

September 10, 2023 11:28 PM GMT

Tesla Inc

Unlocking Tesla's AI Mojo... Enter the Dojo: Upgrade to OW, PT \$400, Top Pick

The autonomous car has been described as the mother of all AI projects. In its quest to solve for autonomy, Tesla has developed an advanced supercomputing architecture that pushes new boundaries in custom silicon and may put Tesla at an asymmetric advantage in a \$10trn TAM.



Morgan Stanley does and seeks to do business with companies covered in Morgan Stanley Research. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of Morgan Stanley Research. Investors should consider Morgan Stanley Research as only a single factor in making their investment decision.

For analyst certification and other important disclosures, refer to the Disclosure Section, located at the end of this report.

+ = Analysts employed by non-U.S. affiliates are not registered with FINRA, may not be associated persons of the member and may not be subject to FINRA restrictions on communications with a subject company, public appearances and trading securities held by a research analyst account.

Contributors



MORGAN STANLEY & CO. LLC

Adam Jonas, CFA

Equity Analyst
+1 212 761-1726

Adam.Jonas@morganstanley.com



MORGAN STANLEY & CO. LLC

Daniela M Haigian

Research Associate
+1 212 761-6071

Daniela.Haigian@morganstanley.com



MORGAN STANLEY & CO. INTERNATIONAL PLC+

Edward Stanley

Equity Strategist
+44 20 7425-0840

Edward.Stanley@morganstanley.com



MORGAN STANLEY & CO. INTERNATIONAL PLC+

Matias Ovrum

Equity Strategist
+44 20 7425-9902

Matias.Ovrum@morganstanley.com



MORGAN STANLEY TAIWAN LIMITED+

Charlie Chan

Equity Analyst
+886 2 2730-1725

Charlie.Chan@morganstanley.com



MORGAN STANLEY & CO. LLC

Joseph Moore

Equity Analyst
+1 212 761-7516

Joseph.Moore@morganstanley.com



MORGAN STANLEY & CO. LLC

Mason Wayne

Research Associate
+1 212 761-6012

Mason.Wayne@morganstanley.com

Contents

- 7 Executive Summary - Our Thesis in 3 Charts
- 9 Welcome to the Dojo
- 20 A Closer Look at the World of Custom Silicon
- 23 Implications for Tesla: Upgrade to OW, Increase PT to \$400, New Top Pick
- 35 Further Industry Implications
- 39 Tesla Valuation Matrix
- 40 Risk Reward - Tesla Inc (TSLA.O)
- 42 Appendix - Dojo in Detail
- 56 Appendix - Glossary

Unlocking Tesla's AI Mojo... Enter the Dojo: Upgrade to OW, PT \$400, Top Pick

Investors have long debated whether Tesla is an auto company or a tech company. We believe it's both, but see the biggest value driver from here being software and services revenue. The same forces that have driven AWS to reach 70% of AMZN total EBIT can work at Tesla, in our view, opening up new addressable markets that extend well beyond selling vehicles at a fixed price. **The catalyst?** Dojo, Tesla's custom supercomputing effort in the works for the past 5 years. *Version 12 of Tesla's full self driving system (OTA by year-end) and Tesla's next AI day (early 2024) are worth watching.*

STOCK RATING

Overweight

INDUSTRY VIEW

In-Line

PRICE TARGET

\$400.00

We believe that Dojo can add up to \$500bn to Tesla's enterprise value, expressed through a faster adoption rate in Mobility (robotaxi) and Network Services (SaaS). The change drives our PT increase to \$400 vs. \$250 previously. We upgrade to Overweight and make Tesla our Top Pick.

What is Dojo? Dojo is a purpose-built supercomputer designed in-house by Tesla to *train* the full-self-driving (FSD) system that sits inside every Tesla vehicle. **Why is Tesla doing it?** Tesla's cars are sensor encrusted robots making life and death decisions in highly unpredictable environments and driving situations. Tesla's ability to improve the efficacy of its full self driving system is limited by the ability to collect and process real world video data from the edge and to train these robots from the experience of its vehicle fleet in service, which is 5mm units today and closer to 50mm by end of decade. **Tesla management has said it needs as much compute power/NVIDIA GPU clusters it can get its hands on and currently it cannot physically secure the amount of chips necessary to train cars. In addition, they believe they can develop a more efficient system for their specific needs while not funding a supplier's 60% gross margin.**

With a highly experienced semiconductor team, Tesla has built a custom AI ASIC chip, that, due to its core function of processing vision-based data for autonomous driving use cases, can operate more efficiently (energy consumption, latency) than the leading cutting-edge general-purpose chips on the market (NVIDIA's A100s and H100s), and at a fraction of the cost. *Dojo is a training computer made up of many thousands of D1 chips housed in an AI data center. It trains the inference engine (FSD chip) that sits within the vehicles at the edge which Tesla has designed in-house for the past 7 years.*

Tesla is not the first tech player to attempt to build a custom silicon system in-house. What's

Tesla Inc (TSLA.O, TSLA US)

Top Pick

Autos & Shared Mobility | United States of America

Stock Rating	Overweight
Industry View	In-Line
Price target	\$400.00
Shr price, close (Sep 7, 2023)	\$251.49
Mkt cap, curr (mm)	\$875,508
52-Week Range	\$313.80-101.81

Fiscal Year Ending

	12/22	12/23e	12/24e	12/25e
EPS (\$)***	4.07	3.26	3.92	5.58
Prior EPS (\$)***	-	-	3.74	4.94
ModelWare EPS (\$)	3.62	2.77	3.43	5.07

Unless otherwise noted, all metrics are based on Morgan Stanley ModelWare framework

*** = Based on consensus methodology
e = Morgan Stanley Research estimates

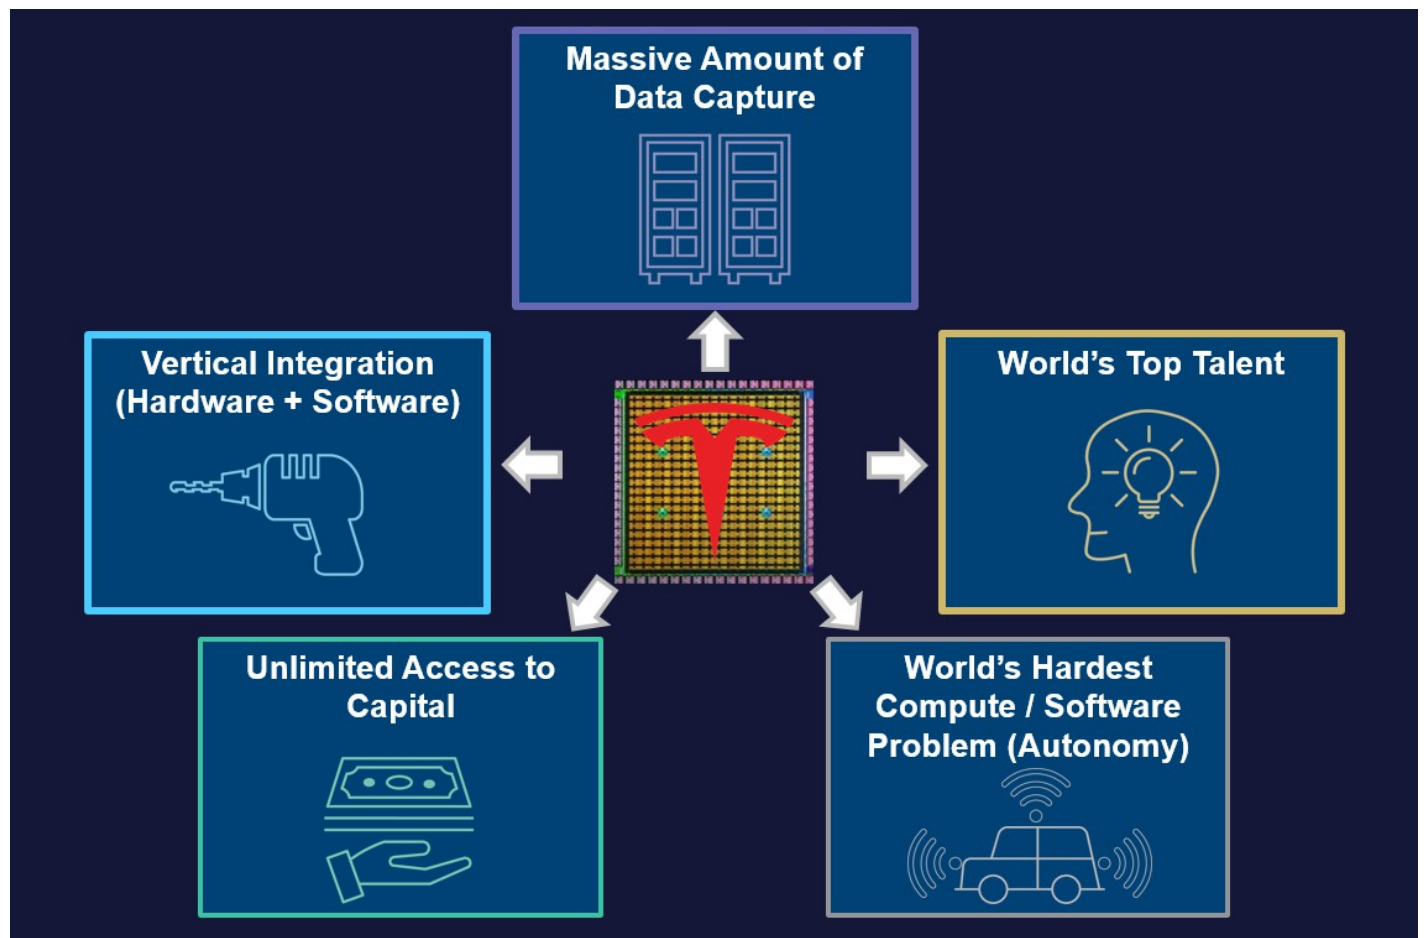
WHAT'S CHANGED

	From	To
Tesla Inc		
Price Target	\$250.00	\$400.00
Rating	Equal-weight	Overweight
Top Pick	RACE.N	TSLA.O

unique about Tesla is the company's longtime experience with advanced driver assist systems (ADAS). It has commercialized a vast network of vehicles that is constantly increasing (400k+ FSDs on the road already collecting data from 300+ million miles traveled). In addition, the company has brought together a world class design team, and has allocated expansive resources towards the autonomy problem. Like other tech platforms, Tesla pursues high vertical integration in key technology domains to enable high iteration and continual improvement while helping to diversify away from over-reliance on 3rd party suppliers that may not be able to provide an optimal solution for Tesla's specific needs. *While it is difficult to explicitly validate the many claims Tesla has made about Dojo's cost and performance, we believe Tesla has a chance of bringing forth a competitive customized solution given the company's innovation track record and capabilities.*

In this report we present a comprehensive "primer" on Dojo involving insights and opinions across Morgan Stanley's AI/Semis research teams. We explore what Dojo is, why Tesla is doing it and how it can impact the business and the stock's valuation. We invite investors to dive into the world of custom silicon at exaFLOP scale solving some of the world's most challenging problems (autonomy) offering a gateway into vast untapped commercial potential... essentially *everything with a camera that can process data to make decisions*. The more we looked at Dojo, the more we realized the potential for underappreciated value in the stock. Like many other large cap tech stocks on your screen, we believe Tesla can reasonably test its all-time highs of \$400 over the next 12 months.


Exhibit 1: Tesla's capabilities and business model can significantly benefit from the development of custom AI tools. It's too big and too specialized an opportunity not to have in-house.



Source: Morgan Stanley Research

For years, we've tried to focus investor attention on the potential of Tesla's leadership in EV hardware (semi-autonomous electric 'robots') to convert vehicle owners into 'subscribers' generating highly recurring (and high margin) revenue. While Tesla has had some modicum of success on this front to date, we believe Tesla's in-house computing efforts have the potential to materially accelerate the network effect and speed of data capture/analysis/learning from the 1 billion miles traveled *per day* we forecast is executed by its global fleet (Tesla + 3rd party) by 2027. *The scale and complexity of the data (the collective global light vehicle fleet travels a distance of nearly 2 light years annually) combined with the high standard of safety requirement make the global mobility market highly relevant territory for the expression of the AI investment theme.*

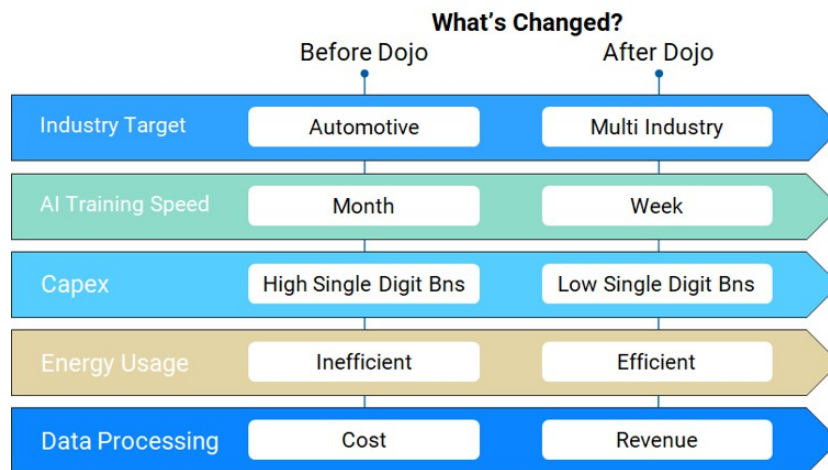
Exhibit 2: We believe Dojo can accelerate Tesla's Auto & SAAS 'Double Flywheel'



Source: Morgan Stanley Research

We collaborated with a range of AI experts at MS to weigh in on Tesla's AI ambitions and to compare and contrast them to those of NVDA and other hyper-scalers. Extending beyond the implications for Tesla, we hope this work sheds light on the potential of the broader custom silicon market.

Exhibit 3: Dojo - Before & After



Source: Morgan Stanley Research

In conjunction with this report, we feel it is prudent to include a degree of optionality for Tesla's AI potential into our revised price target which has increased to \$400 vs. \$250 previously. Our bull case valuation is raised to \$550 (vs. \$450). And our bear case valuation is raised to \$120 (vs. \$90). Within our forecasts and valuation we express the potential of Dojo through our raised assumptions, primarily for Tesla Mobility (autonomous robo-taxis) and Tesla Network Services (SaaS business derived from Tesla vehicles and 3rd party customers) in the form of faster adoption and higher ARPU. We share a comprehensive review of Tesla valuation on our new earnings forecasts which have increased approximately 20% by FY25/26 which we believe fairly underpins our \$400 price target vs. the growth and valuation multiples of the relevant tech peer group.

Executive Summary - Our Thesis in 3 Charts


We believe Dojo can represent the next step-change in market perception of Tesla. Dojo emphasizes 3 of Tesla's core capabilities: 1) speed, 2) performance, and 3) cost. In the near term, we believe Dojo can accelerate the development and monetization of Tesla's software and services business. Longer term, we see scope for Dojo to provide avenues for Tesla's software and hardware capabilities to extend well beyond the auto industry. If Dojo can help make cars 'see' and 'react,' what other markets could open up? Think of any device at the edge with a camera that makes real-time decisions based on its visual field.

Tesla estimates that Dojo can provide 6x cost saving vs current, state of the art, GPU alternatives. On our calculations, when comparing what Tesla would have to spend on equivalent compute from NVIDIA, Dojo has the potential to drive ~\$6.5bn in cost savings for Tesla over the next couple of years to reach the company's stated goal of materially increasing internal computing power by October 2024 (to 100 exaFLOPs). This is achieved by having a purpose built, in-house semiconductor and AI tech stack. Dojo became operational in July of this year and we believe the continued rollout and subsequent company announcements will provide the catalyst for investors to appreciate Dojo's potential. *We note that 6x cost savings is Tesla's claim and we are unable to verify it with specificity given the early stage of Dojo roll-out. We also note that there are other pieces of data provided by Tesla that suggest other implied cost savings outcomes that could differ from the 6x claim. Finally, just because Tesla is making a major effort to commercialize Dojo for its in-house purposes does not mean that the system will ultimately represent the*

best cost/performance alternative on the market longer term give continuous improvement of rival compute technology.

For our Tesla modeling purposes, we focused on the potential for Dojo to deliver autonomy and network services revenues at a faster attach rate with higher average monthly revenue per user (ARPU), driving a material increase to our estimates. We have NOT given Tesla credit for specific cost savings from Dojo vs. its current supercomputing budget. Nor have we given Tesla credit for any non-auto-related revenue streams. With significantly increased computing power and faster processing speeds (latency), Tesla's path to monetizing vehicle software can materialize sooner, and at higher recurring revenue rates. We also for the first time incorporate non-Tesla fleet licensing revenue into our Network Services model as we expect recent charging station cooperation will extend into FSD licensing (discussions ongoing) and operating system licensing. We now forecast Tesla Network Services to reach \$335bn in revenue in 2040 vs \$157bn previously, and expect the segment to represent over a third of total company EBITDA in 2030, doubling to over 60% of group EBITDA by 2040 (vs. 38% previously). This increase is largely driven by the emerging opportunity we see in 3rd party fleet licensing, increased ARPU, with operating leverage driving higher long-term EBITDA margin vs. prior forecast (65% from FY26 onwards, vs. 50% previously). In addition to Network Services, we indirectly ascribe the value of Dojo to our Tesla Mobility robotaxi assumptions (increased long term fleet size and margin), and 3rd Party Battery Business, as we believe the charging and FSD deals will also result in higher hardware attach.

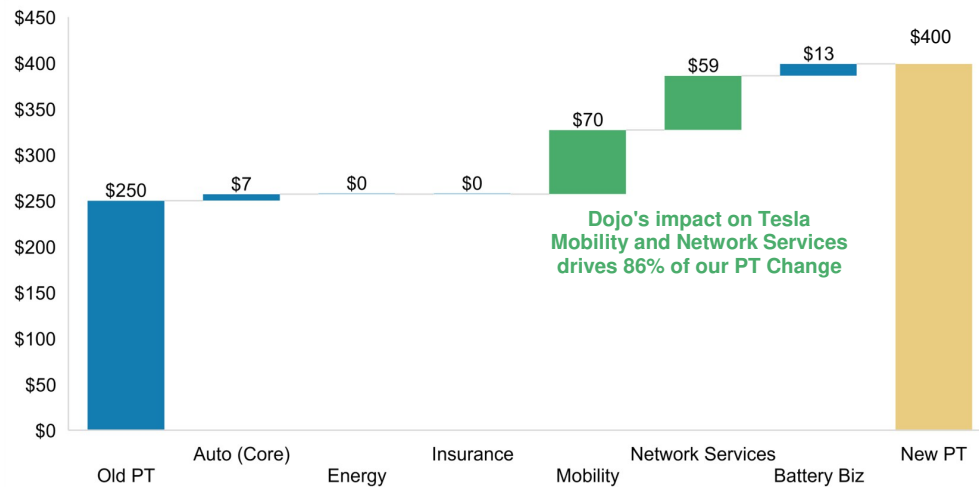
Exhibit 4: Tesla's claims of 6x performance improvements imply multi-billion-\$ cost savings from Dojo



Source: Tesla, Morgan Stanley Research

We made no changes to our assumptions for Tesla Energy or Tesla Insurance. The modest (\$7/share) increase in the value of the core Auto business was mostly related to an increase in our exit EBITDA multiple assumption to 13x from 12x previously. Our near term (FY23/FY24) core Auto assumptions for volume and gross/operating margin were unchanged.

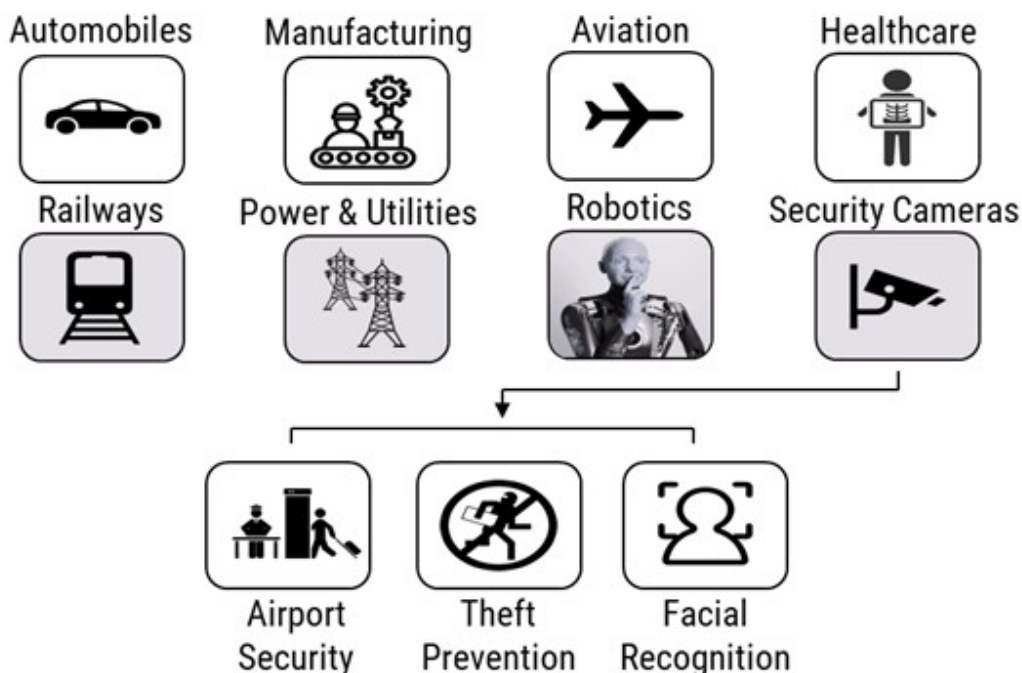
Exhibit 5: Tesla Price Target Bridge: Raised to \$400 from \$250 previously



Source: Morgan Stanley Research

Stretching your thinking. Could success in vehicle autonomy enable Tesla to become Go-To provider for visual data processing across other adjacent markets? Although Dojo is still early in its development, we believe that its applications long-term can extend beyond the auto industry. Dojo is designed to process visual data which can lay the foundation for vision-based AI models such as robotics, healthcare and security. In our view, once Tesla makes headway on autonomy and software, third party Dojo services can offer investors the next leg of Tesla's growth story.

Exhibit 6: Industries that can utilize visual AI



Source: Morgan Stanley Research

Welcome to the Dojo

A *dōjō* translates to "place of the Way" in Japanese. Others may remember this dialogue from *The Karate Kid* (Columbia Pictures, 1984):

Mr. Miyagi: No more fighting.

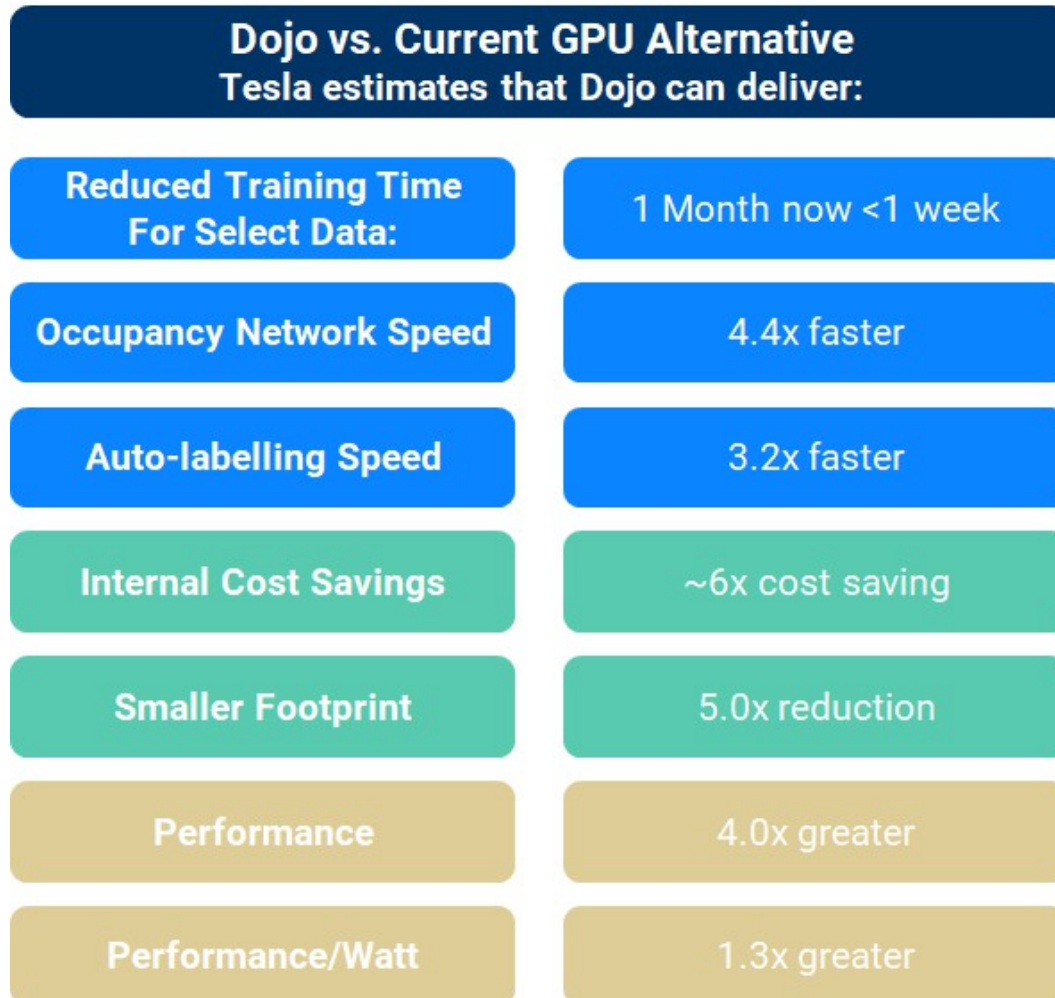
Kreese: This is a karate dojo, not a knitting class. You don't come into my dojo, drop a challenge and leave, old man. Now you get your boy on the mat, or you and I will have a major problem.

Mr. Miyagi: Too much advantage. Your dojo.

Kreese: Name a place.

Mr. Miyagi: Tournament.

Exhibit 7: Estimated Benefits of Dojo compared to Tesla's current GPU system (A100s)



Source: Company Data, Morgan Stanley Research

What is Dojo? Dojo is a purpose-built supercomputer designed in-house by Tesla to train its computer vision systems. Tesla's cars are sensor encrusted robots making life and death decisions in highly unpredictable environments and situations. Tesla's ability to improve the efficacy of its full self driving (FSD) system is limited by the ability to collect, label and process real world video data from the edge and to train these robots from the experience of its fleet in service, which is 5mm units today and closer to 50mm by end of decade. With a highly experienced semiconductor team, Tesla has built a custom AI ASIC chip, that, due to its core function of processing vision-based data for autonomous driving use cases, can operate more efficiently (energy consumption, latency) than the leading cutting-edge general-purpose chips on the market (NVIDIA's A100s and H100s), potentially at a fraction of the cost.

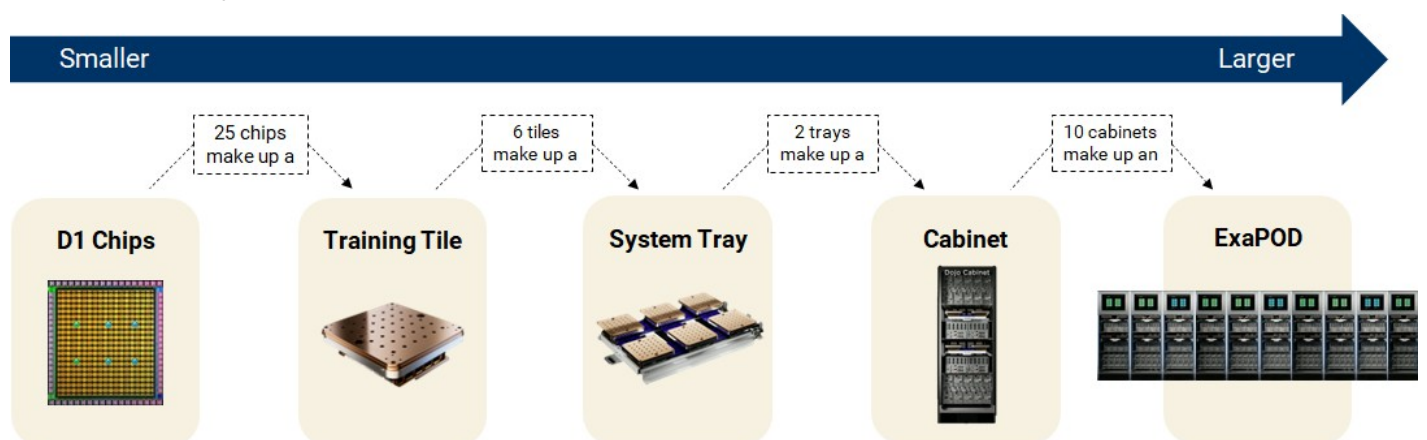
Tesla is not the first tech player to attempt to build a custom silicon system in-house, but given the company's deep understanding of ADAS (pioneer in the EV market), vast network of data that is constantly increasing (400k FSDs on the road already collecting data from 300+ million miles traveled), a world class design team, and expansive resources, in addition to the underlying need to diversify away from over-reliance on NVDA, we believe Dojo may prove competitive in its customized solution. We discuss the custom AI competitive landscape with insights from our Morgan Stanley semiconductor/AI teams in the following section, [A Closer Look at the World of Custom Silicon](#).

A key hurdle to autonomy has been corner-cases, but Elon Musk believes that autonomous driving can be 10x safer than humans (others insist robotaxis must eventually be 10,000x safer than humans), and that Dojo can be the tool to expedite the timeline.

Training for full autonomy is highly complex and is associated with a slew of ethical, legal and regulatory challenges. Models need to be trained on vast data sets to learn responses from the most mundane driving decisions to the most exotic/edge-case scenarios that may confront a driver – the amount of iterations it requires to reach parity with (and eventually exceed) the reaction time of a human driver makes it difficult to sufficiently train a neural net. According to Morgan Stanley's US semis analyst Joe Moore, vision-based training is actually less complex than large language model (LLM) training but nonetheless requires enormous amounts of data (video input from 10 cameras per car constantly on) and vastly different public safety considerations. A more powerful computational training system can meaningfully accelerate the speed at which the autonomous vehicles are trained, creating a shorter timeline to full autonomy – *Tesla estimates that Dojo will reduce the require training time for a typical workload from 1 month to less than a week.*

How is this possible? Dojo Building Blocks: 7nm chips called 'D1' are designed in-house (manufactured by TSMC). D1 chips are organized in a custom architecture and placed into an **ExaPOD** (collection of cabinets with several thousand D1 chips), which features custom-built hardware and software to further increase the chip's efficiency. Our APAC semiconductor team's channel checks suggests that somewhere between 40k - 50k D1 chips have been ordered by Tesla this year. The ExaPOD is then the building block in Tesla's supercomputer. Each complete Dojo ExaPOD supercomputer system will be able to achieve 1.1 exaFLOPs of compute (10^{18} floating operations/sec), making it amongst the **most powerful supercomputers in the world**. Tesla **plans** on building 7 ExaPODs in their Palo Alto data center. Below we map the architecture of an ExaPOD.

Exhibit 8: Tesla Dojo Building Blocks... from chip to ExaPOD



Source: Company Data. Morgan Stanley Research

The first Dojo ExaPOD (online since July 2023) was expected to target Tesla's auto-labeling networks... Oh wait... Tesla FSD V12 may not even need labeling anymore? To this point, for most self driving systems (including Tesla) a key task in processing visual data has been "labeling", which has been mostly a manual process whereby a programmer assigns labels to raw data (stop signs, kitens, potholes) so it can be utilized in training neural nets. Auto-labeling, in conjunction with Tesla's occupancy networks, which use 3D mapping to help the vehicle detect and avoid objects, have reportedly occupied ~50% of Tesla's current computing capacity. However, in its most recent update of FSD (version 12) which is currently in validation stage (scheduled for release to the public by year-end) Tesla essentially 'did away with auto-labeling.' According to Tesla, if 99% of people stop at a stop sign, then over many billions of miles, the computer *learns* that the red octagon to

the right of the road at that junction just IS a stop sign. You don't have to tell the system that it's a stop sign. It just figures it out - rather quickly.

While not claiming perfection, Tesla has described this upcoming FSD version as its 'ChatGPT moment' in terms of delivering a major step change improvement in performance of the system (without labeling, without LiDAR and without HD maps). While the in-house Dojo D1-based system is seen as a critical enabler for this development, Tesla still expects Dojo to operate alongside NVIDIA GPUs. If successful, we would expect to see Tesla rely even further on its own in-house technology.

For further details on the technical aspects of Dojo, see the [Appendix - Dojo in Detail](#).

Exhibit 9: Auto-labeling & occupancy network represent 50% of current GPU cluster usage



Source: Company data

Exhibit 10: Auto-labeling converts data into machine understandable outputs

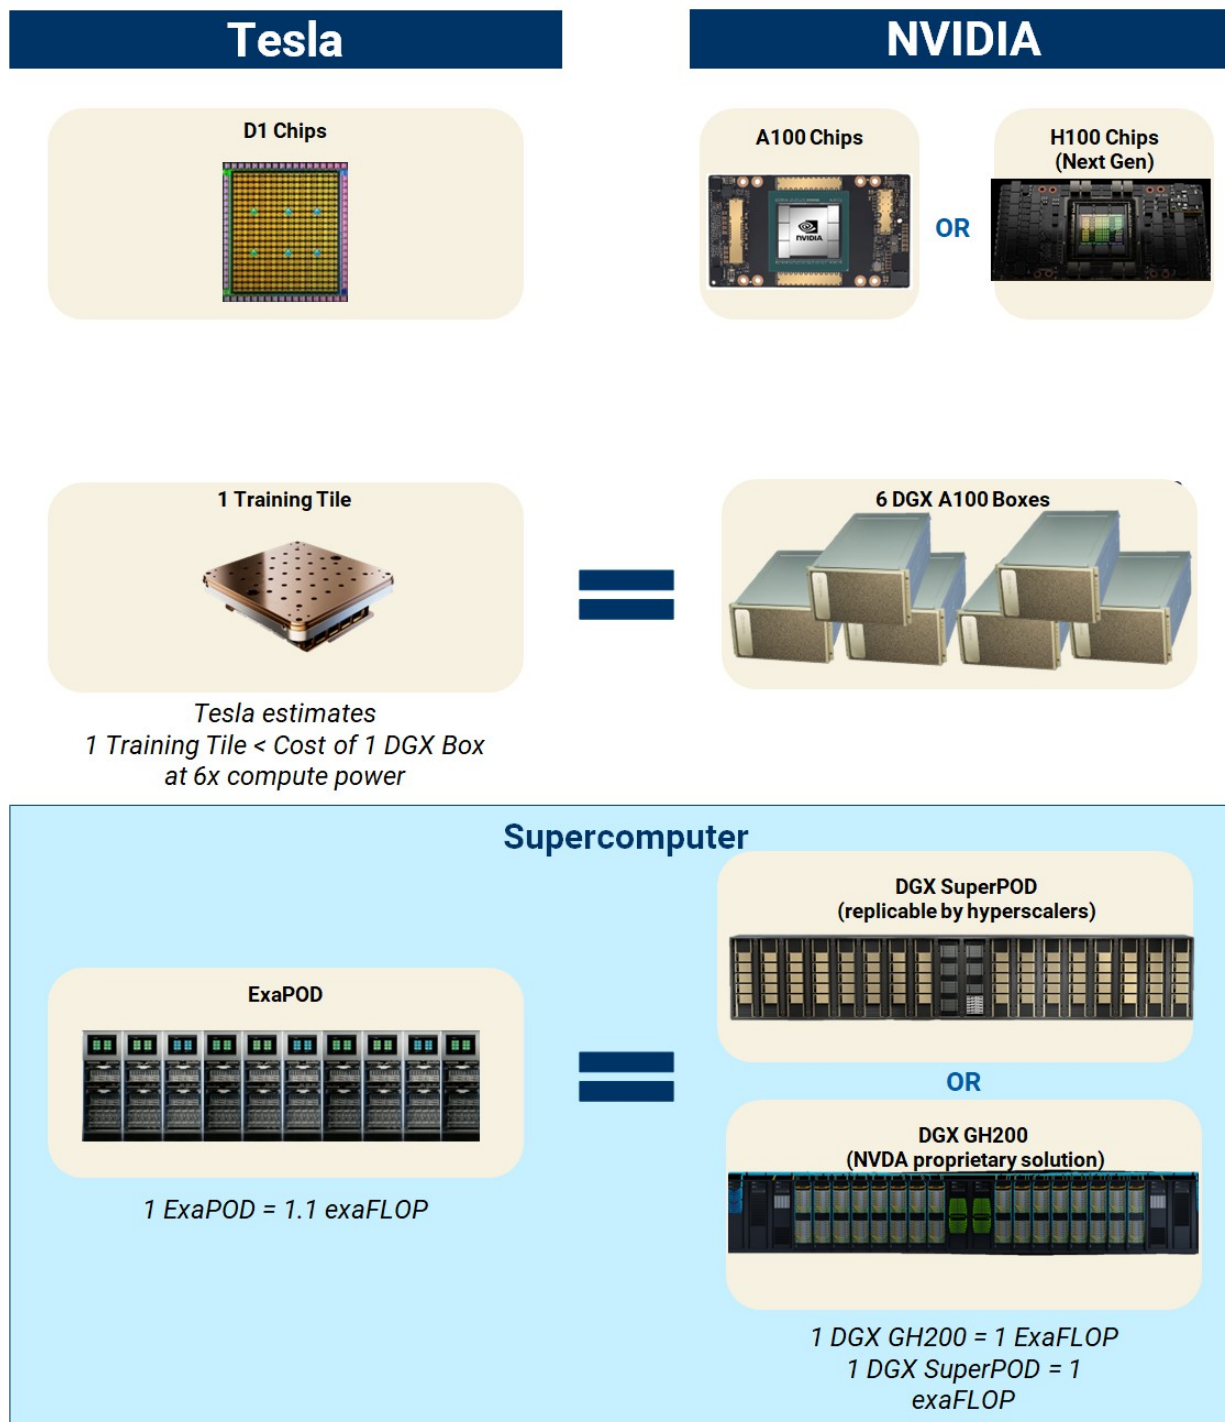


Source: Company data

How Does Dojo Compare to What's Already Out There?

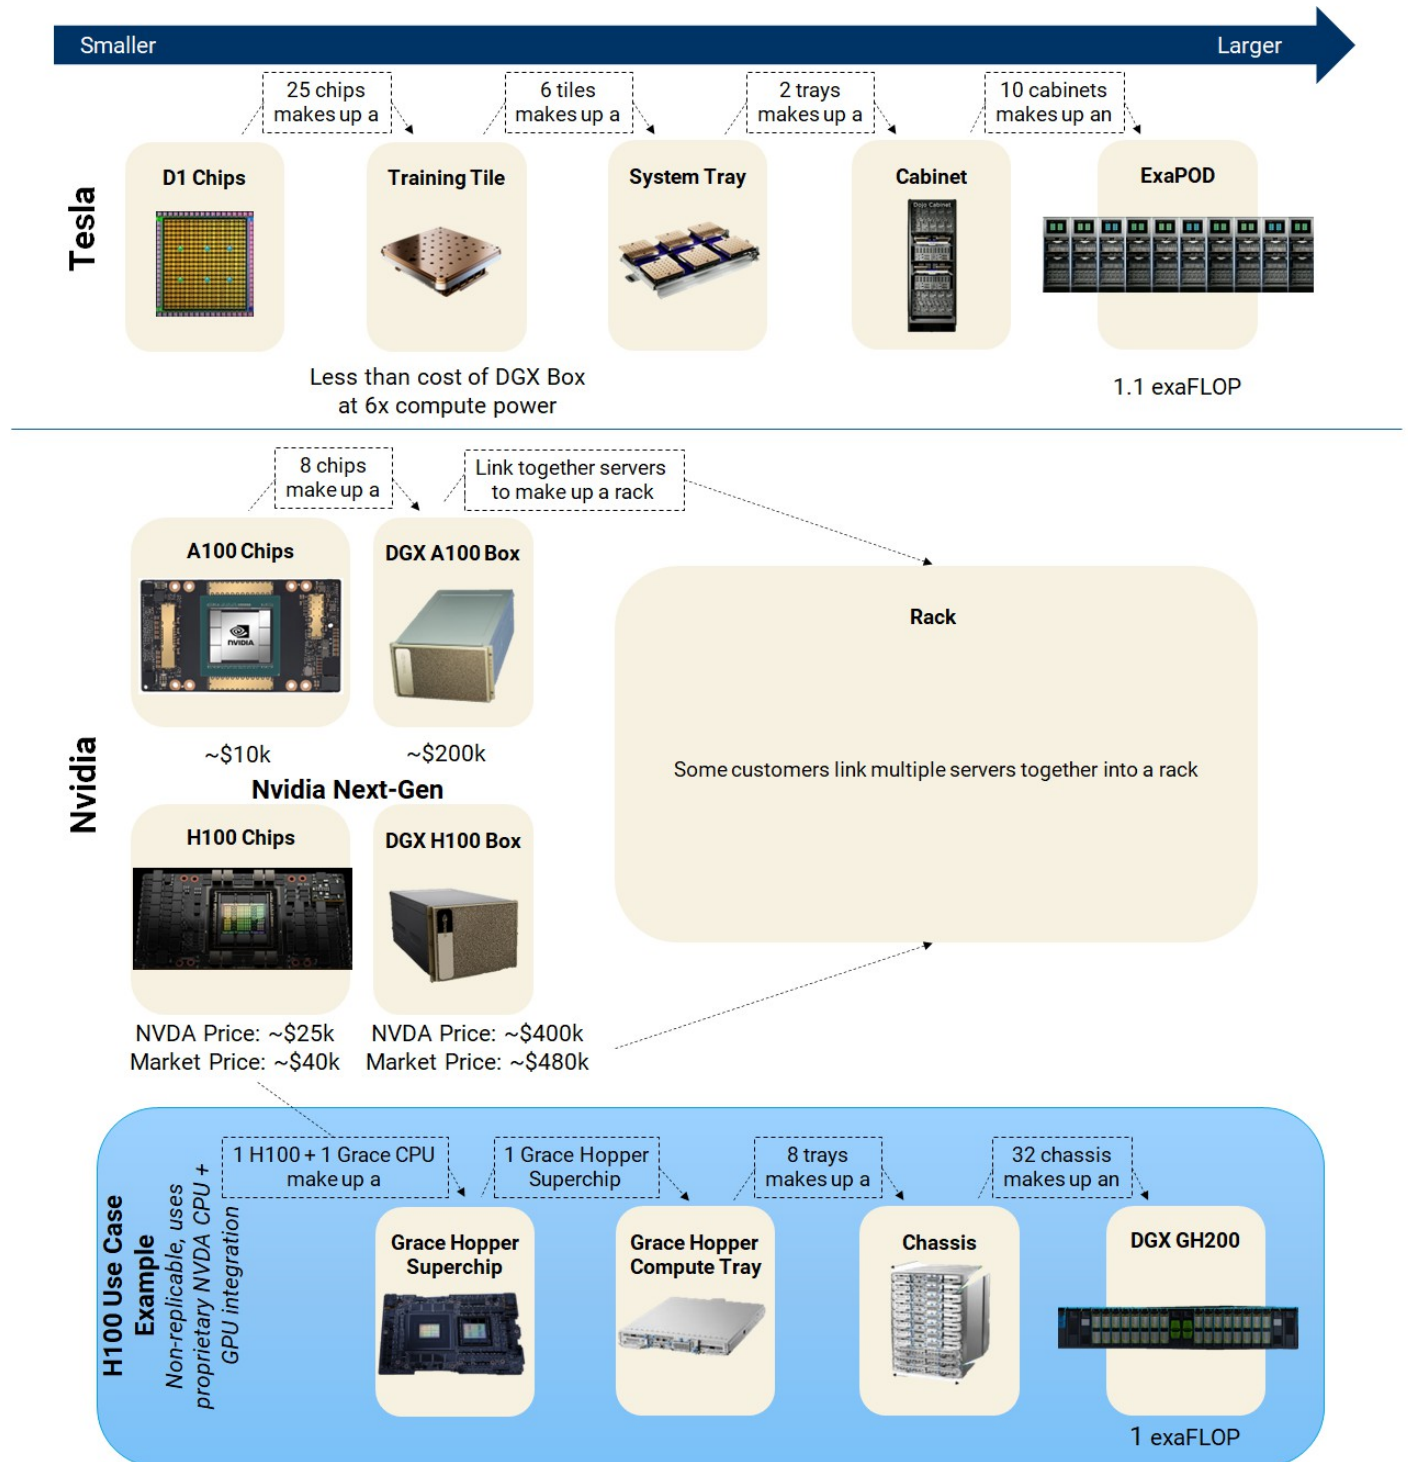
Two diagrams that compare the building blocks of Tesla's Dojo vs. NVIDIA's A100/H100 systems:

Exhibit 11: Compute Power Comparison: Tesla's Dojo vs. NVIDIA's A100/H100



Source: Company Data, Morgan Stanley Research

Exhibit 12: Building Block Comparison: Tesla's Dojo vs. NVIDIA's A100/H100



Source: Morgan Stanley Research

As of September 2022, Tesla had three supercomputers featuring 14,000 NVIDIA A100 GPUs total. 10,000 are reported to be allocated to training purposes and 4,000 for auto-labeling, with Tesla's largest supercomputer system featuring 7,360 A100s. This ranked as the 7th most powerful supercomputer in the world by GPU count. Semiconductors powerful enough to power deep learning of this size are (unsurprisingly) capex intensive:

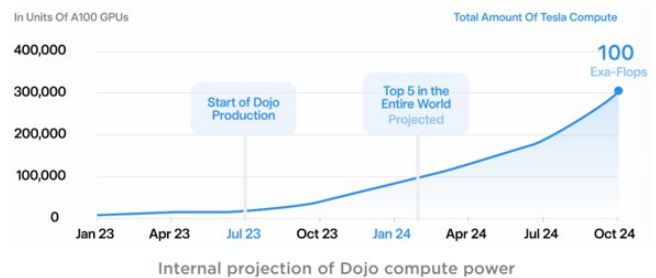
- A fully-loaded DGX A100 system (eight A100s and supporting hardware/software) is approximated to cost ~\$200,000. Following this assumption, Tesla's current largest supercomputer of 7,360 A100s costs ~\$184mn, part of the implied total cost of \$350mn for Tesla's entire 14,000 GPU collection.
- The next-gen NVIDIA DGX H100 is approximated to cost ~\$480,000 and promises up to 9x faster AI training and up to 30x faster AI inference. Each H100 costs ~\$40,000/chip. We note that NVDA's price for H100 is \$25k, but the market price is \$40k given the shortage. Similarly, an H100 DGX costs ~\$400k on NVDA's prices.

Tesla predicts that they will reach 100 exaFLOPs of compute by Q4 2024, up from ~4.5 today. According to Tesla, that's the equivalent of ~300,000 A100 GPUs, which on our estimates would cost \$7.5bn - \$8.0bn. Whether the 100 exaFLOP goal becomes reality by then or not, management believes that Dojo features greater efficiency, scalability, cost-effectiveness, and functionality than any other GPU.

According to Morgan Stanley Taiwan Semis Analyst Charlie Chan, semis design firm Alchip expects Dojo chip volume to reach 40-50k this year (FY23). For comparison, NVIDIA should ship 200-250k H100 chips this year.


On a ranking of supercomputers, Tesla's largest supercomputer was ranked 4th (end of March 2023). Tesla already operates one of the largest publicly known GPU clusters, suggesting that the company was dedicated to building material computing power prior to Dojo. In our view, the move to Dojo signifies Tesla doubling down on computing and software technology as a key competitive advantage for the company going forward.

Exhibit 13: Amount of compute in A100s required to reach Tesla's 100 exaFLOP goal by Q4 2024



Source: Company data

Exhibit 14: Global Supercomputers as of March 2023



Source: State of AI, Morgan Stanley Research

The time is now. Dojo will be phased into the current supercomputing systems, and as of July 2023, the 1st ExaPOD has been scaled into production. According to Musk in June 2023, Dojo has been online and running useful tasks for a few months and Tesla will be buying fewer incremental GPUs in 2023 compared to prior years. Tesla is expected to introduce 6 more ExaPODs into their Palo Alto, CA data center in the near-term, and the next-gen versions of the Dojo hardware and software appear to be well underway.

Tesla's Use Case

When completed, Dojo is intended to be the supercomputer that trains the neural networks in the FSD platform, ADAS applications, and Optimus humanoid.

In January 2023, Tesla reported that 400,000 vehicles were using FSD Beta, up from 160,000 in September 2022. As of July 2023, Tesla reached a milestone of 300+million miles driven using FSD Beta, which is the data inputted to Dojo for AI training. Based on early results with the Dojo-custom Compiler as reported by Tesla, **Dojo can accelerate FSD at a rate that trains workloads that previously took more than one month in less than one week.**


According to management, one Dojo tile (25 D1s) is 30x faster than 24 GPUs, with a given operation taking 5 microseconds on 25 dies vs 150 microseconds on 24 GPUs. The potential efficiency, speed, and cost benefits compared to current GPUs could allow Tesla to materially reduce their autonomy timeline.

Dojo will be used to train Optimus Tesla's humanoid robot. The 'brain' for the humanoid bot will be informed by the same autonomous systems present in Tesla's vehicles. There is natural crossover between autonomous labor (Optimus) and driving (FSD): both require the ability to process raw video input as well as the function to generate 3D maps to inform the user to react to perceived objects (Occupancy Network). As such, the development and refinement of Optimus, like FSD for Tesla vehicles, can be exponentially accelerated by the speed with which Dojo can train its vision-based neural net. We note that we do not ascribe any value to Optimus in our Tesla model at this time. *We discuss Optimus later on in this report, see Further Industry Implications .*

Major potential within the 'Musconomy': We think that there's a possibility that the Dojo system isn't just being built to accelerate and train FSD and Optimus; rather it could be a solution that builds a moat around Tesla and Musk-universe companies with highly advanced ML and AI capabilities. Looking internally at other Musk companies, we believe that Dojo may have the capability to be the core of the Musconomy.

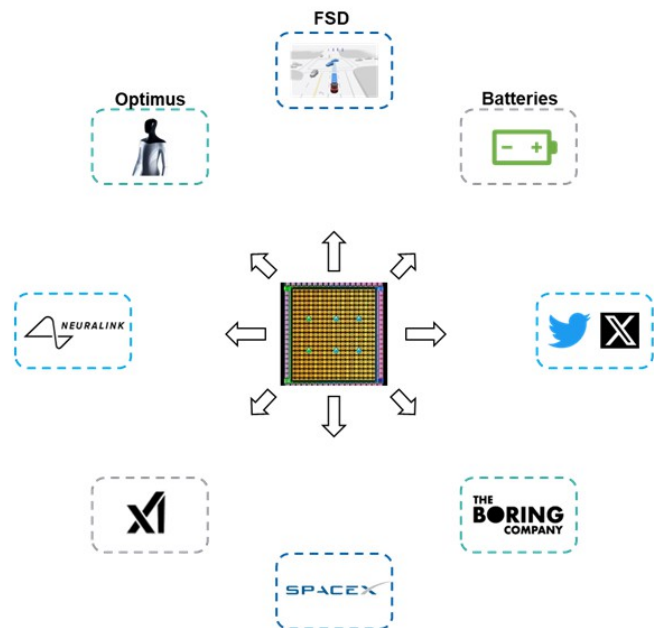
The point we'd make to investors is that Tesla is just one part of a larger area of scientific and commercial interests on which Elon Musk is allocating time, financial resources and talent. We see Tesla as one of the more complex and 'unifying' businesses which can serve as an 'experimental lab' to iterate on advanced technologies with high degree of mission difficulty. Admittedly, it is challenging to put a value on this interconnected/network effect.

Exhibit 15: The difference in training time between GPUs & Dojo for typical workloads



Source: Company data, Morgan Stanley Research

Exhibit 16: Dojo's potential within the Musconomy



Source: Morgan Stanley Research

- **X (Formerly Twitter)** – Investors have wondered where Twitter can fit in the Tesla ecosystem. Twitter can benefit from shifting away from its current compute systems and switching to no/low-cost data computation from the most powerful supercomputer systems in the world, greatly accelerating the platform to new heights. Dojo V2 is anticipated to incorporate the general-purpose AI limitations that V1 currently lacks. Training itself through the LLMs can provide a huge ecosystem of data that feeds into itself. As we've long said, owning and moderating a free speech platform poses incredible moral, political, technological, and regulatory challenges – a supercomputing system has the potential to support the platform. Finally, Twitter's vast collection of data from its hundreds of millions monthly active users provides Tesla with another source of data that can be used to train its systems.
- **SpaceX** – SpaceX's thousands of satellites communicate with each other via Intersatellite Link (ISL), which provides a direct link within the space segment without the need of an intermediate ground segment to relay the data. Since the satellites communicate through lasers, large amounts of data is required for precise measurements needed to connect and hand off. Additionally, immense computational power is required in the field of orbital debris mitigation and avoidance for the satellites need to avoid collisions, where the required metrology can benefit from AI computations. The same way Dojo can be used in autonomous driving to avoid obstacles and create a planned path to use, SpaceX can use Dojo to train their systems to communicate and avoid debris collisions.
- **Batteries** – Analysis of the behavior of magnetic fields and chemical reactions are used to design motors and batteries respectively, using predictive analysis of the thermal properties of such devices. The supply chain, infrastructure, battery, and electric motor are all interconnected with each other and deal with large amounts of data. Tesla has long described the charging network as an extension of the battery and how battery systems can be optimized from the data gathered from the battery fleet on the road and in ESS applications.

Tesla's Expected Cost Savings

Custom silicon solutions upfront costs are high. In the 2Q23 earnings call, Tesla mentioned that they will be spending north of \$1bn on Dojo's R&D over the next year (outlook includes all related expenses). Our global semiconductor analysts estimate the upfront design cost for an AI ASIC project (7nm chip) typically requires an investment of >\$200mn, and it is likely that Dojo has been contributing to Tesla's Capex/R&D spend since 2019. Given Tesla's size, we view investing in custom silicon as a natural choice, similar to many hyperscalers. The silicon utilization rate for custom chips is much higher since it is tailored to a specific workload, and Tesla is expected to see similar cost benefits if able to scale. We discuss the benefits around purpose-built ASICs in greater detail in the following section, [A Closer Look at the World of Custom Silicon](#).

Management believes that one Dojo tile can offer the same amount of machine learning compute as 6 GPU boxes, all at the cost of less than 1 GPU box. Leveraging company assumptions (though we caveat that these may have some definitional inconsistencies when considered relative to other company estimations) and channel checks that suggest that 1 DGX A100 GPU box costs \$200k, it then follows that Tesla's custom Dojo Compiler technology can offer 6x cost savings. If company estimates materialize, Tesla could potentially see cost savings in the billion dollar range as they scale their compute power to reach stated processing capacity in 2024.

- As estimated in the 2Q23 earnings release, it would take 300,000 A100 GPUs to reach the company's 100 exaFLOP target by 4Q24, which at \$200,000 per DGX A100 could cost them \$7.5bn - \$8.0bn.
- Leveraging the company's 6x cost savings estimate, we imply \$1.25bn - \$1.33bn spent on ExaPODs, a ~\$6.5bn reduction in cost.

Whether Tesla will reach their 100 exaFLOP goal, and how (Dojo V2, ExaPOD installation in excess of announcements), and whether these cost saving estimates carry any weight, remain to be seen. We'd like to heavily caveat these cost saving estimates by noting that this is what the company expects to achieve, but that the company has also given a number of different estimates across a number of different investor events, which when considered holistically, may bring to light some definitional inconsistencies as well as a wide range of investor interpretations as a result. One such example is below:

- At the 2022 AI Day, Tesla stated 1 Dojo tile is equivalent to the compute power of 6 boxes. Using this assumption we could calculate that Tesla would need ~65,000 GPU boxes to get to 100 exaFLOPS of computing power (6x the number of tiles required for ~91 ExaPODs).
- In 2Q23, when Tesla set out their goal of reaching 100 exaFLOPS, the company stated that doing so would require 300,000 A100s. With 8 A100 chips per DGX box, we calculate that 37,500 GPU boxes would be required under this assumption.
- Given that compute does not scale linearly, both of these assumptions can co-exist and are not necessarily opposed but we note that there are multiple ways to quantify potential cost savings.
- Our cost analysis in this note builds on the assumption that reaching 100 exaFLOPS with GPUs would require 37,500 but we acknowledge that cost savings *could be higher* depending on the assumptions used. Computing speeds are difficult to measure, given the dependency on calculation accuracy, software/hardware etc. Hence, there is likely no perfect 1:1 comparison between Dojo and other alternatives. This uncertainty makes it challenging to pinpoint the ratio exactly.

Additionally, Tesla is comparing Dojo to NVIDIA's A100 GPU, which is now the company's older technology. NVIDIA has since released a next-generation solution called H100, which is said to have 9x faster AI training and up to 30x faster AI inference.

More importantly, *we do not give Tesla any credit for said cost savings* – our price target increase and changes to our model are driven by speed to market made possible by Dojo and efficacy of the system itself rather than cost savings.


All in all, while it is difficult to explicitly validate the many claims Tesla has made about Dojo's cost and performance, we believe Tesla has a chance of bringing forth a competitive customized solution given the company's innovation track record and capabilities.

Exhibit 17: Tesla estimates that one Dojo tile has the computing power of 6 GPU boxes at less than the cost of a single box...




Source: Tesla, Morgan Stanley Research

Exhibit 18: And our channel checks suggests that 1 DGX A100 (1 GPU Box) costs ~\$200,000...



Source: Morgan Stanley Research

Exhibit 19: ...Making the total savings in reaching the 100 exaFLOP target ~\$6.5bn, if using company estimates




Source: Morgan Stanley Research

Dojo is built for speed and efficiency for the specific purpose - not just a matter of computational horsepower...

On a per dollar basis, Tesla expects Dojo to provide **4.0x better performance, 1.3x better performance/watt, and a 5.0x smaller footprint vs. current alternatives**. AI computing requires a significant amount of energy, and their better anticipated performance/watt highlights Dojo as an energy efficient solution. Additionally, by using significantly fewer systems vs. the current A100 cluster, the supercomputer can be cooled more efficiently. Our semiconductor team emphasizes that the benefits of custom silicon isn't just a battle of the most powerful chip... the benefits to Tesla from Dojo will ultimately be measured in the overall performance of the *entire system* for Tesla's specific vision-based training - defined as how many 9's to the right of the decimal the system can perform on key occupant/pedestrian safety metrics.

Exhibit 20: At the same cost as what's currently available, Dojo boasts great advantages



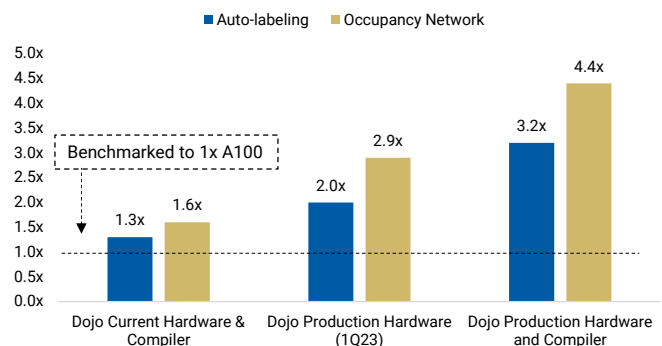
Source: Company data, Morgan Stanley Research

According to **management**, Tesla can replace **6 GPU boxes with a single Dojo tile, which enables the company to reduce network training time from ~one month to ~one week**. Put differently, Tesla can achieve the same throughput on 4 Dojo cabinets as they can with 4,000 GPUs. The first of aforementioned 7 total ExaPODs came

online in July 2023 at the Palo Alto, CA facility and is expected to be 2.5x their current auto-labeling capacity. We note that the state of the art of supercomputing is constantly changing... comparisons of one chip/box/system is complicated on an apples to apples basis and must be seen relative to the trade-offs (cost/energy consumption) and from the perspective of the AI problem being solved (vision vs. LLM, etc).

At its typical workload, Tesla claims its Dojo Compiler Engine (custom platform) outperforms the A100 GPUs currently used in both auto-labeling and occupancy network (on a per die basis). Although preliminary, these results suggest that Tesla can reap significant benefits from the increased efficiency of their platform – even early on. *We note here that Tesla does not provide a performance comparison to NVIDIA's next gen solution, H100.*

Exhibit 21: Benchmarked to an A100 (1.0x), Dojo showed significant improvement in efficiency & speed for a typical workload



Source: Company data, Morgan Stanley Research

Each Dojo tile requires 15kWh of power compared to 39kWh for 6 DGX A100s (GPU boxes), a ~60% improvement. Data center space has become scarce as a result of demand from hyperscalers, and there is a dire need to upgrade the power grid due to the capacity and power density of GPUs (2-3x power per square foot). Dominion Energy, the primary power provider in Loudoun County, VA (world's largest concentration of data centers), **informed** its major customers that power delivery could be delayed until 2026. Corroborating the gravity of the strain on the electrical grid, Tesla **claimed** that while testing Dojo in Oct'22, they tripped the power grid in Palo Alto.

Next generation Dojo will have a broader AI scope. According to a tweet by Musk in June 2023, Dojo V1 is highly optimized for vast amounts of video training while Dojo V2 will incorporate any general-purpose AI limitations that V1 currently faces. Tesla **believes** that an upwards of 10x improvement can be achieved when the next-gen hardware (V2) is developed and implemented, which can enable the company to reach 100 exaFLOPs by 4Q24.

A Closer Look at the World of Custom Silicon

Why Custom AI?

Most of the current supply for AI chips comes from NVIDIA. General purpose GPUs are the default for AI training, but custom semiconductors (ASICs) can offer tailored solutions and drive efficiencies. Among the potential benefits of using a custom AI solution are:

- Reduced rack and chip size
- Reduced power consumption
- Extending the replacement cycle
- Bringing down the cost of hardware
- Increased efficiency


Just one of these aspects can make a custom AI ASIC (application-specific integrated circuit) beneficial to the company operating on it. Budget costs and energy requirements are two major limitations for future AI computing, which will ultimately lead companies with the technical capability and financial motive to seek to create their own energy-efficient and low-cost AI custom chip designs. In recent years, the newer NVIDIA chips' focus have shifted to the larger market of LLMs rather than solely processing visual and video data, which could be one of the motivations for Tesla to develop their own custom solution.

Our Asia Semiconductor team has written extensively on this topic. For further reading please see: [Asia Semiconductors: Tech Diffusion – Fulfilling the surge in AI demand with custom chips \(11 Jun 2023\)](#).

Custom AI chips can provide significant efficiency savings and fixed-cost leverage if they can be scaled. Custom AI ASICs are built

Exhibit 22: As of 2023e, GPUs are projected to account for >90% of cloud AI semiconductor deployment

Cloud semi mix: General-purpose vs. Custom AI, 2023e




Source: Morgan Stanley Research

for predictable specific workloads. NVIDIA solutions like the A100 and H100 are considered the industry's 'gold standard', but are designed for general-purpose tasks: efficiencies can be gained by optimizing for dedicated processing. Modern data centers are faced with capacity challenges – GPUs are very power dense (2-3x power per square foot). These systems are difficult to cool and cost only increases as the systems increase in power. Additionally, scale is limited as marginal improvement reaches a point where size does not provide enough added benefit to outweigh the cost.

Alternatively, AI ASICs are specifically built for an internal system and can operate more efficiently. Often times, companies like Tesla pay for compute or features that are under-utilized (or not utilized at all), and therefore can benefit from significant long-term cost savings with custom AI ASICs if able to scale. It is important to note that upfront chip design cost for an ASIC project is very high, at >\$200mn for a 7nm chip. If a company is able to weather the initial cost of manufacturing in-house chips, they can realize notable fixed-cost leverage.

In 2023, GPUs are projected to account for 92% of cloud AI semiconductor deployment, with 8% allocated to custom AI chips – but our Asia Semiconductors team (led by Charlie Chan) expects AI ASIC chips to capture up to 30% market share by 2027 based on the larger scale of AI computing demand that justifies the cost, and the need for vendor diversification given NVIDIA's bargaining power. Generally, the upfront chip design cost for ASICs is too high at low scale, which is why many major cloud service providers like Meta and Microsoft have been using GPUs until recently. It's expected that hyperscalers will increase their adoption of custom AI ASICs to improve performance per watt and cost of ownership.

Exhibit 23: Custom AI chips are projected to outgrow GPUs through 2027 (CAGR)



Source: Morgan Stanley Research

Beyond cost and efficiency justifications, without diversifying, Tesla's reliance on external GPUs will intensify. If the generative AI trends currently dominating the market continue, demand for NVIDIA systems will remain elevated and procuring sufficient supply of A100s/H100s to meet Tesla's 100 exaFLOP goal will prove difficult. If Tesla were to rely solely on NVDA to reach their stated compute power goal, they alone could comprise 6-11% of NVIDIA's revenue (assuming Tesla replaces their GPUs every 2-3 years). Tesla's incorporation of their in-house FSD system in early 2019 (Hardware 3.0), and now their proprietary training chips, begin to mitigate over-reliance on NVDA. When combining the tailored capabilities of ASICs and better economics that Tesla can realize with fixed-cost leverage given the company's scale, we view custom AI as the best path forward.


Who Else Is Developing Their Own Purpose-Built AI Chips?

Beyond general purpose chip makers like NVIDIA, Cerebras, and Graphcore, many compute-centric companies have begun turning to custom solutions. A key advantage in a custom chip is that since the company knows their workloads, they can tailor the chip and system to serve a direct purpose. While NVIDIA A100s and H100s are the cutting-edge general-purpose GPU for ML training, a chip built to specifically accomplish a certain purpose (*such as utilizing video data to train Tesla's neural network to accelerate FSD*) can be more efficient than the best chips that may have better overall specs. Several companies such as Apple, Microsoft, Google, Amazon, and Meta either develop or [will soon](#) develop their own purpose-built chips. The report from our Asia Semiconductors team estimates that

custom AI ASIC chips will represent \$6.2bn of market value in 2024e, with the combination of Tesla's Dojo and FSD amounting to \$1.2bn.

- **Google:** Following Google's announcement of its first-generation TPU (Tensor Processing Unit) in 2016, now the fourth generation (TPU v4) is going through Broadcom's design service. Google TPU claims 2-3x greater energy efficiency compared to contemporary ML DSAs (machine learning domain-specific accelerators), and more than 3x relative performance improvement over its previous generation, TPU v3.
- **Amazon:** Amazon Web Services (AWS) started its ASIC strategy quite early, with the first announcement in 2018. AWS claimed that its first Inferentia chip delivered up to 2.3x higher throughput and up to 70% lower cost per inference (or up to 50% greater performance per watt) than comparable Amazon EC2 instances (using GP GPUs). Its second-generation chip (manufactured using TSMC's 7nm process, going through Alchip's design service) delivers up to 4x higher throughput and up to 10x lower latency compared to its first-gen Inferentia chip.
- **Microsoft:** Since 2019, Microsoft has been developing its own AI chip named Athena. The chip is designed to handle LLM training on TSMC's 5nm process, and is expected to debut sometime in 2024.
- **Meta:** Meta announced its first-gen AI inference accelerator MTIA v1 this year (2023). According to Meta, this AI ASIC (based on TSMC's 7nm process) can reach up to 2x performance per watt vs. an AI GPU and is targeting 2025 for launch.

Exhibit 24: The custom AI ASIC market is estimated to represent a ~\$6.2bn market value in 2024e



Source: Morgan Stanley Research

Putting this all into context... a few thoughts from MS Semiconductor Analyst Joe Moore:

This is not a new debate, and we have seen custom silicon developments for AI for 10 years. If it were easy to create a better solution than NVIDIA, startups would have done so, or Intel or AMD would have already done so, or custom silicon would already be deployed much more broadly. NVIDIA spends \$8 bn a year on R&D, and has all the advantages of running the same architecture through a higher volume business (gaming) to refine performance and costs. Importantly, that money is not to create a new chip from scratch, that money is iterating on the best in class design. We see a fairly competitive market for inference - though we expect NVIDIA to do well - but we think that for companies to outperform NVIDIA on training, they will need to be very optimized around their own workload.. Conducting due diligence on NVDA - comparable solutions only corroborates this view - the gap to NVDA has historically been difficult to bridge, and where we have seen enthusiasm for new designs, we have often seen that enthusiasm diminish after NVIDIA releases a next generation solution that raises the bar.

That said, NVIDIA's very high price and margins certainly leaves some umbrella to compete, and the scope of their ambitions (to compete in hardware, software, and services, and even to enable Mercedes and JLR to compete with Tesla) is an added incentive to try to replace them.

1) Where we have seen competitive approaches succeed – really Google, thus far – it happened because the customer understood their own workload better than anyone else. Google invented the transformer model over a decade ago, so their 5th generation TPU silicon is fairly optimized around transformers which are suddenly the most important models that there are. *Tesla has been a pioneer in the ADAS space and has a similar understanding of what is needed.*

2) Tesla appears to have a very strong semiconductor design team – this is not a neophyte company relying on ASIC design services. The semiconductor group was originally built by Jim Keller (since departed Tesla), who has broad industry experience in building world class design teams then moving on. Where startups struggle is a lack of resources and/or knowledge to customize the entire stack, with some players who are good at software (e.g. Sima.ai), some who are good at AI hardware (e.g. Graphcore) – Tesla has the ability to tackle both.

Overall, while it is challenging to create a chip that can be unequivocally better than NVIDIA, Tesla's custom Dojo solution may prove competitive in its tailored use case. Tesla is not competing to make a better chip. Tesla is optimizing for a single purpose that can in turn drive an improved total output, at greater efficiency and lower cost. ***NVIDIA used the demanding performance of gaming to develop the world's most powerful GPU chips. Can Tesla use the demands of autonomous cars/FSD to become a global leader in custom AI chips?***

Implications for Tesla: Upgrade to OW, Increase PT to \$400, New Top Pick

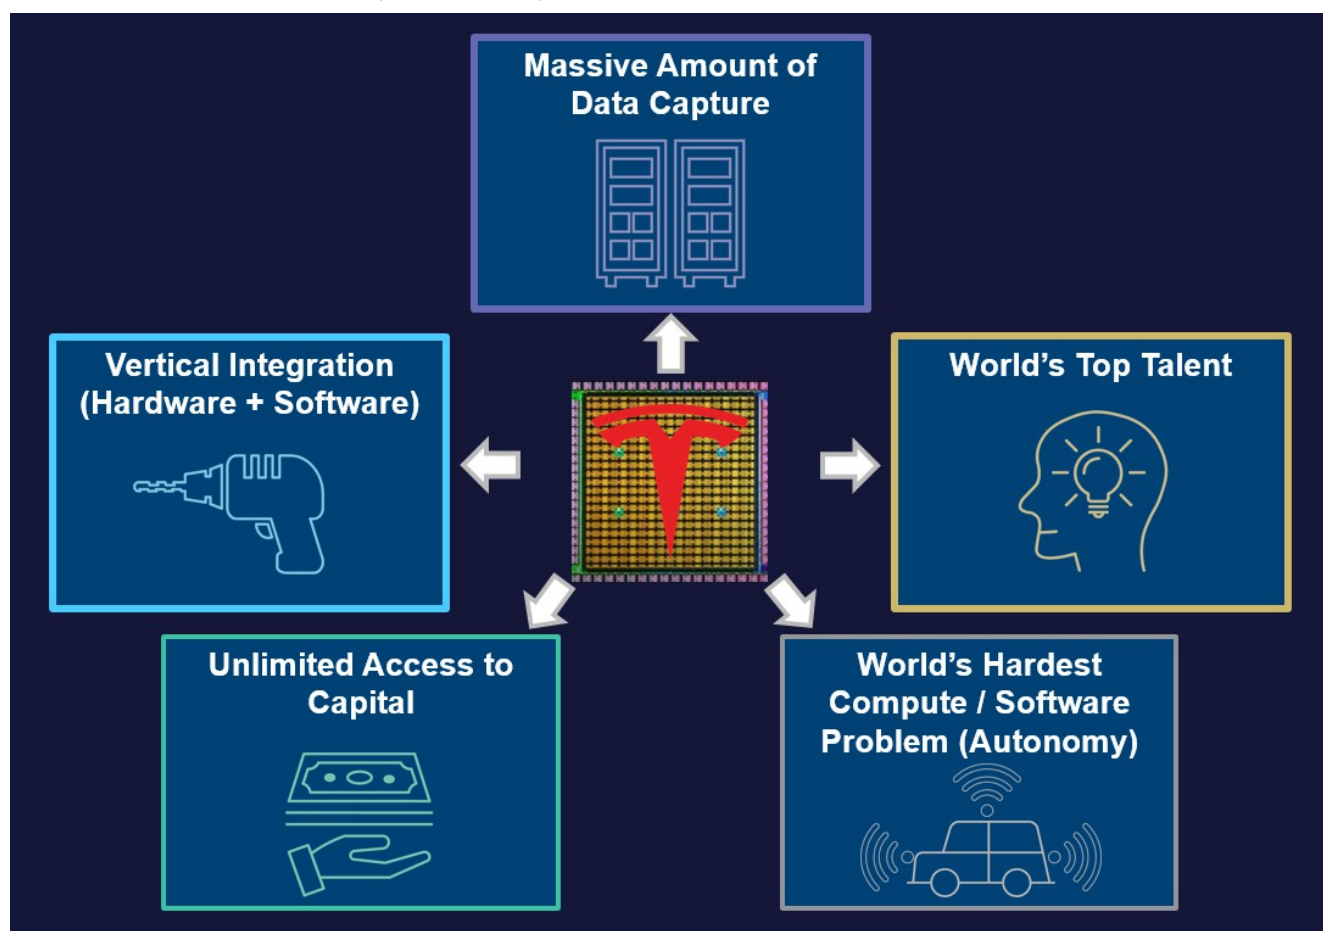
Dojo is the key to unlocking Tesla's double-flywheel effect – integrating and accelerating the synergies between Tesla's Core Auto Flywheel and Tesla's SAAS Flywheel... accelerating time to market and expanding the addressable market.

Following our in-depth Dojo analysis, we believe the AI supercomputer has the potential to efficiently and cost-effectively train Tesla's neural network to help the company achieve competitive asymmetry in the multi-trillion-\$ autonomy market. Tactically, we attached the extra value of Dojo/AI to Network Services (via increased ARPU and added 3rd party licensing), Tesla Mobility (increased long term fleet size and margin), and, to a lesser extent, in our 3rd Party Battery Business model, as we believe the charging and FSD deals will also result in higher hardware attach. We have not made material changes to base case auto, energy, or insurance valuation or assumptions.

We note that while the cost savings associated with Dojo have the potential to be significant, it does not feed directly into our model or impact our price target. Instead, we attach the value of Dojo to Tesla via increased speed to market, performance, and secondary impact on attach rate, licensing, and uptake.

As a result of the newly Dojo-ascribed value to our model, we upgrade Tesla to OW from EW, increase price target to \$400 (vs. \$250 previously), increase Bull Case valuation to \$550 (vs. \$450 previously), and increase Bear Case valuation to \$120 (vs. \$90 previously). Tesla is now our Top Pick.

Exhibit 25: Unlocking Tesla's AI Mojo: Enter the Dojo



Source: Morgan Stanley Research

We see a double-flywheel at work with Tesla, with Dojo as the key accelerant at the intersection of hardware and software iteration.

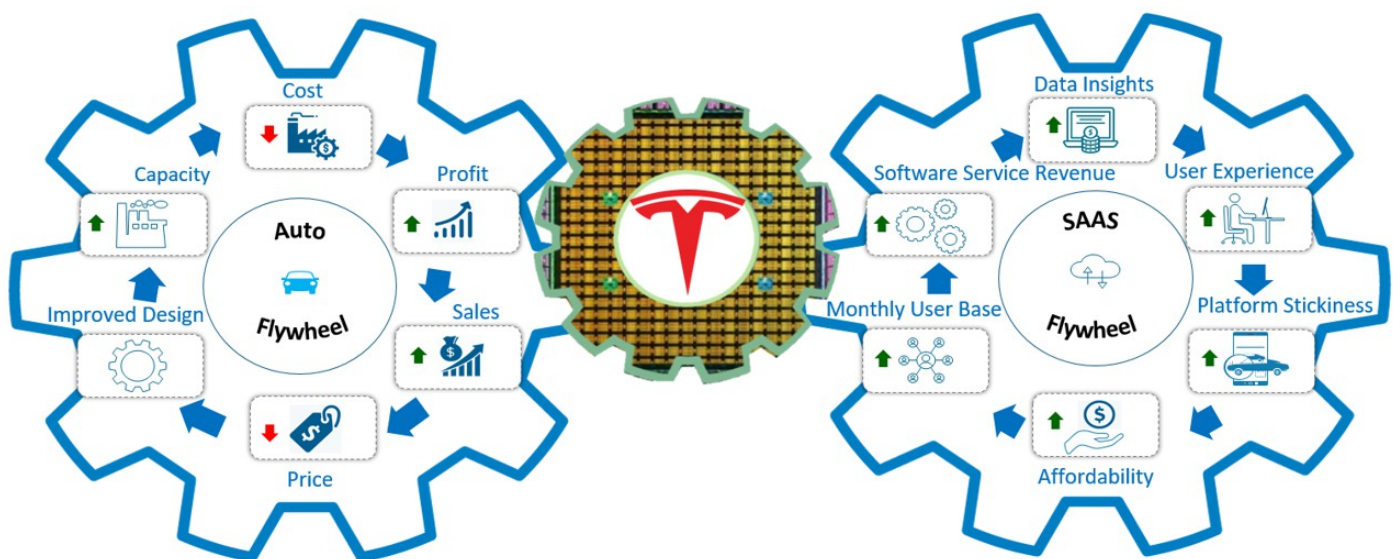
We've long written about the potential of Tesla's leadership in EV hardware (semi-autonomous electric 'robots') to convert vehicle owners into 'subscribers' generating highly recurring revenue. A quick refresher:

- We see Tesla's #1 advantage as scale and cost in EVs. This cost advantage allows Tesla, already the market leader, to lead the industry in price cuts for its products, expanding the addressable market beyond existing capacity. They then add new capacity with design tweaks to further lower cost, supporting profitability and driving the ability to further lower cost... expanding the addressable market.
- As the market expands, Tesla can 'turn on' new features and services for its vehicle fleet... turning owners into 'subscribers.' These services can include everything from charging \$2k for a software update that shaves more than 1 second off your 0 to 60 time, improves your access to available media content, enables various levels of highly automated driving/ FSD, insurance products, access to charging infrastructure and a host of other telematics services. Sound familiar?

- As more and more consumers 'opt in' for a range of these services, Tesla increases its proportion of regularly recurring and high margin (in some cases 100% margin) businesses while further improving the addressable market and the customer experience. This adds stickiness of users to the platform, ballast to the financial profile and top line and margin to the bottom line... further enhancing the company's ability to make continuous improvement to its core auto product (the hardware) to reduce cost, cut price, expand the user base... and so on.

This double-flywheel is only further accelerated and integrated via the compute power of Dojo. While Tesla has had some modicum of success on the double-flywheel front, we believe Tesla's in-house computing efforts could accelerate the network effect and speed of data capture/analysis/learning from the ~1 billion miles traveled per day we forecast is generated by its global fleet (Tesla + 3rd party) by 2027.

Exhibit 26: Dojo-Accelerated Auto Flywheel & SAAS Flywheel



Source: Morgan Stanley Research

Changes to our Tesla Estimates

We attached the extra value of Dojo/AI to Network Services (via increased ARPU and added 3rd party licensing), Tesla Mobility (increased long term fleet size and margin), and 3rd Party Battery Business, as we believe the charging and FSD deals will also result in higher hardware attach.

• Network Services

- We value Network Services at \$119/share, with 23.8mn MAUs at \$180 ARPU by 2030. Licensing fleet 50% probability, vs. \$60/share, 23.8mm MAUs, \$100 ARPU by 2030, and 70% discount previously.
- **Maintained attach rate** at 65% of Tesla end-of-year park, which results in 23.8mn MAUs (connected fleet) in 2030. We have penetration increasing to 70% (and 87.5mn MAUs) by 2040.
- **Increased ARPU to \$180** in 2030 vs. \$100 previously, driven by increases to Upgrades, Content, and Other, while Autonomy, Charging, and Maintenance remains unchanged.

- **Added 3rd Party Licensing** as a result of the charging deals with Ford, GM, Rivian, and Fisker announced in 2023, and 2Q23 earnings call announcement around licensing FSD to an undisclosed OEM. At 11% penetration of the Global EV fleet and a \$36 Monthly ARPU (20% of the \$180 Tesla monthly ARPU) in 2030, we expect Tesla Network Licensing Revenue of \$7,296mn in 2030, ramping up to \$125,250mn in 2040 as penetration of global EV fleet increases to 25.0% and monthly ARPU increases to \$66 (33% of the \$200 Tesla monthly ARPU).
- **EBITDA Contribution.** We increase Network Services EBITDA margin to a run rate of 65% starting in FY26 vs. 50% previously. In addition to a higher margin, the higher combined revenues are as a result of the added licensing revenue and higher ARPU increase total EBITDA as well as Network Services EBITDA as a percent of Tesla Group EBITDA. **By 2040, we expect Network Services EBITDA to contribute ~62% of total EBITDA, vs. ~40% previously.**

Exhibit 27: Tesla Network Services Model Changes

Tesla Network Services Changes	2025			2030			2040		
	New	Old	% Change	New	Old	% Change	New	Old	% Change
Revenue Calculation									
MAUs (Connected Fleet)	5,788,515	5,262,286	0	23,757,833	23,757,833	-	87,459,010	87,459,010	-
% Penetration	55.0%	50.0%	0	65.0%	65.0%	-	70.0%	70.0%	-
Monthly ARPU (\$)	\$ 158.33	\$ 85.00	86.3%	\$ 180.00	\$ 100.00	80.0%	\$ 200.00	\$ 150.00	33.3%
Tesla Network Revenue (\$mm)	10,998	5,368	104.9%	51,317	28,509	80.0%	209,902	157,426	33.3%
Services Revenue/Tesla Group Revenue (%)	6.2%	3.1%	3.1%	12.0%	7.0%	4.9%	20.9%	16.5%	4.4%
Non-Tesla/Licensing Fleet	500,000			16,887,880			158,144,480		
Global EV fleet (mm)	47			154			633		
% Global Penetration	1.1%			11.0%			25.0%		
Monthly ARPU (\$)	\$ 15.83			\$ 36.00			\$ 66.00		
Tesla Network Licensing Revenue	95			7,296			125,250		
Combined Revenue (\$mm)	11,093	5,368	106.7%	58,612	28,509	105.6%	335,152	157,426	112.9%
EBITDA	6,656	2,684	148.0%	38,098	14,255	167.3%	217,849	78,713	176.8%
EBITDA Margin (%)	60%	50%	0	65%	50%	0	65%	50%	0
Services EBITDA/Tesla Group EBITDA (%)	22.0%	10.7%	11.4%	37.3%	21.5%	15.8%	62.0%	37.8%	24.2%

Source: Morgan Stanley Research

Exhibit 28: 2030e Network Services ARPU Breakdown

What's included in Network Services? (by 2030e)

Network Services	Mix %	Rev (\$mm)	\$/Veh Mile	\$/Veh Hour	Month ARPU	Year ARPU	\$/Life of Car
Autonomy	15%	7,919	0.02	0.62	\$ 28	\$ 333	\$ 5,000
Charging	8%	3,982	0.01	0.31	\$ 14	\$ 168	\$ 2,514
Maintenance	9%	4,645	0.01	0.36	\$ 16	\$ 196	\$ 2,933
Upgrades	10%	5,132	0.02	0.40	\$ 18	\$ 216	\$ 3,240
Content	5%	2,566	0.01	0.20	\$ 9	\$ 108	\$ 1,620
Other	53%	27,073	0.08	2.12	\$ 95	\$ 1,140	\$ 17,093
Total	100%	51,317	0.16	4.03	\$ 180	\$ 2,160	\$ 32,400

Source: Morgan Stanley Research

Exhibit 29: Network Services Model


		Tesla Network Services																				
		(\$ million Unless per unit)																				
		2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	2035	2036	2037	2038	2039	2040
Tesla Global Fleet in Service																						
Beg-of-year Parc		865,441	1,330,470	2,213,473	3,438,786	5,163,549	7,435,998	10,524,573	14,430,337	18,918,948	24,184,344	30,076,611	36,550,512	43,282,107	50,369,467	57,849,825	65,782,974	74,151,487	83,060,962	92,540,282	102,641,894	113,422,112
Annual Deliveries		499,647	938,222	1,313,851	1,862,315	2,478,991	3,400,887	4,369,846	5,152,407	6,135,667	7,004,747	7,857,425	8,486,019	9,164,901	9,898,093	10,689,940	11,545,136	12,468,746	13,466,246	14,543,546	15,707,029	16,963,592
% Growth		87.4%	40.3%	41.7%	33.1%	37.2%	28.3%	17.9%	19.1%	14.2%	12.2%	8.0%	8.0%	8.0%	8.0%	8.0%	8.0%	8.0%	8.0%	8.0%	8.0%	8.0%
Scrapage		(34,618)	(63,219)	(88,539)	(137,551)	(206,542)	(312,312)	(463,081)	(663,796)	(870,272)	(1,112,480)	(1,383,524)	(1,754,425)	(2,077,541)	(2,417,734)	(2,776,792)	(3,156,623)	(3,559,271)	(3,986,926)	(4,441,934)	(4,926,811)	(5,444,261)
Scrapage Rate (%)		4.0%	4.0%	4.0%	4.0%	4.0%	4.2%	4.4%	4.6%	4.6%	4.6%	4.6%	4.8%	4.8%	4.8%	4.8%	4.8%	4.8%	4.8%	4.8%	4.8%	4.8%
Tesla End-of-year Parc (units)		1,330,470	2,213,473	3,438,786	5,163,549	7,435,998	10,524,573	14,430,337	18,918,948	24,184,344	30,076,611	36,550,512	43,282,107	50,369,467	57,849,825	65,782,974	74,151,487	83,060,962	92,540,282	102,641,894	113,422,112	124,941,443
Tesla Fleet Analytics																						
Miles per Car		11,000	11,220	11,444	11,673	11,907	12,146	12,388	12,632	12,888	13,146	13,409	13,677	13,951	14,230	14,514	14,805	15,101	15,403	15,711	16,025	16,345
% Growth		2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%
Total Miles Traveled (bn)		15	25	39	60	89	128	179	239	312	395	490	592	703	823	955	1,098	1,254	1,425	1,613	1,818	2,042
Revenue Generating Miles (bn)		5	9	16	27	44	70	107	148	196	253	319	388	464	552	649	757	876	998	1,129	1,272	1,430
Average Vehicle Speed (mph)		25	25	25	25	25	25	25	25	25	25	25	25	25	25	25	25	25	25	25	25	25
Revenue Generating Vehicle Hours (bn)		0.2	0.3	0.6	1.1	1.8	2.8	4.3	6.9	10.1	12.7	15.5	18.6	22.1	26.0	30.3	35.1	39.9	45.2	50.9	57.2	64.0
Implied Revenue/Mile (\$)		\$ 0.11	\$ 0.12	\$ 0.13	\$ 0.14	\$ 0.15	\$ 0.16	\$ 0.16	\$ 0.16	\$ 0.16	\$ 0.16	\$ 0.16	\$ 0.16	\$ 0.16	\$ 0.16	\$ 0.16	\$ 0.15	\$ 0.15	\$ 0.15	\$ 0.15	\$ 0.15	\$ 0.15
Implied Revenue/Hour (\$)		\$ 2.73	\$ 2.99	\$ 3.29	\$ 3.47	\$ 3.70	\$ 3.91	\$ 4.12	\$ 4.10	\$ 4.07	\$ 4.05	\$ 4.03	\$ 3.99	\$ 3.95	\$ 3.92	\$ 3.89	\$ 3.85	\$ 3.81	\$ 3.78	\$ 3.74	\$ 3.71	\$ 3.67
Revenue Calculation																						
Miles (Connected Fleet)		465,665	774,716	1,375,514	2,323,597	3,717,999	5,788,515	8,658,202	11,729,748	15,236,137	19,249,031	23,757,833	28,349,780	33,243,848	38,759,983	44,718,822	51,164,526	58,142,673	64,778,197	71,849,326	79,395,479	87,459,010
% Penetration		35.0%	35.0%	40.0%	45.0%	50.0%	55.0%	60.0%	62.0%	63.0%	64.0%	65.0%	65.5%	66.0%	67.0%	68.0%	69.0%	70.0%	70.0%	70.0%	70.0%	70.0%
Monthly ARPU (\$)		\$ 100.00	\$ 111.67	\$ 123.33	\$ 135.00	\$ 146.67	\$ 158.33	\$ 170.00	\$ 172.50	\$ 175.00	\$ 177.50	\$ 180.00	\$ 182.00	\$ 184.00	\$ 186.00	\$ 188.00	\$ 190.00	\$ 192.00	\$ 194.00	\$ 196.00	\$ 198.00	\$ 200.00
% Growth		11.7%	10.4%	9.5%	8.6%	8.0%	7.4%	6.9%	6.5%	6.1%	5.8%	5.5%	5.2%	5.0%	4.8%	4.6%	4.5%	4.4%	4.3%	4.2%	4.1%	4.0%
Tesla Network Revenue (\$mm)		559	1,038	2,036	3,764	6,544	10,998	17,663	24,281	31,996	41,000	51,317	61,916	73,402	86,511	100,886	116,655	133,961	150,804	168,990	188,644	209,902
% Growth		85.6%	85.6%	84.9%	84.9%	84.9%	84.9%	84.9%	84.9%	84.9%	84.9%	84.9%	84.9%	84.9%	84.9%	84.9%	84.9%	84.9%	84.9%	84.9%	84.9%	84.9%
Services Revenue/Tesla Group Revenue (%)		1.8%	1.9%	2.5%	3.8%	5.0%	6.2%	7.8%	8.9%	9.8%	10.8%	12.0%										
Non-Tesla Licensing Fleet																						
Global EV fleet (mm)		6	10	17	25	35	47	62	79	99	124	154	185	220	259	301	347	397	452	512	573	633
% Global Penetration		0.4%	1.1%	2.0%	4.0%	6.0%	8.0%	10.0%	12.0%	14.0%	16.0%	18.0%	20.0%	22.0%	24.0%	26.0%	28.0%	30.0%	32.0%	34.0%	36.0%	38.0%
Monthly ARPU (\$)		\$ 14.67	\$ 15.83	\$ 17.00	\$ 25.88	\$ 26.63	\$ 26.63	\$ 26.63	\$ 26.63	\$ 26.63	\$ 26.63	\$ 26.63	\$ 26.63	\$ 26.63	\$ 26.63	\$ 26.63	\$ 26.63	\$ 26.63	\$ 26.63	\$ 26.63	\$ 26.63	\$ 26.63
Tesla Network Licensing Revenue		26	95	252	979	1,877	3,564	7,296	10,526	14,602	24,549	32,288	47,530	57,643	71,777	86,573	102,580	119,803	138,258	157,963	178,909	199,990
Combined Revenue (\$mm)		559	1,038	2,036	3,764	6,570	11,093	17,914	25,259	33,873	44,565	58,612	72,442	88,004	111,060	133,174	164,185	191,603	222,580	257,563	296,452	335,152
EBITDA																						
EBITDA Margin (%)		30%	35%	40%	45%	50%	60%	65%	65%	65%	65%	65%	65%	65%	65%	65%	65%	65%	65%	65%	65%	65%
Services EBITDA/Tesla Group EBITDA (%)		4.0%	3.8%	4.7%	11.3%	17.4%	22.0%	26.9%	30.1%	31.6%	34.2%	37.3%	45.2%	58.9%	65.4%	65.5%	64.7%	62.7%	62.5%	62.2%	62.3%	62.0%
Depreciation & Amortization																						
D&A/Capex		78	132	244	352	528	841	1,063	1,253	1,651	2,116	2,679	3,232	3,832	4,516	5,266	6,089	6,993	7,872	8,821	9,847	10,957
% Growth		70%	75%	80%	85%	85%	85%	86%	86%	86%	86%	87%	87%	87%	87%	87%	87%	87%	87%	87%	87%	87%
Tax																						
Tax Rate (%)		25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%
Operating Cash Flow																						
Capex		145	306	672	1,358	2,842	5,202	8,999	12,627	16,926	22,254	29,243	36,123	43,860	55,271	66,239	81,562	95,156	110,476	127,767	146,982	166,126
Capex/Sales (%)		20%	17.0%	15.0%	11.0%	9.5%	9.0%	7.0%	6.0%	6.0%	6.0%	6.0%	6.0%	6.0%	6.0%	6.0%	6.0%	6.0%	6.0%	6.0%	6.0%	6.0%
Free Cash Flow																						
		34	129	366	944	2,221	4,212	7,763	11,710	15,006	19,794	26,164	32,408	39,456	50,080	60,186	74,563	87,117	101,428	117,628	135,664	153,532

Source: Morgan Stanley Research

• Tesla Mobility.

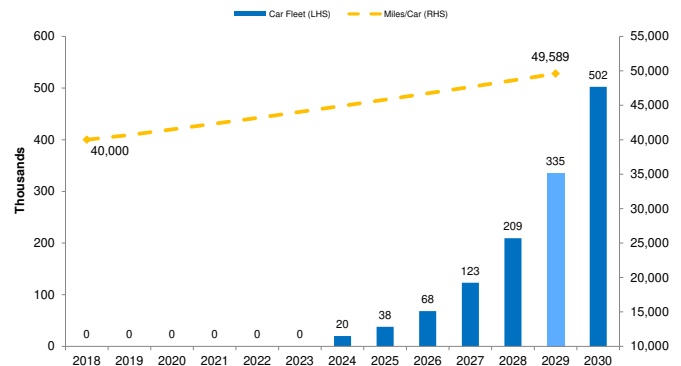
- We value Tesla Mobility at \$81/share on DCF with ~577k fleet by 2030, 18.5% OP margin, \$1.8/mile, and 3% terminal growth rate, vs. \$11/share on DCF with ~500k cars at \$1.7/mile by 2030 and 12.1% OP margin previously.
- **Increased Long-Term Fleet Size** to 577k in 2030 vs. 502k previously. Extended forecasts to 2035, at which point we estimate Tesla's mobility fleet to reach 2.24mn vehicles.
- **Increased Operating Profit Margin** to 18.5% in 2030 vs. 12.1% previously. We expect Tesla to maintain an 18.5% OP margin through 2035.

Exhibit 30: New Fleet and Miles/Car Forecasts




Source: Morgan Stanley Research

Exhibit 31: Prior Fleet and Miles/Car Forecasts




Source: Morgan Stanley Research

Exhibit 32: New Operating Profit Margin Forecast



Source: Morgan Stanley Research

Exhibit 33: Prior Operating Profit Margin Forecast



Source: Morgan Stanley Research

Exhibit 34: Tesla Mobility Model

Tesla Mobility Model: Unlevered DCF (all costs in NOPAT build assumed to be cash)

	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	2035	
Tesla Mobility Fleet	0	0	0	0	0	0	5,000	17,500	50,750	116,725	214,774	360,820	577,313	837,103	1,130,089	1,469,116	1,836,395	2,240,402	
% growth		NM	NM	NM	NM	NM	NM	250.0%	190.0%	130.0%	84.0%	68.0%	60.0%	45.0%	35.0%	30.0%	25.0%	22.0%	
Total Tesla Fleet													37,727,236						
Mobility as a % of Total													1.5%						
Miles/Car	40,000	40,680	41,494	42,323	43,170	44,033	45,795	47,626	50,008	52,508	55,134	57,890	60,206	62,614	64,493	66,427	68,420	70,473	
% growth		1.7%	2.0%	2.0%	2.0%	2.0%	4.0%	4.0%	5.0%	5.0%	5.0%	5.0%	4.0%	4.0%	3.0%	3.0%	3.0%	3.0%	
Total Miles (bn)		0.0	0.0	0.0	0.0	0.2	0.8	2.5	6.1	11.8	20.9	34.8	52.4	72.9	97.6	125.6	157.9		
% growth				NA	NA	NA	264.0%	204.5%	141.5%	93.2%	76.4%	66.4%	50.8%	39.1%	33.9%	28.8%	25.7%		
Billions of Miles/Month		0.0	0.0	0.0	0.0	0.1	0.2	0.5	1.0	1.7	2.9	4.4	6.1	8.1	10.5	13.2			
Rev/Mile (\$)		\$ 2.00	\$ 1.98	\$ 1.96	\$ 1.94	\$ 1.92	\$ 1.90	\$ 1.88	\$ 1.86	\$ 1.85	\$ 1.83	\$ 1.79	\$ 1.75	\$ 1.72	\$ 1.69	\$ 1.65	\$ 1.62		
% growth			-1.0%	-1.0%	-1.0%	-1.0%	-1.0%	-1.0%	-1.0%	-1.0%	-1.0%	-2.0%	-2.0%	-2.0%	-2.0%	-2.0%			
Total Revenue (\$mm)		0	0	0	0	440	1,585	4,779	11,425	21,853	38,163	62,233	91,971	125,328	164,458	207,504	255,535		
% growth				NM	NM	NM	260.4%	201.5%	139.1%	91.3%	74.6%	63.1%	47.8%	36.3%	31.2%	26.2%	23.1%		
OP Margin (%)			-15.0%	-10.0%	-8.0%	-5.0%	-3.0%	0.0%	5.0%	7.5%	12.0%	15.0%	18.5%	18.5%	18.5%	18.5%	18.5%		
Operating Profit (\$mm)		0	0	0	0	-13	0	239	857	2,622	5,724	11,526	17,033	23,211	30,458	38,430	47,326		
Tax Rate (%)			25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%		
NOPAT		0	0	0	0	-10	0	179	643	1,967	4,293	8,644	12,775	17,408	22,843	28,823	35,494		
Mobility as % of Tesla Revs	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.3%	0.9%	2.1%	4.1%	6.3%	9.2%	12.7%						
Mobility as % of Tesla Pretax	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	-0.1%	0.0%	0.7%	2.0%	5.1%	9.4%	15.4%						
Global Miles Traveled (bn)	11,411	11,665	9,916	11,916	12,144	12,382	12,662	12,917	13,248	13,522	13,859	14,169	14,540	14,907	15,272	15,710	16,235	16,752	
Tesla Mobility Share %	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.01%	0.02%	0.05%	0.09%	0.15%	0.24%	0.35%	0.48%	0.62%	0.77%	0.94%	
Utilization (on 24/hr day %)	23%	23%	23%	23%	22%	23%	23%	23%	24%	24%	25%	25%	26%	27%	28%	29%	30%	31%	
Average Speed of Car (Mph)	20	20	21	21	22	22	23	24	24	25	25	26	26	26	26	26	26	26	

Source: Morgan Stanley Research


Tesla 3rd Party Battery Business

- We value Tesla's 3rd party battery business at \$41/share vs. \$28/share previously, with 1.7mn units by 2030, 25% EBITDA margin, and 15x exit EV/ 2040 EBITDA. We previously assumed 1.6mn units by 2030, 25% EBITDA margin, and 24x exit EV/2030 EBITDA. Increased value is driven by higher battery shipments/GWh, as we believe the charging and FSD deals will also result in higher hardware attach.
- **Increased 3rd Party Shipments to 2040.** While we maintain similar, and even slightly more conservative estimates this side of 2030 (we push out the first year of Tesla's 3rd party battery shipments to 2025 vs. 2024 previously), we accel-

erate the pace of shipments from 2030-2040, as we believe the charging and FSD deals will also result in higher hardware attach. In 2035 and 2040, we forecast 5.5mn and 10.2mn shipments, respectively, vs. 4.1mn and 6.8mn previously. Our 2035 GWh supply estimate increased to 439 GWh vs. 325 GWh previously, and our new 2040 estimate is at 815 GWh vs. 543 GWh previously.

- **Maintained 25% EBITDA margin through 2040.** Total EBITDA from 3rd party battery shipments increased to \$23.3bn and \$42.1bn in 2035 and 2040, respectively, vs. \$17.2bn and \$28.0bn previously.

Exhibit 35: Changes to 3rd Party Battery Supply (Shipments and GWh Supply), New vs. Prior Estimates



Source: Morgan Stanley Research

Exhibit 36: Tesla Battery Model

Battery TAM Model	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	2035	2036	2037	2038	2039	2040
Global Battery Supply (GWh)	195	285	377	491	650	824	970	1,109	1,275	1,447	1,672	1,838	2,057	2,309	2,590	2,900	3,181	3,647	4,064	4,354	4,583
% Change	46.4%	32.1%	30.4%	32.3%	26.8%	17.7%	14.3%	15.0%	13.5%	15.6%	9.9%	11.9%	12.2%	12.2%	12.0%	9.7%	14.6%	11.4%	7.1%	5.2%	
\$/kWh (including margin)	\$ 160	\$ 152	\$ 144	\$ 137	\$ 130	\$ 124	\$ 123	\$ 121	\$ 120	\$ 119	\$ 118	\$ 117	\$ 115	\$ 114	\$ 113	\$ 112	\$ 111	\$ 110	\$ 109	\$ 108	\$ 106
% Change	-5.0%	-5.0%	-5.0%	-5.0%	-5.0%	-1.0%	-1.0%	-1.0%	-1.0%	-1.0%	-1.0%	-1.0%	-1.0%	-1.0%	-1.0%	-1.0%	-1.0%	-1.0%	-1.0%	-1.0%	
Global Battery TAM (\$bn)	31	43	54	67	85	102	119	135	153	172	197	214	237	264	293	325	353	400	442	466	488
% Change	39.1%	25.5%	23.9%	25.6%	20.5%	16.5%	13.2%	13.8%	12.4%	14.4%	8.8%	10.8%	11.1%	11.1%	10.8%	8.6%	13.5%	10.3%	6.1%	4.2%	
Tesla Unit Shipments	499,647	936,222	1,313,851	1,862,315	2,478,991	3,400,887	4,368,846	5,152,407	6,135,667	7,004,747	7,857,425	8,486,019	9,164,901	9,898,093	10,689,940	11,545,136	12,468,746	13,466,246	14,543,546	15,707,029	16,963,592
Avg kWh/Unit	80	80	80	80	80	80	80	80	80	80	80	80	80	80	80	80	80	80	80	80	80
Tesla Captive Battery Supply (GWh)	40	75	105	149	198	272	350	412	491	560	629	679	733	792	855	924	997	1,077	1,163	1,257	1,357
Tesla Captive Battery Share (%)	21%	26%	28%	30%	31%	33%	36%	37%	39%	39%	38%	37%	36%	34%	33%	32%	31%	30%	29%	29%	30%
Tesla 3rd Party Shipments	0	0	0	0	10,000	22,000	88,000	396,000	990,000	1,584,000	2,296,800	3,100,680	3,875,850	4,651,020	5,488,204	6,366,316	7,321,264	8,273,028	9,265,791	10,192,370	
% Change					120.0%	300.0%	350.0%	150.0%	60.0%	45.0%	35.0%	25.0%	20.0%	18.0%	16.0%	15.0%	13.0%	12.0%	10.0%		
Non-Tesla Global EV Volume					10,820,497	12,691,688	15,497,882	18,902,699	23,438,387	28,731,809	33,276,499	38,181,380	43,445,411	48,306,795	53,539,257	57,714,782	62,832,792	67,964,643	72,215,371	75,409,356	
Tesla Share of Non-Tesla EV Volume (%)					0.1%	0.2%	0.6%	2.1%	4.2%	5.5%	6.9%	8.1%	8.9%	9.6%	10.3%	11.0%	11.7%	12.2%	12.8%		
Avg kWh/Unit		80	80	80	80	80	80	80	80	80	80	80	80	80	80	80	80	80	80	80	
Tesla 3rd Party Supply (GWh)					1	2	7	32	79	127	184	248	310	372	439	509	586	662	741	815	
Tesla 3rd Party Share (%)					0.1%	0.2%	0.6%	2.5%	5.5%	7.6%	10.0%	12.1%	13.4%	14.4%	15.1%	16.0%	16.1%	16.3%	17.0%	17.8%	
Combined Captive + 3rd Party Share (%)					33.1%	36.2%	37.8%	41.0%	44.2%	45.2%	46.9%	47.7%	47.7%	47.4%	47.0%	47.4%	45.6%	44.9%	45.9%		
Tesla 3rd Party Battery Supply Revenue (\$bn)					0	0	1	4	9	15	21	29	35	42	49	56	64	72	80	87	
Battery Content per Unit (\$)					9,904	9,805	9,707	9,610	9,514	9,419	9,325	9,232	9,139	9,048	8,957	8,868	8,779	8,691	8,604	8,518	
Additional Powertrain Content per Unit (\$)					8,000	8,000	8,000	8,000	8,000	8,000	8,000	8,000	8,000	8,000	8,000	8,000	8,000	8,000	8,000	8,000	
Combined Content (\$/unit)					8,000	17,904	17,805	17,707	17,610	17,514	17,419	17,325	17,232	17,139	17,048	16,957	16,868	16,779	16,691	16,604	
Total Revenue (\$bn)					0	0	2	7	17	28	40	53	66	79	93	107	123	138	154	168	
EBITDA (\$bn)					0.0	0.0	0.1	1.0	3.5	6.9	9.9	13.4	16.6	19.8	23.3	26.8	30.7	34.5	38.5	42.1	
EBITDA Margin (%)					0%	5%	8%	15%	20%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	

Source: Morgan Stanley Research

Changes to PT, Bull Case, and Bear Case

- **Price Target:** We increase our PT to \$400 from \$250 previously. Of the increase, \$7 is related to Core TSLA Auto where we forecast slightly higher EBITDA margins (and still no assumed benefit from IRA), \$70 is from Mobility, and \$58 is from Network Services, and \$13 is from Tesla's 3rd party battery business. At our revised \$400 PT, TSLA trades at 45x 2025 EBITDA and 8x FY25 sales.
- **Core Auto:** We now value Core Auto at \$102/share compared to \$95/share previously. Our unit and ATP forecasts remain unchanged across our forecast horizon, and we now forecast slightly higher EBITDA margins in outer years, partially offset by slightly lower Automotive Gross Profit and slightly higher capex longer term. Our

FY30 EBITDA (GAAP) margin is now at 17.4% vs. 16.6% previously, Auto GP Margin is now at 17.4% vs. 17.7% previously, and Capex as a % of Sales is at 6.9% vs. 6.5% previously. Additionally, we raise our exit EV/EBITDA multiple to 13x from 12x previously. At \$102/share, the automotive business trades at 15.0x our FY25e Auto EBITDA (vs. 14.1x previously).

- **Mobility:** Our valuation for TSLA Mobility now stands at \$81/share vs \$11/share previously. We now assume a fleet of ~577k vehicles by 2030, a 18.5% OP margin, \$1.8/mile, and 3% terminal growth rate. We previously assumed a fleet of ~500k cars at \$1.7/mile by 2030 and 12.1% OP margin. Increased value is driven by higher long term fleet size and margin.

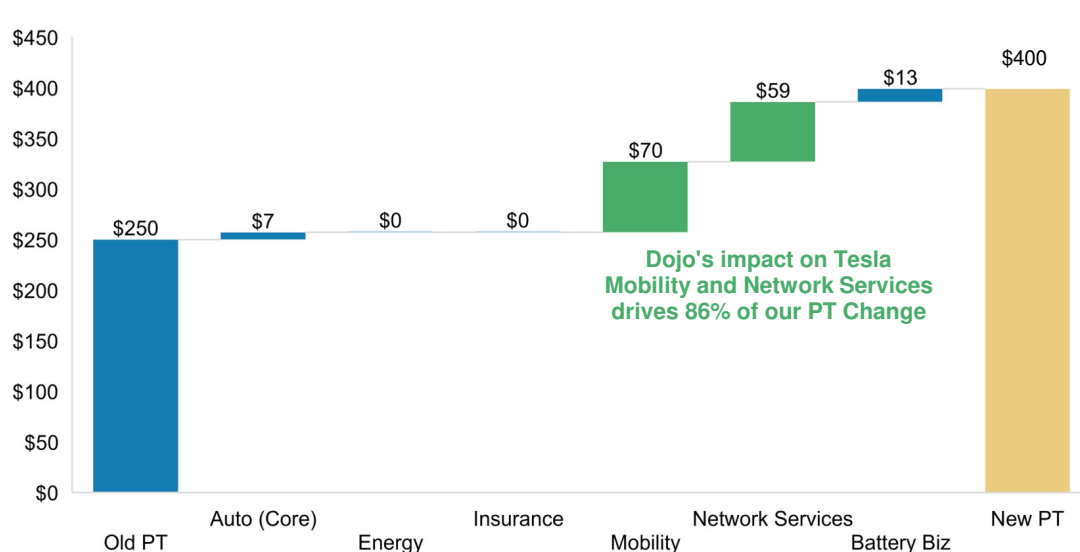
- Network Services:** We value Network Services at \$119/share vs. \$60/share previously, with 23.8mm MAUs at \$180 ARPU by 2030, and added licensing fleet (16.9mn licensing fleet at \$36 ARPU by 2030), at 50% probability. We previously assumed 23.8mm MAUs, \$100 ARPU by 2030, and 70% discount. Increased value driven by higher ARPU and net new 3rd party licensing.
- Battery Biz:** We value Tesla's 3rd party battery business at \$41/share vs. \$28/share previously, with 1.7mn units by 2030, 25% EBITDA margin, and 15x exit EV/ 2040 EBITDA. We previously assumed 1.6mn units by 2030, 25% EBITDA margin, and 24x exit EV/ 2030 EBITDA. Increased value driven by higher battery shipments/GWh, as we believe the charging and FSD deals will also result in higher hardware attach.
- Bull Case:** Our Bull Case valuation now stands at \$550/share vs \$450/share previously. At our revised Bull Case, TSLA trades at 18.5x 2030 EBITDA and 4.4x 2030 sales.
- Bear Case:** Our Bear Case valuation now stands at \$120/share vs \$90/share previously. At our Bear Case, TSLA trades at 3.9x 2030 EBITDA and 0.9x 2030 sales.

Exhibit 37: Price Target Methodology

Price Target Methodology		
Bull Case		
Tesla Auto (Core)	\$150	10 million units by 2030 with 20% EBITDA margin
Tesla Energy	\$65	35% 20yr revenue CAGR, 30% gross margin by 2030
Tesla Insurance	\$16	50% penetration, 20% underwriting margin by 2030
Tesla Mobility/ride-sharing	\$111	1 million car fleet by 2030, 20% OP margin
Tesla Network Services (net)	\$146	25mm MAU at \$200 ARPU by 2030. Net of eliminations
EV P-train 3rd party	\$61	3.0 million units by 2030, 25% EBITDA margin and 25x exit EV/EBITDA
Total	\$550	
Base Case		
Tesla Auto (Core)	\$102	7.9 million units by 2030 at 17.4% EBITDA margin
Tesla Energy	\$48	27% 20yr revenue CAGR, 23.2% gross margin by 2030
Tesla Insurance	\$9	15% penetration and 12% underwriting margin by 2030
Tesla Mobility/ride-sharing	\$81	~600k fleet by 2030, 18.5% OP margin, \$1.8/mile, 3% terminal growth rate
Tesla Network Services (net)	\$119	23.8mm MAU at \$180 ARPU by 2030. Licensing fleet, 50% probability
EV P-train 3rd party	\$41	1.7 million units by 2030, 25% EBITDA margin and 25x exit EV/EBITDA
Total	\$400	
Bear Case		
Tesla Auto (Core)	\$60	5.5 million units by 2030 with 13% EBITDA margin
Tesla Energy	\$15	12% 20yr revenue CAGR, 20% gross margin by 2030
Tesla Insurance	\$0	Assumed \$0 value in bear case
Tesla Mobility/ride-sharing	\$11	100k car fleet by 2030, 15% OP margin
Tesla Network Services	\$34	15mm MAU, \$80 ARPU by 2030
EV P-train 3rd party	\$0	Assumed \$0 value in bear case
Total	\$120	

Source: Morgan Stanley Research

Exhibit 38: Price Target Bridge



Source: Morgan Stanley Research

Valuation

At current levels, the stock is trading at a discount to its 3Y and roughly in line with its 5Y historical valuation (on a forward year EV/EBITDA basis)...


Exhibit 39: Tesla Forward Year EV/EBITDA



Source: Factset, Morgan Stanley Research. Data as of 9/7/23.


...and we think there's room to run. At our price target of \$400, we value Tesla at 91x FY23e EBITDA and 144.2x FY23e earnings. On 2025 numbers, we value the company at 45.4x on an EV/EBITDA basis and 78.9x on a P/E basis.

Exhibit 40: Tesla Forward Year EV/EBITDA



Source: Factset, Morgan Stanley Research. Data as of 9/7/23.

Exhibit 41: Tesla Forward Year PE




Source: Factset, Morgan Stanley Research. Data as of 9/7/23.

Relative Valuation

On a relative basis, Tesla's current valuation is roughly in line with FAANG/software peers on FY26 EBITDA, per our estimates. At its current price of ~\$251 (and relative to MSe), Tesla is trading at a similar multiple to NFLX, at a 4 turn premium to MSFT, ~6 turn premium to GOOGL/AMZN, and 2 turn discount to AAPL. That being said, at a comparable multiple to Big Tech, we believe you also unlock the transformative earnings power of Dojo, which, in our view, can enable a long term revenue CAGR substantially higher vs. Big Tech (in some cases, in excess of 2x consensus FAANG revenue CAGR). *Put simply, we believe Tesla is undervalued at its current price given the revenue opportunity and earnings power we expect around Dojo.*

Exhibit 42: Consensus '26 EV/EBITDA vs Long Term Growth



Source: FactSet, Morgan Stanley Research. Note: GAAP EBITDA (burdened by SBC) is represented. Consensus estimates are used, except where noted. Multiples based on EV as of 9/7/23.

Exhibit 43: Tesla FY26 MSe vs. Consensus Estimates

Tesla FY26 MSe vs. Cons	2026		
	Mse	Cons	% Difference
\$mn			
Total Automotive Revenue	184,275	155,433	18.6%
Services, Energy Storage and Generation	41,534	32,754	26.8%
Total Revenue	225,809	188,186	20.0%
Automotive gross profit	33,587	35,716	(6.0%)
<i>Gross margin (ex.ZEV Credit)</i>	<i>17.2%</i>	<i>24.4%</i>	<i>(7.1%)</i>
Services & Other gross profit	16,184	4,778	238.7%
<i>Gross margin</i>	<i>39.0%</i>	<i>14.6%</i>	<i>24.4%</i>
Total Gross Profit	49,770	42,149	18.1%
<i>Gross margin</i>	<i>22.0%</i>	<i>22.4%</i>	<i>(0.4%)</i>
R&D (burdened by SBC)	7,903	5,467	44.6%
SG&A (burdened by SBC)	9,710	6,826	42.2%
Total Operating Profit (GAAP)	32,157	31,507	2.1%
<i>Operating Margin</i>	<i>14.2%</i>	<i>16.7%</i>	<i>(2.5%)</i>
Net Income	28,177	27,201	3.6%
EPS	7.51	7.40	1.5%
Free Cash Flow	25,732	22,185	16.0%
Capex	11,967.89	10,534.90	13.6%

Source: Factset, Visible Alpha, Morgan Stanley Research

Exhibit 44: Big Tech Morgan Stanley Price Target vs. All-Time Highs


	MS Price Target	All Time High	% Gap
NVDA	\$ 630	\$ 494	27.6%
AAPL	\$ 215	\$ 196	9.4%
GOOGL	\$ 155	\$ 150	3.4%
META	\$ 375	\$ 382	(1.9%)
TSLA	\$ 400	\$ 410	(2.4%)
AMZN	\$ 175	\$ 187	(6.2%)
NFLX	\$ 450	\$ 692	(34.9%)

Source: Factset, Morgan Stanley Research

Relative to MBLY and NVDA – Paying for Quality Long-Term Growth

While trading at parity/slight premium to NVDA on near term estimates, we believe growth in the latter half of the decade via Dojo synergies justifies the valuation. When looking at multiples for these three names, MBLY is the most expensive on FY24 EBITDA, currently trading at 48.7x, while TSLA is slightly more expensive than NVDA, at 41.1x and 39.5x, respectively. Holding EV constant, by FY26 TSLA and MBLY multiples converge, with TSLA at 19.8x 2026 EBITDA and MBLY at 19.6x. Also by FY26, TSLA begins to trade at a noticeable discount to NVDA, and by 2027, the gap widens, with TSLA trading at a ~5 turn discount to NVDA, on our numbers.

This inflection in relative valuation is a function of our view on TSLA: We forecast the company to generate significant Dojo synergies and EBITDA uplift in outer years as an exponentially improved FSD is embedded in a greater share of Tesla's fleet, plus increased licensing revenues, and higher ARPU across Network Services. We expect Network Services to represent about a third (32%) of total company EBITDA by 2028 on our revised forecasts, up from 17% on our prior forecasts. We see incremental EBITDA contribution accelerate to 2040, where we have the segment contributing just over 60% of total company EBITDA, up from 38% previously. We also bake in greater value from Tesla Mobility via increased long term fleet size and OP margin, and 3rd Party Battery Business, as we believe the charging and FSD deals will also result in higher hardware attach.

Exhibit 45: Forward EV/EBITDA Multiples at Current Price

Source: Morgan Stanley Research. Note: We hold EV static based on FY23 balance sheet estimates. EV/EBITDA burdened by SBC. Multiples based on MS estimates, and EV as of 9/7/23.


We highlight that MS Semi Analyst Joe Moore sees NVIDIA's AI capabilities/exposure as a key driver of the stock's multiple premium, and that investor sentiment is clustered around AI rather than at the company's auto exposure/autonomous driving business.

In the below scatterplot, we compare TSLA, MBLY, and NVDA on FY28 EBITDA margin vs. FY28 revenue CAGR. We note that MBLY is in a much earlier stage of their growth trajectory: SuperVision is the next leg of growth for MBLY, and we don't expect to see significant revenue from that product until 2025, driving the higher Revenue CAGR. Additionally, NVDA's materially higher EBITDA margin represents the entire company, while Auto is a lower margin business. We'd also highlight Tesla's underlying need to diversify away from over-reliance on NVDA – Tesla believes they can develop a more efficient system for their specific needs while not funding a supplier's 60% gross margin.

As Tesla begins to unlock Dojo synergies in the back half of the decade and beyond 2030, we expect to see meaningful EBITDA margin expansion. We forecast Network Services to deliver a 65% EBITDA margin, and that it will represent 62% of Tesla's total EBITDA in 2040. We can thus imply a 35% total company EBITDA margin in FY40e, up from 15% in FY23e and 24% FY30e.

As Network Services, at meaningfully higher margins powered by Dojo, begins to compose a greater share of total company EBITDA, we expect to see a mix shift in EBITDA margin up beyond 2030, similar to AWS' role at AMZN today.

Exhibit 46: '22-'28 Revenue CAGR vs 2030 EBITDA Margin



Source: Morgan Stanley Research. Note: EBITDA margin and Revenue CAGR based on Morgan Stanley Research estimates. EBITDA is burdened by SBC.

Exhibit 47: Network Services Unlocks Meaningful Total Company EBITDA Margin Expansion Long Term


Implied 2040 Group EBITDA Margin	2040
Tesla Network Revenue (not incl. licensing)	209,902
<i>Services Revenue/Tesla Group Revenue (%)</i>	<i>20.9%</i>
Implied Group Revenue	1,006,372
Combined Network Services Revenue (incl. licensing)	335,152
Group Revenue	1,006,372
<i>Combined Services Revenue/Tesla Group Revenue (%)</i>	<i>33.3%</i>
Network Services EBITDA	217,849
<i>EBITDA Margin (%)</i>	<i>65%</i>
<i>Services EBITDA/Tesla Group EBITDA (%)</i>	<i>62.0%</i>
Implied Group EBITDA	351,631
Implied Group EBITDA Margin	34.9%

Source: Morgan Stanley Research

The AWS Parallel

In the early 2000s, Amazon set out to increase developer productivity by decreasing the time spent on building foundational website features for each new project, and thus AWS was born. AWS, an internal cloud service, gave developers a shared software infrastructure that helped drive efficiency for developers. As the bookseller and online marketplace business grew, the company realized that the technology used to power their core business could be monetized by being rented to 3rd parties. Today, AWS has grown to contribute ~70% of total Amazon EBIT (on our Internet team's 2023e numbers), generating 2.5x the EBIT as the retail business, at 20% of the revenue.

Exhibit 48: Amazon EBIT Contribution from AWS and Retail Businesses




Source: Company Data, Morgan Stanley Research

The same forces that have driven AWS to reach 70% of AMZN total EBIT today can work at Tesla, in our view, opening up new addressable markets that extend well beyond selling vehicles at a fixed price. We forecast Network Services to reach 62% of total TSLA EBIT by 2040.

Our intent is not to say that Dojo will replace AWS. However, we do believe that AWS provides a powerful example of how leveraging core capabilities, in this case, specialized data processing and software, can create significant value drivers.

Exhibit 49: Network Services as a % of TSLA EBIT vs. AWS as a % of AMZN EBIT



Source: Company Data, Morgan Stanley Research

Further Industry Implications

Can Other OEMs & AI Companies Benefit From Dojo?

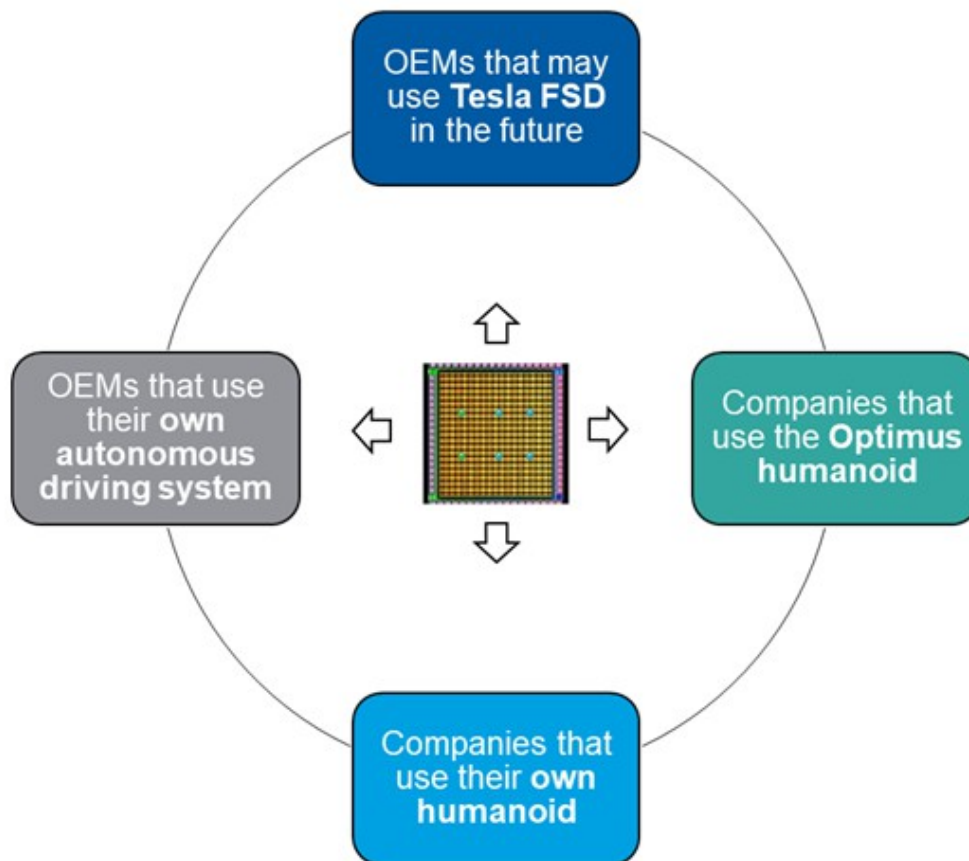
Before getting into Dojo's long term potential, we note that we do not believe that current cloud providers will be a target for Tesla in the short to medium term. We view the potential of a Tesla 'Cloud' as a long term Moonshot, sitting below solving autonomy internally on the company's near term priority list. At the same time, we do recognize that Tesla is building material internal capabilities that over time, could be monetized in an 'AWS-type' business model.

According to Senior Hardware Director Ganesh Venkataramanan, Tesla's priority is to internally implement Dojo, but that it eventually may be available to non-Tesla users. In the 1Q23 earnings call, Musk mentioned that the future of Dojo could be something like

Amazon Web Services (AWS), giving it the potential to be "a sellable service that we would offer to other companies". He can also be quoted calling Dojo a "long-shot bet" that could pay off in the "multi-hundred billion" dollar potential outcome if it develops as the company hopes.

Rather than selling their chips and supercomputer system, the company appears open to potentially supplying cloud services in the future. This could potentially span beyond vehicles to across any vision-focused industry. Within the wave of selling AI compute as a service, in March 2023, NVIDIA released their DGX Cloud, an all-inclusive AI supercomputer in the cloud – starting at \$36,999 a month. The rental includes access to eight H100s or A100s and is close to double the price (\$20,000 per month) of Microsoft Azure's A100 option with 96 CPUs.

Exhibit 50: AWS-esque potential beyond the Muskonomy



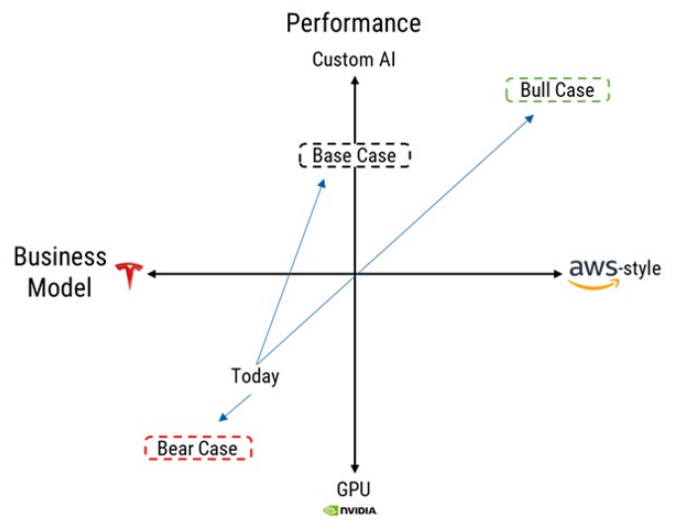
Source: Morgan Stanley Research

If another OEM were to incorporate Tesla's FSD system into their vehicles, Dojo's scalability materially increases. Musk recently mentioned in the 2Q23 earnings that Tesla has been in talks with an OEM about a potential implementation of FSD into their vehicles, which can directly benefit from Dojo's training. Even if other OEMs were to develop their own autonomous systems, Dojo is purposely tailored to turn video and visual data into useful outputs for autonomous driving. Its designed use case places it more favorably for transportation-based neural network training than the current selection of preferred general purpose GPUs. The chip is designed to advance and accelerate autonomous driving, which can potentially be taken advantage of in any scenario that involves large and complex video recognition data.

If Dojo can report efficiencies in data acceleration, Tesla has potential to be among the best vision and video processing ML training system on the market. Similar to how NVIDIA dominates the GPU market and Qualcomm dominates the mobile semiconductor market, Tesla may have the ability to be the top player in the machine learning training market that utilizes *visual data* - and the benefits can go beyond companies producing autonomous vehicles. While a host of industries could benefit from visual AI acceleration, we offer our thoughts on industries that we view as examples of where Dojo could possibly be implemented to accelerate safety and efficiency via visual-data AI applications:

- **Automobiles:** The application for visual AI within the automotive space extends from ADAS to full autonomous driving. All OEMs may be able to benefit from Dojo as it may be able to train autonomous vehicles at a greater efficiency and cost than all other supercomputer offerings.
- **Robotics:** Any company that either uses or develops their own humanoid robot may be able to benefit from Dojo's accelerated visual training. Internally, Dojo is being developed to train Optimus and has potential to be used to train other humanoid bots as it can be the cutting-edge solution for visual processing. On the 4Q21 earnings call, Elon Musk stated that "*Optimus ultimately will be worth more than the car business and worth more than FSD*".
- **Aviation/Air Mobility:** Computer vision plays a key role in improving safety of airplanes through object identification and early detection for potential hazards in real-time. In addition to the current usage, we could see the development of autonomous or remotely supervised planes that will utilize data similar to how an autonomous car would.

Exhibit 51: Potential scenarios for Dojo

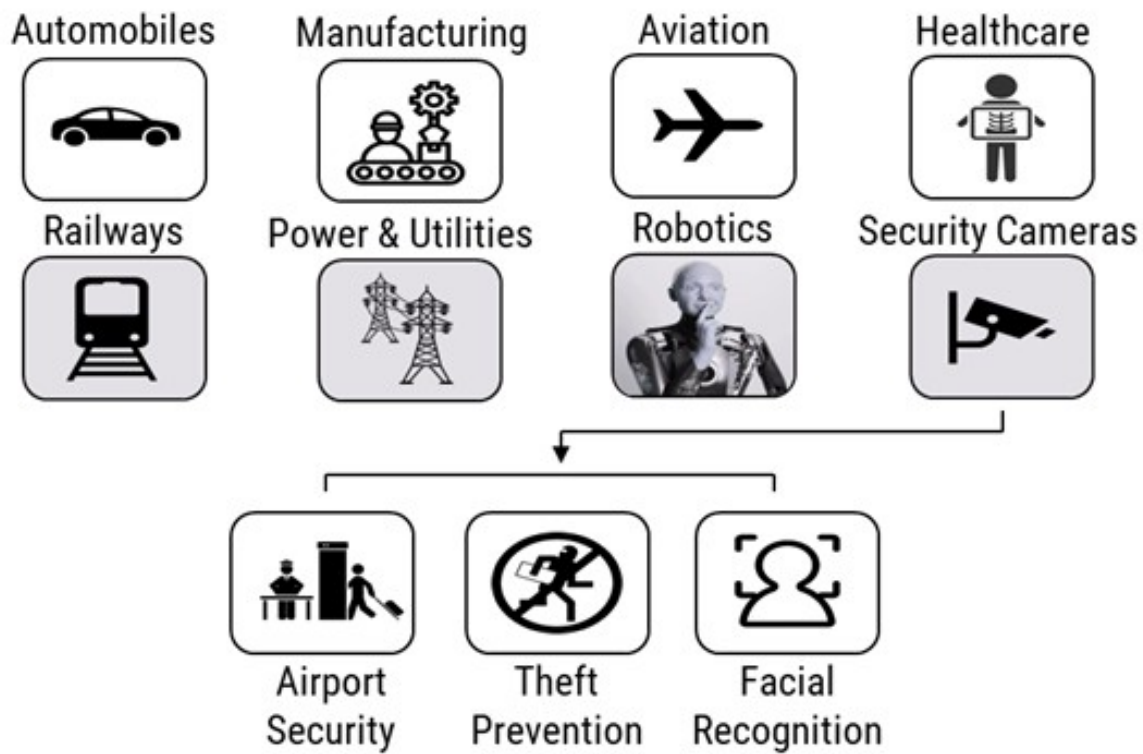


Source: Morgan Stanley Research

- **Railways:** The AI applications for railways range from autonomous trains to enhanced safety. Railway computers are used in driver assistance systems for driving, breaking control, and collision protection systems, which can reduce energy consumption by 15%. Through predictive maintenance via evaluating telemetry data, work can be planned at an earlier stage which can reduce downtime and railway maintenance costs.
- **Security & Cameras:** The application of AI in video surveillance can help identify people, vehicles, objects, and behaviors in real time. It can be used to detect emergency alert situations, automate the location of people or vehicles (for traffic management systems), and prevent crime through facial recognition and automated reactions. In addition to that, AI can monitor and process multiple visual feeds at once beyond the capabilities of a human, which can reduce costs and the need for an operator. AI applications for camera improvements can be utilized by just about every industry, ranging from identifying theft and managing stock in retail to monitoring factories in manufacturing.
 - **Facial Recognition:** AI can help with facial recognition, such as identifying patterns like height, width of face, color of face, and unique features of the face. With these tools, cameras can automatically help identify faces in a crowd and can match faces together from multiple frames in real time.

- **Airport Security:** AI can significantly improve airport security through the identification of passengers, automate the baggage check-in process, and can speed up the security process. The hassle surrounding security checks would be reduced due to cameras that can use facial recognition to match a passenger's face to their boarding pass and passport. This can help shorten the time required for security checks, reduce queues, and help airports become more efficient.
- **Healthcare:** Visual data can play a key role in the automation and improvement of medical technology and disease detection. Medical imaging, which includes X-rays, MRIs, and Ultrasounds, is one area that can see improvement in accuracy from visual AI. A second area that can see improvements is in robotic surgeries, which uses small cameras and robotic controls to perform minimally invasive operations. In addition to that, AI can be used to predict rare diseases based on medical records that can potentially lead to predictive solutions.

Exhibit 52: Industries that can utilize visual AI in the near future



Source: Morgan Stanley Research

A quick aside on Optimus... Dojo's accelerated visual training is used to enhance the neural net internally, which in turn trains Optimus. While AI Day and the March Investor Day garner investor hype around robotics, we look forward to seeing whether the Tesla Bot is much more than a gimmick...whether we could expect either significant cost-savings in manufacturing down the line, a new revenue stream, or both. In our view, Dojo has the potential to significantly reduce time to market and cost saving realization from Optimus applications, in the long-term.

Elon Musk does mention that production at scale is a main priority, and "is why Tesla engineering has transitioned to focus heavily on designing the machine that makes the machine – turning the factory itself into a product." Successful robotic assistance in the production line could result in systematic cost reductions and alleviate labor shortages long-term. We remind investors that there are far bigger forces at work here on the interplay of labor demographics, education, immigration, union organization and other factors. Energy transition and on-shoring industrial manufacturing significantly accelerate the pay-backs, trade-offs and social implications of human replacement behind the wheel, in the mine, at the warehouse and on the factory floor.

That being said, we do not ascribe any value to Optimus, either as a 'line item' or via potentially realized cost savings in our Tesla model (and would discourage investors from doing the same) at this time.

See relevant reading:

- [Tesla Inc: Tesla AI Day: What's the Tesla Bot Really About? \(28 Sep 2022\)](#)
- [Tesla Inc: Giga Austin: Tesla Leading the 'Race to the Bottom' on EV Costs \(2 Mar 2023\)](#)

Risks & Caution:

- **Regulatory adoption of autonomous vehicles:** Even if Dojo can accelerate the timeline of autonomous vehicles and humanoid bots, the world may not be ready to adapt it yet. There will need to be significant regulations in place to handle the technological advancements, which could add time to when we begin to see Dojo's results in our lives.
- **Hype > Results:** Tesla has previously made promises regarding timelines, efficiencies, and results that haven't necessarily panned out yet. As with all disruptive technology, there is a possibility that Dojo turns into an expensive science project that doesn't reach the heights originally anticipated. In that sense, Tesla would likely lean on NVIDIA GPUs for compute and AI training, continuing to fight for availability within the scarce market.
- **Efficiency/cost benefits still hold at scale?** A common part of why companies shy away from custom AI ASICs is due to the challenges with scaling. It often takes longer to hurdle the initial design investment costs than companies expect, as Google's TPU technology (released in 2016) only started to be scaled for external use in recent years. Unless Tesla can scale to their intended size, it may be longer before they begin to see cost benefits.

Tesla Valuation Matrix

Exhibit 53: Tesla Valuation Matrix

Tesla Valuation Matrix						
Inputs		2022-Q4				2022 CAGR to 2030
Current Share Price (\$)		\$291.26				
Diluted Shares Outstanding (mm) - YE22		3,471				
Market Cap		1,010,963				
Net Debt (YE-22)		(19,086)				
Enterprise Value		991,877				
	2022	2023E	2024E	2025E	2030E	
Sales	81,462	100,254	131,265	176,403	429,310	23%
Operating Income	13,656	10,098	14,245	21,446	81,225	
% Margin	16.8%	10.1%	10.9%	12.2%	18.9%	
EBITDA	17,360	15,049	20,825	30,194	102,112	25%
% Margin	21.3%	15.0%	15.9%	17.1%	23.8%	
EPS	\$3.62	\$2.77	\$3.43	\$5.07	\$17.35	22%

Share Price (\$)	Market Value (\$MM)	Enterprise Value (\$MM)	Price/Sales (x)			EV/Sales (x)			EV/EBITDA			P/E		
			2023E	2025E	2030E	2023E	2025E	2030E	2023E	2025E	2030E	2023E	2025E	2030E
\$50	\$173,550	\$154,464	1.7x	1.0x	0.4x	1.5x	0.9x	0.4x	10.3x	5.1x	1.5x	18.0x	9.9x	2.9x
\$70	\$242,970	\$223,884	2.4x	1.4x	0.6x	2.2x	1.3x	0.5x	14.9x	7.4x	2.2x	25.2x	13.8x	4.0x
\$90	\$312,390	\$293,304	3.1x	1.8x	0.7x	2.9x	1.7x	0.7x	19.5x	9.7x	2.9x	32.5x	17.8x	5.2x
\$110	\$381,810	\$362,724	3.8x	2.2x	0.9x	3.6x	2.1x	0.8x	24.1x	12.0x	3.6x	39.7x	21.7x	6.3x
\$120	\$416,520	\$397,434	4.2x	2.4x	1.0x	4.0x	2.3x	0.9x	26.4x	13.2x	3.9x	43.3x	23.7x	6.9x
\$150	\$520,650	\$501,564	5.2x	3.0x	1.2x	5.0x	2.8x	1.2x	33.3x	16.6x	4.9x	54.1x	29.6x	8.6x
\$170	\$590,070	\$570,984	5.9x	3.3x	1.4x	5.7x	3.2x	1.3x	37.9x	18.9x	5.6x	61.3x	33.5x	9.8x
\$190	\$659,490	\$640,404	6.6x	3.7x	1.5x	6.4x	3.6x	1.5x	42.6x	21.2x	6.3x	68.5x	37.5x	11.0x
\$210	\$728,910	\$709,824	7.3x	4.1x	1.7x	7.1x	4.0x	1.7x	47.2x	23.5x	7.0x	75.7x	41.4x	12.1x
\$230	\$798,330	\$779,244	8.0x	4.5x	1.9x	7.8x	4.4x	1.8x	51.8x	25.8x	7.6x	82.9x	45.4x	13.3x
\$250	\$867,750	\$848,664	8.7x	4.9x	2.0x	8.5x	4.8x	2.0x	56.4x	28.1x	8.3x	90.2x	49.3x	14.4x
\$270	\$937,170	\$918,084	9.3x	5.3x	2.2x	9.2x	5.2x	2.1x	61.0x	30.4x	9.0x	97.4x	53.3x	15.6x
\$290	\$1,006,590	\$987,504	10.0x	5.7x	2.3x	9.9x	5.6x	2.3x	65.6x	32.7x	9.7x	104.6x	57.2x	16.7x
\$310	\$1,076,010	\$1,056,924	10.7x	6.1x	2.5x	10.5x	6.0x	2.5x	70.2x	35.0x	10.4x	111.8x	61.2x	17.9x
\$330	\$1,145,430	\$1,126,344	11.4x	6.5x	2.7x	11.2x	6.4x	2.6x	74.8x	37.3x	11.0x	119.0x	65.1x	19.0x
\$350	\$1,214,850	\$1,195,764	12.1x	6.9x	2.8x	11.9x	6.8x	2.8x	79.5x	39.6x	11.7x	126.2x	69.0x	20.2x
\$370	\$1,284,270	\$1,265,184	12.8x	7.3x	3.0x	12.6x	7.2x	2.9x	84.1x	41.9x	12.4x	133.4x	73.0x	21.3x
\$390	\$1,353,690	\$1,334,604	13.5x	7.7x	3.2x	13.3x	7.6x	3.1x	88.7x	44.2x	13.1x	140.6x	76.9x	22.5x
\$400	\$1,388,400	\$1,369,314	13.8x	7.9x	3.2x	13.7x	7.8x	3.2x	91.0x	45.4x	13.4x	144.2x	78.9x	23.1x
\$410	\$1,423,110	\$1,404,024	14.2x	8.1x	3.3x	14.0x	8.0x	3.3x	93.3x	46.5x	13.7x	147.9x	80.9x	23.6x
\$430	\$1,492,530	\$1,473,444	14.9x	8.5x	3.5x	14.7x	8.4x	3.4x	97.9x	48.8x	14.4x	155.1x	84.8x	24.8x
\$450	\$1,561,950	\$1,542,864	15.6x	8.9x	3.6x	15.4x	8.7x	3.6x	102.5x	51.1x	15.1x	162.3x	88.8x	25.9x
\$470	\$1,631,370	\$1,612,284	16.3x	9.2x	3.8x	16.1x	9.1x	3.8x	107.1x	53.4x	15.8x	169.5x	92.7x	27.1x
\$490	\$1,700,790	\$1,681,704	17.0x	9.6x	4.0x	16.8x	9.5x	3.9x	111.7x	55.7x	16.5x	176.7x	96.7x	28.2x
\$510	\$1,770,210	\$1,751,124	17.7x	10.0x	4.1x	17.5x	9.9x	4.1x	116.4x	58.0x	17.1x	183.9x	100.6x	29.4x
\$530	\$1,839,630	\$1,820,544	18.3x	10.4x	4.3x	18.2x	10.3x	4.2x	121.0x	60.3x	17.8x	191.1x	104.6x	30.5x
\$550	\$1,909,050	\$1,889,964	19.0x	10.8x	4.4x	18.9x	10.7x	4.4x	125.6x	62.6x	18.5x	198.3x	108.5x	31.7x
\$570	\$1,978,470	\$1,959,384	19.7x	11.2x	4.6x	19.5x	11.1x	4.6x	130.2x	64.9x	19.2x	205.6x	112.4x	32.9x
\$590	\$2,047,890	\$2,028,804	20.4x	11.6x	4.8x	20.2x	11.5x	4.7x	134.8x	67.2x	19.9x	212.8x	116.4x	34.0x
\$610	\$2,117,310	\$2,098,224	21.1x	12.0x	4.9x	20.9x	11.9x	4.9x	139.4x	69.5x	20.5x	220.0x	120.3x	35.2x
\$630	\$2,186,730	\$2,167,644	21.8x	12.4x	5.1x	21.6x	12.3x	5.0x	144.0x	71.8x	21.2x	227.2x	124.3x	36.3x
\$650	\$2,256,150	\$2,237,064	22.5x	12.8x	5.3x	22.3x	12.7x	5.2x	148.6x	74.1x	21.9x	234.4x	128.2x	37.5x


Source: Morgan Stanley Research

Risk Reward – Tesla Inc (TSLA.O) Top Pick


Beyond EVs: Internet of Cars and Network/Software Services Optionality

PRICE TARGET **\$400.00**

Our PT of \$400 is comprised of 6 components: **(1)** \$102/share for core Tesla Auto business on 7.9mm units in 2030, 9.0% WACC, 13x 2030 exit EBITDA multiple, exit EBITDA margin of 17.4%. **(2)** Tesla Mobility at \$81 on DCF with ~600k cars at \$1.8/mile by 2030. **(3)** Tesla as a 3rd party supplier at \$41/share. **4)** Energy at \$48/share, **5)** Insurance at \$9/share, & **6)** Network Services at \$119, 23.8mm MAUs, \$180 ARPU by 2030, 50% discount.



RISK REWARD CHART AND OPTIONS IMPLIED PROBABILITIES (12M)




Source: Refinitiv, Morgan Stanley Research, Morgan Stanley Institutional Equities Division. The probabilities of our Bull, Base, and Bear case scenarios playing out were estimated with implied volatility data from the options market as of {{date}} 07 Sep, 2023. All figures are approximate risk-neutral probabilities of the stock reaching beyond the scenario price in either three-months' or one-years' time. View explanation of Options Probabilities methodology [here](#)

OVERWEIGHT THESIS

- **A Dojo-Enabled Double-Fly-Wheel.** We believe TSLA can leverage its EV cost leadership to expand user base and generate a higher % of revenue from recurring/high-margin software & services. Dojo is the key accelerant at the intersection of hardware and software.
- **Network Services in focus.** We forecast TSLA's services EBITDA to account for 37% of total EBITDA by 2030 & 62% by 2040. Includes: FSD, infotainment, upgrades, charging, maintenance, etc.
- **Attractive R/R.** Incl. Services, Energy & Mobility in our forecast, at \$400, Tesla trades at ~13x 2030 EBITDA and ~3x 2030 sales. Bear Case \$120 & Bull Case \$550.
- **Growth:** We forecast TSLA to sell 7.9mm units by 2030 and grow revenue at a 25% 8yr CAGR.

Consensus Rating Distribution



Risk Reward Themes

- Disruption:** Positive
- Secular Growth:** Positive
- Electric Vehicles:** Positive

View descriptions of Risk Rewards Themes [here](#)

BULL CASE	\$550.00	BASE CASE	\$400.00	BEAR CASE	\$120.00
~18.5x 2030e EV/EBITDA		~13.4x 2030e EV/EBITDA		~3.9x 2030e EV/EBITDA	
For the core auto business, we assume TSLA is able to deliver 10mm units by 2030 with ~20% EBITDA margin, which implies a value of ~\$150/share. We value TSLA Mobility / Rideshare at \$111/share. For Energy, \$65/share (35% 20yr rev CAGR). \$61/share allocated for TSLA as an EV powertrain & battery supplier (assumes 3mm units at 25% EBITDA margin and 25x exit EV/EBITDA at 2030). Insurance is valued at \$16/share. TSLA Network Services valued at \$146/share, on 25mm connected MAUs at \$200 Monthly ARPU.		Our \$400 PT is comprised of 6 components (1) \$102/share for core Tesla Auto business on 7.9mm units in 2030, 9.0% WACC, 13x 2030 exit EBITDA multiple, exit EBITDA margin of 17.4%. (2) Tesla Mobility at \$81/share on DCF with ~600k cars at \$1.8/mile and 18.5% OP margin by 2030. (3) Tesla as a 3rd party supplier at \$41/share. 4) Energy at \$48/share, 5) Insurance at \$9/share, & 6) Network Services at \$119/share, 23.8mm MAUs, \$180 ARPU by 2030, 50% discount for tech, execution, and competitive risks		Our \$120 bear case ascribes \$60/share for automotive which assumes 5.5mm units by 2030 at a 13% EBITDA margin. Other value is ascribed to Tesla Mobility at \$11/share on a 100k car fleet and 15% OP margin by 2030, Tesla Network Services at \$34/share (15mm MAUs at \$80 ARPU by 2030) and Tesla Energy at \$15/share. Tesla as a 3rd party supplier valued at \$0/share, \$0/share value for insurance.	

Risk Reward – Tesla Inc (TSLA.O)


KEY EARNINGS INPUTS

Drivers	2022	2023e	2024e	2025e
Total Deliveries	1,313,851	1,862,315	2,478,990	3,400,887
Total Revenue (\$, mm)	81,462	100,254	131,265	176,403
Auto Gross Margin (%)	28.5	19.5	18.2	17.9
Free Cash Flow (\$, mm)	7,566	4,468	13,206	19,044
Net Debt (Cash) (\$, mm)	(19,086)	(23,766)	(36,972)	(56,016)

INVESTMENT DRIVERS

- Sales, China Market Announcements
- Pricing Adjustments
- Berlin/Austin Giga Ramp
- Fremont Battery Pilot Ramp
- Cybertruck & New Model announcements
- Emerging Competition (from traditional OEMs, startups, & large tech firms)
- Services disclosure

GLOBAL REVENUE EXPOSURE



Source: Morgan Stanley Research Estimate
View explanation of regional hierarchies [here](#)

MS ALPHA MODELS

5/5 BEST	24 Month Horizon	5/5 MOST	3 Month Horizon
-------------	---------------------	-------------	--------------------

Source: Refinitiv, FactSet, Morgan Stanley Research; 1 is the highest favored Quintile and 5 is the least favored Quintile

RISKS TO PT/RATING

RISKS TO UPSIDE

- Disclosure on service revs
- Increased FSD attach rate
- Cost milestones on new battery
- New model intro (Cybertruck, multivan, Semi)
- 3rd party battery win
- Geographic penetration & new capacity

RISKS TO DOWNSIDE

- Competition: legacy OEMs/Chinese players/big tech
- Execution risk: multiple factory ramps
- Market does not recognize Dojo-enabled services op, lower than expected attach rate & RPU
- China risk
- Dilution
- Valuation


OWNERSHIP POSITIONING

Inst. Owners, % Active	45.8%	<div style="width: 45.8%;"></div>
HF Sector Long/Short Ratio	1.5x	<div style="width: 1.5;"></div>
HF Sector Net Exposure	9.3%	<div style="width: 9.3%;"></div>

Refinitiv; MSPB Content. Includes certain hedge fund exposures held with MSPB. Information may be inconsistent with or may not reflect broader market trends. Long/Short Ratio = Long Exposure / Short exposure. Sector % of Total Net Exposure = (For a particular sector: Long Exposure - Short Exposure) / (Across all sectors: Long Exposure - Short Exposure).

MS ESTIMATES VS. CONSENSUS

FY Dec 2024e



Source: Refinitiv, Morgan Stanley Research


Appendix - Dojo in Detail

Steve Jobs [quoted](#) Alan Kay at Apple's iPhone launch back in 2007, saying "people who are really serious about software should make their own hardware". It's evident by the resources dedicated to Dojo that Tesla is serious about their future beyond just selling vehicles.

Top-of-the-industry solution. Comparing an individual Dojo D1 chip to other industry standard ML training chips, Tesla may be creating one of the most powerful supercomputers in the world, all while building each chip with no virtual memory, less cores, and transistors than competing chips. The Dojo system contains 1.3TB of SRAM and 13TB of high-bandwidth DRAM through the Dojo Interface Processors.

From 2010 – 2020, the amount of annual global data generated went from 2 zettabytes to 64 zettabytes, with a projected 181 zettabytes in 2025. Over the same period, 80% of data growth has been unstructured, or machine learning focused ("Software 2.0") and between 2015 – 2020, compute used to train the largest models increased by 300,000x (doubled every 3.5 months). As data gets more complex, for example when implementing real world data, traditional systems and chip architectures struggle to process it. However, machine learning techniques enable the processing of new data, traditional data, AND real-world data, such as driving in streets made for humans, and robotics interfacing for human environments.

Exhibit 54: Global data volume per year (zettabytes)



Source: Company data, Morgan Stanley Research

The process for learning features AI, ML, and Deep-Learning (DL).

- AI refers to tasks that require near-human intelligence in applicable world settings.
- ML, a subset of AI, is geared for specific tasks by learning from data and making predictions.
- DL, a subset of ML, uses deep learning neural network architecture to process higher-level data features.

Traditional compute systems are used to input data, which programmers then hand-code logic through programs into traditional computers, leading to useful outputs. ML systems, however, take input and output data and insert it through learning computers to receive trained logic. Afterwards, trained logic is inserted back into the learning computers to calculate useful outputs. Both learning computer processes are part of training, while the second is used for inference. Training refers to the learning of a new capability from existing


data, while inference applies this capability to new data through application and makes predictions to produce actionable results.

Dojo is an ideal long-term solution due to gaps in chip scaling and processing visual data.

The build of a standard system goes from chip to package, package to board, board to box, box to rack, and rack to data center. Throughout the process, a great deal of bandwidth is lost, and the latency increases multi-fold. A low latency computer network is optimized to process a very high volume of data messages with minimal delay, which is sought after when computing data in real-time. Measured by the amount of energy per 64-bit operation, the traditional hierarchy of power is very taxing on the system. Systems mitigate this by using AI chips that are on reticle (about the size of a quarter) sized dies. Key problems with modern data centers are that GPUs take up great capacity and are very power dense (2-3x power per square foot). These systems have cooling difficulties and cost increases as the systems get more powerful, and a point is


reached where scaling up in size does not provide enough added benefit to outweigh the cost. Lateral power delivery is used, which causes voltage to go down (leading to an increase in amperes needed), requiring more power and resulting in higher costs. Performance and power consumption go hand-in-hand – as compute systems get more advanced, the amount of power needed increases.

Exhibit 55: Global Power Trend (measurement of TDP in watts)



Source: Company data, Morgan Stanley Research

Exhibit 56: Lateral power delivery challenges: as systems get bigger, voltage goes down, leading to higher costs and power required



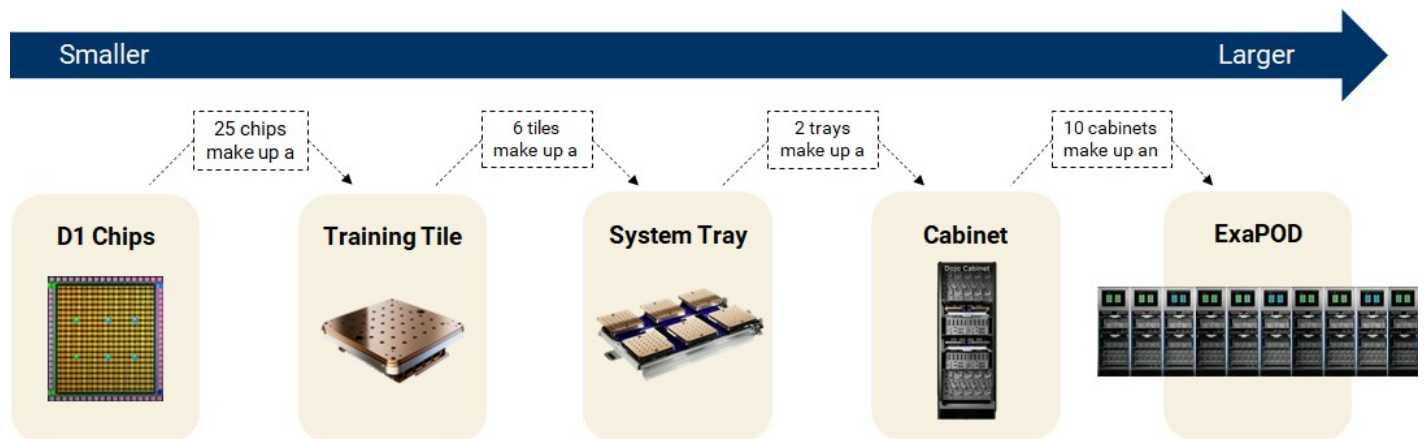
Source: Company data, Morgan Stanley Research

Traditional systems use Internet Optimized data center arches, which have high latency, high packet size, low bandwidth, and is HPU/CPU centric. In Tesla’s case, the cost and delay from supporting enough GPUs to continue to improve and expedite the autonomous learning process isn’t sustainable nor efficient. In the long-term, Tesla needs a system that can be highly scalable, have low latency, high bandwidth, and can be utilized at high efficiency to support, develop, and train their FSD platform (and beyond). This is where Dojo fits in.

Dojo's Hardware

Hardware problem? A problem that generative-AI and ML training systems are facing is that the advancement of hardware is well behind the software. Tesla’s solution to this problem is developing their own innovative hardware structure, piece-by-piece, to provide the ideal architecture to run Tesla’s computations. Its architecture flows from CPU (nodes, cores) to D1 Die to Training Tile to System Tray to Cabinet to ExaPOD system. There are 354 computing cores per D1 Die, 25 D1 Dies per Tile, 6 Tiles per System Tray, 2 Trays per Cabinet, and 10 Cabinets per ExaPOD, resulting in 1,062,000 cores and ~1.1 exaFLOPs of compute.

Exhibit 57: The Dojo Building Blocks: From D1 to ExaPOD



Source: Morgan Stanley Research

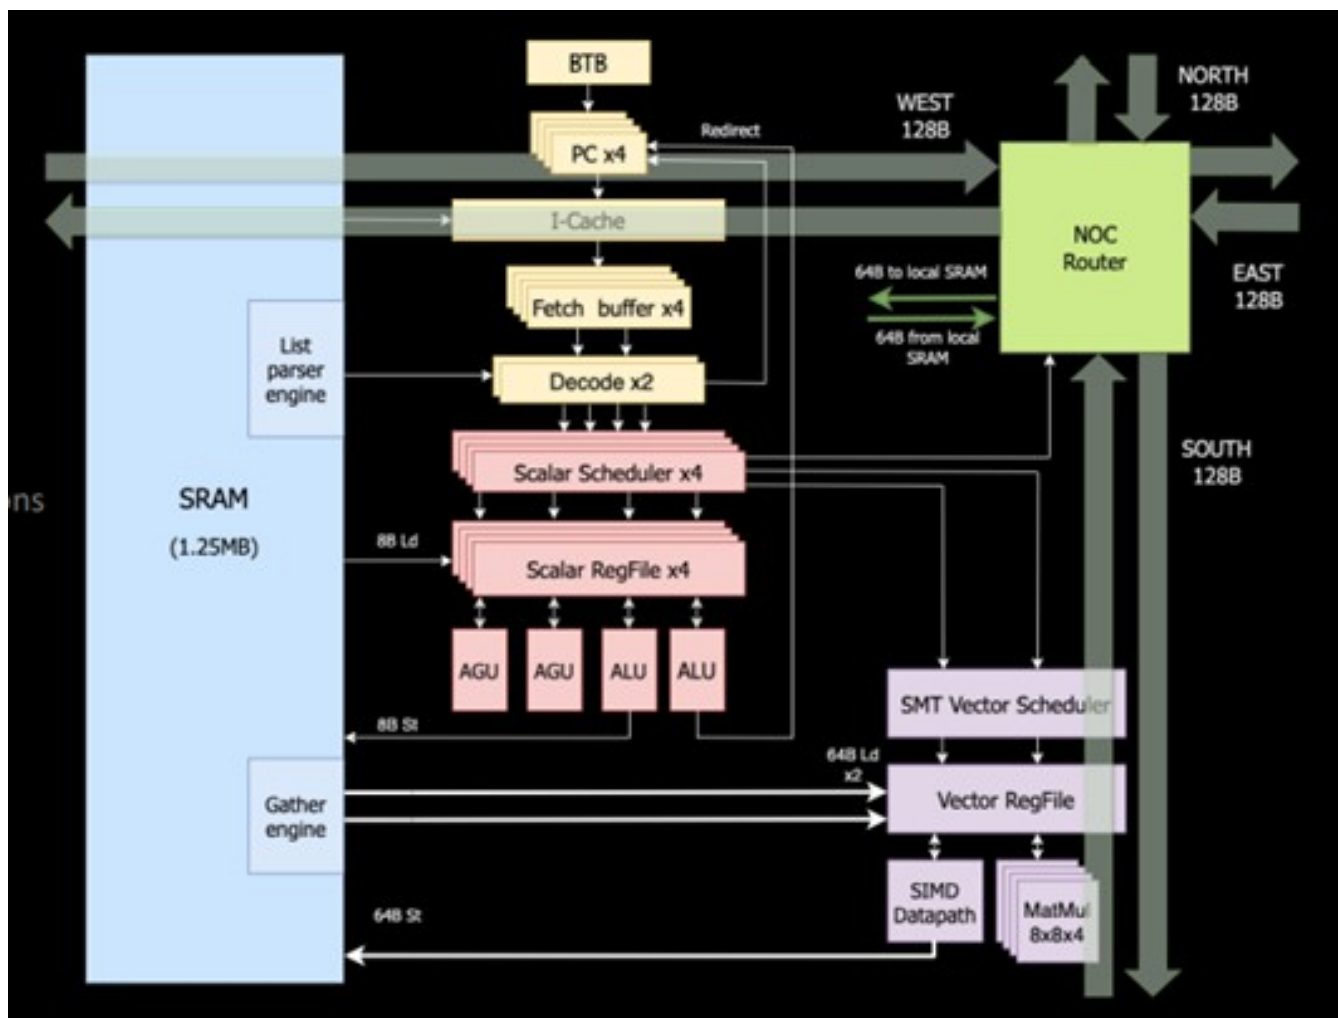
Unique vertical architecture. The Dojo chip itself has a unique architecture compared to conventional supercomputer designs. Typically, modern GPUs have many more cores than the Dojo chip (354 cores) does, as seen in NVIDIA’s A100 and H100 (6,912 and 16,895 respectively), the Cerebras WSE-2 (850,000+), Google’s TPU v4 (4,096) and Graphcore’s MK2 IPU (1,472). Dojo instead combines each of these cores, which are full-fledged computers themselves, together in a 10-cabinet ExaPOD that features a total 1,062,000 cores. Unlike other chips, its vertically stacked infrastructure and lack of virtual memory uniquely caters to Tesla’s AI needs of data-transfer capacity and speed. Tesla mentioned at AI Day 2021 that at the same cost as what’s currently available, Dojo will have 4.0x better performance, 1.3x better performance/watt (energy savings), and a 5.0x smaller footprint.

Musk and NVIDIA VP of Automotive Danny Shapiro can both be quoted mentioning that Tesla’s A100s typically use the TF32 compute format, but Dojo uses a FP16 mixed precision format (known as

BF16), allowing them to take advantage of a higher compute system, resulting in less Dojo chips needed to replace the A100 equivalent of compute power.

The Dojo D1 chip is comprised of small nodes, each containing purpose-built 64-bit CPUs with superscalar cores. Built on TSMC’s 7nm technology, each node acts like an individual computer, with a dedicated CPU, local memory, and communicative I/O SerDes interfaces, meaning that each core can operate independently and isn’t dependent on shared caches or register files. However, the cores do not support virtual memory like other chips would, as management believes it takes up too much space, slows down bandwidth and increases latency. Instead, each core has 1.25MB of SRAM that acts as the chip’s main internal memory, with a load speed of 400GB/sec and store speed of 270GB/sec.

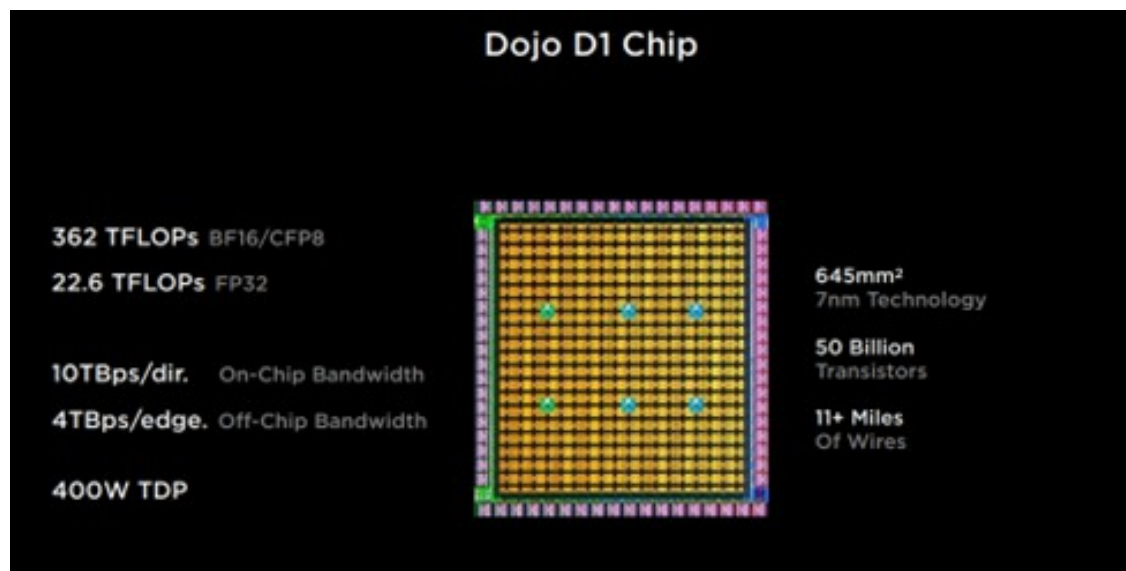
Exhibit 58: A node in more detail



Source: Company data

System wide DDR4 memory works like high-bandwidth bulk storage. The SRAM has a unique list parser that feeds a pair of decoders (changes code into signals), that then feeds the vector register file (contains operands for vector computations and the results), which together transfers information across nodes. Each core is connected to a Network on Chip (NoC) router that communicates with other cores in all four directions. The cores have an integer unit that borrows instructions from commonly-used RISC-V architecture (reduced instruction set computer), with most of the vector math being implemented from scratch. The Dojo Compiler provides the chip with explicit core-to-core data transfer instructions rather than allowing interrupts that would normally push the original task aside and store it into memory. The instructions usually move data to/from an external SRAM of memories of other cores within the Dojo system.

Exhibit 59: Dojo D1 Chip's architecture




Source: Company data

The nodes have a modular design and are arranged into a 2-D, 18x20 array on a Die, with 354 processing D1 cores available. The nodes run at 2GHz, and totals 440MB of SRAM and delivers 362 teraFLOPs at BF16/CFP8 and 22 teraFLOPs at FP32 (there is no FP64 support in the vector units). The D1 chips are placed on a training tile, which connects them together through mesh wiring and is the cooling and power source to the stack. Tesla chose to use a vertical layout of the system because it provides seamless connection to neighboring dies.

Dojo Training Tile: Each Training Tile contains 25 known-good, tested, D1 chips (packaged into a 5x5 2-D array) with each Tile supporting 36 TB/sec of bandwidth via 40 I/O chips. The Tile itself has 11GB of SRAM memory, but the Dojo Interface Processor (DIP) and I/O Core Processor system packs 32GB of shared high bandwidth memory (HBM). The Tiles also support ethernet interfaces beyond the system to hop between Tiles and cores quicker than traditional systems. Utilizing the Tesla Transport Protocol (TTP) and TTP over ethernet (TTPoE), it takes just 4 hops (devices that data travels through) compared to the typical 30 hops to go end-to-end. The TTP is a completely custom protocol used to communicate across the entire accelerator, while the TTPoE enables extending communications across the accelerator.

Exhibit 60: Training Tile power and bandwidth



Source: Company data


The Tiles consume 15kW of power delivery and features electrical, thermal, and mechanical uses. Using vertical power delivery and cooling, the Tiles are water cooled and are designed so multiple Tiles can be interconnected without additional power/cooling design. Each tile uses a TSMC integrated fan-out 'system on wafer', providing shorter interconnects with improved thermal and electrical performance compared to conventional packages. As Dojo was built with a 'no limits philosophy', there is an unlimited amount of Tiles that can be connected to each other, however bandwidth would decrease if the number got too high.

Exhibit 62: 20 Dojo Interface Processors are attached underneath the system tray



Source: Company data

Exhibit 61: The Training Tile's vertical stack



Source: Company data

Dojo Interface Processor. Six Tiles are aggregated into a System Tray, which is integrated with a Host Interface. The Host Interface includes 512 x86 cores providing a Linux-based user environment, and each Tile has one host with 160GB of shared DRAM. A System Tray is needed to realize Tesla's vision of using a single accelerator, which seamlessly connects Tiles together within/between Cabinets at tight spacing, aiding to uniform communication. Trays are densely integrated, at 75mm in height and are built to support 135kgs. An individual system tray can be compared to 3-4 fully loaded high performance racks, with 54 petaFLOPs at BF16/CFP8, and 2000A of power, mechanical, and thermal at 52VDC.


Exhibit 63: The Host Interface connects directly to the Dojo Interface Processors



Source: Company data


Although DRAM isn't used in the D1 chip, it's used in the DIP. Each DIP card has 32GB of RAM, and 20 cards are used across four host servers to reach 640GB of DRAM. The DIP does the processing that a graphics card would normally do, and each Tile has five cards per Tile edge, offering 160GB/sec of bandwidth to host servers and 4.5TB/sec to the tile. DIPs sit on the edge of tile arrays and are hooked into the mesh wiring. Dojo Host Interfaces provide ingest processing and connects to interface processors through PCIe (peripheral component interconnect express) and offers video decoder support for video training. Host systems power the DIPs and perform various system management functions. The Host Interfaces feature 32GB of high-bandwidth DRAM, 900TB/sec TTP bandwidth, 50GB/sec ethernet bandwidth, and 32 GB/sec GEN4 PCIe bandwidth. High-radix z-plane connectivity allows for further shortcuts across the compute plane, leading to 640GB high-bandwidth DRAM, 1TB/sec ethernet bandwidth, and 18TB/sec aggregate bandwidth to Tiles.

Exhibit 64: Communication mechanisms within the Tray



Source: Company data

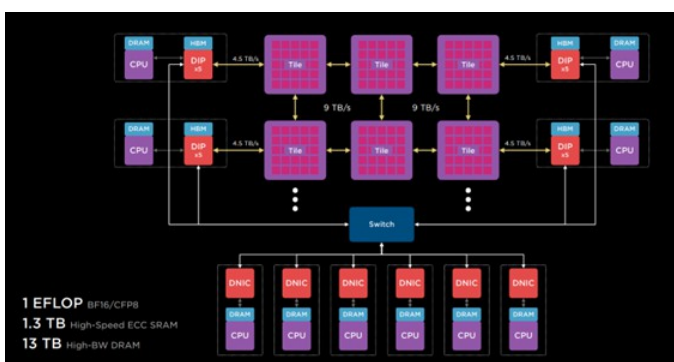
Exhibit 65: Z-plane connectivity allows for 4 hops between data-points vs. 30 hops



Source: Company data


Dojo Cabinet & ExaPOD. Each Cabinet holds two System Trays that are vertically stacked to minimize the distance and communication time between each other. Vertical integration addresses all workload bottlenecks, including data loading, bandwidth, and latency. The Dojo ExaPOD system holds 10 Cabinets, which is equivalent to 120 Tiles, 3,000 D1 chips, and 1,062,000 usable cores, resulting in 1.1 exaFLOPs of AI compute.

Exhibit 66: Complete Dojo system's connectivity with each other



Source: Company data

Exhibit 67: One (out of two) System Trays per Cabinet




Source: Company data

Exhibit 68: Each Cabinet has power sources below/above the System Tray



Source: Company data

Exhibit 69: Complete Dojo ExaPOD system

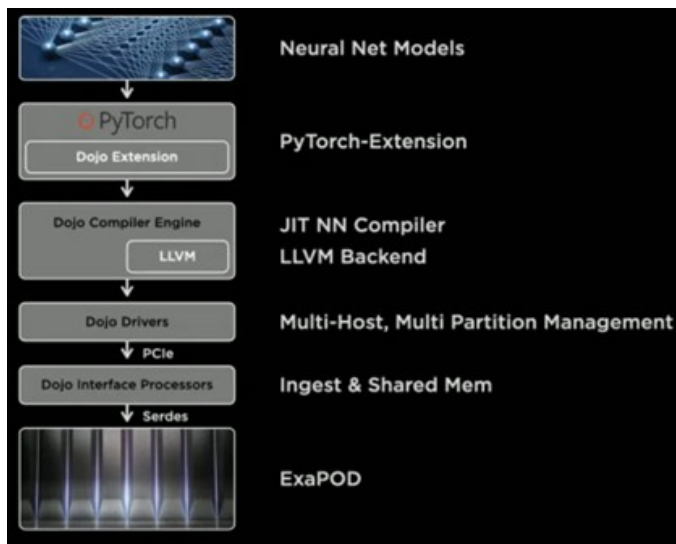


Source: Company data

Dojo's Software

The Dojo hardware interconnects to its software stack through its compiler friendly ISA, which defines how the CPU is controlled by the software, its flexible StateMachine model of computation for ML layers, its fault tolerance systems, and its fire and forget communication protocols. Tesla uses the Dojo PyTorch extension, ensuring the same user-level interfaces that machine learning scientists are used to, rather than using C, C++, or CUDA. From there, the Dojo Compiler Engine acts as the Just-In-Time (JIT) neural net compiler with a LLVM backend, which generates code on-the-fly so it can be used for subsequent execution. This connects to the Dojo Drivers (multi-host, multi-partition management system), which is connected to the DIPs through the PCIe, and is ultimately connected to the ExaPOD.


Exhibit 70: Dojo V1 software stack



Source: Company data

Why use its own custom Compiler? Residual Network (ResNet) models are DL models commonly used for computer vision applications and are easily scalable by replicating the one-accelerator process. For larger vision models, batch sizes that fit in the single accelerator are often smaller than the batch norm surface, so they typically run on multiple accelerators. This goes against the purpose of Dojo, as multiple accelerators can lead to latency bound communication issues, and manually working around it is insufficient for an autonomous machine.

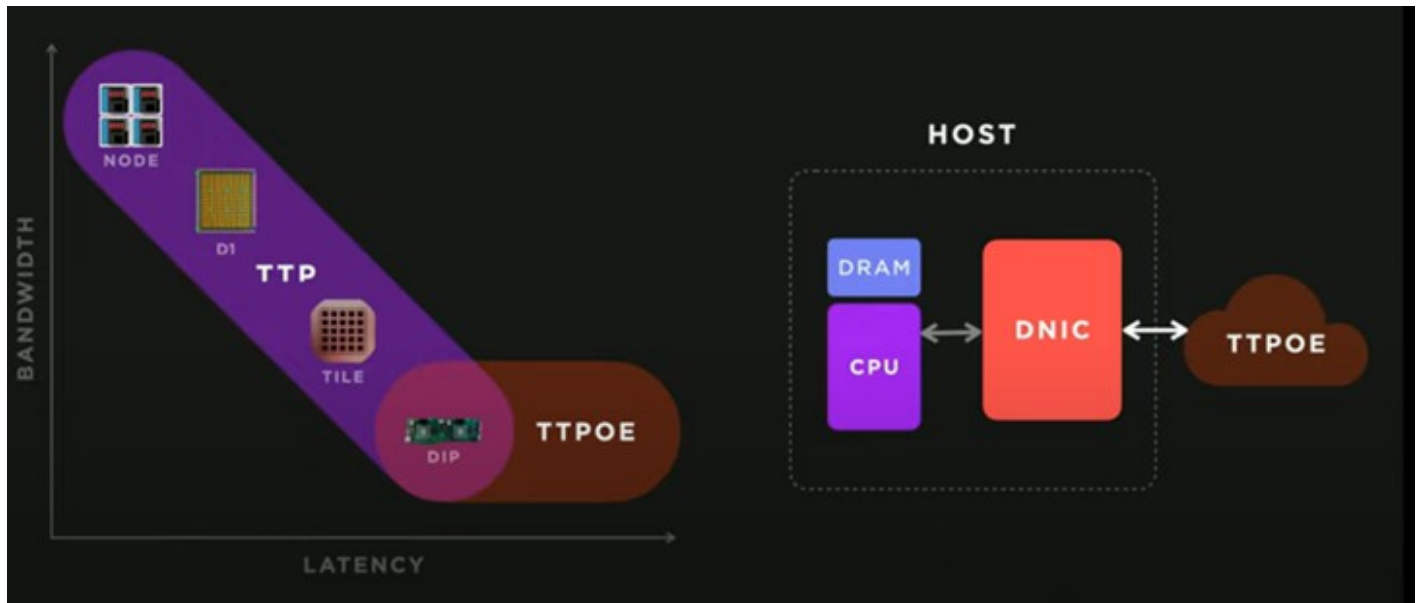
Exhibit 71: Compared against each other for batch norm results, the Dojo D1 was 30x faster than A100s



Source: Company data, Morgan Stanley Research

The Dojo Compiler Engine is custom-built to allow the model to work at high utilization. The compiler's job is to extract utilization from the hardware, and ingest pipelines make sure that data can be fed at throughput high enough for the hardware to never starve. Its high-density integration was built to accelerate compute and latency bound portions like batch norm, and bandwidth bound portions like gradient All-reduce (data-parallel distributed algorithm) and MPI All-gather (gathers and distributes data to all tasks). Using the system's disaggregated structure, a slice of the Dojo mesh can be carved to run any model as long as the portion is large enough to fit the model's batch norm. The system maximizes utilization by extracting parallelism model data, performing JIT replication, overlapping compute and data transfers when needed, and recomputing when successful. The introduction of the data network identification codes (DNIC), which allows for remote direct memory access (DMA) over the TTPoE and enables remote compute for pre/post processing allowed for the hardware occupancy within the Dojo system to skyrocket from 4% to 97% (expected to reach 100% in the near future).

Exhibit 72: The TTP utilizes the TTPoE to connect to the host (memory & CPU)




Source: Company data

Dojo Processing Units, or DPUs, are virtual devices that can be sized according to the application’s needs. The compiler performs mapping onto the DPU, which doesn’t require any human interference and uses aspects like chaining, hybrid partitioning, placement, and memory allocation. Chaining refers to the attachment of each step directly to previous one to speed up reaction time. Hybrid partitioning uses data parallel, model parallel, and graph parallel, which allows a table or index to be subdivided into smaller pieces, called partitions. Placement combines pieces of several chips into one uniform picture on one chip, and memory allocation utilizes distributed tensors (algebraic object), recomputes, and spill/fills based on the input.

How Does The Software & Hardware Tie Together?


There are aspects beyond compute that share the goal of creating an efficient and scalable system. Communication reaches speeds of terabytes per second, memory has the capacity for GBs/TBs, disaggregation moves with workloads, networking topologies include TTP/TTPoEs to increase efficiency speed, and it’s all done with the implication to seamlessly scale the system. The supercomputer for ML features new integration, which enables high-bandwidth and performance, uniform high-bandwidth, which enables full exploitation of parallelism by software, and vertically integrated I/O, and addresses all workload bottlenecks including data loading.

Exhibit 73: The role of each system in the disaggregated, scalable system

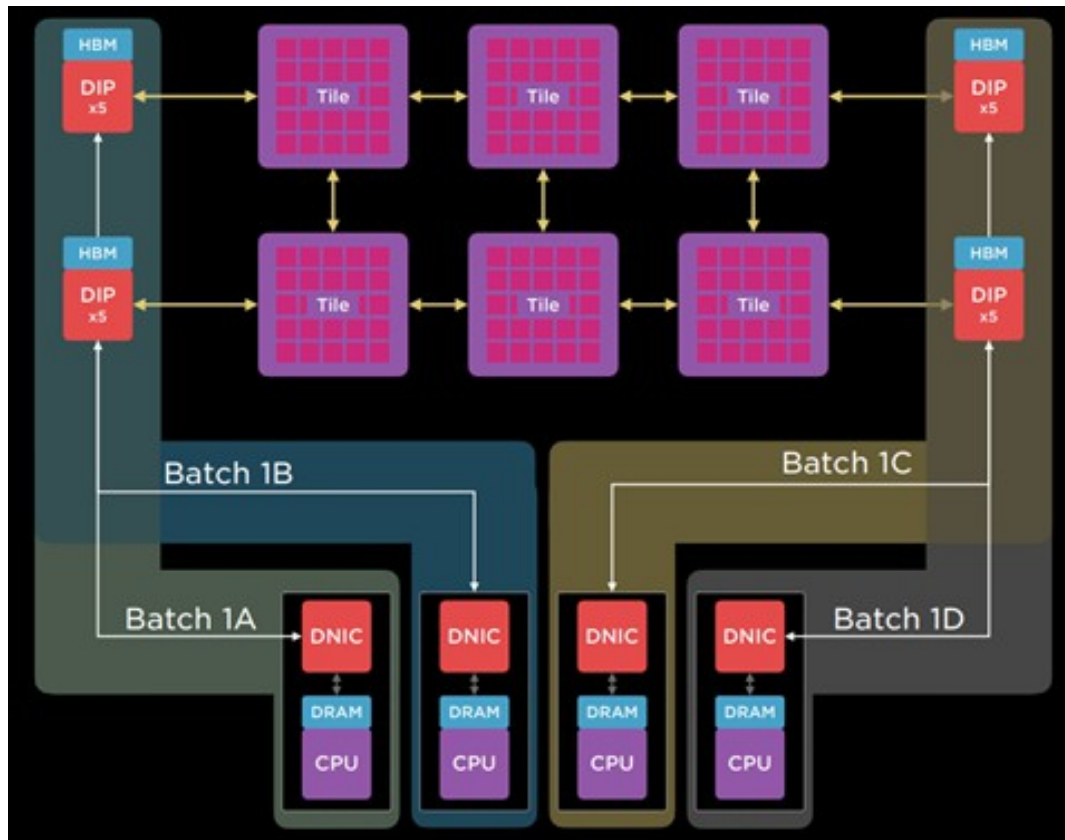


Source: Company data

Exhibit 74: Dojo Training Matrix: how each Tile works together to fix and solve



Source: Company data

Exhibit 75: Disaggregation model of data tier loading through a system

Source: Company data

Dojo Timeline

What has happened since 2019? While the public hasn't been regularly updated on Dojo's progress, what we do know is primarily from recent Tesla AI Days, management's tweets, and the 2Q23 earnings call.

- April 2019:** During Tesla's autonomy day, Elon Musk [mentions](#) that Tesla was working on a "super powerful training computer". He says that the goal of Dojo will be to perform unsupervised training at a large-scale level using vast amounts of video data.
- August 2019:** Musk [tweets](#) "will Dojo be the difference", alluding towards the potential significance of the supercomputer, even at an early stage.
- August 2020:** Musk makes a series of tweets about Dojo, with one claiming V1.0 to be "about a year away".
- September 2020:** More Musk tweets are made, mentioning that Dojo uses Tesla-made chips (rather than a GPU cluster) and a computer architecture that is optimized for neural net training. He can be quoted saying "I think it will be the best in the world".
- August 2021:** Tesla hosts AI Day and unveils Dojo's hardware and software in greater detail.
- August 2022:** Tesla [releases](#) two presentations at the Hot Chips 34 conference ahead of their planned AI Day that feature high-level information about Dojo. Senior Hardware Director Ganesh Venkataramanan introduces that there is a large gap between what's needed to accelerate AI and what is available, and that Dojo can fill the void.
- Musk replies to a tweet about Dojo, saying that Tesla won't need to buy as many incremental GPUs next year.
- September 2022:** Tesla hosts AI Day and provides further updates on Dojo's architecture and progress.
- June 2023:** A twitter account from Tesla is created in May 2023 called "Tesla AI", and a series of tweets are made in June with videos of the autonomous technology in action, prompted by their neural networks. Tweets mention that these models will learn from a huge set of data, and that it will be trained on "enormous amounts of compute" using Dojo in the future. A chart is released that projects that Dojo will start producing (one ExaPOD will be completely brought

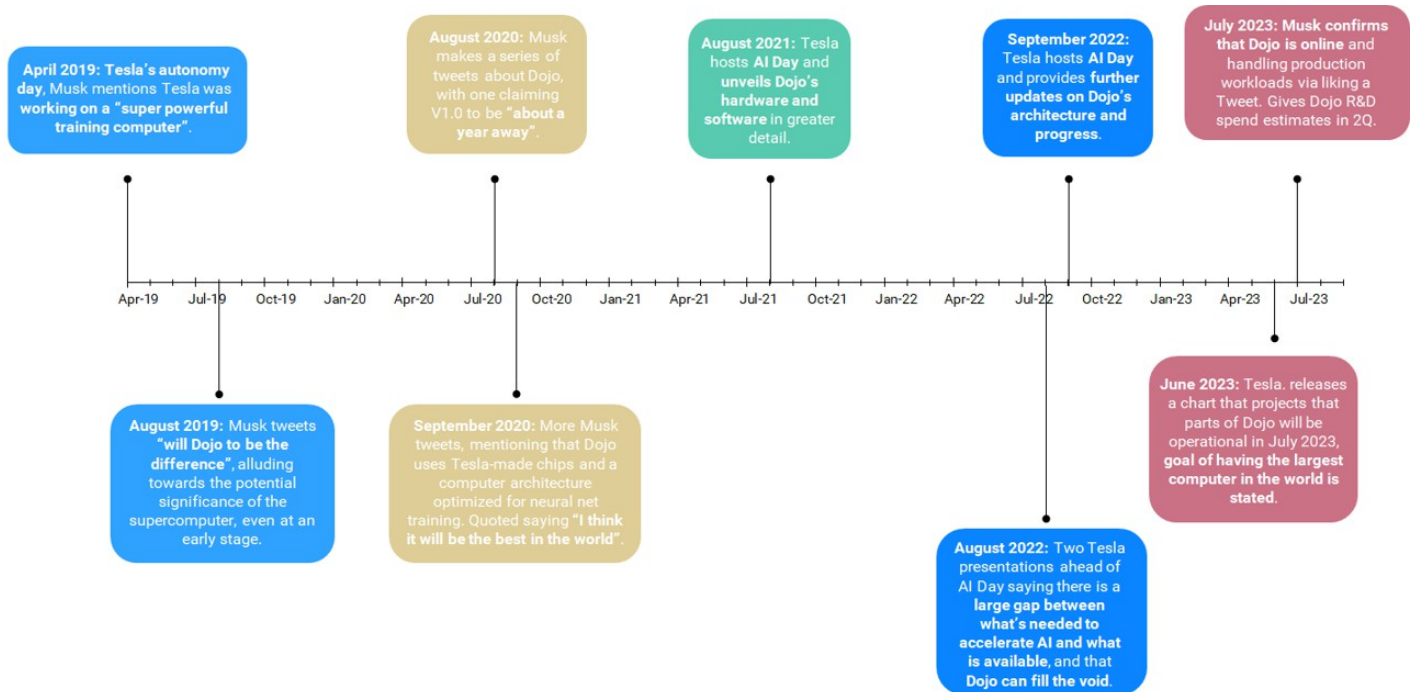
online) in July 2023, Tesla will have a top 5 compute power in the world by January 2024, and Tesla will reach 100 exaFLOPs of compute by October 2024.

Musk replies to a tweet on a post from the “Tesla AI” account stating that Dojo has been “online & running useful tasks for a few months”.

- **July 2023:** Musk confirms that Dojo is online and handling production workloads via liking a Tweet.

Tesla reports 2Q earnings and provides updates on Dojo in the earning call, including R&D/CapEx costs over the next year of \$1bn+.

Exhibit 76: Tesla Dojo Timeline

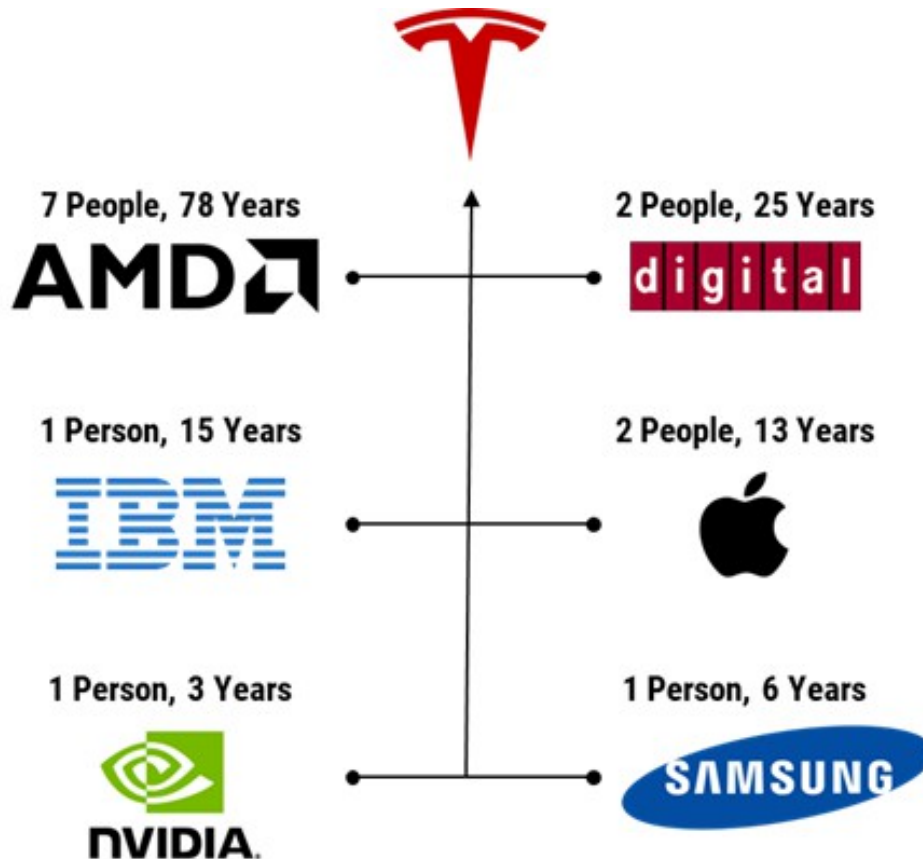


Source: Morgan Stanley Research

Meet The Team

The team hired to run the Dojo project shows that Tesla has been committing themselves to this supercomputer. Tesla brought in several experienced professionals that have years of expertise in each leg of the chip (hardware & software). Looking at some of the names on the list, many members were directly involved or responsible for the advancement of industry-changing technology.

Exhibit 77: The team has a combined 250+ years of relevant hardware & software experience between 12 of the notable engineers working on Dojo



Source: Morgan Stanley Research

Some key members of the tenured Dojo team include:

- **Ganesh Venkataramanan** is the Senior Director of Autopilot Hardware & “Head of Dojo” at Tesla. He has been at Tesla since 2016, joining as the director of autopilot hardware and since 2018, has overseen the Dojo, Silicon, Systems, and Firmware/Software departments. Ganesh worked for AMD’s CPU group for 14+ years, overseeing 200+ engineers as a Senior Director of Design Engineering. He was a member of the team that created many generations of x86 and ARM cores. Notable achievements include being involved in creating the first x86-84 chip, the first Dual-Core x86, and the Zen Core. In addition to his time at AMD, he previously spent 3.5 years at Analog Devices.
- **Peter Bannon** is the VP of Low Voltage & Silicon Engineering and

leads the team that created the FSD system that currently sits inside Tesla vehicles. Prior to his arrival to Tesla in 2016, he was the lead architect on the first 32b ARM CPU used in the iPhone 5 and built the team that created the first 64b RAM processor in the iPhone 5s. He has been designing computer systems for almost four decades through 7 years at Apple, 1 year at Intel, 4 years at PA Semi, and 16 years at Digital Equipment Corporation.

- **Bill Chang** is a Principal System Engineer for Dojo & Autopilot at Tesla and has been the Chief Engineer responsible for system architecture and design of AI training systems since 2020. He has 20+ years of experience in technology and product engineering through 6 years at Apple and 15 years at IBM. At Apple, he was a Senior Engineering Program Manager where he managed SoC architecture and definition.

- **Ashok Elluswamy** is the Director of Autopilot Software at Tesla and has been working there since 2014. He is responsible for leading the autonomy software team to create large-scale automatic ground truth pipelines geared to train neural networks with massive amount of data. Prior to Tesla, he had stints at VW Electronic Research Lab and WABCO Vehicle Control Systems.
- **Rajiv Kurian** is a Principal Engineer at Tesla and has been responsible for “All things Dojo” since 2020. He previously worked at Waymo for 2 years on hardware architecture and the acceleration of autonomous driving workloads, and hardware and software design for motion planning. Prior to joining Tesla, Rajiv has compiled 10 years of engineering experience.
- **Emil Talpes** is a Hardware Engineer at Tesla and has been with the company since 2016. Emil previously worked at AMD for 9+ years as a Senior Member of Technical Staff where he was a member of the architecture team for the K12 ARM core, and the Bulldozer, Piledriver, Steamroller x86 cores. On each team, he was responsible for the decode unit and out-of-order execution engine.
- **Debjit Das Sarma** is a Principal Autopilot Hardware Architect at Tesla where he is responsible for the architecture of FSD, neural network inference, and training chips. Prior to working at Tesla starting in 2016, he was a Fellow and Chief Architect at AMD during his 15-year stint there. He was responsible for the architecture and design of several high-performance computer processors including the Ryzen processor, which acts as the brain of many laptops and desktops used today.
- **Doug Williams** is an Autopilot Hardware Engineer at Tesla where he is involved in computer systems, architecture, NoC, and serial communication protocols since 2017. He previously spent 12+ years at AMD as a PMTS Design Engineer, where he worked on computer architecture, brand prediction, instruction fetching, security and virtualization, and cache hierarchies.
- **Eric Quinnell** is a Senior Staff Engineer at Tesla that works primarily on Dojo/AI computing. Prior to joining Tesla in 2020, he spent 7 years at AMD, 6 at Samsung R&D Center, and 1 at Arm. He worked on the Bobcat and Jaguar x86 systems at AMD as a floating-point architect.
- **Bill McGee** is a Principal Hardware Engineer on the Dojo Software Team at Tesla and has worked there since 2016. He spent 18 years working at AMD as a Senior Fellow and Chief Engineer Server SoC, where he led several problem-solver teams on most AMD CPU designs (K8, Bulldozer, Kaveri, Excavator). He also has experience as a Principal Designer at Digital Equipment Corporation, where he spent 9 years.
- **Anton Lawrenda** is a Staff Autopilot Hardware Engineer at Tesla, where he’s been focusing on Dojo Silicon Development since 2020. Prior to joining in 2018, he spent 3.5 years at AMD as a Senior ASIC/Layout Design Engineer.

Competitive Landscape - GPUs & Select Custom Chips

The Dojo chip is comparable to other ML research and production chips like the Google TPU v4, the NVIDIA A100 and H100, Cerebras WSE-2, and Graphcore MK2 IPU. Each chip serves a similar purpose in the system, with the A100 and H100 being used to allow organizations to build large-scale machine learning infrastructure, as Tesla has done with their 14,000 GPUs (and 7,360 A100 count supercomputer). With many general purpose GPUs, the larger you scale, the less cost effective and compute efficient the supercomputer gets, which could be one major reason why Tesla decided to work on their own chip. NVIDIA’s chips are designed to be general compute chips (GP GPUs) that can easily and efficiently be implemented into a company’s own systems, while the TPU is Google’s custom supercomputer platform. Since Dojo is tailored specifically for Tesla’s predictable workloads, operations, and computational systems, the product is advantageous compared to outsourced chips.

- **NVIDIA A100:** Powered by the ampere, the A100 is optimized for data science and AI workflows in applications across the ML training environment. The use cases for the A100 range from AI model development and inference, video/image decoding, natural language processing, and augmented fault and error detection. At its release in 2020, the A100 was the most powerful computer accelerator in the IT industry and had 11x higher throughput than the 2019 NVIDIA V100. The 7nm chip ranks higher than the Google TPU v3/v4 and Graphcore MK2 IPU in both compute power (312 teraFLOPs for TF32 and 624 for FP16) and TDP (400W) and supports up to 80GB of on-chip memory with 54 billion transistors. OpenAI trained and ran ChatGPT on A100s and will be transferring to H100s on its Azure supercomputer. Eight A100s are paired together with its supporting system to form a DGX A100, which costs ~\$200,000.

Exhibit 78: Specs: Dojo D1 vs NVIDIA A100

	Tesla D1	NVIDIA A100
Solution	Data center training	Data center training
Type	ASIC	GPU
Year of release	2019	2020
Process Node	TSMC 7nm	TSMC 7nm
Die size (mm ²)	645	826
Transistor Count	50	54.2
TDP (watt)	400	400

Source: Company Data, Morgan Stanley Research

- **NVIDIA H100:** Released in Q3 2022, the next-gen model of the A100, the H100, is powered by the NVIDIA hopper GPU and is optimized for developing, training, and deploying generative AI, large language models (LLM), and recommender systems. Its use cases

are in deep learning, high-performance computing, AI inference, computer vision, and computational biology. It is able to [achieve](#) the performance of 2.2x-3.3x A100s as the 4nm chip ranks above than competing chips at 3,026 teraFLOPs and up to 700W in TDP, with 80GB of on-chip memory and 80 billion transistors. Meant to be an upgrade from their A100 offering, NVIDIA's H100 promises up to 9x faster AI training and up to 30x faster AI inference. It's designed to work seamlessly with NVIDIA's NVLink interconnect technology, which allows for high-bandwidth interconnection between GPUs. Eight H100s are paired together with its supporting system to form a DGX H100, which costs ~\$480,000.

- **Google TPU v4:** Released in 2020, the purpose-built TPUs are optimized for training large and complex deep learning models that feature many matrix calculations. These systems are used for ML research and production workloads across language models, recommender systems, and generative AI. The TPU was the first supercomputer chip to deploy a reconfigurable OCS, allowing it to dynamically reconfigure their interconnect topology to improve scale, availability, utilization, modularity, deployment, security, power, and performance, leading to a cheaper and lower power-consuming system. The 7nm TPU v4 chip can compute up to 275 teraFLOPs, has 4,096 cores per pod, a TDP of 275W, and 22 billion transistors per chip.
- **Cerebras WSE-2:** Announced in April 2021, the WSE-2 chip features a first-in-its-kind structure, with the chip being one massive chip rather than several chips paired together. The chip is 'designed for AI' since it's independently programmable and optimized for tensor-based operations that underpin neural network training and inference. It's approximated that the CS-2 system that the WSE-2 chips sit in costs ~\$7 million, a large capital cost for a large chip. The unique design eliminates the communication slowdown and inefficiencies of connecting hundreds of small devices via wires and cables, like how the Dojo training tile functions. As the world's largest computer chip at 46,225mm² (56x larger than the next largest GPU – the A100) it's not as versatile in terms of being implemented into any type of supercomputer system. The 7nm chip boasts 850,000+ cores, 15kW of TDP, 2.6 trillion transistors, 40GB of on-chip memory, and reaches compute power of 503 teraFLOPs.
- **Graphcore MK2 IPU:** Released in 2021, the MK2 chip utilizes an IPU, which is a programmable network device that intelligently manages system-level infrastructure resources by securely accelerating functions in a data center. The IPU is a unique kind of parallel processor and is used to accelerate machine intelligence. The 7nm MK2 IPU has a compute power of 250 teraFLOPs, has 1,272 cores, a TDP of 300W, 900 MiB of on-chip memory, and 59 billion transistors.

Appendix - Glossary

64-bit operating system: operating system designed to work in a computer that processes 64 bits at a time

Accelerator: computer hardware that specifically handles AI requirements in order to speed up work

AI computing: math-intensive process of calculating machine learning algorithms, typically using accelerated systems and software

Ampere: NVIDIA's microarchitecture used for the A100 chip

Artificial Intelligence (AI): the simulation of human intelligence processes by machines (computer systems)

ASIC: application-specific integrated circuit – a chip customized for a particular use (in this case, an AI ASIC is used for AI applications)

Auto-labeling: automatically converts video, IMU, GPS, and odometry data into a machine-identifiable language within an autonomous system. Tesla claims it no longer needs to explicitly 'label' data as the system will take large amounts of raw video data as an input with driving decisions (steering and pedal angles) as the output.

Bandwidth: the capacity at which a network can transmit data

Batch normalization (norm): technique to standardize the inputs to a network

BF16: custom 16-bit floating point format that is a cross between FP16 and FP32 – like FP16, it has a smaller memory footprint, enabling faster training and inference time than FP32

CUDA: parallel computing platform and programming model created by NVIDIA

Cabinet: rack that holds two system trays - each cabinet has 12 training tiles, 300 D1 dies, 106,200 cores

Cerebras WSE-2: Cerebras' large, single wafer, chip solution for deep learning computer systems

CFP8: floating point 8 – a natural progression for accelerating deep learning beyond the 16-bit format

Chaining: style of programming that attaches each step directly to the previous one to remove the intermediate step between it (via invoking multiple method calls occurring sequentially)

Cooling system: removes excess heat, maintains operating temperature, and brings the engine to the correct temperature that it works most efficiently in

Core: small CPU or processor built into a larger CPU/chip

D1 Die/Chip: Tesla's chip designed specifically for artificial intelligence machine learning training

Data parallel model: tasks are assigned to processes that are parallel to each other and each task performs similar operations on different data

Data parallelism: a consequence of single operations that is applied on multiple data items

Data ingestion: process of obtaining and importing data for immediate use or storage in a database

DDR4 memory: double data rate fourth generation memory – memory that achieves higher speed and efficiency due to increased transfer rates and decreased voltage

Decoder: circuit that converts signals from one form to another

Deep learning (DL): method of AI that teaches computers to process data similarly to a human brain

Disaggregation: decoupling data center resources of memory, compute, storage so each can be scaled and provisioned independently (ratios move with workloads) – components are divided into subsystems

DMA: direct memory access – the process of transferring data without the involvement of the processor itself; is often used for transferring data to/from input/output devices

DNIC: data network identification code – four-digit number designed to provide identification of individual public data networks, often intended to identify and permit automated switching of data traffic to particular networks

Dojo Compiler Engine: Tesla's custom-built program that helps translate a programming language's source code into a machine language code

Dojo Drivers: conducts multi-host, multi partition management

Dojo Host Interface: provides power to the Dojo Interface Processor, ingest processing, connects to interface processors through PCIe, and offers video decoder support for video training

Dojo Interface Processor (DIP): provides both connectivity to the outside world and shared memory that helps feed data to the training tiles

Dojo Processing Unit (DPU): a virtual device that the Compiler performs mapping onto that can be sized according to the application's needs

DRAM: dynamic random-access memory – memory used for data/program code needed by a computer processor; is commonly used where low-cost and high-capacity memory is required

Edge computing: distributed computing paradigm that brings computation and data storage closer to the sources of data – saves bandwidth and improves response time

Ethernet: system for connecting several computer systems to form a local area network

exaFLOP: performance measure for a supercomputer that can calculate 10¹⁸ floating point operations/sec

ExaPOD: the complete Dojo system that holds 10 cabinets - 4 system trays, 120 training tiles, 3000 D1 dies, 1,062,000 cores, equating 1.1 exaFLOPs

Fanout system/wafer: integrated circuit packaging technology that provides a smaller package footprint and improved thermal/electrical performance, allowing a high number of contacts without increasing the die size

Floating point format: computer number format represented by wide dynamic range of numeric values

FP16: half-precision floating point format that occupies 16-bits in computer memory and has a smaller memory footprint, enabling faster training and inference time than FP32

FP32: full-precision floating point format that occupies 32-bits in computer memory

FP64: double-precision floating point format that occupies 64-bits in computer memory

Google TPUv4: Google's supercomputer chip solution for deep learning computers

GP GPU: general-purpose graphics processing unit – handles computation for computer graphics that can be utilized in several different use cases

Gradient All-reduce: popular synchronous data-parallel distributed algorithm often used in PyTorch DistributedDataParallel

Graph parallel: data-parallel computation applied to graph data

Graphcore MK2 IPU: Graphcore's processor specifically designed for machine learning

High-bandwidth memory (HBM): high speed computer memory

High-radix z-plane connectivity: high-radix routers allow shortcuts across the compute plane by lowering networking diameter while providing high bandwidth and path diversity

Hopper: NVIDIA's microarchitecture used for the H100 chip

Hops: number of devices (usually routers) that a piece of data travels through

Hybrid partitioning: enables partitions to reside both in database files and external sources

Hyperscale: ability of an architecture to scale appropriately as increased demand is added to the system; hyperscalers refer to large cloud service providers

I/O Compute (input/output): describes any operation, program, device that transfers data in/out of a computer

I/O Processor: specialized processor that loads and stores data into memory along with the execution of I/O instructions

IMU: inertial measurement unit – describes a collection of measurement tools through accelerometers, gyroscopes, magnetometers

Inference: process of taking a model and deploying it onto a device which will then process incoming data to look for and identify what it's trained to recognize

Ingest pipeline: allows a system to perform common transformations on data before indexing (removing fields, extracting values from text)

Intelligence Processing Unit (IPU): processor specifically designed for machine learning and AI applications

ISA: instruction set architecture – part of the abstract model of a computer that defines how the CPU is controlled by the software

Just-In-Time (JIT) Compiler: executes computer code that involves compilation during the execution of a program at run time (in real time) rather than before execution

Known-good die (KGD): a tested die that either met or exceeded quality, reliability, functional specifications

Language learning models (LLM): type of machine learning model trained to conduct a probability distribution over words

Latency: the time data takes to transfer across a network (delay in network communication)

Lateral power delivery (LPD): modules are placed adjacent to the processor

Linux-based environment: open-source/community developed operating system that is supported on most major computing platforms (x86, ARM, SPARC)

LLVM backend: target-independent code generator that can create output for several types of CPUs

Load speed: measure of the amount of computation work that a computer system performs

Local memory: portion of memory that is designated exclusively to an individual compute unit

Machine learning (ML): subset of AI that performs specific tasks by learning from data & making predictions

Mapping: uses the speed and versatility of computer graphics to display spatial data

Memory allocation: reserves virtual/physical compute space for a specific purpose

Mixed precision: the combined use of different computational methods (example: FP16 & FP32)

MPI All-gather: gathers data from all tasks and distributes the combined data to all tasks

Natural language processing (NLP): ML technology that gives computers the ability to interpret, manipulate, and comprehend humans

Network on Chip (NoC): network-based communications subsystem on a microchip that typically sits between modules in a system on a chip (SoC)

Neural network: method of AI that teaches computers to process data in a similar way to human brain

Node: connection point among network devices that can send/receive data from one endpoint to another

NVIDIA A100: NVIDIA's GPU that allows organizations to build large-scale ML infrastructure

NVIDIA H100: NVIDIA's GPU geared towards large amounts of data, making it a good choice for real-time inference applications

Occupancy network: uses 3D mapping to detect obstacle detection with usages in collision avoidance

Optical circuit switch (OCS): optical networking technology that dynamically reconfigures the interconnect topology to improve scale, availability, utilization, modularity, security, power, and performance

Odometry: uses data from motion sensors to estimate change in position over time

Optimus: Tesla's general-purpose humanoid bot

Parser: program (part of the compiler) that takes code from a pre-processor and breaks it into smaller pieces and analyzes it so other software can understand it

Packet: small segment of a larger message

Parallel processing: computing method of running two or more processors (CPUs) to handle separate parts of a task

Partition: division of a disk into logical sections that is treated as a separate operating unit & file systems

PCIe: standardized interface that connects a system to one or more high-speed components

PyTorch (Dojo Extension): framework for building deep learning models (written in Python)

Register: storage space for units of memory used to transfer data by the CPU for data processing

ResNet: residual network – deep learning model used for computer vision applications

Reticle: around 800mm², which is the common maximum size for chips

RISC-V: commonly used ISA that software can be ported, hardware can be developed, and processors can be built to support it

Recommender system: class of machine learning that uses data to help predict, narrow down, and find what people are looking for among exponentially growing number of options

SerDes: serializer/deserializer - facilitates data to compensate for limited input/output, converts parallel data to serial (one bit at a time) data

Shared cache: hardware/software component that stores data for future requests that can be accessed by multiple cores

Spill: the movement of some variables to/from RAM when there's not enough registers to hold them all

SRAM: type of RAM that retains data bits in memory as long as power is being supplied

StateMachine: behavioral model that allows dynamic flow to states depending on values from previous states of user input

Store speed: minimum sustained speed that memory can be stored

Supercomputer: computer with a high level of performance, commonly measured in floating point operations per second

Superscalar: method of parallel computing used in many processors where the CPU manages multiple instruction pipelines to execute several instructions concurrently during a clock cycle

System Tray: the tray that holds 6 training tiles – each tray has 150 D1 dies, 53,100 cores

System on a Chip (SoC): type of integrated circuit design that combines many or all high-level function elements of an electronic device onto a single chip (rather than using separate components mounted to a motherboard)

Tensors: fundamental data structure used by machine and deep learning algorithms

Tensor Processing Unit (TPU): Google's custom-developed application specific integrated circuit used to accelerate machine learning workloads

Teraflop: measurement of 10¹² floating point operations per second

Tesla Transport Protocol (TTP): Tesla's interconnect over PCIe that is a completely custom protocol used to communicate across the entire accelerator

Tesla Transport Protocol over Ethernet (TTPoE): enables extending communications over standard ethernet, provides native hardware support with little overhead Thermal Design Power (TDP)

TF32: computing format used by NVIDIA A100s with same numerical range as FP32 but has 10 bits instead of 23 (similar to FP16)

Thermal Design Power (TDP): the power consumption under maximum theoretical load in watts

Throughput: the volume of work or information flowing through a system

Total Cost of Ownership (TCO): comprehensive assessment of information technology or other costs across enterprise boundaries over time

Trained logic: use of computers to establish facts – the output of learning systems

Training: learning a new capability from existing data (fed by a developer in a data center typically)

Training Tile: the tile that holds 25 D1 chips and has thermal, electrical, mechanical functions – each tile has 25 D1 dies and 8,850 cores

Transistor: semiconductor device used to amplify or switch electrical signals and power

Vector math: one-dimensional array that typically stores numbers

Vertical power: delivers high current at low processor core voltage and minimizes space between systems resulting in faster communication and data transfers

Virtual memory: storage allocation that allows secondary memory to be addressed as if it were part of the main memory

VRM: voltage regulator module – converter that provides the microprocessor and chipset the appropriate supply voltage

x86 core: type of ISA for computer processors

Zettabyte: a measurement of digital storage capacity equal to 1021 bytes, or 1 billion terabytes

Disclosure Section

The information and opinions in Morgan Stanley Research were prepared by Morgan Stanley & Co. LLC, and/or Morgan Stanley C.T.V.M. S.A., and/or Morgan Stanley Mexico, Casa de Bolsa, S.A. de C.V., and/or Morgan Stanley Canada Limited. As used in this disclosure section, "Morgan Stanley" includes Morgan Stanley & Co. LLC, Morgan Stanley C.T.V.M. S.A., Morgan Stanley Mexico, Casa de Bolsa, S.A. de C.V., Morgan Stanley Canada Limited and their affiliates as necessary.

For important disclosures, stock price charts and equity rating histories regarding companies that are the subject of this report, please see the Morgan Stanley Research Disclosure Website at www.morganstanley.com/researchdisclosures, or contact your investment representative or Morgan Stanley Research at 1585 Broadway, (Attention: Research Management), New York, NY, 10036 USA.

For valuation methodology and risks associated with any recommendation, rating or price target referenced in this research report, please contact the Client Support Team as follows: US/Canada +1 800 303-2495; Hong Kong +852 2848-5999; Latin America +1 718 754-5444 (U.S.); London +44 (0)20-7425-8169; Singapore +65 6834-6860; Sydney +61 (0)2-9770-1505; Tokyo +81 (0)3-6836-9000. Alternatively you may contact your investment representative or Morgan Stanley Research at 1585 Broadway, (Attention: Research Management), New York, NY 10036 USA.

Analyst Certification

The following analysts hereby certify that their views about the companies and their securities discussed in this report are accurately expressed and that they have not received and will not receive direct or indirect compensation in exchange for expressing specific recommendations or views in this report: Charlie Chan; Adam Jonas, CFA; Joseph Moore; Edward Stanley.

Global Research Conflict Management Policy

Morgan Stanley Research has been published in accordance with our conflict management policy, which is available at www.morganstanley.com/institutional/research/conflictolicies. A Portuguese version of the policy can be found at www.morganstanley.com.br

Important Regulatory Disclosures on Subject Companies

The analyst or strategist (or a household member) identified below owns the following securities (or related derivatives): Daniela M Haigian - Apple, Inc.(common or preferred stock), Meta Platforms Inc(common or preferred stock); Adam Jonas, CFA - Alphabet Inc.(common or preferred stock), Amazon.com Inc(common or preferred stock); Matias Ovrum - Amazon.com Inc(common or preferred stock); Edward Stanley - NVIDIA Corp.(common or preferred stock).

As of August 31, 2023, Morgan Stanley beneficially owned 1% or more of a class of common equity securities of the following companies covered in Morgan Stanley Research: Alphabet Inc., Amazon.com Inc, American Axle & Manufacturing Holdings Inc, Apple, Inc., Aptiv Plc, Asbury Automotive Group Inc, AutoNation Inc., BorgWarner Inc., Carvana Co, Ferrari NV, Fisker Inc, Ford Motor Company, FREYR Battery SA, General Motors Company, Group 1 Automotive, Inc, Harley-Davidson Inc, Lithia Motors Inc., Lucid Group Inc, Meta Platforms Inc, Netflix Inc, NVIDIA Corp., Quantumscap Corp, Rivian Automotive, Inc., Tesla Inc, Visteon Corporation.

Within the last 12 months, Morgan Stanley managed or co-managed a public offering (or 144A offering) of securities of Amazon.com Inc, Apple, Inc., Ford Motor Company, FREYR Battery SA, General Motors Company, Meta Platforms Inc, Mobileye Global Inc, Quantumscap Corp, Rivian Automotive, Inc., Tesla Inc.

Within the last 12 months, Morgan Stanley has received compensation for investment banking services from Alphabet Inc., Amazon.com Inc, Apple, Inc., Aptiv Plc, Avis Budget Group Inc, Ford Motor Company, General Motors Company, Meta Platforms Inc, Mobileye Global Inc, Quantumscap Corp, Rivian Automotive, Inc., Tesla Inc.

In the next 3 months, Morgan Stanley expects to receive or intends to seek compensation for investment banking services from Alphabet Inc., Amazon.com Inc, American Axle & Manufacturing Holdings Inc, Apple, Inc., Aptiv Plc, AutoNation Inc., Avis Budget Group Inc, BorgWarner Inc., Carvana Co, Ferrari NV, Ford Motor Company, General Motors Company, Hertz Global Holdings Inc, Lear Corporation, Lithia Motors Inc., Lucid Group Inc, Magna International Inc., Meta Platforms Inc, Mobileye Global Inc, Netflix Inc, NVIDIA Corp., Quantumscap Corp, Rivian Automotive, Inc., Tesla Inc, Visteon Corporation.

Within the last 12 months, Morgan Stanley has received compensation for products and services other than investment banking services from Alphabet Inc., Amazon.com Inc, American Axle & Manufacturing Holdings Inc, Apple, Inc., Aptiv Plc, Avis Budget Group Inc, BorgWarner Inc., Ferrari NV, Ford Motor Company, General Motors Company, Harley-Davidson Inc, Hertz Global Holdings Inc, Lear Corporation, Meta Platforms Inc, Netflix Inc, NVIDIA Corp., Rivian Automotive, Inc., Tesla Inc, Visteon Corporation.

Within the last 12 months, Morgan Stanley has provided or is providing investment banking services to, or has an investment banking client relationship with, the following company: Alphabet Inc., Amazon.com Inc, American Axle & Manufacturing Holdings Inc, Apple, Inc., Aptiv Plc, AutoNation Inc., Avis Budget Group Inc, BorgWarner Inc., Carvana Co, Ferrari NV, Ford Motor Company, FREYR Battery SA, General Motors Company, Hertz Global Holdings Inc, Lear Corporation, Lithia Motors Inc., Lucid Group Inc, Magna International Inc., Meta Platforms Inc, Mobileye Global Inc, Netflix Inc, NVIDIA Corp., Quantumscap Corp, Rivian Automotive, Inc., Tesla Inc, Visteon Corporation.

Within the last 12 months, Morgan Stanley has either provided or is providing non-investment banking, securities-related services to and/or in the past has entered into an agreement to provide services or has a client relationship with the following company: Alphabet Inc., Amazon.com Inc, American Axle & Manufacturing Holdings Inc, Apple, Inc., Aptiv Plc, Avis Budget Group Inc, BorgWarner Inc., Ferrari NV, Fisker Inc, Ford Motor Company, General Motors Company, Harley-Davidson Inc, Hertz Global Holdings Inc, Lear Corporation, Meta Platforms Inc, Netflix Inc, NVIDIA Corp., Rivian Automotive, Inc., Tesla Inc, Visteon Corporation.

An employee, director or consultant of Morgan Stanley is a director of Alphabet Inc., General Motors Company. This person is not a research analyst or a member of a research analyst's household.

Morgan Stanley & Co. LLC makes a market in the securities of Adient PLC, Alphabet Inc., Amazon.com Inc, American Axle & Manufacturing Holdings Inc, Apple, Inc., Aptiv Plc, Asbury Automotive Group Inc, AutoNation Inc., Avis Budget Group Inc, BorgWarner Inc., Carmax Inc, Ford Motor Company, General Motors Company, Group 1 Automotive, Inc, Harley-Davidson Inc, Lear Corporation, Lithia Motors Inc., Magna International Inc., Meta Platforms Inc, Mobileye Global Inc, Netflix Inc, NVIDIA Corp., Penske Automotive Group, Inc, Rivian Automotive, Inc., Sonic Automotive Inc, Tesla Inc, Visteon Corporation.

The equity research analysts or strategists principally responsible for the preparation of Morgan Stanley Research have received compensation based upon various factors, including quality of research, investor client feedback, stock picking, competitive factors, firm revenues and overall investment banking revenues. Equity Research analysts' or strategists' compensation is not linked to investment banking or capital markets transactions performed by Morgan Stanley or the profitability or revenues of particular trading desks.

Morgan Stanley and its affiliates do business that relates to companies/instruments covered in Morgan Stanley Research, including market making, providing liquidity, fund management, commercial banking, extension of credit, investment services and investment banking. Morgan Stanley sells to and buys from customers the securities/instruments of companies covered in Morgan Stanley Research on a principal basis. Morgan Stanley may have a position in the debt of the Company or instruments discussed in this report. Morgan Stanley trades or may trade as principal in the debt securities (or in related derivatives) that are the subject of the debt research report.

Certain disclosures listed above are also for compliance with applicable regulations in non-US jurisdictions.

STOCK RATINGS

Morgan Stanley uses a relative rating system using terms such as Overweight, Equal-weight, Not-Rated or Underweight (see definitions below). Morgan Stanley does not assign ratings of Buy, Hold or Sell to the stocks we cover. Overweight, Equal-weight, Not-Rated and Underweight are not the equivalent of buy, hold and sell. Investors should carefully read the definitions of all ratings used in Morgan Stanley Research. In addition, since Morgan Stanley Research contains more complete information concerning the analyst's views, investors should carefully read Morgan Stanley Research, in its entirety, and not infer the contents from the rating alone. In any case, ratings (or research) should not be used or relied upon as investment advice. An investor's decision to buy or sell a stock should depend on individual circumstances (such as the investor's existing holdings) and other considerations.

Global Stock Ratings Distribution

(as of August 31, 2023)

The Stock Ratings described below apply to Morgan Stanley's Fundamental Equity Research and do not apply to Debt Research produced by the Firm.

For disclosure purposes only (in accordance with FINRA requirements), we include the category headings of Buy, Hold, and Sell alongside our ratings of Overweight, Equal-weight, Not-Rated and Underweight. Morgan Stanley does not assign ratings of Buy, Hold or Sell to the stocks we cover. Overweight, Equal-weight, Not-Rated and Underweight are not the equivalent of buy, hold, and sell but represent recommended relative weightings (see definitions below). To satisfy regulatory requirements, we correspond Overweight, our most positive stock rating, with a buy recommendation; we correspond Equal-weight and Not-Rated to hold and Underweight to sell recommendations, respectively.

Stock Rating Category	Coverage Universe		Investment Banking Clients (IBC)			Other Material Investment Services Clients (MISC)	
	Count	% of Total	Count	% of Total IBC	% of Rating Category	Count	% of Total Other MISC
Overweight/Buy	1345	37%	273	43%	20%	606	39%
Equal-weight/Hold	1686	46%	301	47%	18%	719	47%
Not-Rated/Hold	3	0%	0	0%	0%	1	0%
Underweight/Sell	594	16%	67	10%	11%	218	14%
Total	3,628		641			1544	

Data include common stock and ADRs currently assigned ratings. Investment Banking Clients are companies from whom Morgan Stanley received investment banking compensation in the last 12 months. Due to rounding off of decimals, the percentages provided in the "% of total" column may not add up to exactly 100 percent.

Analyst Stock Ratings

Overweight (O). The stock's total return is expected to exceed the average total return of the analyst's industry (or industry team's) coverage universe, on a risk-adjusted basis, over the next 12-18 months.

Equal-weight (E). The stock's total return is expected to be in line with the average total return of the analyst's industry (or industry team's) coverage universe, on a risk-adjusted basis, over the next 12-18 months.

Not-Rated (NR). Currently the analyst does not have adequate conviction about the stock's total return relative to the average total return of the analyst's industry (or industry team's) coverage universe, on a risk-adjusted basis, over the next 12-18 months.

Underweight (U). The stock's total return is expected to be below the average total return of the analyst's industry (or industry team's) coverage universe, on a risk-adjusted basis, over the next 12-18 months.

Unless otherwise specified, the time frame for price targets included in Morgan Stanley Research is 12 to 18 months.

Analyst Industry Views

Attractive (A): The analyst expects the performance of his or her industry coverage universe over the next 12-18 months to be attractive vs. the relevant broad market benchmark, as indicated below.

In-Line (I): The analyst expects the performance of his or her industry coverage universe over the next 12-18 months to be in line with the relevant broad market benchmark, as indicated below.

Cautious (C): The analyst views the performance of his or her industry coverage universe over the next 12-18 months with caution vs. the relevant broad market benchmark, as indicated below. Benchmarks for each region are as follows: North America - S&P 500; Latin America - relevant MSCI country index or MSCI Latin America Index; Europe - MSCI Europe; Japan - TOPIX; Asia - relevant MSCI country index or MSCI sub-regional index or MSCI AC Asia Pacific ex Japan Index.

Important Disclosures for Morgan Stanley Smith Barney LLC & E*TRADE Securities LLC Customers

Important disclosures regarding the relationship between the companies that are the subject of Morgan Stanley Research and Morgan Stanley Smith Barney LLC or Morgan Stanley or any of their affiliates, are available on the Morgan Stanley Wealth Management disclosure website at www.morganstanley.com/online/researchdisclosures. For Morgan Stanley specific disclosures, you may refer to www.morganstanley.com/researchdisclosures.

Each Morgan Stanley research report is reviewed and approved on behalf of Morgan Stanley Smith Barney LLC and E*TRADE Securities LLC. This review and approval is conducted by the same person who reviews the research report on behalf of Morgan Stanley. This could create a conflict of interest.

Other Important Disclosures

Morgan Stanley & Co. International PLC and its affiliates have a significant financial interest in the debt securities of Alphabet Inc., Amazon.com Inc, Apple, Inc., Aptiv Plc, Asbury Automotive Group Inc, AutoNation Inc., Avis Budget Group Inc, BorgWarner Inc., Ford Motor Company, General Motors Company, Lear Corporation, Lithia Motors Inc., Magna International Inc., Meta Platforms Inc, Netflix Inc, Penske Automotive Group, Inc, Rivian Automotive, Inc., Tesla Inc.

A member of Research who had or could have had access to the research prior to completion owns securities (or related derivatives) in the Ford Motor Company, NVIDIA Corp.. This person is not a research analyst or a member of research analyst's household.

Morgan Stanley Research policy is to update research reports as and when the Research Analyst and Research Management deem appropriate, based on developments with the issuer, the

sector, or the market that may have a material impact on the research views or opinions stated therein. In addition, certain Research publications are intended to be updated on a regular periodic basis (weekly/monthly/quarterly/annual) and will ordinarily be updated with that frequency, unless the Research Analyst and Research Management determine that a different publication schedule is appropriate based on current conditions.

Morgan Stanley is not acting as a municipal advisor and the opinions or views contained herein are not intended to be, and do not constitute, advice within the meaning of Section 975 of the Dodd-Frank Wall Street Reform and Consumer Protection Act.

Morgan Stanley produces an equity research product called a "Tactical Idea." Views contained in a "Tactical Idea" on a particular stock may be contrary to the recommendations or views expressed in research on the same stock. This may be the result of differing time horizons, methodologies, market events, or other factors. For all research available on a particular stock, please contact your sales representative or go to Matrix at <http://www.morganstanley.com/matrix>.

Morgan Stanley Research is provided to our clients through our proprietary research portal on Matrix and also distributed electronically by Morgan Stanley to clients. Certain, but not all, Morgan Stanley Research products are also made available to clients through third-party vendors or redistributed to clients through alternate electronic means as a convenience. For access to all available Morgan Stanley Research, please contact your sales representative or go to Matrix at <http://www.morganstanley.com/matrix>.

Any access and/or use of Morgan Stanley Research is subject to Morgan Stanley's Terms of Use (<http://www.morganstanley.com/terms.html>). By accessing and/or using Morgan Stanley Research, you are indicating that you have read and agree to be bound by our Terms of Use (<http://www.morganstanley.com/terms.html>). In addition you consent to Morgan Stanley processing your personal data and using cookies in accordance with our Privacy Policy and our Global Cookies Policy (http://www.morganstanley.com/privacy_pledge.html), including for the purposes of setting your preferences and to collect readership data so that we can deliver better and more personalized service and products to you. To find out more information about how Morgan Stanley processes personal data, how we use cookies and how to reject cookies see our Privacy Policy and our Global Cookies Policy (http://www.morganstanley.com/privacy_pledge.html).

If you do not agree to our Terms of Use and/or if you do not wish to provide your consent to Morgan Stanley processing your personal data or using cookies please do not access our research. Morgan Stanley Research does not provide individually tailored investment advice. Morgan Stanley Research has been prepared without regard to the circumstances and objectives of those who receive it. Morgan Stanley recommends that investors independently evaluate particular investments and strategies, and encourages investors to seek the advice of a financial adviser. The appropriateness of an investment or strategy will depend on an investor's circumstances and objectives. The securities, instruments, or strategies discussed in Morgan Stanley Research may not be suitable for all investors, and certain investors may not be eligible to purchase or participate in some or all of them. Morgan Stanley Research is not an offer to buy or sell or the solicitation of an offer to buy or sell any security/instrument or to participate in any particular trading strategy. The value of and income from your investments may vary because of changes in interest rates, foreign exchange rates, default rates, prepayment rates, securities/instruments prices, market indexes, operational or financial conditions of companies or other factors. There may be time limitations on the exercise of options or other rights in securities/instruments transactions. Past performance is not necessarily a guide to future performance. Estimates of future performance are based on assumptions that may not be realized. If provided, and unless otherwise stated, the closing price on the cover page is that of the primary exchange for the subject company's securities/instruments.

The fixed income research analysts, strategists or economists principally responsible for the preparation of Morgan Stanley Research have received compensation based upon various factors, including quality, accuracy and value of research, firm profitability or revenues (which include fixed income trading and capital markets profitability or revenues), client feedback and competitive factors. Fixed Income Research analysts', strategists' or economists' compensation is not linked to investment banking or capital markets transactions performed by Morgan Stanley or the profitability or revenues of particular trading desks.

The "Important Regulatory Disclosures on Subject Companies" section in Morgan Stanley Research lists all companies mentioned where Morgan Stanley owns 1% or more of a class of common equity securities of the companies. For all other companies mentioned in Morgan Stanley Research, Morgan Stanley may have an investment of less than 1% in securities/instruments or derivatives of securities/instruments of companies and may trade them in ways different from those discussed in Morgan Stanley Research. Employees of Morgan Stanley not involved in the preparation of Morgan Stanley Research may have investments in securities/instruments or derivatives of securities/instruments of companies mentioned and may trade them in ways different from those discussed in Morgan Stanley Research. Derivatives may be issued by Morgan Stanley or associated persons.

With the exception of information regarding Morgan Stanley, Morgan Stanley Research is based on public information. Morgan Stanley makes every effort to use reliable, comprehensive information, but we make no representation that it is accurate or complete. We have no obligation to tell you when opinions or information in Morgan Stanley Research change apart from when we intend to discontinue equity research coverage of a subject company. Facts and views presented in Morgan Stanley Research have not been reviewed by, and may not reflect information known to, professionals in other Morgan Stanley business areas, including investment banking personnel.

Morgan Stanley Research personnel may participate in company events such as site visits and are generally prohibited from accepting payment by the company of associated expenses unless pre-approved by authorized members of Research management.

Morgan Stanley may make investment decisions that are inconsistent with the recommendations or views in this report.

To our readers based in Taiwan or trading in Taiwan securities/instruments: Information on securities/instruments that trade in Taiwan is distributed by Morgan Stanley Taiwan Limited ("MSTL"). Such information is for your reference only. The reader should independently evaluate the investment risks and is solely responsible for their investment decisions. Morgan Stanley Research may not be distributed to the public media or quoted or used by the public media without the express written consent of Morgan Stanley. Any non-customer reader within the scope of Article 7-1 of the Taiwan Stock Exchange Recommendation Regulations accessing and/or receiving Morgan Stanley Research is not permitted to provide Morgan Stanley Research to any third party (including but not limited to related parties, affiliated companies and any other third parties) or engage in any activities regarding Morgan Stanley Research which may create or give the appearance of creating a conflict of interest. Information on securities/instruments that do not trade in Taiwan is for informational purposes only and is not to be construed as a recommendation or a solicitation to trade in such securities/instruments. MSTL may not execute transactions for clients in these securities/instruments.

Morgan Stanley is not incorporated under PRC law and the research in relation to this report is conducted outside the PRC. Morgan Stanley Research does not constitute an offer to sell or the solicitation of an offer to buy any securities in the PRC. PRC investors shall have the relevant qualifications to invest in such securities and shall be responsible for obtaining all relevant approvals, licenses, verifications and/or registrations from the relevant governmental authorities themselves. Neither this report nor any part of it is intended as, or shall constitute, provision of any consultancy or advisory service of securities investment as defined under PRC law. Such information is provided for your reference only.

Morgan Stanley Research is disseminated in Brazil by Morgan Stanley C.T.V.M. S.A. located at Av. Brigadeiro Faria Lima, 3600, 6th floor, São Paulo - SP, Brazil; and is regulated by the Comissão de Valores Mobiliários; in Mexico by Morgan Stanley México, Casa de Bolsa, S.A. de C.V. which is regulated by Comisión Nacional Bancaria y de Valores. Paseo de los Tamarindos 90, Torre 1, Col. Bosques de las Lomas Floor 29, 05120 Mexico City; in Japan by Morgan Stanley MUFG Securities Co., Ltd. and, for Commodities related research reports only, Morgan Stanley Capital Group Japan Co., Ltd; in Hong Kong by Morgan Stanley Asia Limited (which accepts responsibility for its contents) and by Morgan Stanley Bank Asia Limited; in Singapore by Morgan Stanley Asia (Singapore) Pte. (Registration number 199206298Z) and/or Morgan Stanley Asia (Singapore) Securities Pte Ltd (Registration number 200008434H), regulated by the Monetary Authority of Singapore (which accepts legal responsibility for its contents and should be contacted with respect to any matters arising from, or in connection with, Morgan Stanley Research) and by Morgan Stanley Bank Asia Limited, Singapore Branch (Registration number T14FC0118); in Australia to "wholesale clients" within the meaning of the Australian Corporations Act by Morgan Stanley Australia Limited A.B.N. 67 003 734 576, holder of Australian financial services license No. 233742, which accepts responsibility for its contents; in Australia to "wholesale clients" and "retail

clients" within the meaning of the Australian Corporations Act by Morgan Stanley Wealth Management Australia Pty Ltd (A.B.N. 19 009 145 555, holder of Australian financial services license No. 240813, which accepts responsibility for its contents; in Korea by Morgan Stanley & Co International plc, Seoul Branch; in India by Morgan Stanley India Company Private Limited having Corporate Identification No (CIN) U22990MH1998PTC115305, regulated by the Securities and Exchange Board of India ("SEBI") and holder of licenses as a Research Analyst (SEBI Registration No. INH000001105); Stock Broker (SEBI Stock Broker Registration No. INZ000244438), Merchant Banker (SEBI Registration No. INM000011203), and depository participant with National Securities Depository Limited (SEBI Registration No. IN-DP-NSDL-567-2021) having registered office at 18th Floor, Tower 2, One World Center, Plot- 841, Jupiter Textile Mill Compound, Senapati Bapat Marg, Lower Parel, Mumbai 400013, India Telephone no. +91-22-61181000; Compliance Officer Details: Mr. Anil Shenoy, Tel. No.: +91-22-61181000 or Email: Anil.Shenoy@morganstanley.com; Grievance officer details: Mr. Anil Shenoy, Tel. No.: +91-22-61181000 or Email: msic-compliance@morganstanley.com; in Canada by Morgan Stanley Canada Limited; in Germany and the European Economic Area where required by Morgan Stanley Europe S.E., authorised and regulated by Bundesanstalt fuer Finanzdienstleistungsaufsicht (BaFin) under the reference number 149169; in the US by Morgan Stanley & Co. LLC, which accepts responsibility for its contents. Morgan Stanley & Co. International plc, authorized by the Prudential Regulation Authority and regulated by the Financial Conduct Authority and the Prudential Regulation Authority, disseminates in the UK research that it has prepared, and research which has been prepared by any of its affiliates, only to persons who (i) are investment professionals falling within Article 19(5) of the Financial Services and Markets Act 2000 (Financial Promotion) Order 2005 (as amended, the "Order"); (ii) are persons who are high net worth entities falling within Article 49(2)(a) to (d) of the Order; or (iii) are persons to whom an invitation or inducement to engage in investment activity (within the meaning of section 21 of the Financial Services and Markets Act 2000, as amended) may otherwise lawfully be communicated or caused to be communicated. RMB Morgan Stanley Proprietary Limited is a member of the JSE Limited and A2X (Pty) Ltd. RMB Morgan Stanley Proprietary Limited is a joint venture owned equally by Morgan Stanley International Holdings Inc. and RMB Investment Advisory (Proprietary) Limited, which is wholly owned by FirstRand Limited. The information in Morgan Stanley Research is being disseminated by Morgan Stanley Saudi Arabia, regulated by the Capital Market Authority in the Kingdom of Saudi Arabia, and is directed at Sophisticated investors only.

The information in Morgan Stanley Research is being communicated by Morgan Stanley & Co. International plc (DIFC Branch), regulated by the Dubai Financial Services Authority (the DFSA), and is directed at Professional Clients only, as defined by the DFSA. The financial products or financial services to which this research relates will only be made available to a customer who we are satisfied meets the regulatory criteria to be a Professional Client. A distribution of the different MS Research ratings or recommendations, in percentage terms for Investments in each sector covered, is available upon request from your sales representative.

The information in Morgan Stanley Research is being communicated by Morgan Stanley & Co. International plc (QFC Branch), regulated by the Qatar Financial Centre Regulatory Authority (the QFCRA), and is directed at business customers and market counterparties only and is not intended for Retail Customers as defined by the QFCRA.

As required by the Capital Markets Board of Turkey, investment information, comments and recommendations stated here, are not within the scope of investment advisory activity. Investment advisory service is provided exclusively to persons based on their risk and income preferences by the authorized firms. Comments and recommendations stated here are general in nature. These opinions may not fit to your financial status, risk and return preferences. For this reason, to make an investment decision by relying solely to this information stated here may not bring about outcomes that fit your expectations.

The trademarks and service marks contained in Morgan Stanley Research are the property of their respective owners. Third-party data providers make no warranties or representations relating to the accuracy, completeness, or timeliness of the data they provide and shall not have liability for any damages relating to such data. The Global Industry Classification Standard (GICS) was developed by and is the exclusive property of MSCI and S&P.

Morgan Stanley Research, or any portion thereof may not be reprinted, sold or redistributed without the written consent of Morgan Stanley.

Indicators and trackers referenced in Morgan Stanley Research may not be used as, or treated as, a benchmark under Regulation EU 2016/1011, or any other similar framework.

The issuers and/or fixed income products recommended or discussed in certain fixed income research reports may not be continuously followed. Accordingly, investors should regard those fixed income research reports as providing stand-alone analysis and should not expect continuing analysis or additional reports relating to such issuers and/or individual fixed income products.

Morgan Stanley may hold, from time to time, material financial and commercial interests regarding the company subject to the Research report.

Registration granted by SEBI and certification from the National Institute of Securities Markets (NISM) in no way guarantee performance of the intermediary or provide any assurance of returns to investors. Investment in securities market are subject to market risks. Read all the related documents carefully before investing.

INDUSTRY COVERAGE: Autos & Shared Mobility

COMPANY (TICKER)	RATING (AS OF)	PRICE* (09/07/2023)
Adam Jonas, CFA		
Adient PLC (ADNT.N)	U (03/17/2021)	\$38.10
American Axle & Manufacturing Holdings Inc (AXL.N)	O (02/28/2022)	\$7.72
Aptiv Plc (APTV.N)	E (11/28/2022)	\$100.68
Asbury Automotive Group Inc (ABG.N)	U (05/20/2022)	\$219.20
AutoNation Inc. (AN.N)	U (01/17/2023)	\$152.44
Avis Budget Group Inc (CAR.O)	O (06/20/2023)	\$197.66
BorgWarner Inc. (BWA.N)	O (05/15/2023)	\$39.75
Carmax Inc (KMX.N)	O (07/10/2018)	\$81.59
Carvana Co (CVNA.N)	U (07/26/2023)	\$47.63
Ferrari NV (RACE.N)	O (05/09/2019)	\$298.79
Fisker Inc (FSR.N)	U (01/25/2023)	\$6.23
Ford Motor Company (F.N)	O (10/05/2022)	\$11.96
FREYR Battery SA (FREY.N)	O (06/28/2023)	\$6.35
General Motors Company (GM.N)	O (05/01/2023)	\$32.57
Group 1 Automotive, Inc (GPI.N)	U (05/20/2022)	\$259.09
Harley-Davidson Inc (HOG.N)	O (03/21/2023)	\$33.51
Hertz Global Holdings Inc (HTZ.O)	E (12/06/2021)	\$15.76
Lear Corporation (LEA.N)	E (02/28/2022)	\$139.77
Li-Cycle Holdings Corp. (LICY.N)	U (01/25/2023)	\$4.56
Lithia Motors Inc. (LAD.N)	U (02/09/2021)	\$297.81
Lucid Group Inc (LCID.O)	U (09/13/2021)	\$6.00
Magna International Inc. (MGA.N)	O (10/14/2021)	\$57.70
Mobileye Global Inc (MBLY.O)	E (07/24/2023)	\$35.83
Penske Automotive Group, Inc (PAG.N)	U (11/17/2021)	\$158.95
Quantumscape Corp (QS.N)	U (11/09/2022)	\$6.96
Rivian Automotive, Inc. (RIVN.O)	O (12/05/2021)	\$23.42
Sonic Automotive Inc (SAH.N)	U (11/17/2021)	\$51.84
Tesla Inc (TSLA.O)	O (09/08/2023)	\$251.49
Visteon Corporation (VC.O)	E (06/01/2022)	\$136.66

Stock Ratings are subject to change. Please see latest research for each company.

* Historical prices are not split adjusted.

Morgan Stanley

© Morgan Stanley 2023

The Americas

1585 Broadway
New York, NY 10036-8293
United States
Tel: **+1 (1) 212 761 4000**

Europe

20 Bank Street, Canary Wharf
London E14 4AD
United Kingdom
Tel: **+44 (0) 20 7 425 8000**

Japan

1-9-7 Otemachi, Chiyoda-ku
Tokyo 100-8104
Japan
Tel: **+81 (0) 3 6836 5000**

Asia/Pacific

1 Austin Road West
Kowloon
Hong Kong
Tel: **+852 2848 5200**