COBANK
COOPERATIVE. CONNECTED. COMMITTED.

February 2022

Self-Driving EVs Could Overwhelm the Data Center Market. Can the Industry Respond?

Jeff Johnston
Lead Economist, Communications

Inside…

Key Points:

- Auto manufacturers are making precedent-setting investments as they pivot away from internal combustion engine (ICE) vehicles to electric vehicles (EVs). These investments are laying the ground work on the path to self-driving EVs.
- Some of the obvious challenges to self-driving EV adoption are safety, social acceptance, and regulatory restrictions. But another major challenge is the impact these vehicles will have on the data center market.
- Self-driving EVs are expected to generate unthinkable amounts of data that will have a profound impact on the markets for data center storage and computation.
- While technology developments like efficient hyper-scale data centers and on-board processing may provide long-term solutions, there's an enormous chasm today between the existing infrastructure and what is needed to support widespread adoption of self-driving EVs.
- The amount of capital required to build the data center infrastructure for self-driving EVs is so massive that it begs the question: Without major technological advancements in compute and storage processes, will the industry be able to handle the deluge of data self-driving EVs will generate?

Introduction

As auto manufacturers pivot their business model away from internal combustion engine (ICE) vehicles and towards electric vehicles (EVs), the stage is being set for millions of self-driving EVs to buzz through the streets. And while the timing is debatable, conventional wisdom suggests it's not a matter of if but when it will become a reality. Beyond the obvious challenges of safety, social acceptance, and regulatory restrictions, other factors must be considered.

Self-driving EVs are expected to have a profound impact on the "datasphere" (the total data created in a year), the likes of which the industry has never seen. There is an enormous chasm between the existing data center infrastructure and what is needed to support wide spread adoption of self-driving EVs. In the meantime,

EXHIBIT 1: Autonomous Driving Levels

Level 0	Level 1	Level 2	Level 3	Level 4	Level 5
The driver is fully responsible and permanently carries out all the aspects of the driving tasks	The driver can delegate either steering or accelerating/ braking to the system	The driver must permanently monitor the system	In certain situations, the driver can turn attention away from the road, but must always be ready to take full control again	The driver can transfer complete control to the system but can take control at any time	No driver needed
No driver assistance systems	The system will perform one of the driving tasks	The system will perform several of the driving tasks	The system can autonomously control the vehicle on defined routes	The system is able to perform all driving tasks	The system controls the vehicle autonomously under all conditions
Manual Driving →					**Automation**

Source: tesmanian.com

data center energy efficiency has leveled off over the last few years, which creates a less than ideal situation of flattening efficiencies and surging demand.

In this report we look at the impact this technology will have on the data center market. It appears that barring profound technology advancements in storage and computing processes, and massive investments in renewable energy, the industry might not be able to handle the deluge of data that is expected to come.

EV Investments

EV sales and investments are important leading indicators for self-driving technology.

Legacy auto manufacturers are making precedent-setting investments in their EV business in response to customer demand and climate trends. For example, General Motors (GM) plans to spend $35 billion and to reach 1 million annual EV unit sales by 2025. GM shared its vision of a world of zero crashes, zero emissions and zero congestion several years ago, along with plans to stop selling ICE vehicles by 2035. Ford plans to sell 600,000 EV units by 2023 and will invest $18 billion to build new plants in Tennessee, Kentucky and Michigan. And Toyota said that EVs will represent about one-third of its sales by 2030 as it rolls out 30 new models. Then there is the market leader, Tesla, which reported selling 936,000 units in 2021 with a goal of selling 20 million per year by 2030.

So clearly the industry is already moving quickly towards EVs. The U.S. market trails the European and Chinese markets where EV sales represented 10.3% and 14.3% respectively in 2021. By contrast, EV sales in the U.S. represented only 5% in 2021. Based on U.S. auto manufacturers' investments and the expected ramp in domestic sales, this EV adoption gap between the U.S. and other markets should start to close.

Self-driving EVs

Self-driving EV technology has been advancing for quite some time. According to predictions several years ago, we'd have Level-5 robotaxis in every major U.S. city by now *(Exhibit 1).* But the market has fallen short of the most optimistic forecasts due to technical, social, safety, and regulatory issues. However, we are now starting to see local jurisdictions allow limited use of self-driving technologies for ride share services. For example, Las Vegas announced that starting this year, it will allow Lyft and Motional to deploy a limited fleet of self-driving taxis (with nobody behind the wheel) with a full-fledged commercial launch slated for 2023. And Elon Musk, CEO and founder of Tesla, said that he expects Tesla to achieve Level 4 autonomy this year.

An abrupt switch to driverless cars is unlikely. Rather, we see gradual adoption as auto manufacturers move up the autonomous driving scale, starting with commercial applications then on to consumer use. Non-Tesla manufacturers are just beginning to launch EVs and still need to collect a significant amount of driving data to fine-tune their self-driving algorithms. Tesla is way ahead of the competition and its lead is expected to increase over the coming years. The general consensus today is that most EVs will reach Level 4 autonomy by 2030.

Data Generation

Impact projections from widespread adoption of Level 4 EVs on the data center market are mind blowing *(Exhibit 2).* If these projections are in the ballpark of reality, do we have the digital infrastructure to handle what is coming?

According to the Automotive Edge Computing Consortium (AECC), self-driving EVs may eventually need to offload as much as 5,000GBs per hour of operation – to put that into perspective, in 2020 the average person worldwide generated about 150GBs per day. Holon Investments estimates that the global self-driving fleet could reach 400 million by 2035. Extrapolating these numbers implies the total global datasphere will surge from 64 zettabytes (ZB) in 2020 to an eye-popping 10,000 – 15,000ZBs in 2035 just from self-driving EVs.

EXHIBIT 2: Datasphere Forecast

2020	2035
Datasphere **64 ZBs** • Storage Capacity **2.1 ZBs** • Cost to Build 2.1 ZBs of Capacity **$2 trillion**	Datasphere Forecast (from self-driving EVs) **10,000 to 15,000 ZBs** **Cost to Build Needed Capacity** **????**

Source: Holon Investments

EXHIBIT 3: Global Data Storage Capability*

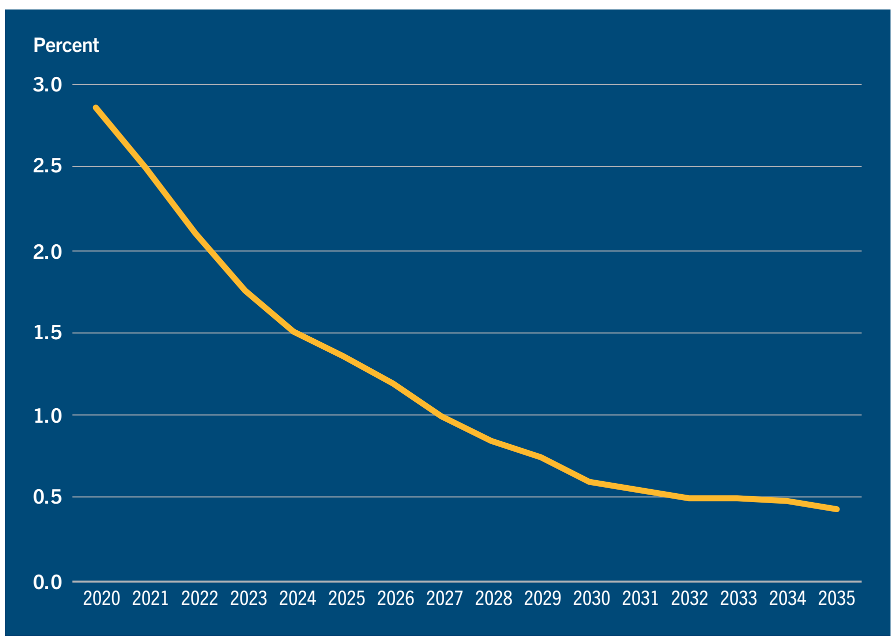

Source: Holon Investments

**If we were to increase our storage capacity 100 fold (to 210ZBs) by 2035, only about .4% of the datasphere would be stored.*

EXHIBIT 4: Data Center Average Annual Power Usage Effectiveness

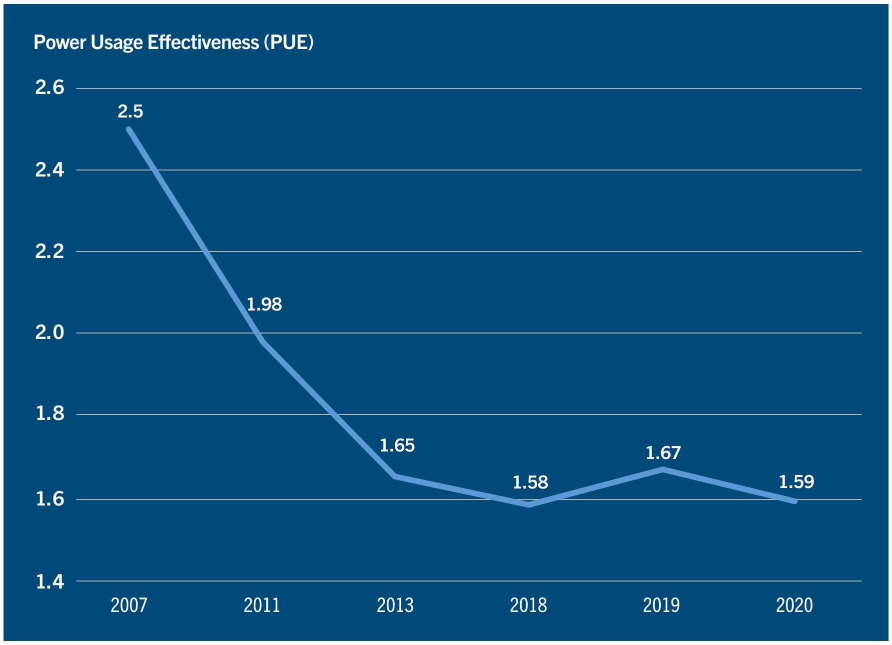

Source: Statista

New types of storage technologies and architectures including edge computing will help manage the deluge of data. Filecoin, a decentralized storage network made up of a large number of diverse storage providers and developers, could be another arrow in the industry's quiver. It's also likely that as autonomous driving algorithms mature, onboard processing will increase, meaning less data will need to be offloaded to a data center. But even with these developments, the chasm between the data we generate and store today versus where we will be in 10 years is bigger than anything the data center market has seen, by orders of magnitude.

Most of the data that self-driving EVs generate will not be offloaded to a data center. The AECC estimates that 30% of video will be uploaded to help refine models and train algorithms, while about 2% of the video captured will be uploaded to retain audit trails for accidents or traffic incidents.

Data Storage

In 2020, globally we stored just under 3% of all data created, with a total storage capacity of 2.1ZBs. If we were to increase our storage capacity 100 fold (to 210ZBs) by 2035, only about .4% of the datasphere would be stored *(Exhibit 3).* A 100 fold increase in storage capacity represents a 15-year 36% CAGR, a good bit higher than the market intelligence firm IDC's annual forecast of 18% (from 2020 to 2024). So, will the industry have enough storage capacity to support self-driving EVs if we can only store .4% of the data created? And considering that it has taken over $2 trillion to build 2.1ZBs of storage capacity, is there a (data center) business case to support mass adoption of self-driving EVs? At this point there appears to be more questions than answers.

Energy Costs

With energy costs representing about 40% of a data centers' total operating expenditures, it's important to explore this area of the business in the context of exponential growth to the datasphere.

The good news is that despite the increase in internet traffic and storage demand, data centers have done an excellent job managing their energy expenses. For example, from 2010 to 2018, data storage increased by 25-fold, with energy consumption increasing by only a factor of three. Data center computation was up 650% while associated energy usage was up just 25%. In total, energy consumption for data center storage and computation only increased by 6% from 2010 to 2018.

These efficiencies were, in part, due to Moore's law, which predicts a 25% annual energy decline per processing unit. However, these gains don't last forever, which could present cost challenges. In fact, the power usage effectiveness (PUE) for data centers has been relatively flat in recent years *(Exhibit 4).* A flattening in energy efficiency coupled with an explosion in

EXHIBIT 5: Renewable Energy Cost Declines, 2010-2019

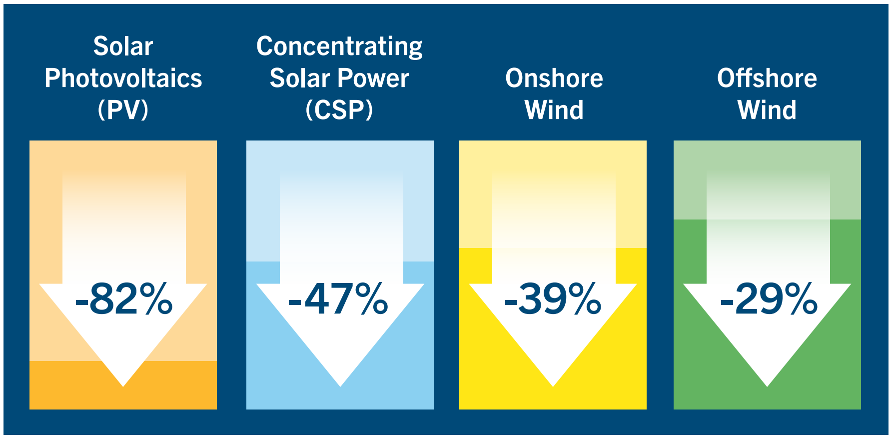

Source: IRENA, 2020

datacenter demand is not ideal for data center operators. Self-driving EVs are expected to represent the lion's share of future datacenter energy usage given their disproportionate increase in data storage requirements

One possible solution for energy efficiency is the growth in hyper-scale data centers. According to Amazon Web Services, hyper-scale data centers only require 16% of the power compared to on-premise infrastructure. And of course, leveraging renewables is an important part of the equation. Renewable energy costs have declined sharply over the last several years *(Exhibit 5).* A study from the International Renewable Energy Agency (IRENA) showed that more than half of the renewable capacity added in 2019 achieved lower electricity costs than new coal-fired capacity.

The large hyper-scale data center providers have been investing aggressively in renewables. In 2017 Google announced that it achieved 100% renewable energy across all its operations. Amazon is committed to powering its entire operations with renewables by 2025. While a commitment to renewable energy – and its potential cost efficiencies – is great news, it is important to note that major infrastructure investments are needed to continue to scale renewable energy. And our aging grid is making energy transport more expensive and less efficient.

Summary

The impact of self-driving EVs on the total global datasphere is stunning, as estimates show the total data created in a year will surge from 64ZBs in 2020 to an eye-popping 10,000 – 15,000ZBs in 2035 – just from self-driving EVs.

Variables such as efficient renewable energy, hyper-scale data centers, edge computing, and on-board processing could certainly affect the impact self-driving EV technology will have on the data center market. So, is there a business case that justifies the massive investments required in capacity/computation resources to support the flood of data that is likely to come? At this point, profound technological advancements will be needed to bridge the enormous gap that exists today.

We have a long way to go before self-driving EVs are adopted at scale, and technology will look a lot different when that happens. But even if the current estimates are on the high side, the industry will still likely face major challenges in how data is stored and processed.

References

https://www.ktnv.com/news/las-vegas-selected-as-first-city-for-new-fully-driverless-taxi-service

https://fortune.com/2021/11/10/motionals-driverless-taxi-fleet-will-expand-worldwide-after-vegas-2023-launch-ceo-says/

https://datacenterfrontier.com/rolling-zettabytes-quantifying-the-data-impact-of-connected-cars/

https://holon.investments/how-a-deluge-of-data-from-electric-vehicles-will-create-a-global-data-storage-capacity-crisis/

https://cleantechnica.com/2021/10/08/teslas-goal-20-million-annual-car-sales-by-2030/

https://www.sciencedirect.com/science/article/pii/S0306261921003019

https://energypost.eu/the-nexus-between-data-centres-efficiency-and-renewables-a-role-model-for-the-energy-transition/

https://www.nytimes.com/2021/12/14/business/toyota-electric-vehicles.html

https://plants.gm.com/media/us/en/gm/ev.detail.html/content/Pages/news/us/en/2021/jun/0616-gm.html

https://techjury.net/blog/how-much-data-is-created-every-day/#gref

CoBank's Knowledge Exchange Division welcomes readers' comments and suggestions.
Please send them to KEDRESEARCH@cobank.com.