Hyperscalers Amazon and Facebook Drive First Quarter Storage Revenue

In a reversal of 2019's trends, hyperscale providers have sharply increased their spending in Q1 2020, providing some much-needed buoyancy to the overall storage market according to IDC's latest Enterprise Storage System Tracker. While the overall market results in Q1 2020 included an 18.1% drop in total storage capacity shipped and a drop in original design manufacturer (ODM) capacity shipped of 20.2%, it also included an increase in ODM revenues due to increased sales to hyperscalers. This is likely due to the purchase of partially-populated, latest-generation storage systems by the hyperscale companies from the ODMs.

OEM Storage Systems Hits and Misses

Counterbalancing this, enterprise external storage capacity increased by 3%, while revenues fell 8.2%. The total all-flash array market generated $2.8B in revenue, a slight increase, while the hybrid flash array market (HFA-disk +SSD) was down 11.5% at $2.5B. Disk arrays declined 18%. Uncertainty, pandemic, stay at home

orders seemed to have only a limited impact on Huawei, Pure Storage, and IBM. All three vendors reported revenue gains, which can largely be attributed to their focus on all-flash arrays. Huawei in particular saw an increase in year-to-year revenue of 17.7%. As a group, these three vendors went from an 11.8% share if the storage market in Q1 2019 to 14.1% in Q1 2020.

Storage industry giants Dell, NetApp, HPE, and Hitachi all slipped in calendar Q1 2020 revenue. Dell still dominates the worldwide enterprise external OEM storage systems market, with 33.2% market share but took a big hit with revenues down 8.2%.

Top 5 Companies, Worldwide Enterprise External OEM Storage Systems Market, First Quarter of 2020 *(Revenues are in US$ millions)*

Company	1Q20 Revenue	1Q20 Market Share	1Q19 Revenue	1Q19 Market Share	1Q20/1Q19 Revenue Growth
1. Dell Technologies[a]	$2,162.0	33.2%	$2,355.9	33.2%	-8.2%
2. NetApp	$715.7	11.0%	$894.9	12.6%	-20.0%
3. HPE/New H3C Groups[b]	$646.2	9.9%	$778.2	11.0%	-17.0%
4. Hitachi	$430.3	6.6%	$470.5	6.6%	-8.5%
T5. IBM*	$332.1	5.1%	$320.0	4.5%	3.8%
T5. Pure Storage*	$311.7	4.8%	$289.5	4.1%	7.7%
T5. Huawei*	$270.7	4.2%	$230.0	3.2%	17.7%
Rest of Market	$1,652.9	25.3%	$1,766.5	24.9%	-6.4%
Total	**$6,521.7**	**100.0%**	**$7,105.4**	**100.0%**	**-8.2%**

Source: IDC Worldwide Quarterly Enterprise Storage Systems Tracker, June 9, 2020.

How Much Trunk Space Will Onboard Storage Consume in Autonomous Vehicles?

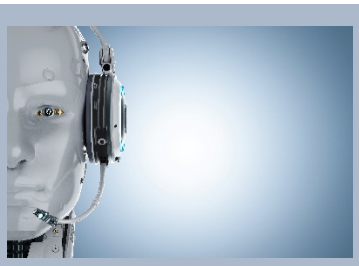

As our cars (and other vehicles) become smart/smarter, the use of high-performance onboard storage is increasing rapidly. But how much storage does an autonomous vehicle (AV) need? It turns out (unsurprisingly) that the amount of storage depends on the amount of automation and autonomy in the vehicle. The Society of Automotive Engineers defines 6 levels of vehicle autonomy:

1) L0- None- (or mostly none) Driver controls everything except engine fuel injection system
2) L1- Driver Assistance- A process of driving is automated such as cruise control
3) L2- Partial Automation- More driver functions are automated
4) L3- Conditioned Automation-Driving is entirely automated but the driver/passenger is prepared to take over if needed

5) L4- High Automation- No driver/passenger control, but restricted use

6) L5- Full Automation- True driverless cars

Each subsequent level of automation requires collection and storage of more data, including from outside the car collected via cameras and radar and data collected from within the car, including engine output, exhaust, suspension spring rates. However, experts disagree on exactly how much storage is needed at each level. Quantum has provided estimates between 5TB and 20TB for autonomous test vehicles. At CES, Seagate said Renovo's AV software can generate 32TB/day of data collected and stored at the edge before migrated to the cloud. John Hayes, CEO of Ghost, puts the number closer to 3.2TB/day. The range of these projections is problematic because they require design choices and costs that factor into overall feasibility to be within a spending comfort range for consumers.

Design option for data storage range from centralized to distributed, to hybrid options. Factoring in the real-time function needs such as instant braking and steering coupled with gaps in communications network coverage and a 32TB/day estimate of generated data, a 30-day storage period to cover extended non-uplink time would require 960TB capacity. For speed and reliability, disk storage is not realistic; flash storage would be required. All of this combine to thousands of dollars in additional cost- conservatively, around $30k.

A wrong estimate of data storage needs means failure, aggravation, and cost. Tesla upgraded to a MCUv2 with 32GB flash card in 2018. Prior models used a MCUv1 firmware that was 300MB but grew to require 1GB because car logging information filled the card up leaving no cells for use by the card's wear-leveling function. From 2010-2018 Tesla Model S and X cars the MCU, which controls 99% of the car functionality, failed in under three years because they ran out of storage space.

As SSDs are integrated into the vehicle, the challenge is creating fast, accessible, reliable, data, while addressing changing temperatures and vibration, and the ability to operate independently and in conjunction with an entire system.

"Now, as the core evolves, we've started working directly with OEMs on the car's partitioning of storage," says Verhoeven. "If you have an infotainment cluster on the one hand, an ADAS cluster on the other hand, and in the future you'll have an event recorder for all kinds of legal or insurance reasons in each of these systems, optimizing the way you use storage will be of great value."

"Any evolution in sophistication requires new data center infrastructures and edge technologies to extract immediate intelligence from captured data, not only to enhance the driving experience, but to also make it safer." Huibert Verhoeven, SVP and GM, Automotive, Mobile, and Emerging Markets, Western Digital.

AV System Storage Must Account for Cybersecurity and Training System Data Demands

In addition to the storage needed above for AV operation, onboard cybersecurity solutions (which can include anti-malware software, threat databases, and security incident logging data) require significant storage not included in the capacity estimates of the previous article. In many ways, AVs (especially commercial AVs) will be very attractive to attack by cybercriminals and other cybersecurity actors. NVIDIA's 2018 Safety Report for Self-Driving Cars summarizes their commitment to addressing cybersecurity in AV systems:

> "Our cybersecurity team works with the Automotive Information Sharing and Analysis Center (Auto-ISAC), NHTSA, SAE, and the Bureau of Industry and Security (Department of Commerce). We contribute to Automatic Identification System (Department of Homeland Security), Federal Information Processing Standards (Federal Information Security Management Act), and Common Criteria standards or specifications. In addition, we use the SAE J3061 cybersecurity process as a guiding principle and leverage processes and practices from other cybersecurity-sensitive industries. We also participate in the SAE J3101 standard development, which ensures the necessary building blocks for cybersecurity are implemented at the hardware chip level. And we review platform code for security conformance and use static and dynamic code analysis techniques."

Given the rapidity in which cybersecurity threats evolve and adapt to the environment, the amount of storage needed for cybersecurity attack countermeasures will likely be something that grows significantly over time.

The other factor that will impact storage in AVs, and in the artificial intelligence (AI) systems that support them, will be the data storage needs for training. AI applications, based on massive datasets, are still challenged by real-time application under varying hazards, conditions, and circumstances. AI is capable but not maximal. Validation programs for AV neural networks require at least 20 petabytes of training data. Keith Rieken, Solutions Manager, Artificial Intelligence, Pure Storage, explains that the objective is achieving the ideal reaction to a situation, not just 'any' reaction, "That's the challenge we see today—how do we sufficiently train those networks in order to achieve full autonomy?"

If you want to look deeper into these challenges – including how to construct systems to analyze training data for AV – please join us for our July webinar (register here).

John Hayes, CEO of self-driving car technology company Ghost, "Driving is not a rules problem; it is a learning problem- most of driving is learning what is important and what can be ignored."

Excerpt from Medium…John Hayes, CEO, Ghost

"In robotics, you write a bunch of rules, put them in a car and drive around to see if they work. It's essentially 'guess and check' — your models are based on whatever assumptions an engineer might make about driving. When it breaks, you solve the edge cases, writing exceptions by hand. Then you repeat the process. This is limiting in many ways — first, your assumptions are made up, subject to human error or bias. And second, your assumptions are difficult to change — with every exception built on top of every rule, it becomes extremely expensive to adjust your initial assumptions, akin to starting over.

In contrast, a learning system is an iterative process designed to discover the ideal program to drive a car. Instead of starting by making assumptions about how to drive, you start by observing how people actually drive in the real world. This data — capturing what people see and what they do next — serves as ground truth for your models. You then use machine learning techniques to discover which features in the environment actually impact driving decisions, and how people best navigate the world safely. This process is highly iterative, testing a broad selection of potential inputs and measuring model performance against more real-world data. You are essentially developing a driving model in reverse, starting with the right answer, and then using math to discover how people get to that answer.

The obvious advantage here is flexibility. When engineers try to write driving rules in robotics, they are guessing. No one knows exactly what features or inputs in a scene influence our driving behavior. With learning systems, the guesswork is removed — you can discover the most important inputs *and* the correct behavior by observing lots of real-world driving. A system that can rapidly hypothesize, train and measure is no longer constrained by a list of rules; any array of signals (e.g. objects, classifiers, sign recognizers, road

rules from HD maps) can be tested to find the optimal mix. This is probably the single greatest benefit of the new machine learning age — allowing computers to find patterns in the world that people have failed to fully describe with linear programming. The flexibility to test a lot of variables leads to a superior product."

Upcoming G2M Research Webinars for the Rest of 2020

As our industry continues to be "virtual" during the rest of 2020, webinars can be a good way to stay up-to-date. G2M has several webinars scheduled for this year on hot topics in our industry. Interested in attending our webinars? Register by clicking on the dates of interest. Interested in Sponsoring a webinar? Contact **G2M** for a prospectus.

July 21:	AV, Self-Driving Cars, and Advanced Storage
August:	**VMware and NVMe-oF™ - Providing Storage Networking At Lightning Speeds *(new)***
Sept 15:	Edge Computing/Storage – Get (and Keep) Your Data Off Of My Cloud
Oct 20:	AI and Storage Use Cases in Healthcare
Nov 17:	NVMe-oF™ - Using Telemetry to Improve Network Latency

G2M RESEARCH

Check back for additional 2020 company-specific, conference, and other webinars (to be posted soon).
Let us know if there are any enterprise storage topics you would like to see covered this year or next.
Our first quarter 2021 webinar schedule will be released soon.

Upcoming 2020 Virtual Enterprise Storage Events

- International Supercomputer, June 22-25, free but register for access to videos post conference
- HPE Discover, starts June 23, spans several weeks with 100+ speakers
- Microsoft Ignite, September 21-25, more info soon
- VMworld US, September 28- October 2, may have regional gatherings this Fall as well
- Flash Memory Summit, October, more info soon

Enterprise Storage Newsfeed

New approach to DNA **data storage** makes **system** more dynamic, scalable
Science Daily

"Most of the existing DNA **data storage systems** rely on polymerase chain reaction (PCR) to access stored files, which is very efficient at copying ..

Pure **Storage** All-Flash Software Gets Important Updates
Channel Futures (blog)

... and data migration hassles. They don't want to worry about having people learn all kinds of specialized skills to run their **data storage systems**."

HPE Further Blurs The **Storage** Line Between On Premises And The **Cloud**
The Next Platform

... networking, and **storage** hardware and software like operating **systems**, virtualization, **cloud management** and network security to enterprises and ..

StorONE touts Optane **Flash** Array
Blocks and Files

StorONE, the data storage startup, has crafted a super all-**flash** array by twinning Optane and QLC **Flash SSDs** in a 2-tier, 40TB, dual Xeon SP server ..

Kingston Digital ships 2TB KC2500 **NVMe** PCIe SSD
ZAWYA

2 **NVMe**™ PCIe SSD for desktop, workstations and high-performance computing (HPC) systems. The new, higher capacity gives customers the flexibility ...

Memory Vendors Move Up the Stack to Optimize Storage
Eetasia.com (press release)

As NVMe continues to mature and help storage systems get the most from **flash**-based **SSDs**, memory vendors are looking a little higher up the stack ...

Mission-critical computing and HCI: the time has come
Blocks and Files

Intel® Optane™ **SSDs** fit the bill here, because Intel® Optane™ is based on different technology to the NAND **flash** found in most other **SSDs**. Current ..

ZADAK Announces First PCIe SSD: The Spark RGB M.2, **NVMe** Up to 2 TB
AnandTech

ZADAK, a company that up until now has primarily been known for its memory modules, has just announced its first-ever PCIe 3.0 SSD. The ZADAK ...

Hybrid drive vs. **SSD**: What's best for your organization?
TechTarget

SSHDs include a small amount of **flash** in combination with the core hard disk architecture. The high-speed **SSD** portion of the drive acts as a cache ...

Samsung 980 Pro **SSD** likely to be released within the next two months
TechGenyz

The Samsung 980 Pro **SSD** was introduced at the CES 2020 which took ... TLC NAND **flash**, something that is widely used by the rest of the industry.

HPE updates Primera and Nimble arrays
Blocks and Files

The Primera array gets **NVMe** support via Primera OS v4.2. HPE said all-flash Primera already delivers 75 per cent of I/O within 250 µs latency.

Worldwide Enterprise External OEM **Storage Systems** Market Revenue Declined 8.2% During the ...
Business Wire (press release)

According to the IDC Worldwide Quarterly Enterprise **Storage Systems** ... individuals and employees leverage **cloud**-based collaboration tools, and ..

Splunk Strengthen its Technologies Platform With **Cloud** and Machine Learning Capabilities
Data & Storage Asean (press release) (blog)

“With Splunk's **cloud** and platform **solutions**, organizations can minimize ... popular mobile **device management** (MDM) providers such as MobileIron .

The mainstream all-**flash** array is an enterprise dead-end, these three startups say

Blocks and Files

A thin **SSD** layer acts as an intervening store between disk and the DRAM. There are three controllers per array, which enables parallel access to the ..

Sabrent Rocket Q NVMe 8TB **SSD** Review (Page 1)

TweakTown

As you may imagine, fitting 8TB of **flash** on an M.2 x 2280 PCB along with DRAM and **SSD** controller is not easy to do. You need 1TB per **flash** ...

Cloudian Object **Storage**: Product Overview and Insight

eWeek

Object **storage** addresses this challenge with a distributed **system in** which nodes may be deployed wherever needed. Seamless **cloud** integration: Most ..

Building Bulletproof Bioinformatics **Storage**

The Next Platform

This pre-processed **data** is sent to two other HPC sites via an OpenStack ... The CLIMB team's current Dell-integrated **storage system** layers Red Hat ..

IBM Wheels And Deals With Solution Edition Booster Pack

IT Jungle

Add it all up, and this SAS **flash SSD** Power S914 costs $28,035 at list price. By moving to a pair of 1.6 TB NVM-Express **flash** cards, for a total of 3.2 ...

Your occasional **storage** digest featuring Druva, WekaIO with STACs more benchmarks, and more

Blocks and Files

The Kx kdb+ 3.6 database **system** was distributed across 14 HPE Proliant XL170r Gen10 servers, with **data stored** in a cluster of 18 HPE Proliant ...

Storage class memory could benefit VDI storage, remote work

TechTarget

Compared to **flash SSDs**, the Optane DC drives offer faster and more reliable performance, with latencies typically less than 10 microseconds (µs). In ...

Scality claims big savings for hospital **data storage**

Blocks and Files

A chart in the report shows the average costs for the Scality system and the customers' equivalent **storage systems** not using Scality. Costs are split ...

Highlights of the day: YMTC said to enter **SSD** brand business

Digitimes

China-based NAND **flash** maker YMTC reportedly is looking to cross into brand **SSD** business in the third quarter and will focus on supplying the ...

KIOXIA Adds Thin Provisioning to KumoScale Software Suite, Increases Cloud Storage Efficiency

insideHPC

... Corporation, has added thin provisioning to its KumoScale™ storage software based on NVM Express™ over Fabrics (**NVMe**-oF™) technology.

IBM Joins the Active Archive Alliance

EnterpriseAI

IBM's active archival solution, the IBM **Cloud** Object **Storage System**, ... end users with active archive solutions for intelligent **data** management.”

WekaFS on the HPE Apollo 2000 Gen10 **System** Sets 17 New Benchmark Records

HPCwire

Optimized for NVMe and the hybrid **cloud**, Weka handles the most demanding **storage** challenges in the most **data**-intensive technical computing ...

How Mayo Clinic manages patient **data** privacy, consent in licensing deals with tech companies

Becker's Hospital Review

One of the health **system's** partnerships is with Google, which provides **data storage** and research support. Mayo said it may provide a small number of ...

NVMe SSDs Ready to Embrace EDSFF Form Factors

Eetasia.com (press release)

Just as **NVMe** freed SSDs from relying on legacy technologies designed for spinning disks, the Enterprise & Data Center SSD Form Factor (EDSFF) ...

Argonne's New Menu of **Data Storage** Software Helps Scientists Realize Findings Earlier

HPCwire

It's nearly impossible to build a single **data storage system** that would satisfy both ... This research used the Theta and Cooley systems at the Argonne ...

Cloudian Adds Object-Based **Storage** Platform Monitoring Tool

DevOps.com

Reliance on object-based **storage systems** to manage unstructured **data** within local **data** centers is still a relatively new phenomenon. **Cloud** storage ...

Samsung Announces New NAND **Flash** Facility to Address Future Data Center and Mobile Demands

WebWire (press release)

Samsung Electronics, the world leader in advanced memory technology, today announced plans to expand its NAND **flash** production capacity in ..

Samsung aims to kill server CPUs with this special new **SSD**

TechRadar India

When it comes to **flash** memory, the most common configuration in data centers worldwide is JBOF (or just a bunch of **flash**), which uses NVMe **SSD** ...

Huawei launches OceanStor Pacific series **storage system**

InfotechLead.com (press release) (blog)

Data does not need to be migrated between multiple **storage systems**, improving service processing efficiency by 25 percent and reducing space by 20 ...

6 converged infrastructure vendors' product options compared

TechTarget

Dell EMC VxBlock 1000 is built for use in private **cloud** or hybrid **cloud** environments based on VMware. The **system** uses Dell EMC **storage** and **data** ...

SK Hynix Returns to Top 5 NAND **Flash** Producers

BusinessKorea

... server **SSD** sales surged in the first quarter to the point of accounting for 40 percent of its NAND **flash** sales. The company's NAND **flash** shipments ...

Infinidat adds business offerings, support for **NVMe** over Fabrics

ITWeb

NVMe over Fabrics is designed to connect hosts to storage across a network fabric using the **NVMe** protocol. On the business side, Infinidat is offering ...

Fujitsu Launches Flexible, Scale-out Software-Defined **Storage** ETERNUS **Data** Services Platform ...

Data & Storage Asean (press release) (blog)

The solution enables flexible and rapid **storage** expansion, reducing **management** burden and costs throughout the **system** lifecycle. To make the most ...

Will **NVMe** become the universal block storage access protocol?

Blocks and Files

Facebook's virtual OCP summit earlier this month hosted two provocative presentations that suggested **NVMe** could become a universal block access ...

How 6 new midrange storage arrays from top vendors stack up

TechTarget

The use of **NVMe**, storage class memory (SCM) and flash means most midrange array users will have way more performance than they could ever use ...