurately

How it works?

- **Content Scanning:** ACR systems scan audio tracks in UGC, analyzing sound patterns to detect copyrighted material.
- **Music Recognition:** Using advanced algorithms, it can recognize songs even when mixed with background noise, voices, or other audio.
- **Metadata Matching:** Identified content is matched with a database of music metadata (artist name, song title, ISRC identifier).
- **Real-Time Reporting:** The system generates automated reports detailing the identified content, including play duration, date, time, and associated metadata.
- **Actionable Insights:** These reports are used to clear copyright issues, ensure compliance, and calculate royalties.

Use Cases:

- **UGC Platforms (e.g., YouTube, TikTok):** Automatically detect copyrighted songs in user-uploaded videos.
- **Streaming Services (e.g., Spotify, SoundCloud):** Identify and manage music usage in live streams or podcasts to avoid copyright infringements.
- **Social Media Platforms (e.g., Instagram, Facebook):** Scan uploaded content for copyrighted background music to ensure compliance with global copyright laws

- **Intelligent Transcription:** An advanced transcription service that goes beyond simply converting spoken words into text. It uses cutting-edge technologies such as Automatic Speech Recognition (ASR) and voice AI to transcribe conversations in real-time while understanding the context, intent, and meaning behind the words

- **Tech** - It uses Automatic Speech Recognition (ASR)+ Contextual Understanding+ ML

Use Cases:

- **Customer Service:** By transcribing live conversations in real-time, customer support teams can quickly understand the issue and provide a solution
- **Meetings and Conferences:** In business settings, Intelligent Transcription helps teams capture key discussions, action points, and decisions, even in noisy environments
- **Content Creation and Media:** Intelligent Transcription is valuable for content creators who need to transcribe interviews, podcasts, or media content with speed and accuracy, capturing key moments and theme



SoundHound AI

Market Cap
\$ 4,608 Mn

Revenue
\$ 67 Mn

EBITDA
\$ -88 Mn

PAT Margin
\$ -110 Mn

As of Sept 2024

Product Offerings

Apps

- Houndify - An app for developers:** An AI voice platform that allows **businesses and developers to integrate voice recognition** capabilities into their applications
 - **Technology:**
 - **CaiLAN (Conversational AI Language)** - A structured way to connect different pieces of information. Ex - Find a nearby coffee shop with free Wi-Fi that is open now
 - **Speech-to-Meaning® Technology** - This patented technology enables Houndify to process speech in real time, extracting meaning and context as soon as words are spoken
 - **ASR & NLU**
 - **Use Cases:**
 - **Automotive Industry** - In-Car Voice Assistants. Ex - Drivers can use voice commands to control navigation, play music, set reminders, check the weather, or get directions without taking their hands off the wheel
 - **Smart Home and IoT** - Smart Device Control & Personalized Routines (i.e **Good night” to turn off lights, lock doors**)
 - **Healthcare** - Virtual Health Assistants (i.e Schedule an appointment with Dr. Lee and remind me to take my medication at 8 PM)
 - **Retail and E-commerce** - Voice Shopping Assistants (**Ex. Show me red dresses available in size medium, and tell me the delivery time if I order now**)
 - **Hospitality** - In-Room Assistants for Hotels (**Ex. Turn on the air conditioning and order breakfast for 8 AM**)
 - **Media and Entertainment** - Voice-Controlled Streaming (**Ex. Play the latest episode of my favourite show and turn on subtitle**)
 - **Clients:** Hyundai, Mercedes-Benz, KIA, Netflix, Snap
- SoundHound Music App:** A consumer-facing application that provides **music recognition, allowing users to identify songs by simply humming or playing a tune.** The app also offers personalized music recommendations.
 - **Tech** - ASR, **Sound and Audio Fingerprinting**, NLU, Deep Neural Networks and Machine Learning
 - **Features:**
 - **Identifying Songs by Humming or Singing** - This feature is especially helpful for users who only remember part of a melody but not the lyrics or title
 - **Getting Song Lyrics in Real Time** - Displays synchronized lyrics as the song plays
 - **Personalized Music Recommendations:**
 - **Building a Music Library of Favorites**
 - **Song Identification on the Go**
 - **Client:** Spotify



SoundHound AI

Market Cap
\$ 4,608 Mn

Revenue
\$ 67 Mn

EBITDA
\$ -88 Mn

PAT Margin
\$ -110 Mn

As of Sept 2024

Product Offerings- APP

Competition:

- **Google Cloud Speech-to-Text:** Convert Speech into Text in more than 120 languages and variants
- **Deepgram:** Deepgram's voice AI power apps with real-time speech-to-text and text-to-speech APIs
- **Otter.ai:** A transcription software company that uses artificial intelligence (AI) to create written transcripts of speech
- **Twilio:** A cloud communications company that provides tools for developers and businesses to add communication features to their applications

Acquisition:

- **Amelia** – A leading conversational AI platform
- **SYNQ3** - The largest voice AI provider for restaurants

Partnership

- **Perplexity** - SoundHound partnered with Perplexity to bring cutting-edge online LLMs to SoundHound Chat AI
- **Connex2X** - An innovative aftermarket connected vehicle company, partnered with SoundHound to integrate leading voice AI technology into its connected vehicle products



SoundHound AI

Market Cap
\$ 4,608 Mn

Revenue
\$ 67 Mn

EBITDA
\$ -88 Mn

PAT Margin
\$ -110 Mn

As of Sept 2024

Financials

Particulars	FY19	FY20	FY21	FY22	FY23	TTM
Revenue (In \$ million)	7.70	13.00	21.20	31.10	45.90	67.30
EPS Growth Rate	-	68.83%	63.08%	46.70%	47.59%	46.62%
EBITDA (In \$ million)	-59.90	-60.00	-59.80	-101.60	-61.70	-88.40
EBITDA Margin	-777.92%	-461.54%	-282.08%	-326.69%	-134.42%	-131.35%
PAT (In \$ million)	-64.50	-74.40	-79.50	-116.70	-88.90	-110.10
PAT Margin	-837.66%	-572.31%	-375.00%	-375.24%	-193.68%	-163.60%
EPS (In \$)	-0.28	-0.32	-0.35	-0.51	-0.39	-0.36
EPS Growth Rate	-	15.35%	6.85%	46.79%	-23.82%	-7.44%
Price (31st March)	0.00	0.00	7.50	1.21	23.95	15.02
Historic P/E	0.00	0.00	-21.63	-2.38	-61.77	-41.85
Current P/E (TTM)			-41.85			
Current PE/ROE			-			
Shareholder funds (In \$ million)	20.10	-0.70	-63.70	-36.60	28.20	296.00
Minority Interest(In \$ million)	0.00	0.00	0.00	0.00	0.00	0.00
Debt (In \$ million)	6.30	16.60	73.30	44.10	90.20	43.80
Cash (In \$ million)	38.20	43.70	21.60	9.20	95.30	135.60
No. of shares (In million)	11.60	11.80	67.30	157.30	229.30	306.80
Closing Price	0.00	0.00	7.50	1.21	2.10	15.02
Market Cap (In \$ million)	0.00	0.00	504.75	190.33	481.53	4,608.14
Enterprise Value (In \$ million)	0.00	0.00	556.45	225.23	476.43	4,516.34

	Ratios					
D/E	0.31	-23.71	-1.15	-1.20	3.20	0.15
PE/ROE	-	-	-	-	-	-
ROCE	-	-	-	-	-	-
ROE	-	-	-	-	-	-
EV/Sales	0.00	0.00	26.25	7.24	10.38	67.11

	FY19	FY20	FY21	FY22	FY23	TTM
R&D	48	54	59	76	51	63
% of revenue	621%	418%	279%	246%	112%	93%



IBM

Market Cap
\$ 2,08,316 Mn

Revenue
\$ 61,867 Mn

EBITDA
\$ 11,951 Mn


PAT Margin
\$ 6,397 Mn

As of Sept 2024

Product Offerings

Granite developed by IBM

- Granite is a family of foundation models developed by IBM, specifically designed for business applications. They are trained on diverse, business-relevant datasets, including Internet data, academic content, legal documents, and financial reports. This ensures the models are familiar with industry-specific jargon and context. Additionally, these models prioritize trust and governance, incorporating IBM's AI ethics principles to mitigate risks like bias and ensure data quality throughout the training process.
- The models were trained on over 12 trillion tokens from 12 natural languages and 116 programming languages. Granite models are also environmentally efficient and support businesses in creating customized AI solutions without compromising on data privacy, as IBM ensures users retain ownership of their data and models.
- The IBM Granite 3.0 models are available as open-source under the permissive Apache 2.0 license, making them accessible for customization and deployment in various use cases. Additionally, the instruct variants of these models are offered for commercial use on IBM's watsonx.ai platform, along with integrations in platforms like NVIDIA NIM and Google Cloud's Vertex AI Model Garden. This ensures flexibility for enterprise applications and edge deployments.
- **Granite-13b-chat-v2:** This model is optimized for conversational use cases, such as virtual assistants. It builds on its predecessor with improved contextual understanding and user interaction capabilities.
- **Granite-13b-instruct-v2:** Designed for tasks like text classification, extraction, and summarization. It excels at identifying mentions, summarizing long texts, and extracting relevant information from data.
- **Granite-20b-multilingual:** Supports English, German, Spanish, French, and Portuguese. It is geared for multilingual applications, including translation and question-answering tasks.



Average scores across benchmarks included in Hugging Face's RAGBench Leaderboard

- **Granite-8b-japanese:** Specializes in Japanese language tasks, including translation between Japanese and English, and performs well in extraction and classification.
- **Granite Code Models:** These are instruction-tuned models for coding tasks such as code generation, debugging, and language translation. The model sizes range from 3 billion to 34 billion parameters and support a wide variety of programming languages like Python, Java, C++, and more.



IBM

Market Cap
\$ 2,08,316 Mn

Revenue
\$ 61,867 Mn

EBITDA
\$ 11,951 Mn

PAT Margin
\$ 6,397 Mn

As of Sept 2024

Product Offerings

WATSONX.AI– One-stop platform to simplify the development process. Offering enterprise-grade AI development for business by managing the full cycle starting from Data generation to building an AI focused on using RAG for more accurate responses. Supports Developers by offering pre-built SDK, IDEs and APIs.

Pricing: \$0.1 per 1 million tokens and \$5.22 per hour.

Capabilities

- **Enterprise search using Retrieval-Augmented Generation (RAG)** – Optimize enterprise search by improving semantic search accuracy and relevance across multiple documents, tables, images and data inputs, providing effective responses based on near real-time data feeds.
- **Conversational AI chatbots** – Deploy gen AI-powered chatbots using RAG patterns for 24x7 customer support. These conversational AI chatbots, powered by watsonx.ai studio’s proprietary and open-source foundation models, engage customers naturally and intuitively.
- **AI-powered employee training** – Use gen AI to rapidly onboard new employees, using LLMs to connect relevant data to a central knowledge repository. Conversational search enables employees to quickly access accurate, up-to-date information on product-related or HR questions, reducing the need for manual research and minimizing the risk of errors, thereby streamlining the onboarding process and empowering employees to hit the ground running.
- **Conversational Assistant** – Allows the company to build, train and deploy the models on Watsonx.ai, which can be integrated into the existing chatbots to deliver contextual responses.

- Scalability

99%

Improvement in turnaround time of journey testing.

[Learn more about Vodafone](#) →

75%

Faster hiring processes anticipated.

[Learn more about Silver Egg Technology](#) →

>40%

Reduction in cloud-related operating costs.

[Learn more about Water Corporation](#) →

- **Pattern Analyses** – Extract insights from both structured and unstructured data. Gives prediction by analyzing patterns.
- **Content Generation** – To build models to generate various content types, such as ideas for marketing and sales campaigns, emails, blogs, social media posts, automated reports, and scripts.
- **Offers Multiple foundation model options**- Watsonx.ai integrates many Hugging Face open-source libraries, and models from third-party providers like Meta, Mistral and their self-trained IBM Granite series.
- **Scale-able Faster Application development**– Enhance developer productivity with AI-recommended code based on natural language inputs or existing source code.

Pricing: \$0.1 per 1 million tokens and \$5.22 per hour.



IBM

Market Cap
\$ 2,08,316 Mn

Revenue
\$ 61,867 Mn

EBITDA
\$ 11,951 Mn

PAT Margin
\$ 6,397 Mn

As of Sept 2024

Product Offerings

WATSONX.DATA

- Offers a comprehensive solution for managing and utilizing enterprise data for AI and analytics. It is built on an open lakehouse architecture that combines unstructured and structured data.
- Interoperability – Allows multiple data formats.
- 50% cost-effective by decreasing the workload cost.
- **Pricing:** Supporting services \$3 per hour.
- **Clients-** IBM watsonx.data enhances **Cogniware's** Argos platform by reducing implementation time by 45%, speeding data processing by 60%, and enabling seamless integration of diverse data sources for more efficient investigations.

WATSONX.GOVERNANCE

- The solution is designed to ensure responsible, transparent, and compliant use of AI across an organization's workflows. It provides tools for managing the lifecycle of AI and ML models, from development to deployment, while ensuring alignment with governance and regulatory requirements.
- **Clients** – IBM is enhancing the US Open's digital experience by integrating AI-generated insights and content through its watsonx platform, improving fan engagement and automating editorial workflows
- IBM's watsonx.governance enables **Deloitte** to establish trustworthy AI foundations by supporting regulatory compliance, risk management, and lifecycle governance for AI applications.

Watson NLP

- It can communicate in real-time with the help of generative natural language.
- **Clients-** IBM is helping its client by leveraging IBM Cloud Pak for AIOps to enhance IT operations with AI-driven insights, improving efficiency, reducing time-to-resolution, and enabling better decision-making.
- **Pricing:** \$125 per month for 250,000 API calls per month

Watson Assistant


- Assistance in tasks by process automation.
- Pricing:** Starting at \$3,000 per month.

Watson Text to speech

- It can convert the text into voice.
- **Clients-** IBM is helping its client, **CodeObjects**, enhance customer service efficiency and reduce costs during disasters by leveraging Watson Text to Speech as part of the InsurBot.ai AI-powered assistant.

Watson restaurant order

- Maximize drive-thru revenue by increasing the efficiency
- Features
 - Vehicle Detection
 - Order initiation
 - Audio processing
 - Order display
 - Order completion
 - Crew notification



IBM

Market Cap
\$ 2,08,316 Mn

Revenue
\$ 61,867 Mn

EBITDA
\$ 11,951 Mn

PAT Margin
\$ 6,397 Mn

As of Sept 2024

Product Offerings

Watson Code assistant

- It can provide code assistance to developers

Watson Speech to text

- It can convert the voice into text.

IBM Z Anomaly Analytics

- Advanced solution designed for monitoring and identifying anomalies in IBM Z mainframe environments. It uses ML to analyze operational log and metric data, creating a baseline of "normal" behavior. Deviations from this baseline are flagged as potential issues, enabling IT operations teams to proactively address anomalies before they escalate into critical problems.
- **ML-Based Detection:** The solution models expected behavior based on historical data, allowing it to detect issues without relying on static thresholds. This approach adapts to variations in workload patterns, such as differences between weekday and weekend operations.
- **Comprehensive Analysis:** IBM Z Anomaly Analytics integrates data from various subsystems, including Db2, CICS, and MQ, to monitor KPIs. It employs multivariate analysis, reducing false alarms by considering the broader operational context.
- **Unified Insights:** Anomalies detected in both log messages and metrics are presented in a consolidated user interface, enabling easier correlation and diagnosis. This also includes event notifications with detailed evidence, which can integrate into broader IT operations platforms like IBM Cloud Pak for Watson AIOps.
- **Operational Efficiency:** By proactively identifying potential issues, the solution reduces the mean time to resolution (MTTR) for IT operations teams, minimizing system downtime and operational disruptions.


IBM Watson Discovery

- Platform designed for intelligent search and content analytics, enabling businesses to extract insights from structured and unstructured data.
 - **Customizable AI Models:** Allows businesses to train models specific to their industry or domain using supervised learning and active learning techniques.
 - **Smart Document Understanding (SDU):** Helps users label and analyze text, including headers, tables, and other document structures.
 - **Optical Character Recognition (OCR):** Extracts information from images of text, even in suboptimal conditions.
 - **Enterprise-grade Security:** Ensures data privacy with features like encryption and compliance with GDPR.

Use case:

- **Legal and Contract Analysis:** Automates legal research and contract management, helping firms save time and improve accuracy.
- **Insurance and Financial Services:** Accelerates processes like risk assessment and document analysis.
- **Healthcare and Life Sciences:** Helps researchers and healthcare professionals extract insights from clinical and scientific data.
- **Education and Public Sector:** Automates insights from documents to streamline operations and decision-making.

Pricing: Starting \$500



IBM

Market Cap
\$ 2,08,316 Mn

Revenue
\$ 61,867 Mn

EBITDA
\$ 11,951 Mn

PAT Margin
\$ 6,397 Mn

As of Sept 2024

Product Offerings

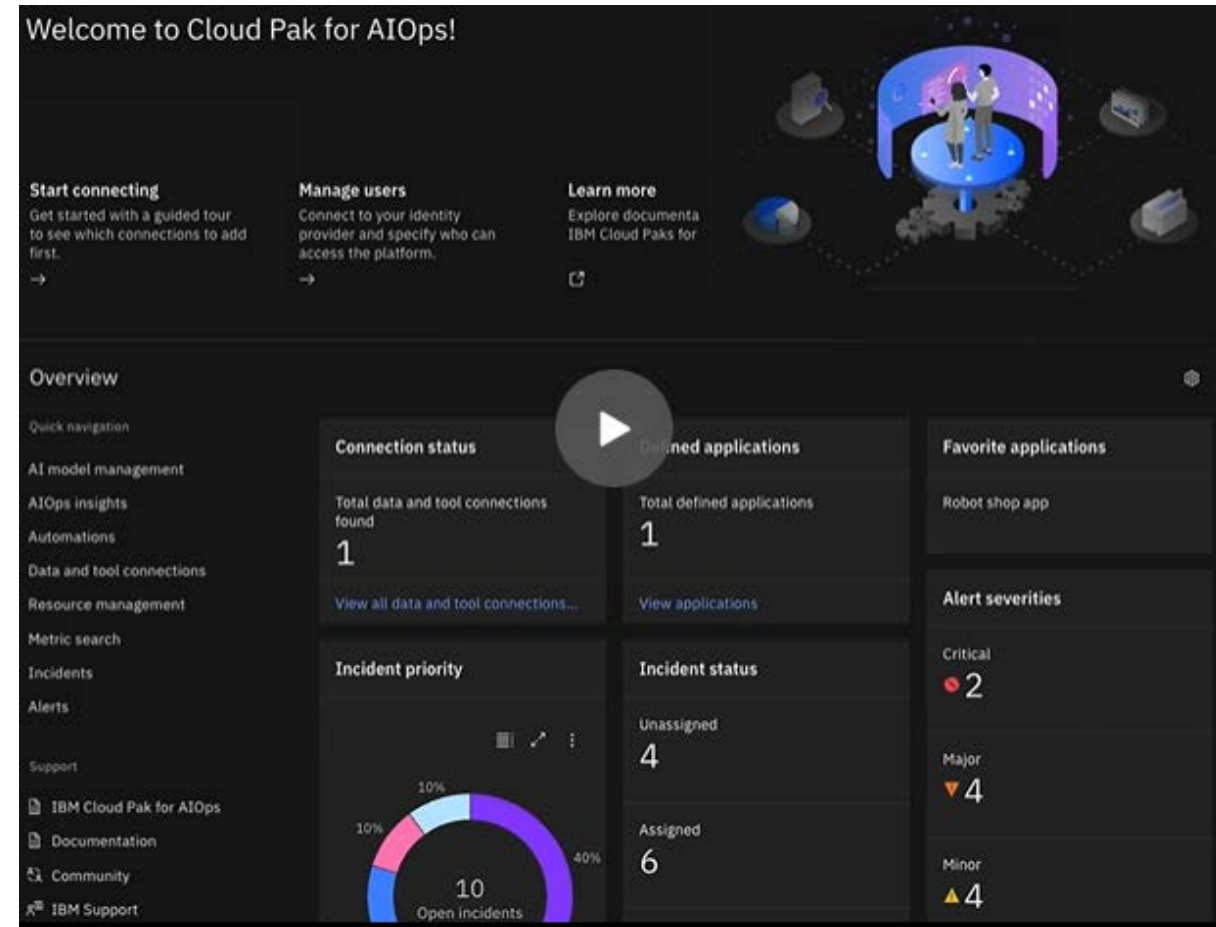
IBM Cloud Pak for AIOps

- Platform designed to enhance IT operations by automating incident detection, root cause analysis, and resolution. It integrates AI and ML with observability tools to manage complex hybrid cloud environments efficiently.
- Event Correlation and Analysis – Uses AI to group related alerts and events, reducing noise and helping IT teams focus on critical incidents.
 - Predictive Insights
 - Incident Automation
 - End-to-End Observability
 - Hybrid and Multi-Cloud Support:
 - Supports deployment across on-premises, public cloud, or hybrid environments.

Benefits:

- Improved Uptime: Reduces mean time to detect (MTTD) and mean time to resolve (MTTR) incidents.
- Cost Efficiency: Automates manual tasks, freeing up resources for strategic projects.
- Scalability: Handles the complexity of modern, distributed IT environments.

Clients- IBM Cloud Pak for Watson AIOps enables **Electrolux** to automate IT operations, enhance event correlation, reduce resolution times, and support sustainability goals through optimized resource usage.





IBM

Market Cap
\$ 2,08,316 Mn

Revenue
\$ 61,867 Mn

EBITDA
\$ 11,951 Mn

PAT Margin
\$ 6,397 Mn

As of Sept 2024

Financials

Particulars	FY19	FY20	FY21	FY22	FY23	TTM
Revenue (In \$ million)	56,498	54,195	56,571	59,879	61,123	61,867
EPS Growth Rate	-	-4.08%	4.38%	5.85%	2.08%	1.22%
EBITDA (In \$ million)	13,598	11,357	12,000	6,475	14,291	11,951
EBITDA Margin	24.07%	20.96%	21.21%	10.81%	23.38%	19.32%
PAT (In \$ million)	9,431	5,590	5,743	1,639	7,502	6,397
PAT Margin	16.69%	10.31%	10.15%	2.74%	12.27%	10.34%
EPS (In \$)	10.23	6.06	6.23	1.78	8.14	6.85
EPS Growth Rate	-	-40.73%	2.74%	-71.46%	357.72%	-15.76%
Price (31st March)	121.46	136.69	129.21	134.21	163.46	223.18
Historic P/E	11.88	22.55	20.75	75.51	20.09	32.56
Current P/E (TTM)			32.56			
Current PE/ROE			0.98			
Shareholder funds (In \$ million)	20,841.00	20,597.00	18,901.00	21,944.00	22,533.00	24,448.00
Minority Interest(In \$ million)	144.00	129.00	95.00	77.00	80.00	82.00
Debt (In \$ million)	68,158.00	65,088.00	55,140.00	54,013.00	59,935.00	60,126.00
Cash (In \$ million)	8,868.00	13,788.00	7,250.00	8,738.00	13,441.00	13,702.00
No. of shares (In million)	892.80	896.60	904.60	912.30	922.10	933.40
Closing Price	121.46	136.69	129.21	134.21	160.08	223.18
Market Cap (In \$ million)	1,08,439.49	1,22,556.25	1,16,883.37	1,22,439.78	1,47,609.77	2,08,316.21
Enterprise Value (In \$ million)	1,67,873.49	1,73,985.25	1,64,868.37	1,67,791.78	1,94,183.77	2,54,822.21

	Ratios					
D/E	3.27	3.16	2.92	2.46	2.66	2.46
PE/ROE	0.26	0.83	0.68	10.11	0.60	1.24
ROCE	15.25%	13.23%	16.19%	8.52%	17.31%	14.12%
ROE	45.25%	27.14%	30.38%	7.47%	33.29%	26.17%
EV/Sales	2.97	3.21	2.91	2.80	3.18	4.12

	FY19	FY20	FY21	FY22	FY23	TTM
R&D	5,910	6,262	6,488	6,567	6,775	7,260
% of revenue	10%	12%	11%	11%	11%	12%



NVIDIA

Market Cap
\$ 34,83,527 Mn

Revenue
\$ 1,13,269 Mn

EBITDA
\$ 72,741 Mn

PAT Margin
\$ 63,074 Mn

As of Sept 2024

Product Offerings

[1] Generative AI Solutions:

- NVIDIA's full-stack AI platform integrates hardware, software, and APIs to empower businesses to harness the transformative potential of generative AI
- It offers enterprise-ready APIs from the NVIDIA API Catalog to enable developers to prototype and deploy applications seamlessly."
- Combines software and hardware (DGX Systems, A100 GPUs) to optimize AI workflows
- Provides tools, frameworks, and training resources for developers to build and refine AI models

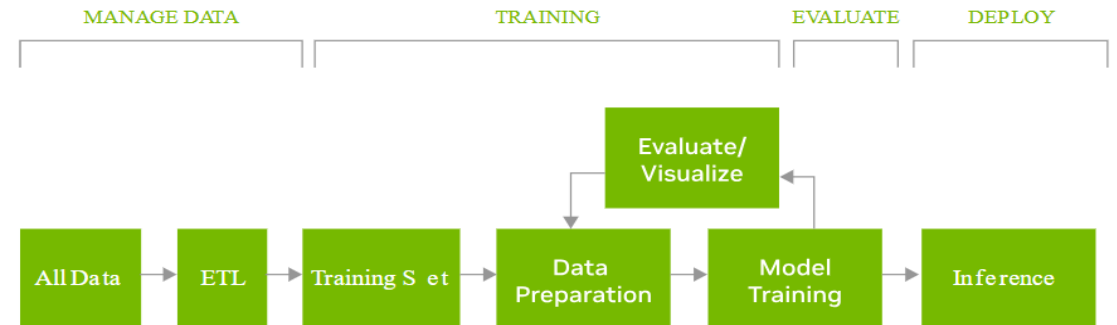
Use Cases:

- **Automotive:** streamlining vehicle design and manufacturing, accelerating self-driving technology development
- **Healthcare:** Expediting drug discovery with AI-enabled molecular analysis, enabling patient record management
- **Telecommunications:** Optimizing networks and predicting hardware failures
- **Financial Services:** personalized customer service, risk mitigation, and automated document processing, regulatory compliance.
- **Retail:** Generating personalized product recommendations and marketing campaigns
- **Media & Entertainment:** content creation and personalization

Solutions:

- Accelerated Pipelines: Assembly time reduced from 4–6 weeks to a few clicks.
- Data Integration: Processed 5TB of multimodal data for actionable insights.
- Scalable Experimentation: Saved 7–10 months of setup time using DGX Cloud.

Machine Learning to Deep Learning, All on GPU





NVIDIA

Market Cap
\$ 34,83,527 Mn

Revenue
\$ 1,13,269 Mn

EBITDA
\$ 72,741 Mn

PAT Margin
\$ 63,074 Mn

As of Sept 2024

Product Offerings

[2] AI Inference Solutions:

- AI inference solutions involve deploying pre-trained AI models to make real-time predictions or generate insights from new data.
- NVIDIA leads in AI inference solutions by offering a full-stack approach combining software, and tools optimized for diverse workloads

Tech:

- **NVIDIA DGX Cloud:** A virtual supercomputer in the cloud that helps companies build and run advanced AI projects without needing to buy and manage expensive hardware themselves. It provides powerful GPUs (graphics processing units) and software tools, making it easy for businesses to train AI models and use them in real-world applications
- **Triton Inference Server:** You can simply upload your trained AI model to Triton, and it takes care of everything: running the model, handling user requests, and giving results quickly.

- **Customers:** Amdocs, Snapchat, American Express, Encord, GE HealthCare, Infosys, NIO, Yahoo Japan
 - Amdocs: Developed a generative AI platform for telcos using NVIDIA DGX Cloud and NIM microservices, improving latency and accuracy while reducing costs.
 - Snapchat: Used Triton Inference Server for optical character recognition and apparel shopping experiences, scaling effectively while reducing time to production.

[3] Conversational AI Solutions:

- Conversational AI systems are commonly used in applications like chatbots, AI virtual assistants, and digital humans, enabling real-time, personalized, and interactive conversations between humans and machines.

Solutions:

- **NVIDIA Riva:** A platform designed for building high-performance conversational AI applications with a focus on speech recognition and speech synthesis
 - Multilingual Speech Recognition: Riva supports accurate transcription of spoken language into text, supporting many languages and dialects, making it suitable for diverse global applications
 - Speech Translation: Translating speech in real-time, enabling multilingual conversations
 - Speech Synthesis (Text-to-Speech): Creating lifelike, expressive synthetic voices for conversational agents
- **NVIDIA NeMo:** A platform for building and deploying LLMs, including generative AI for text and multimodal tasks. It allows developers to create AI systems capable of understanding and generating text in a conversational manner
- **Generative AI Deployment:**
 - **NVIDIA GPUs:** Conversational AI platforms like Riva and NeMo are optimized for NVIDIA's powerful GPUs, leveraging CUDA and Tensor Cores to accelerate training and inference, ensuring low latency, and real-time interactions.
 - **Use Cases:** Digital Humans, Content Generation, AI Virtual Assistant, Agent Assist, AI Translation, AI Robot
- **Customers:** Infosys, Kore.ai, Intelligent Voice, BotPress, Data Monsters, Morningstar



NVIDIA

Market Cap
\$ 34,83,527 Mn

Revenue
\$ 1,13,269 Mn

EBITDA
\$ 72,741 Mn

PAT Margin
\$ 63,074 Mn

As of Sept 2024

Product Offerings

[4] Vision AI Solutions (Metropolis)

- An advanced platform that empower developers with tools to develop video analytics and deployment

Use Cases:

- **Smart Cities & Spaces:**

- Intelligent Traffic Systems: In a smart city like Singapore, traffic cameras combined with Metropolis AI could automatically adjust traffic lights to avoid congestion or accidents, improving road safety and reducing fuel consumption. The system can also identify traffic violations (like running red lights) and alert authorities in real time.
- Parking and Access Control: Metropolis use AI-powered cameras to monitor parking spaces and automatically detect vehicle entry or exit. The system can offer real-time availability of parking spots via an app, improving the convenience for users and reducing time spent searching for parking

- **Retail & Logistics**

- Customer and store analytics: data generated from PoS, cameras, and sensors help to determine demographics, buying preferences, queue analytics, etc. Retailer using this data run promotions and customer preferences.
- Store simulation to enhance store layout and merchandising
- Stockout and inventory management: scan items to check stock levels and alert associates to restock, correct shelf location, and even adjust prices when needed.

- **Industrial and Manufacturing**

- Improve industrial inspection through Automated Optical Inspection
- Remote asset monitoring
- Safety and compliance

- **Healthcare**

- Temperature sreening
- Patient monitoring
- PPE detection



NVIDIA

Market Cap
\$ 34,83,527 Mn

Revenue
\$ 1,13,269 Mn

EBITDA
\$ 72,741 Mn

PAT Margin
\$ 63,074 Mn

As of Sept 2024

Product Offerings

[5] Cybersecurity AI (Morpheus)

- NVIDIA Morpheus uses GPU acceleration with generative AI framework designed to help enterprises build, customize, and scale cybersecurity applications more efficiently and cost-effectively

What is problem

- Traditional Cybersecurity systems run on CPU Based Cyber Security
- On an average threat are detected 287 days after they originally infected the network
- 20% on average threats go missing

- It uses GPU acceleration with generative AI to scan all the traffic in real-time across the entire network, including data center, edge, gateway, and centralized compute
- GPU acceleration at scale—up to 600X faster than CPU-only solutions—reducing time to detect from weeks to minutes, while drastically lowering operational costs.

Use Cases -

- Digital Fingerprinting: Morpheus can uniquely fingerprint users, services, accounts, and machines within an organization. It continuously analyzes behavioral patterns and flags anomalies that may indicate potential security threats
- Ransomware Detection: Morpheus enables real-time detection and mitigation of ransomware attacks by analyzing files, network traffic, and system behavior to identify patterns associated with ransomware activity
- Phishing Detection: By analyzing email content and metadata, Morpheus can identify phishing attempts and other forms of social engineering. Using generative AI, improves the detection of new, evolving phishing tactics
- Fraudulent Transaction and Identity Detection: It helps prevent fraudulent transactions in real-time by analyzing patterns in financial data and user behavior to flag suspicious activities before they can cause damage



NVIDIA

Market Cap
\$ 34,83,527 Mn

Revenue
\$ 1,13,269 Mn

EBITDA
\$ 72,741 Mn

PAT Margin
\$ 63,074 Mn

As of Sept 2024

Financials

Particulars	FY20	FY21	FY22	FY23	FY24	TTM
Revenue (In \$ million)	10,918	16,675	26,914	26,974	60,922	1,13,269
EPS Growth Rate	-	52.73%	61.40%	0.22%	125.85%	85.92%
EBITDA (In \$ million)	3,227	5,819	11,215	7,121	34,480	72,741
EBITDA Margin	29.56%	34.90%	41.67%	26.40%	56.60%	64.22%
PAT (In \$ million)	2,796	4,332	9,752	4,368	29,760	63,074
PAT Margin	25.61%	25.98%	36.23%	16.19%	48.85%	55.69%
EPS (In \$)	0.11	0.17	0.39	0.18	1.19	2.54
EPS Growth Rate	-	54.94%	125.12%	-55.21%	581.32%	112.60%
Price (31st March)	6.11	13.48	28.00	16.00	53.14	140.11
Historic P/E	54.50	77.61	71.61	91.36	44.53	55.23
Current P/E (TTM)			55.23			
Current PE/ROE			0.80			
Shareholder funds (In \$ million)	12,204.00	16,893.00	26,612.00	22,101.00	42,978.00	65,899.00
Minority Interest(In \$ million)	0.00	0.00	0.00	0.00	0.00	0.00
Debt (In \$ million)	2,643.00	7,597.00	11,831.00	12,031.00	11,056.00	10,225.00
Cash (In \$ million)	10,897.00	11,561.00	21,208.00	13,296.00	25,984.00	38,487.00
No. of shares (In million)	24,720.00	25,100.00	25,350.00	25,070.00	24,940.00	24,862.80
Closing Price	6.11	13.48	28.00	16.00	52.25	140.11
Market Cap (In \$ million)	1,51,039.20	3,38,348.00	7,09,800.00	4,01,120.00	13,03,115.00	34,83,526.91
Enterprise Value (In \$ million)	1,42,785.20	3,34,384.00	7,00,423.00	3,99,855.00	12,88,187.00	34,55,264.91

	Ratios					
D/E	0.22	0.45	0.44	0.54	0.26	0.16
PE/ROE	2.38	3.03	1.95	4.62	0.64	0.58
ROCE	21.74%	23.76%	29.17%	20.86%	63.81%	95.56%
ROE	22.91%	25.64%	36.65%	19.76%	69.24%	95.71%
EV/Sales	13.08	20.05	26.02	14.82	21.14	30.50

	FY19	FY20	FY21	FY22	FY23	TTM
R&D	2,829	3,924	5,268	7,339	8,675	11,665
% of revenue	26%	24%	20%	27%	14%	10%

Sources

- ▶ <https://aws.amazon.com/>
- ▶ <https://gemini.google.com/?hl=en-IN>
- ▶ <https://www.soundhound.com/>
- ▶ <https://www.ibm.com/in-en>
- ▶ <https://www.nvidia.com/en-in/>
- ▶ <https://seekingalpha.com/>

DISCLAIMER

- *ValPro has taken due care in the preparation of this document. All information/views provided herein, have been prepared from publicly available information and proprietary database of ValPro and is believed to be reliable at the time of publication. ValPro reserves its rights to withdraw or amend information contained in this document at any time based on developments and/or new information becoming available to it.*
- *ValPro, and ValPro's shareholders, employees, directors, agents and representatives ('Affiliates') do not warrant the completeness or accuracy of the information contained herein and shall in no circumstance whatsoever be liable or responsible for any inaccuracies, omissions, mistakes or errors in this paper and for any economic or any other loss or damage (direct or indirect), incurred or suffered by the reader/ user or any other person, arising out of or in connection with such information/views. Any person relying on information contained herein is advised to undertake their own independent diligence, analysis and verification.*
- *Information contained herein is meant to be for the recipient only. By accepting the information, the recipient is implicitly undertaking to keep the contents thereof as confidential. Dissemination of such information to any third party without express written consent of ValPro is prohibited. If so requested by ValPro in writing, the recipient shall be obligated to return and/or destroy this document/information retained by it in its records whether in physical or digital form.*
- *This document does not constitute an offer, invitation or recommendation to buy, sell, subscribe for or issue any securities or invest in any business or asset or a solicitation of any such offer or invitation. ValPro shall not be liable for any decision made by the recipient or any third party on the basis of this information.*
- *ValPro and its Affiliates, may have a position or vested interest in the company or businesses covered by this document or may be representing such companies or businesses in relation to their business in transactions. By receiving this information, the recipient acknowledges the fact of the aforesaid possible interest of ValPro and/or its Affiliates in the circulation of this note and waives any claims in this regard.*
- *By receiving this information, the recipient acknowledges that any dispute arising herefrom shall be subject to the laws of India and shall be resolved by a sole arbitration appointed in accordance with the Arbitration and Conciliation Act 1996. The venue of arbitration shall be New Delhi and language of arbitration shall be English.*

CONTACT US

EMAIL: info@valpro.co.in

ADDRESS: A-125 NEETI BAGH, NEW DELHI – 110049, INDIA