Cloud Automation Software

Tech Decoded | Datacenter is Back in Vogue (Vol #1) - Compute, Storage, Silicon

CONCLUSION

In our latest *Tech Decoded* series, we are providing a multi-volume overview around Datacenter solutions across Cooling, Compute, GPUaaS, Networking, Power, Security, Storage, & Silicon. While the shift towards cloud / IaaS has been a secular change in the space, focus has shifted back to datacenter solutions given the needed investments by hyperscalers and enterprises around AI and datacenter modernization. This *Tech Decoded* volume across the +$500B Datacenter opportunity focuses on the Compute, Silicon/Chips, Storage, and Backup & Recovery segments. The report discusses some of the datacenter basics, market opportunities, and trends across these sub-categories. From an investment view, we see AMD, AVGO, MRVL, NTNX, NVDA, PSTG, and RBRK as some of our favorite ways to play the themes around each of the sub-categories currently, while also noticing a handful of privates that investors should be monitoring given their differentiated technologies.

- **Datacenters are "Cool" Again:** Datacenters house the processing, storage, and networking capabilities needed to run applications and services. There have been multiple evolutions of the datacenter, and with the advancements of AI, there is another one upon us. Datacenter drivers include sites growing +1% CAGR, mix-shift towards IaaS but a hybrid environment overall, and refresh of on-premise solutions as organizations modernize their stacks for AI and non-AI purposes. This has created an "arms-race" over the years by large vendors looking to offer consolidated infrastructure options across the +$500B TAM (11% CAGR).

- **Compute & Silicon:** Gartner estimates Server spending will increase from $187B in 2024 to >$300B by 2028, with CSPs allocating +70% of server hardware budgets to AI-servers over the next two-years. Spending intentions by enterprises on Compute are the highest we've seen in years, with edge computing, Blackwell GPUs, and shift in legacy CPUs to GPUs and share underneath some trends to monitor. There is a lot of private innovation going on within chips, including at Cerebras, Groq, and SambaNova. *Most exposed names: AVGO (OW, $250 PT), MRVL (OW, $120 PT), NVDA (OW, $175 PT), AMD (OW, $180 PT), CSCO (N, $57 PT).*

- **Primary Storage** is anticipated to increase from $41B to $53B by 2028, as SSD-arrays and HCI gain share underneath. Enterprise spending intent appears elevated heading into 2025. Data growth, architecture shifts towards software-defined storage, SSD capacity and pricing, AI requirements, and STaaS are some themes to watch here. Additionally, Pure Storage has seen hyperscalers evaluate storage-software for AI-datacenters. While there are multiple privates to watch, some up-and-coming names to be aware of include Vast, Weka, Qumulo, Scale, Nasuni, and Wasabi. *Most exposed names: PSTG (OW, $76 PT), NTNX (OW, $83 PT).*

- **Backup & Recovery:** "Secondary Storage" or Backup & Recovery is a space we discussed earlier this year, with a TAM increasing to >$20B by 2028E from $16B today that is only helped by recent IT outages. The space has shifted into the "Gen-Z" age as enterprises upgrade their backup systems in order to support cyber-resiliency, multi-cloud usage, protect against ransomware, data compliance, and unstructured data and insights. Scaled privates to be cognizant of include Veeam, Cohesity (+Veritas now), and Druva. *Most exposed names: RBRK (OW, $72 PT), CVLT (N, $154 PT).*

James E. Fish
Sr. Research Analyst, Piper Sandler & Co.
612 303-6509, james.fish@psc.com

Harsh V. Kumar
Sr. Research Analyst, Piper Sandler & Co.
901 302-1646, harsh.kumar@psc.com

Quinton Gabrielli
Research Analyst, Piper Sandler & Co.
612 303-0306, quinton.gabrielli@psc.com

Robert Aguanno
Research Analyst, Piper Sandler & Co.
212 284-9345, robert.aguanno@psc.com

Braeden Kuether
Research Analyst, Piper Sandler & Co.
503 275-7725, braeden.kuether@psc.com

Related Companies:	Share Price:
AMD	125.02
AVGO	240.23
CSCO	58.52
CVLT	166.44
MRVL	112.25
NTNX	66.14
NVDA	130.39
PSTG	65.76
RBRK	72.97

INDUSTRY RISKS

Macroeconomic, IT Spending, Competition, ASPs, Architectural Shifts, & Valuations

Tech Decoded

Tech Decoded | Datacenter is Back in Vogue (Vol #1) - Compute, Storage, Silicon

James Fish
Sr. Research Analyst
Cloud Automation Software
+1 612 303-6509
James.fish@psc.com

Quinton Gabrielli
Research Analyst
Cloud Automation Software
+1 612 910-7533
Quinton.gabrielli@psc.com

Braeden Kuether
Technology Product Manager
+1 605 214-3696
braeden.kuether@psc.com

Harsh Kumar
Sr. Research Analyst
Semiconductors
+1 901 302-1646
Harsh.kumar@psc.com

Robert Aguanno
Research Analyst
Semiconductors
+1 212 284-9345
Robert.aguanno@psc.com

Table of Contents

Tech Decoded Report Links	Date
Future of Work: Opportunity Brewing in the Private Landscape	9/17/2024
B2B FinTech: Reshaping Products, Payments and Money Itself	7/16/2024
IoT Security - Dangerous When "Left to [Their] Own Devices"	2/22/2024
Cloud, Seeds & Machines Underpin AI Revolution in Agriculture	9/26/2023
Future of Work | Unlocking Workforce Potential Through Cloud & AI	7/25/2023
Generative AI | The Next $100B+ Tech Revolution	4/14/2023
Connected Software for a Better Built World	3/13/2023
New Innovations to Supercharge Vertical Software	2/21/2023
Automation 2.0 | Redefining Infrastructure and CX Operations	10/21/2022
Automation 2.0 | A New Era Redefining Operations & Efficiency	9/8/2022
IAM Renaissance | "Who are you? Who, who, who, who?"	5/23/2022

1. Back to Basics
What is a Datacenter?

What is a Datacenter?

A datacenter is a physical location that houses the processing, storage, and networking capabilities needed to run software applications and services such as web applications, databases, and virtual machines. The first datacenter was built in 1945 at the University of Pennsylvania for the purpose of the Electronic Numerical Integrator and Computer (ENIAC) that helped with artillery firing, weather forecasting, and weapon designs. There are four tiers of datacenters today: onsite datacenters, colocation, hyperscale / cloud datacenters, and edge datacenters (sometimes referred to as micro-datacenters or points-of-presence). The four tiers of datacenters deal with uptime and capabilities. Datacenters are actually typically owned and operated by large organizations (including cloud-players) or by co-location companies that lease out the space and provide network capacity, power, and cooling equipment vs. the tenants bringing their own IT equipment.

While there have been multiple evolutions of the datacenter, such as the change from mainframe computing to virtualized computing, arguably the biggest evolution to-date was the rise of public cloud providers in the 2010s that gave organizations the ability to access on-demand resources that scale with usage and significantly accelerated the pace of development and growth of software applications. However, with the recent advancements around artificial intelligence and machine learning, datacenters are once again undergoing a new evolution today resulting in “AI Datacenters” and “Datacenter Modernization” that is driving investment and investor interest back into the space.

Datacenter Types

1. **On-Premise:** Fully-owned & managed by the company, resulting in more customization but higher capex / upfront cost and limited flexibility.
2. **Colocation:** Large datacenter facilities (ex: Equinix, CoreSite, Digital Realty) that companies can lease space at for a fixed monthly cost for IT equipment.
3. **Hyperscaler:** Enterprises rent the underlying space and infrastructure, creating an “OpEx” spend over capex. This can be higher cost at scale but provides for the greatest flexibility.
4. **Edge:** Small datacenters that are closest to the end-user / network edge. This provides the best latency but limited scale.

Datacenter Tiers

1. **Tier-1:** 99.671% uptime. Basic capacity for IT systems’ support that has uninterruptible power, dedicated 24/7 cooling, and a backup power generator.
2. **Tier-2:** 99.749% uptime. Better maintenance + cooling.
3. **Tier-3:** 99.982% uptime. Greater data resiliency / redundancy.
4. **Tier-4:** 99.995% uptime. Contains isolated systems in-case of disruption with fully redundant systems.

Source: Piper Sandler Research, Amazon

Datacenter Stack Basics

Datacenter Equipment	• **Compute / Servers:** Physical machines that process data for applications/services that can be in a blade or rack configuration. CPU-based servers handle general purpose computing tasks, while GPU-based servers support accelerated computing for AI, 3D rendering, analytics, etc. • **Storage:** Where the data for applications / services sit in order to be processed. This can come in multiple forms (file, block, object) and types (array, software-defined, hyper-converged). • **Networking:** Comprised of systems (switches, routers, load balancers, etc.) that leverage chips, optics, cables, and operating systems that provide data movement between other IT equipment and connectivity to endpoints / end-users. • **Cyber-Security:** Includes firewalls, gateways, Web-App Firewalls, and other on-premise cyber-security solutions that prevent bad-actors and allow good-actors into applications & services based on policies / rules. • **Backup & Recovery:** Duplicates data and stores the data in a secure place in case of loss (hardware failure, software failure, natural disasters, cyber-attacks, human errors, etc.)
Datacenter Software & Services	• **Infrastructure Management / Hypervisor:** Software or firmware that allows an organization to run virtual machines that allow one endpoint / server to host multiple 'guests' by sharing resources like processing power and memory. • **Security Software:** Includes endpoint, virtual firewalls or gateways, Security Incident & Event Management (SIEM), data loss prevention, vulnerability management, micro-segmentation, etc. • **Network Performance Management:** Helps organizations oversee, secure, and optimize the network. • **Other Services:** Includes managed hosting, Infrastructure-as-a-Service, and GPU-as-a-Service.
Datacenter Utilities	• **Telco & Bandwidth:** Carriers provide network connectivity, redundancy, and bandwidth connections. • **Power Supply & Backup Generators:** Power is provided by utility companies through a main electrical service entrance. Datacenters utilize three-phase power (AC) that optimizes Power Usage Efficiency (PUE) and lowers wiring costs. Backup generators are needed on-site for power redundancy. • **HVAC & Cooling:** Cooling and climate control systems allow datacenter equipment to function in an optimal environment, making processing more efficient, extending the lifetime of components, and reducing errors / malfunctions. Cooling can leverage specialized HVAC systems or liquid-methods. • **Fire Suppression Systems:** Utilize dry-pipe sprinklers (pressured air / nitrogen) or a clean-agent. • **Physical Building Security Systems:** Fences, 911-alarms, video surveillance, personnel, etc. • **Real Estate:** Where the datacenter (whether owned or leased) sits.

Source: Piper Sandler Research

Mainframe to x86 to Containers | Impact on Datacenters

While cloud computing was pioneered in 2006 by AWS and was arguably the biggest impact to computing to-date, there have been three other major changes in the datacenter to be mindful of: 1) Mainframe computing; 2) Virtualized Computing; 3) Container-based infrastructure.

Mainframe Computing

Relevancy: 1940s-1990s

Computer used by large organizations for mission-critical applications and large data and transaction processing. Mainframes are known for their high reliability, availability, serviceability, and security capabilities.

Major Vendors*:* IBM, CA (Broadcom), BMC, Rocket Software, UNICOM, Micro Focus (OpenText), ASG, Unisys.

Virtualization

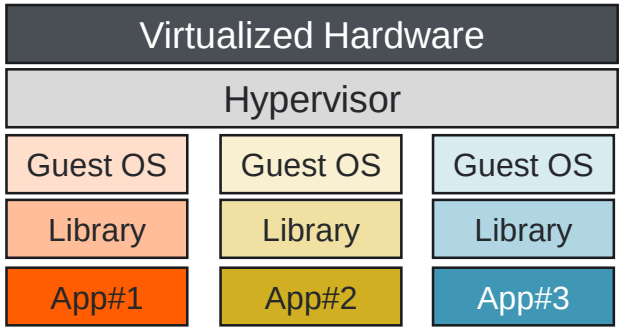

Relevancy: 1980s-2020s

Virtualization of hardware can occur across desktops, storage, & compute. The software is separated from the underlying hardware, helping to better leverage hardware across apps, and given data and compute was becoming more distributed.

Major Vendors*:* VMware, Red Hat (IBM), Citrix (CSG), Nutanix, Cisco, Microsoft, & Oracle.

Containers

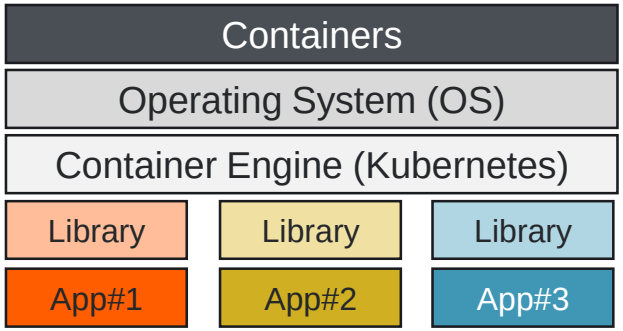

Relevancy: 2010s-Today

Containers are lightweight packages of application code, specific programming languages, & libraries that virtualize the OS and run anywhere, allowing for teams to deploy software efficiently and quickly at a very large scale.

Major Vendors*:* Google Kubernetes, Red Hat OpenShift (IBM), VMware Tanzu (Broadcom), Docker, Amazon EKS, Azure KS, D2iQ (Nutanix), DigitalOcean, Rancher (SUSE).

Source: Piper Sandler Research, Broadcom, Google Kubernetes

The Movement of Bytes | Example of Data's Journey in a Datacenter

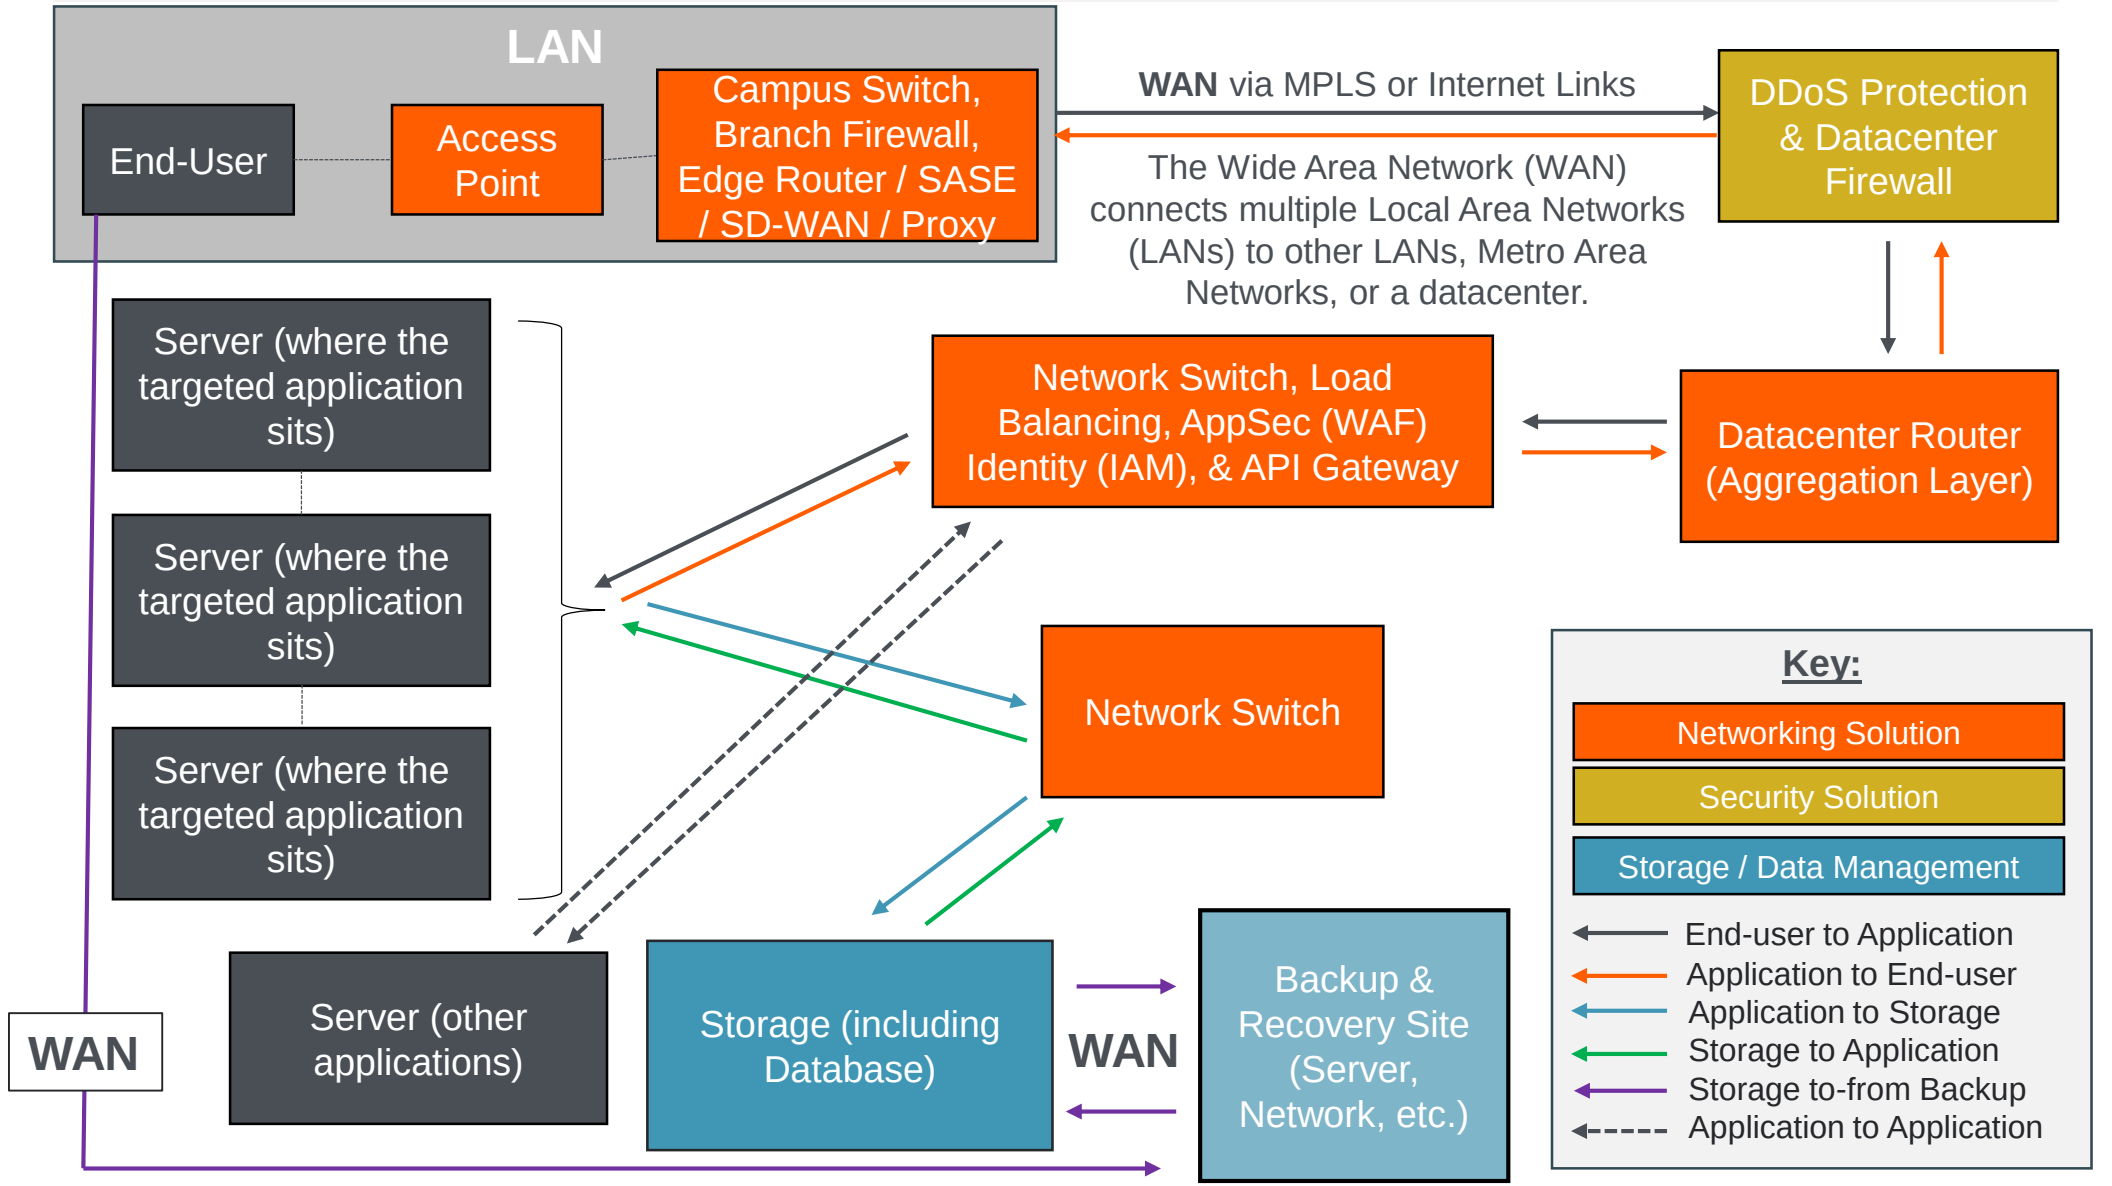

Source: Piper Sandler Research; Solutions can be multi-part (ex: SASE is both networking & security).

The Movement of Bytes | Datacenter – North v. South, East v. West

North-South Traffic

East-West Traffic

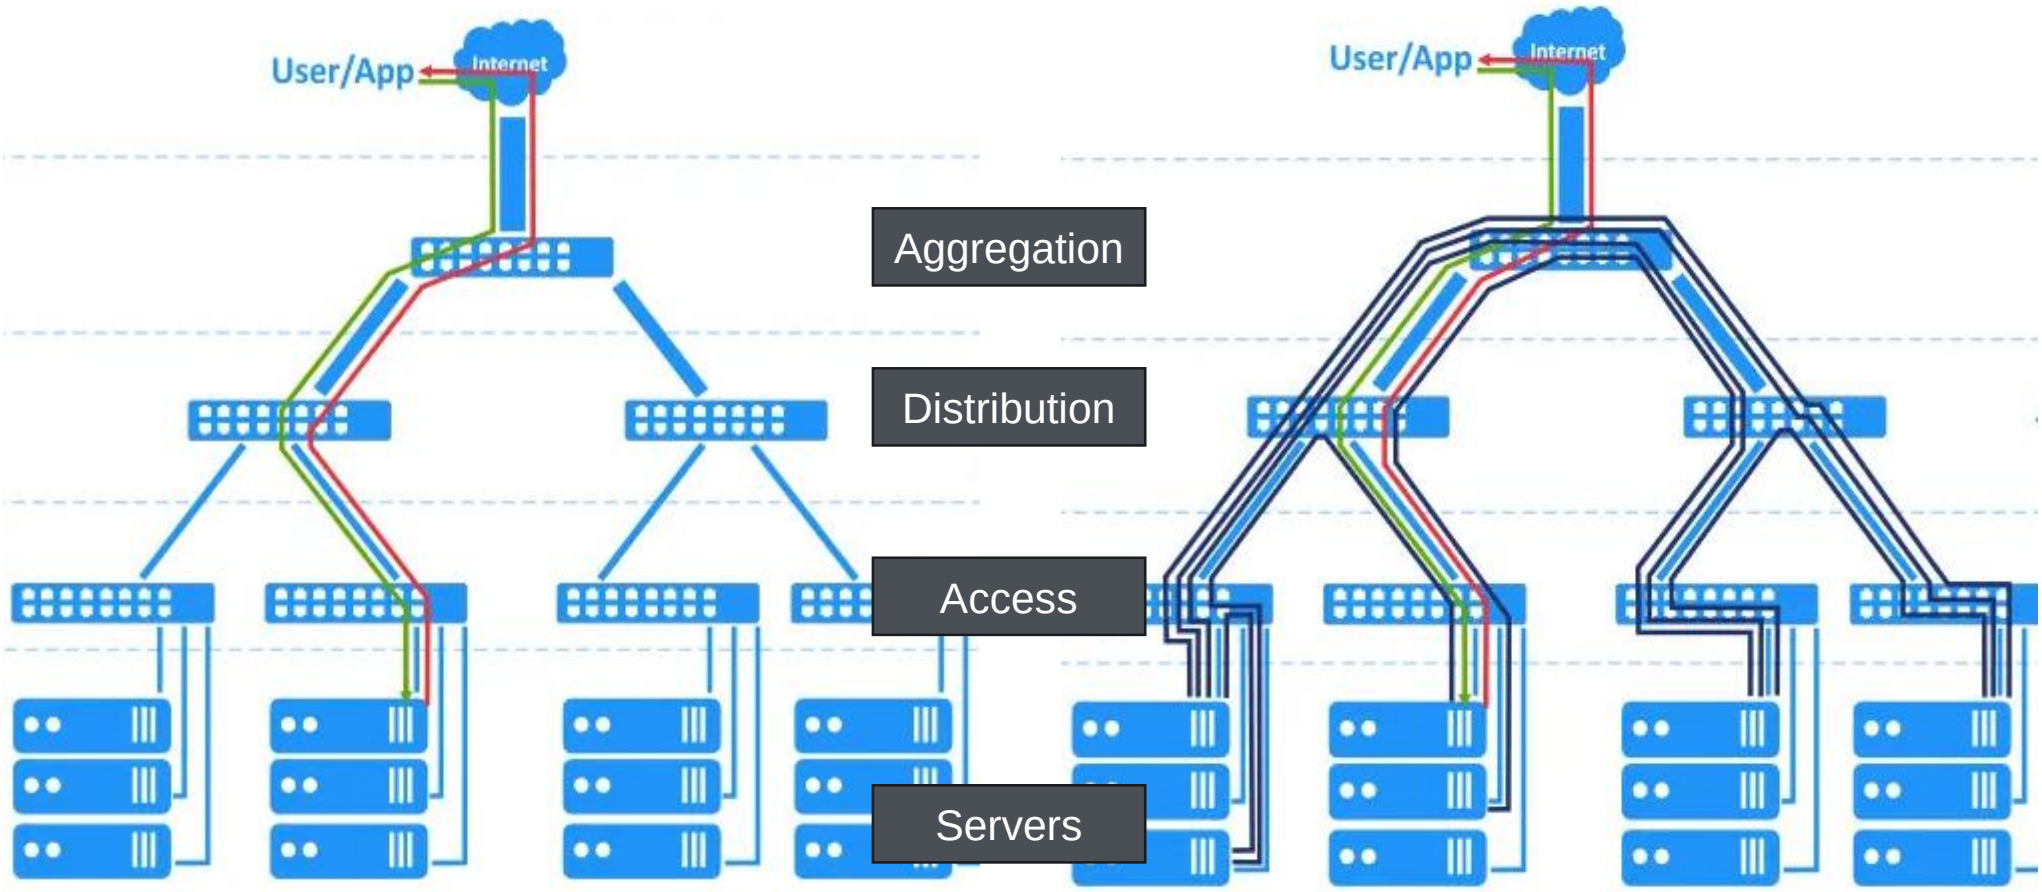

Source: Piper Sandler Research, Cisco

2. Datacenter Trends & Catalysts

Datacenter Sites | +1.3% CAGR for Larger Sites Thru '27

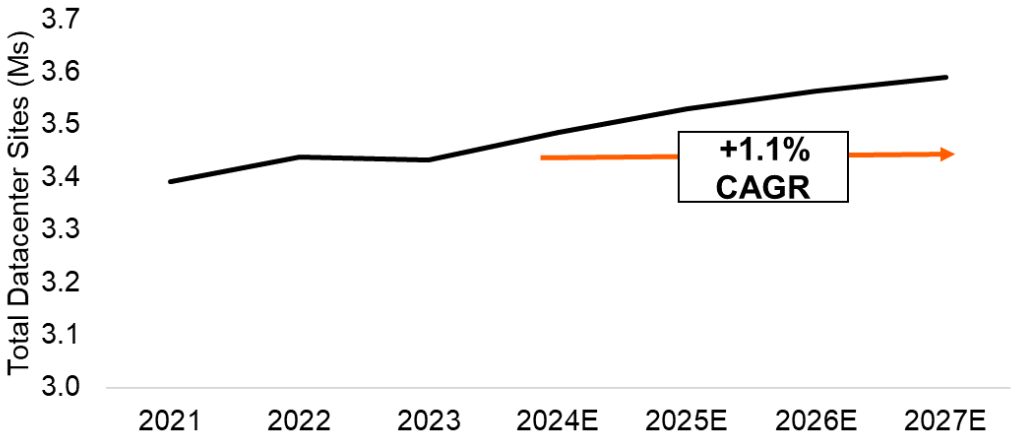

3.43M
Total Datacenter Sites in 2023

3.59M
Total Datacenter Sites in 2027

23K
Total Mid-to-Large Sites in 2023

24K
Total Mid-to-Large Sites in 2027

3.41M
Total Single-to-Small Sites in 2023

3.57M
Total Single-to-Small Sites in 2027

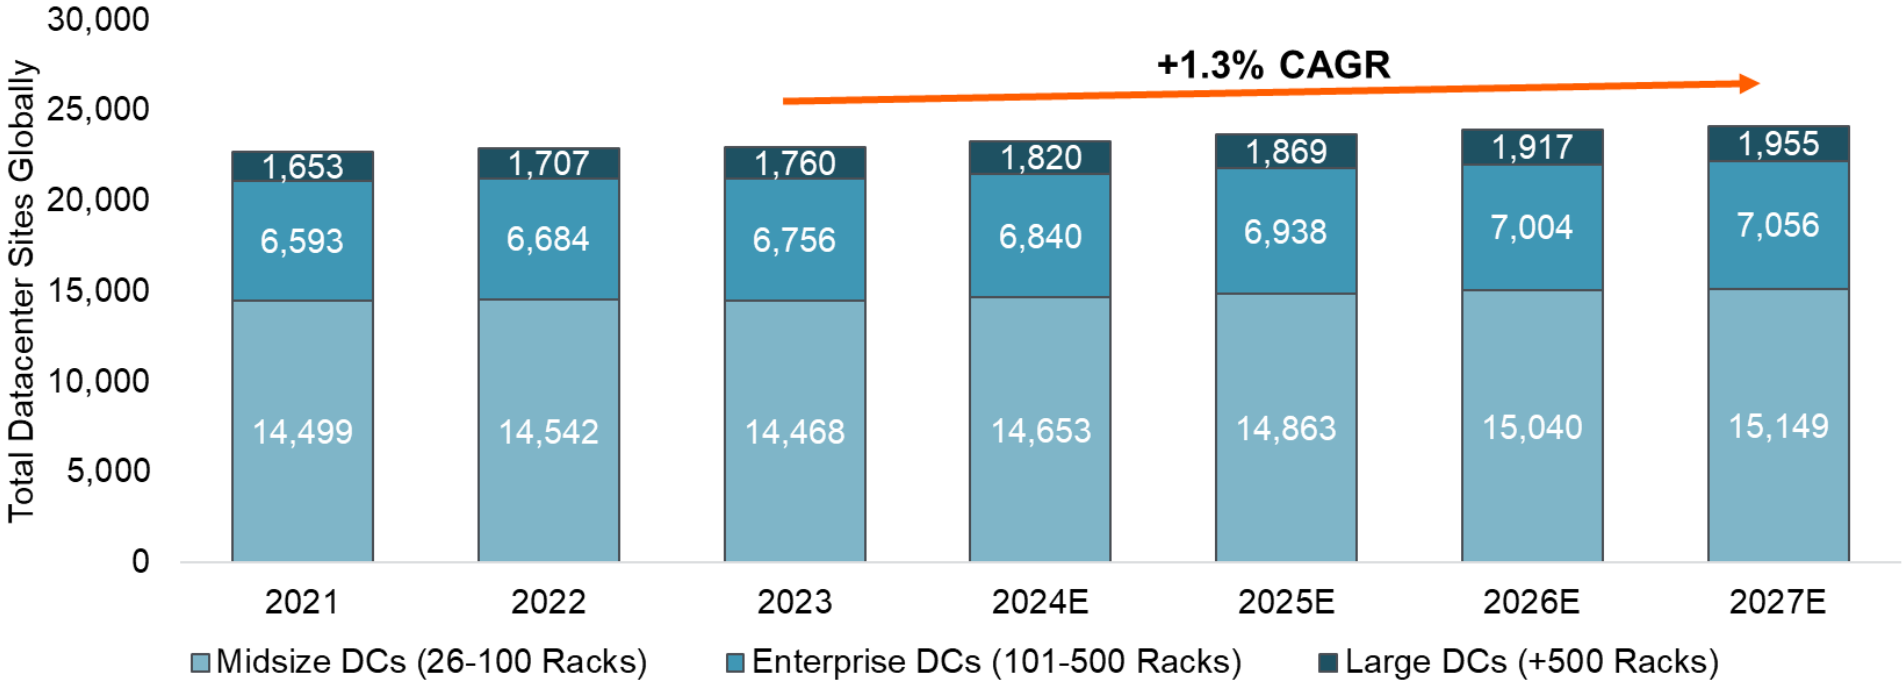

Source: Piper Sandler Research, Gartner

On-Premise vs. Cloud IT Spending

On-Prem the Largest Spending Cohort, Though Public Cloud Taking Share of Overall Budgets

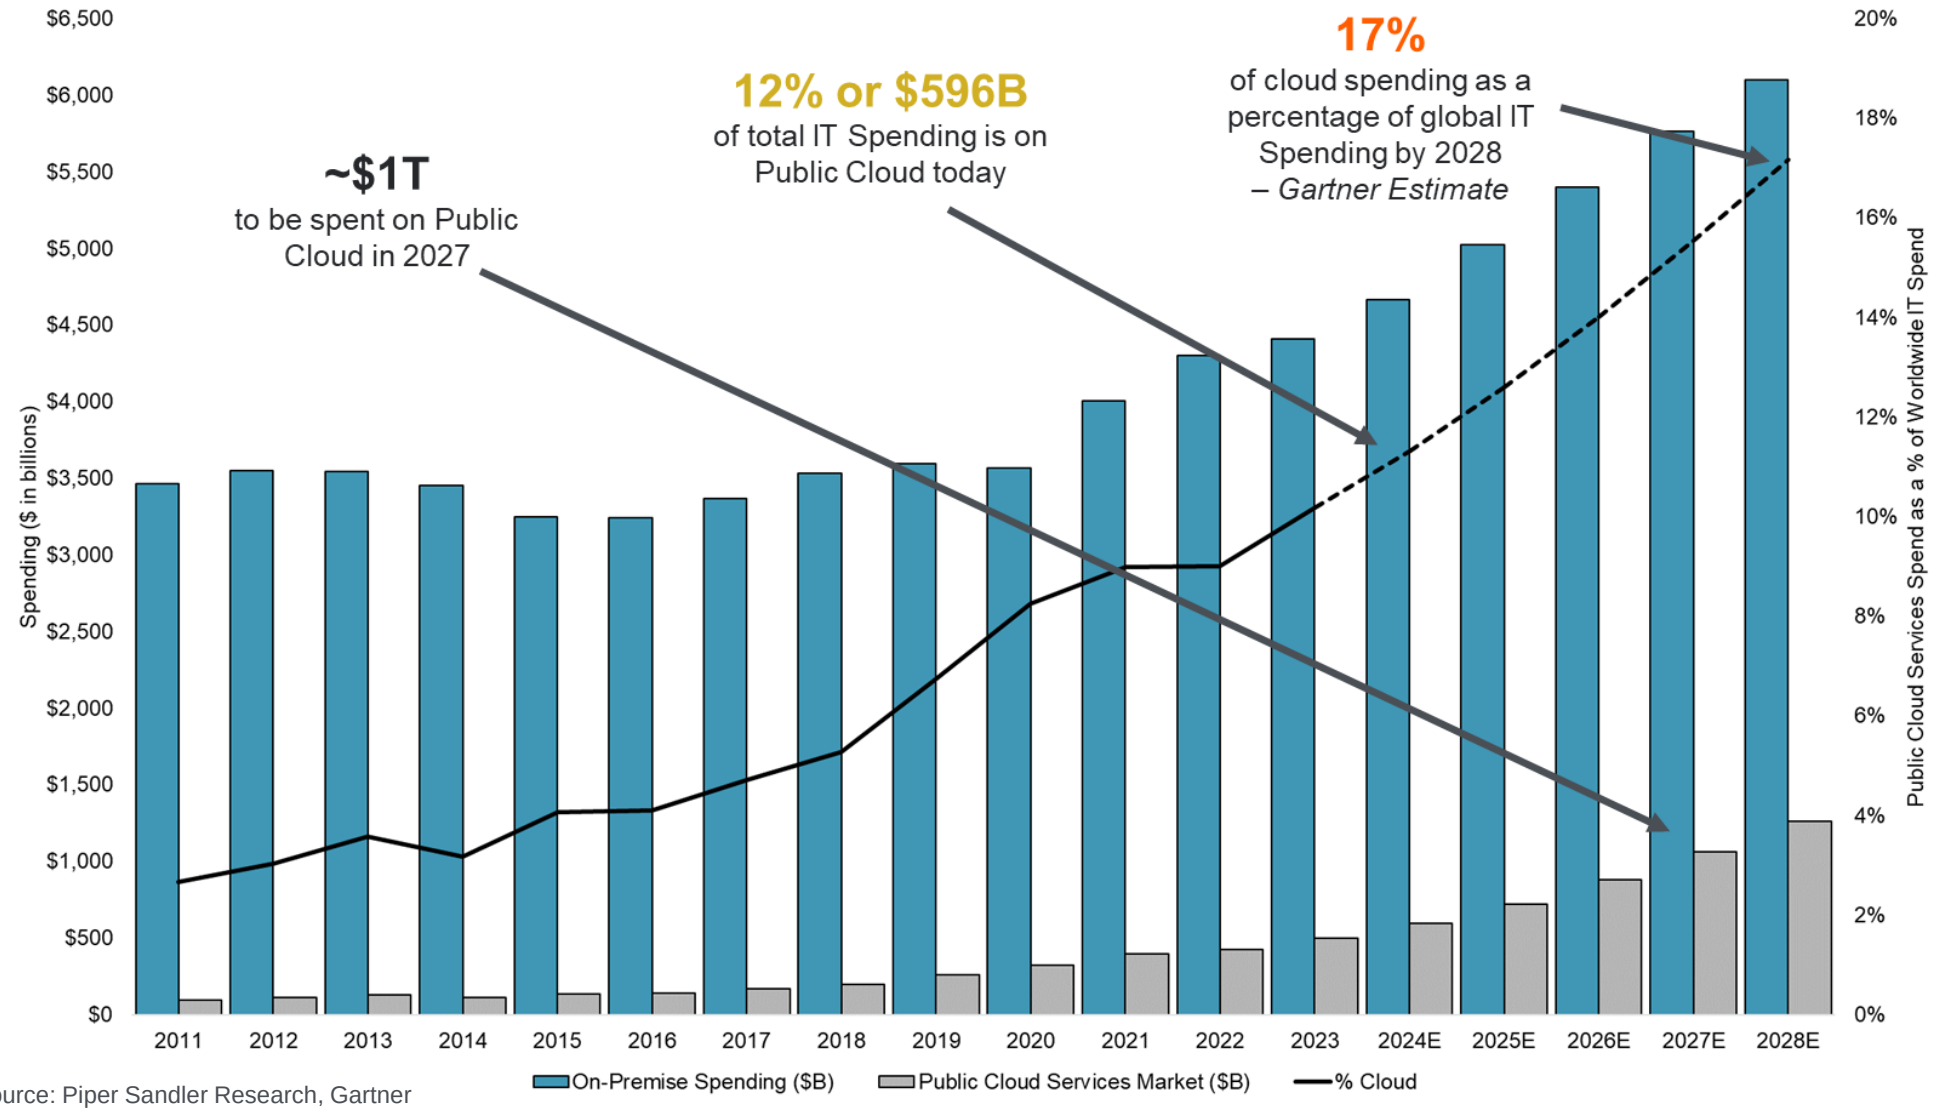

Source: Piper Sandler Research, Gartner

CIOs Shifting Towards Public Cloud, But Hybrid Here to Stay

On average, CIOs estimate they are currently running just 29% of applications in the Public Cloud today (with apps evenly spread across On-Premise, Public Cloud and Private Cloud environments). On-Premise is expected to see the largest decline in share (-2%) from 2024 to 2026, with Private Cloud (~flat%) and Public Cloud (+2%) seeing the largest increases.

We believe the increased preference for Public Cloud environments will continue to drive Cloud IaaS Provider Capex dollars, though we note the lower expected mix shift corroborates our view that hybrid and on-prem infrastructure is here to stay in a major way through at least this decade.

Cloud Shift Delayed vs. Prior Survey, But Cloud Still the Increasing Preference

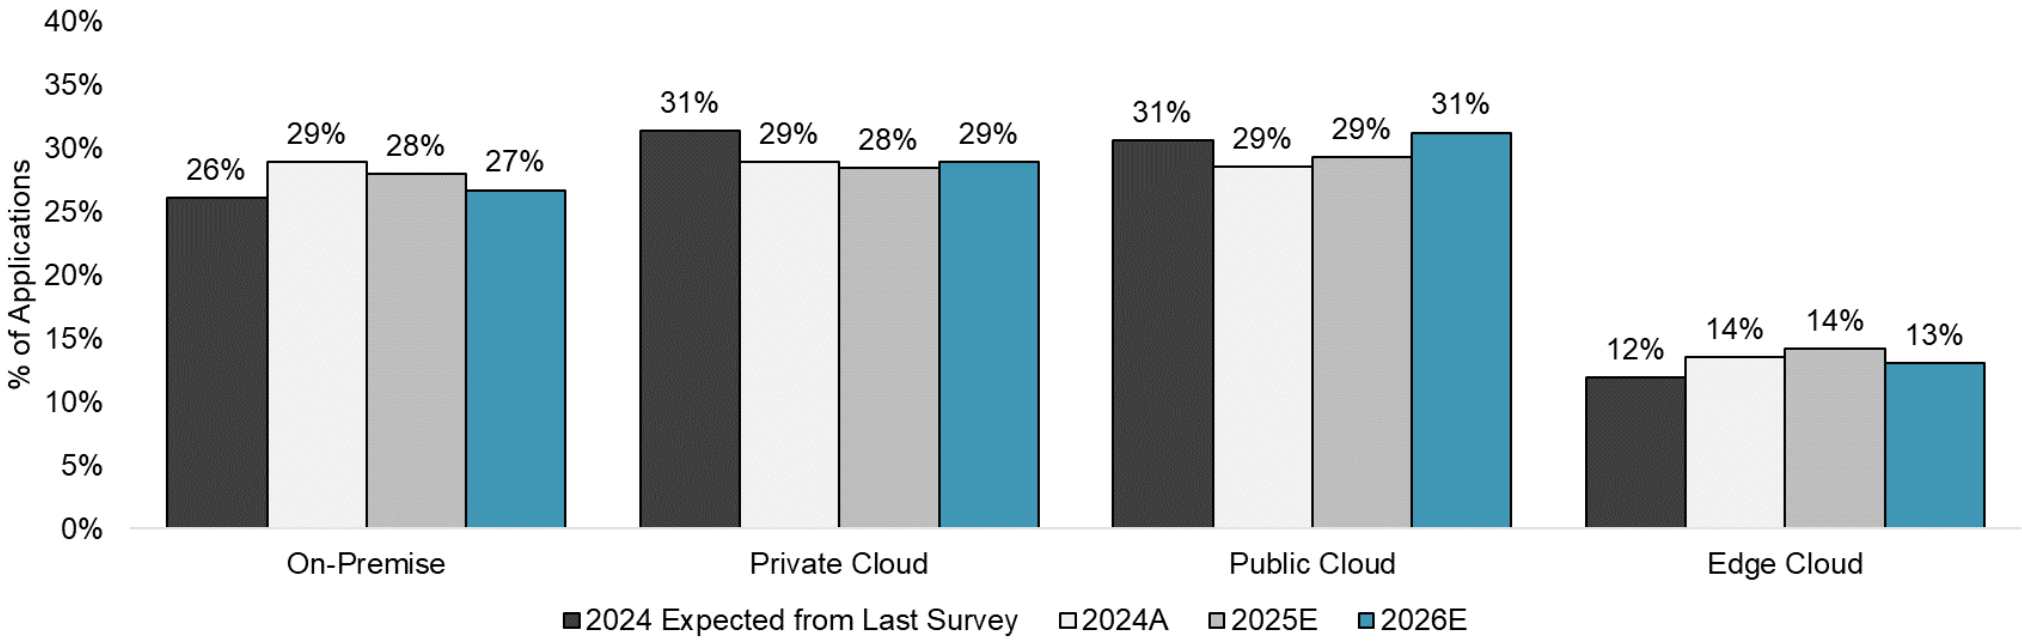

Source: Piper Sandler Research

Cloud Capex Spending & Refresh Opportunities

+$330B of Capex

is expected in 2024 from Cloud Providers, which is expected to grow 18% in 2025. Cloud Service Provider (CSP) spending makes up the majority of Total Service Provider capex, and is growing faster than its Telco-Service Provider counterparts (+21% in 2025 vs. SP at +5%)

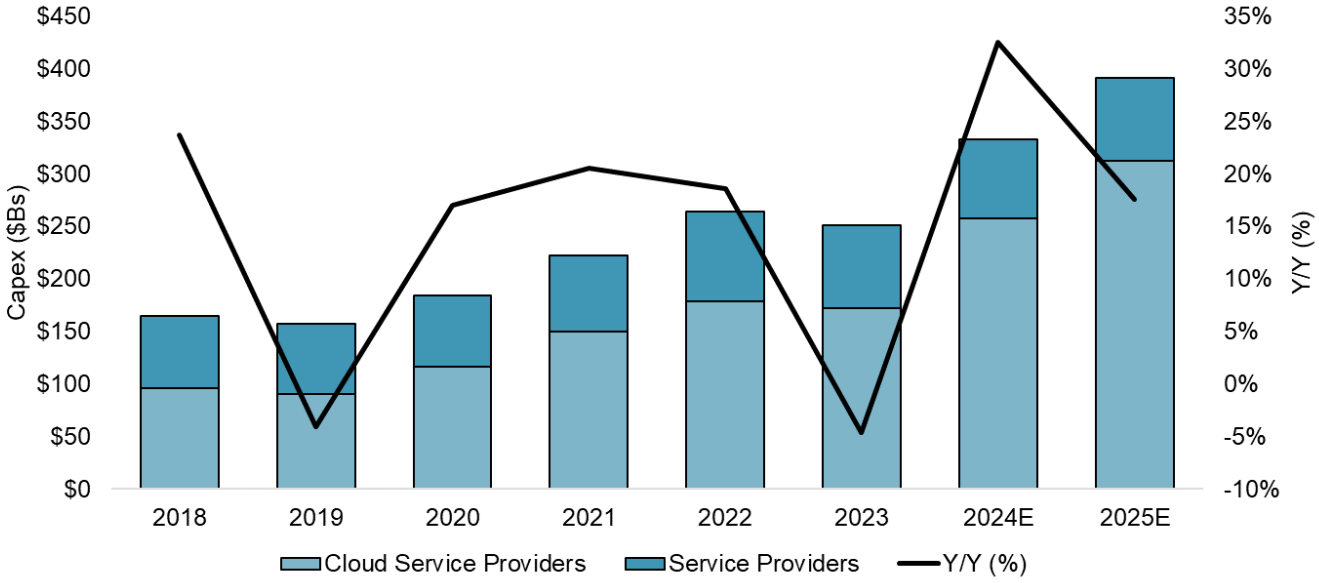

Our recent CIO survey called out a **potential 2025 refresh of on-premise systems**, though 2026 is likely the larger refresh opportunity. Datacenter firewalls are anticipated to see the highest amount of refresh in 2025, while Servers and Datacenter Switches are more delayed.

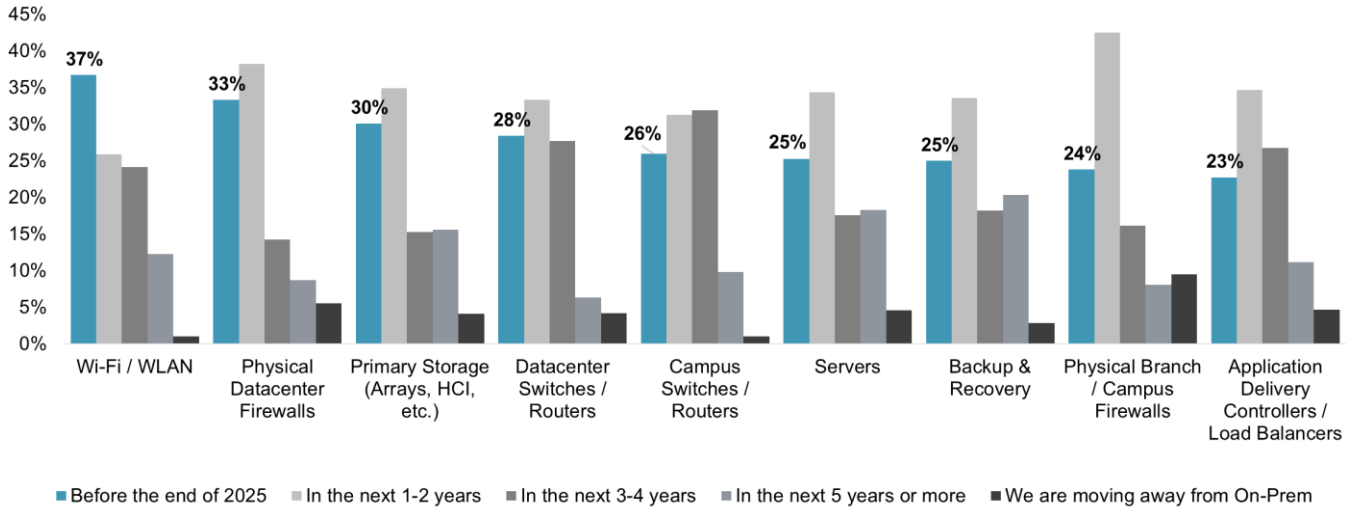

Source: Piper Sandler Research, Gartner; Enterprise Capex = Russell 2000 Capex + S&P 500 capex, excluding Amazon, Google, Microsoft, and Meta.

U.S. Government Investing in Critical IT Infrastructure

Planned IIJA Investments ($B)

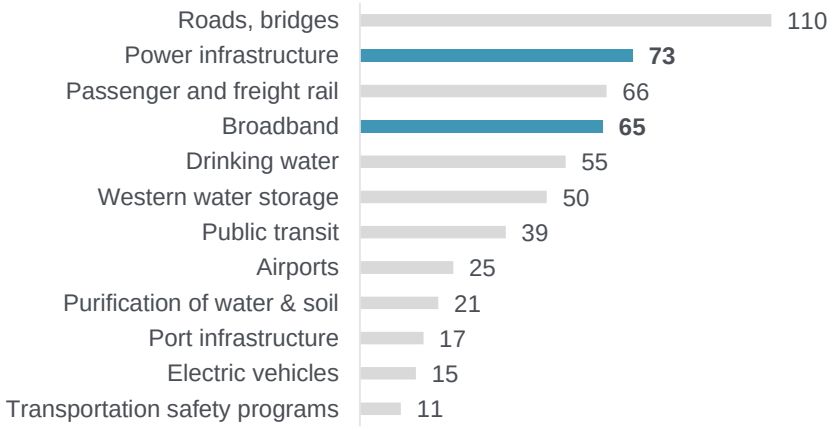

$73B
Investment in power infrastructure

$65B
Investment in broadband

Funding Allocation from CHIPS Act

Company	Amount	Factory Locations
Intel	$8.5B	Arizona, New Mexico, Oregon
TSMC	$6.6B	Arizona
Samsung	$6.4B	Texas
Micron	$6.14B	New York
Global Foundries	$1.5B	New York, Vermont
Microchip Technology	$162M	Colorado, Oregon
Polar Semiconductor	$120M	Minnesota
BAE Systems	$35M	New Hampshire

$53B
Investment in U.S. production of semis

$1.5B
Investment in U.S. wireless supply chains

The Infrastructure Investment and Jobs Act (IIJA) allocates $138B across power infrastructure (including renewables) and broadband as the government looks to address rising Datacenter energy consumption, and expand the public's access to high-speed internet as products and services become increasingly digitized. Similarly, the CHIPS act is investing $53B to bring semiconductor manufacturing back to the U.S., mitigating significant supply-chain risk arising from a complete dependence on Taiwan.

In September, the White House convened leaders from hyperscalers, AI companies, Datacenter operators, and utilities to discuss U.S. leadership in AI, and announced a number of new initiatives:

- AI Datacenter Infrastructure Task Force
- Faster Datacenter Permitting
- DoE AI Datacenter Engagement Team
- DoE supporting projects to repurpose closed coal sites.

Source: Piper Sandler Research, NPR, U.S. Department of Commerce

Datacenter Modernization | Upgrades Across the Stack

The Datacenter is modernizing across the full stack to keep up with increasing internet/software demand, meet modern application requirements, improve energy efficiency (offsetting growing consumption levels), reduce complexity and burden on IT and Datacenter operations, and enhance security posture / infrastructure management across hybrid, multi-cloud environments.

Legacy Datacenter Stack	*Modern Datacenter Stack*	*Benefits*
CPU Servers	Mix of CPU & GPU Servers	Improved capabilities and efficiency on high-performance workloads.
CPU Infrastructure Workloads	DPU / SmartNIC Infrastructure Workloads	Offloading infrastructure workloads to purpose-built hardware eliminates CPU overhead.
10 - 100Gb/s Switches	400 - 800Gb/s Switches	Higher throughput removes bottlenecks for high-performance workloads.
Routing, Load Balancing, Security	SASE / SD-WAN	Zero-trust security, consistent policy, greater visibility, and ease-of-management across hybrid, multi-cloud environments.
3-Tier Storage	Software Defined Storage / Hyperconverged Infrastructure	More scalable for rapid growth environments, enables higher levels of automation.
HDD	SSD / all-flash	Removes storage bottlenecks for high-performance applications, better space efficiency.
Traditional Backup & Recovery	“Gen-Z” Backup with Data Protection & Resiliency	Zero downtime, rapid recovery with zero data loss, improved compliance and security posture.
Legacy HVAC	Hot / Cold Aisles, Liquid Cooling	Significantly lower power consumption drives cost reductions.
Non-Renewable Energy	Renewables	Shift to purpose-built renewable projects reduces cost and strain on public utilities.

Source: Piper Sandler Research

And Gen AI Models Are Still Getting Better…

Concerns about hitting a plateau in model performance have been growing over the past year, as the rate of model releases has slowed compared to the initial rush post ChatGPT, and model size has not necessarily yielded proportional gains in performance (ex: Llama 3.1 405B only 6 points higher on quality index than Llama 3.1 70B). We continue to believe that Gen AI is still in the early innings, and developers will unlock new methods to solve increasingly complex problems, reduce model hallucination, and improve speed. OpenAI's recently launched o1-preview model shows a stark improvement in performance (albeit at a significantly higher cost), and the company emphasizes allowing the model to "think" for a longer time on more complex problems. We think o1 could represent a new frontier of Gen AI models that use a collection of AI agents in a "chain of thought" to deliver a more accurate solution on the most complex problems.

Artificial Analysis Quality Index

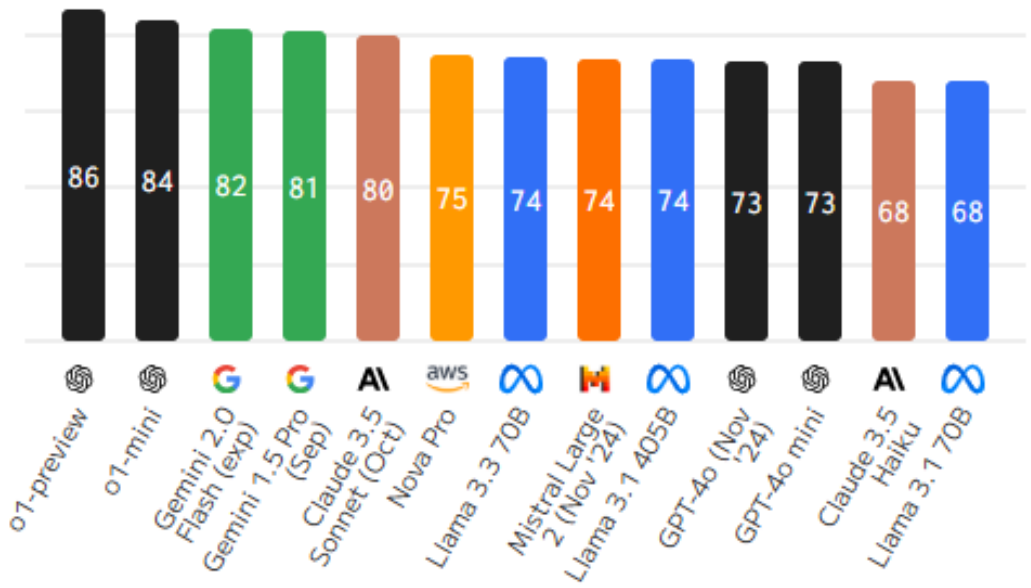

OpenAI

Model	Quality Score	Release date
GPT-3.5 Turbo	52	Nov 6 2023
GPT-4 Turbo	74	Nov 6 2023
GPT-4o	73	May 13 2024
o1-preview	86	Sep 17 2024
o1	??	??

Source: Piper Sandler Research, Artificial Analysis

…and AI is Gaining Steam in Other Areas

Uptake of traditional AI and new emerging AI use-cases should continue to support the transition to the cloud and growth in Datacenter. Vertical applications of AI stand to create new revenue opportunities and reduce costs across transportation, pharmaceuticals, manufacturing, defense, agriculture, etc. To accelerate these AI initiatives, C-suites have adopted a more aggressive posture on AI-related investments, including increased spending on the underlying infrastructure, co-locations, and cloud services.

Self-Driving

Waymo One available in San Francisco, Phoenix, with plans to expand to Los Angeles and Austin.

Aurora's Driver as a Service model could deliver TCO benefits to the ~$1T U.S. trucking market.

Robotaxi event…

- *Very large TAM*
- *New transportation opportunities*
- *Significantly improved safety*
- *Lower cost structure long-term*

Drug Discovery

Using AI, Pfizer reduced computational time to design the Covid-19 drug Paxlovid by 80%-90%.

Schrödinger offers a computational platform to accelerate the discovery of new therapeutics and materials.

Exscientia designed the first AI-based cancer treatment with a prospective clinical trial showing improvement.

- *Very large TAM*
- *New drugs and clinical therapies*
- *Better approval success rates*
- *Faster time-to-market*

Manufacturing

Predictive maintenance systems at Frito-Lay plants added 4,000 hours a year of manufacturing capacity.

BMW uses AI to identify and correct misplaced studs, leading to cost savings and more efficient labor force.

Siemens uses AI and digital twins to increase production and lower carbon footprint 50% at its Erlangen factory.

- *Very large TAM*
- *Improved quality control*
- *Less downtime from maintenance*
- *Work shifted to high value tasks*

Source: Piper Sandler Research, Company Websites, Forbes, Wall Street Journal, CNBC

The Rise of GPU as a Service

Gartner estimates 15% of all AI end-user spending will be on IaaS platforms by 2028. This is through both a combination of AI enabled software and services, but also due to the rise in GPU-as-a-Service (GPUaaS). Today, the largest cloud vendors aren't directly monetizing their GPU buildouts (as they are utilizing the fleet of GPUs for internal model building and tuning), leaving room for private players like Coreweave and Lambda to service the upper-end of the market. Public players like DigitalOcean and Akamai-Linode do have basic functionality here, but it is more targeted towards the lower-end of the market.

SMB / Midmarket Focused Vendors

DigitalOcean
Akamai
linode
VULTR
CLOUDFLARE Workers

Enterprise Focused Vendors

CoreWeave
Lambda
APPLIED DIGITAL
NORTHERN DATA

Hyperscalers

aws
Microsoft
Google
Meta
ORACLE

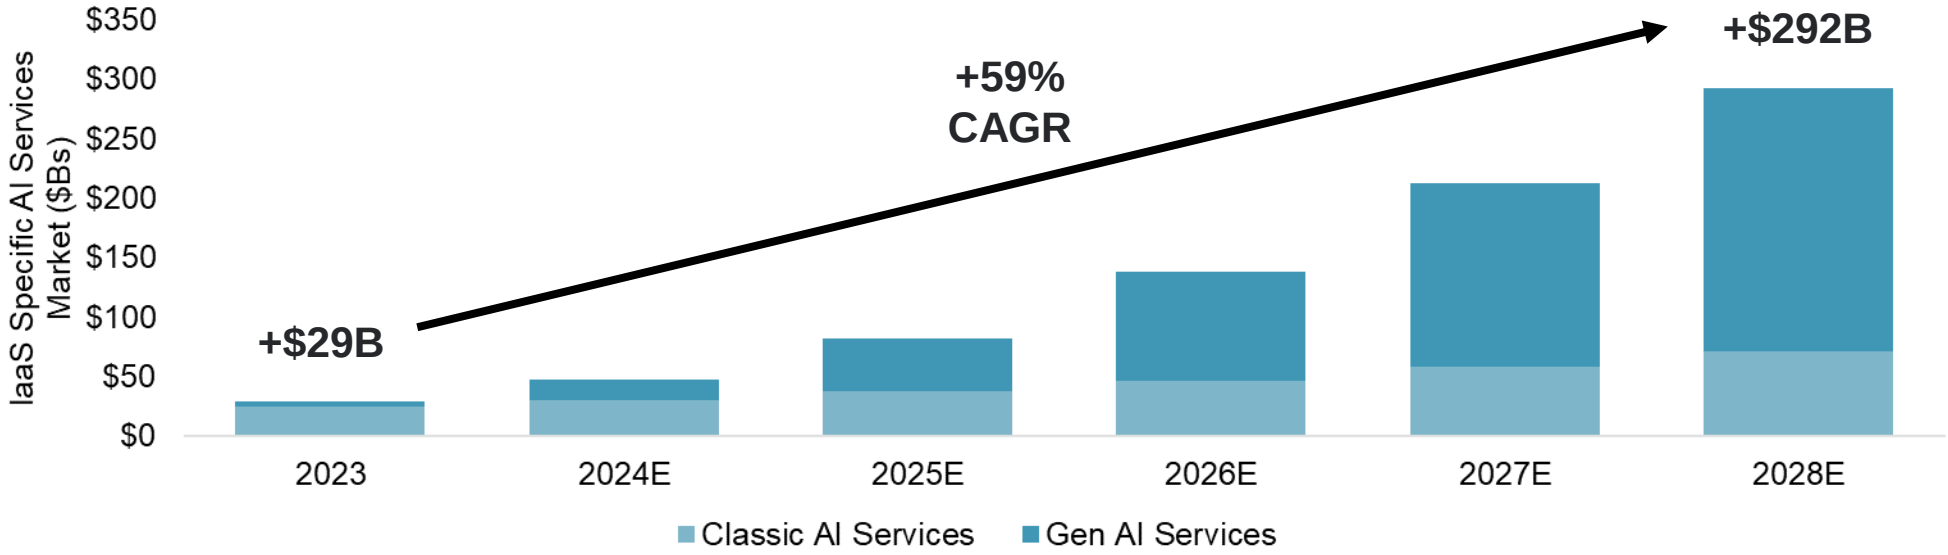

Source: Piper Sandler Research, Company Website, Gartner

Consolidation | Desire to "Do All" Again

We are increasingly seeing the largest vendors in the space accelerate the "do-it-all" mentality across Networking, Servers, Storage and the underlying Chips / Silicon (essentially owning exposure across the datacenter value chain). Vendors like Huawei already have functional capabilities across the stack, while vendors like Arista have seen the most success focusing solely on one sub-segment (networking). We expect the increased focus on datacenter investments would continue to drive a portfolio mentality, and increase consolidation in the space (especially with an Executive branch that may be more open to consolidation).

Company	DC Networking	Servers	Storage	Chips / Silicon
ARISTA	✓			
CISCO	✓	✓	/	✓
DELL	/	✓	✓	
Hewlett Packard Enterprise	✓*	✓	✓	
HUAWEI	✓	✓	✓	✓
NVIDIA	✓			✓

Source: Piper Sandler Research (✓* - via JNPR acquisition, / = limited portfolio), Company Websites,

Datacenter "Arms Race" Recent Examples

BROADCOM	vmware	→	*Cloud computing, networking, and security platform; leading VM solution*	→ **$69B** Acquisition price on 11/22/2023
IBM	HashiCorp	→	*Leading IaC solution with full infrastructure lifecycle management*	→ **$6.4B** Acquisition price on 4/24/2024
AMD	zt Systems	→	*Complex compute, storage, and accelerator solutions provider*	→ **$4.9B** Acquisition price on 8/19/2024
LUMENTUM	cloudlight	→	*Optical Transceiver Products*	→ **$750M** Acquisition price on 11/07/2023
NVIDIA	run:ai	→	*GPU orchestration software*	→ **$700M** Acquisition price on 4/24/2024
CISCO	Acacia	→	*Optics*	→ **$4.5B** Acquisition price on 3/1/2021
Hewlett Packard Enterprise	JUNIPER NETWORKS	→	*Network switching & routing vendor* *Network switching & routing vendor*	→ **$14B** Acquisition price on 1/09/2024

Source: Piper Sandler Research, TechCrunch, Reuters, Company Press Releases

Typical Constraints | Land

In looking at the typical constraints of datacenter sizing, land tends to be one of the most talked about "bottlenecks", as large-scale datacenters can be anywhere from 20k to +10M square feet. Gartner estimates there are more than 3.4M datacenters today (with the mix of "Large" continuing to grow). Utilizing a weighted average, Gartner estimates over 163M square feet are **already** in use by datacenter sites, though we believe this is overly conservative considering the massive size of the largest DC's

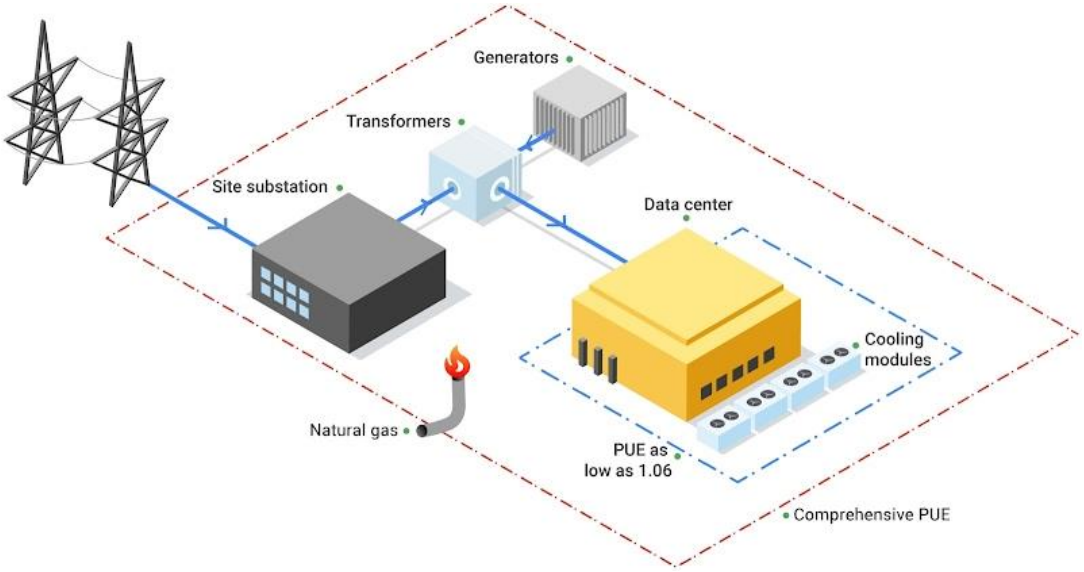

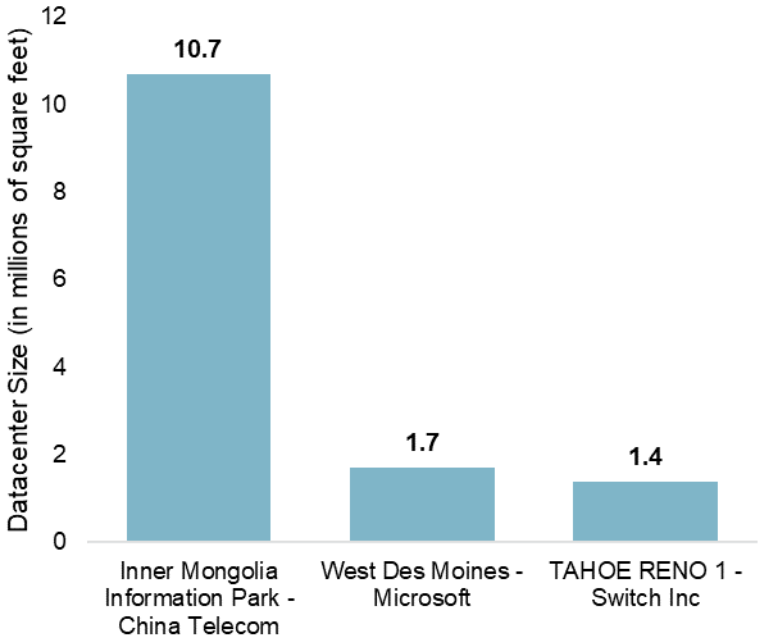

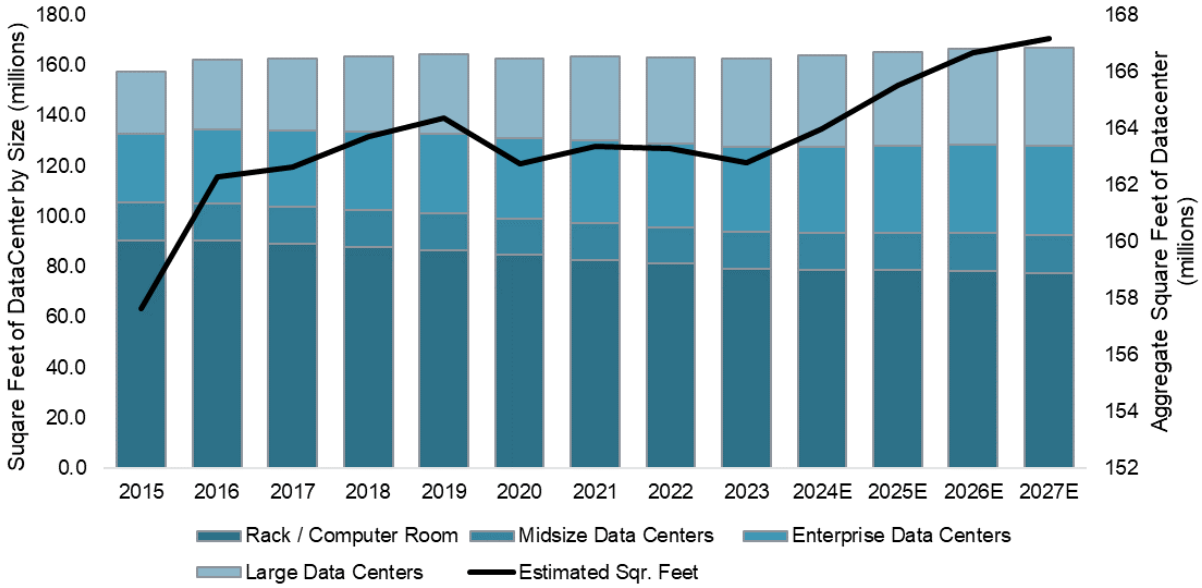

Source: Piper Sandler Research, Gartner, Company Reports

Typical Constraints | Power

The International Energy Agency estimates that data centers, crypto, and AI consumed about 460 TWh of electricity worldwide in 2022 (2% of total electricity demand), with projections this could more than double to more than 1,000 TWh in 2026. The IEA estimates computing accounts for ~40% of electricity demand of a datacenter, cooling requirements to achieve stable processing efficiency make up another 40%, while the remaining 20% comes from other associated IT equipment.

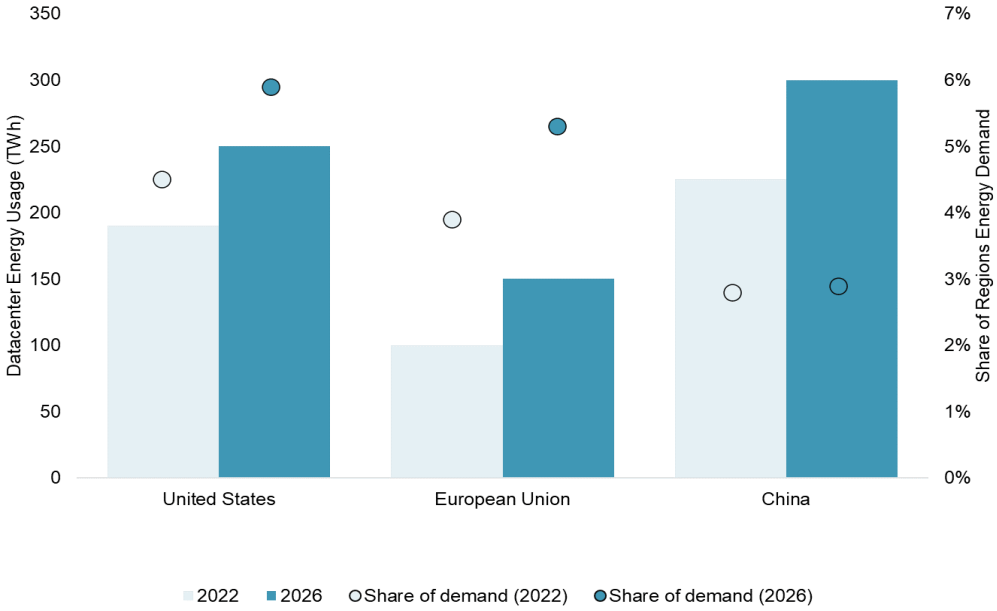

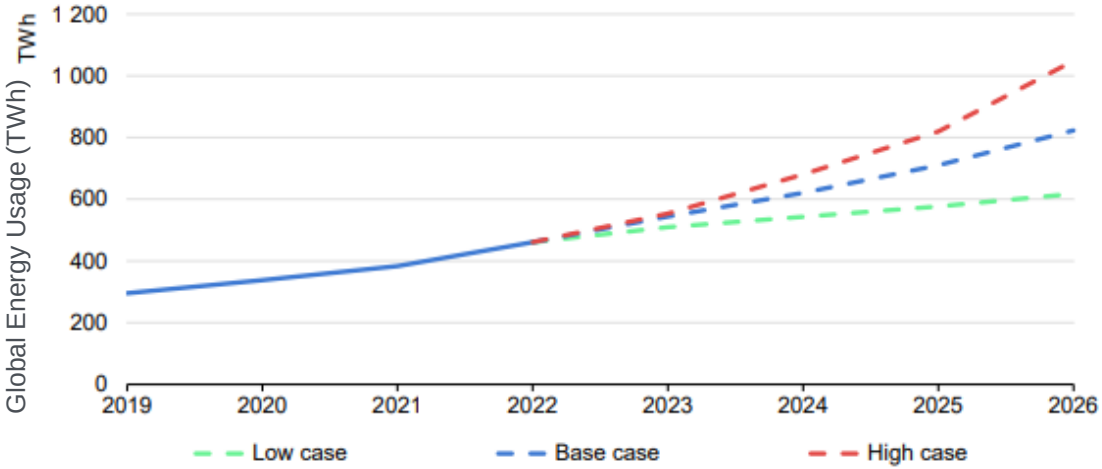

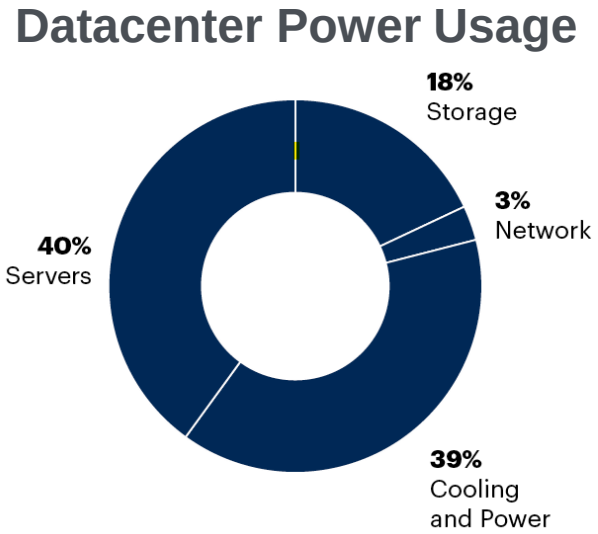

Source: Piper Sandler Research, US Energy Information Administration, Gartner

Typical Constraints | Cooling

Datacenter cooling is an increasing focus for most large scale operators, as Microsoft estimates that for every 1 megawatt hour of electrical energy consumed, 1 MWh of heat energy is produced. Because systems within the datacenter will fluctuate performance with the heat, Datacenter operators try to keep the temperature within 64.4-80.6 degrees Fahrenheit. To reduce heat, datacenters can divert and transfer heat generated, or cool the underlying systems with chilled water or air. Increasingly, operators are looking to utilize liquid cooling for high-density server-cooling problems due to water's 3,000x better conduction of heat relative to air (so less energy is required for the same heat removal).

Liquid & Immersion Cooling Still Early

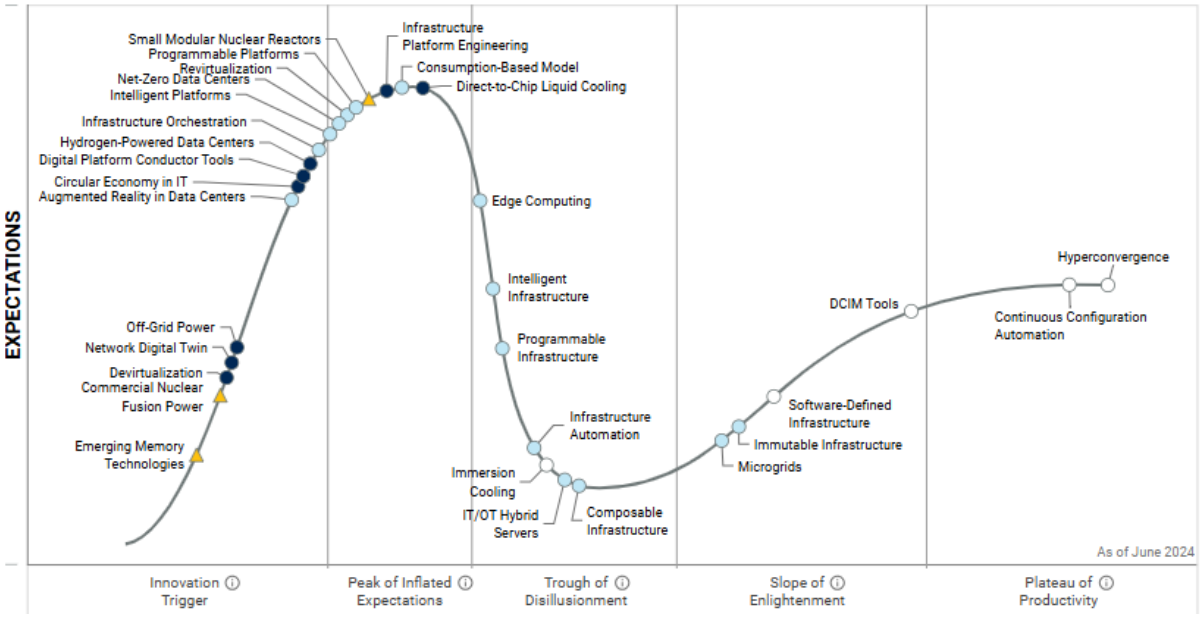

Water Saving Technologies

- Thermal Storage Technologies
- Side Stream Filtration
- Reverse Osmosis Treatment Systems
- Direct Liquid Cooling Systems

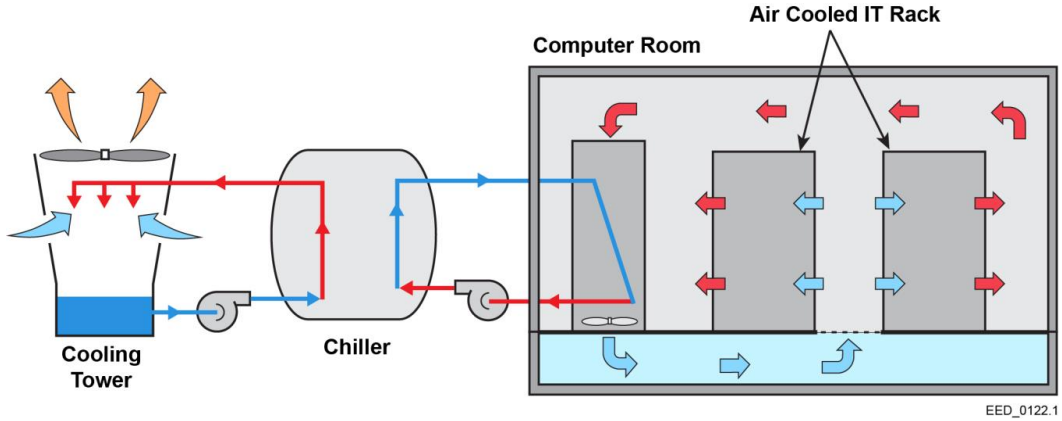

Source: Piper Sandler Research, Department of Energy, Gartner, Microsoft

3. +$500B of Datacenter Solution Spending by 2028

Datacenter Solution TAM | +$500B by 2028

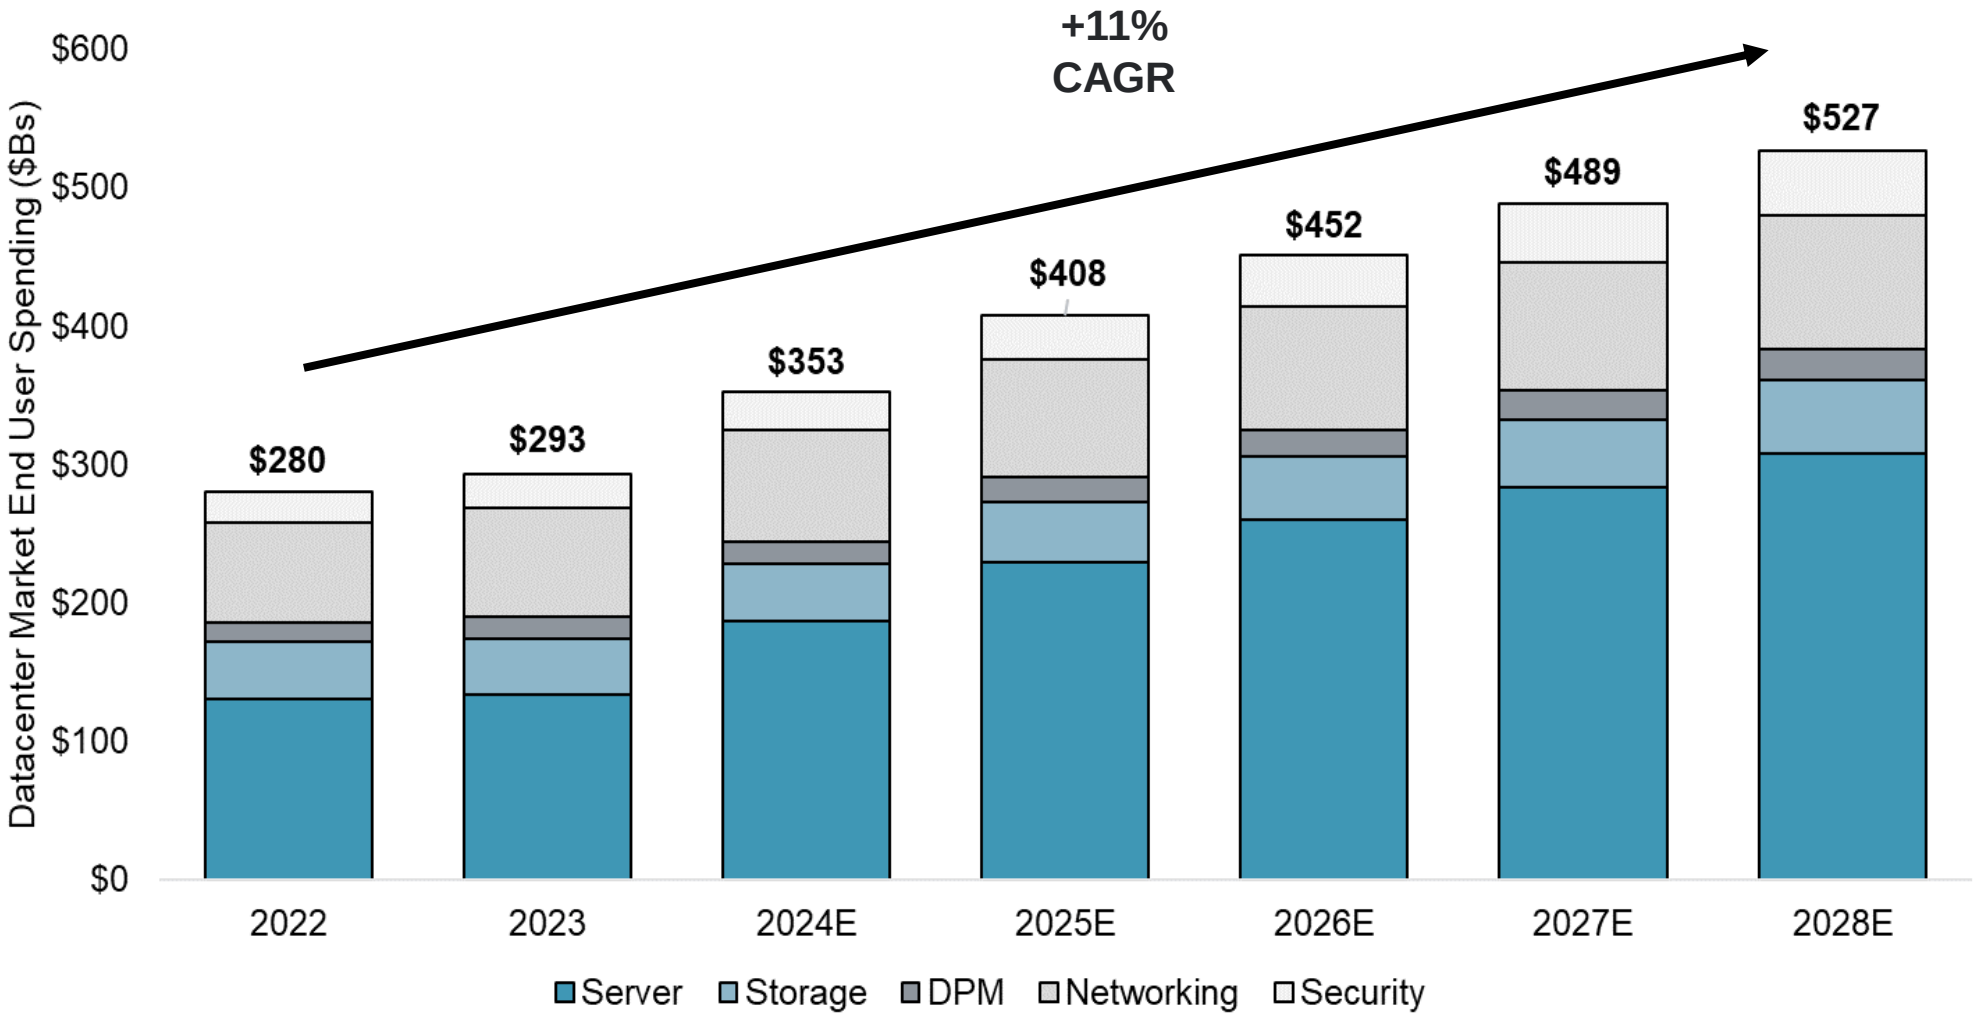

Source: Piper Sandler Research, Gartner

Server TAM | +$300B by 2028 as AI Takes Majority of Share

Gartner estimates Server spending will be ~$187B in 2024, growing at a 17% CAGR to over $300B by 2028. Importantly, Gartner expects Service Providers will allocate **>70%** of their server hardware budgets to AI servers in 2025 and 2026, accounting for the majority of spending. This AI server spending will be substitutive to server hardware budgets for more than 80% of enterprise buyers during that time frame.

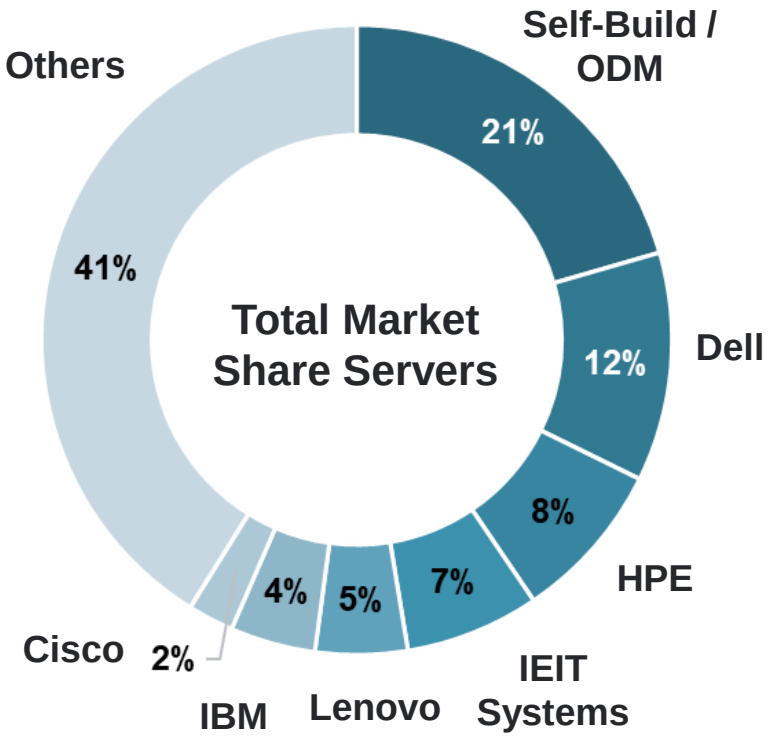

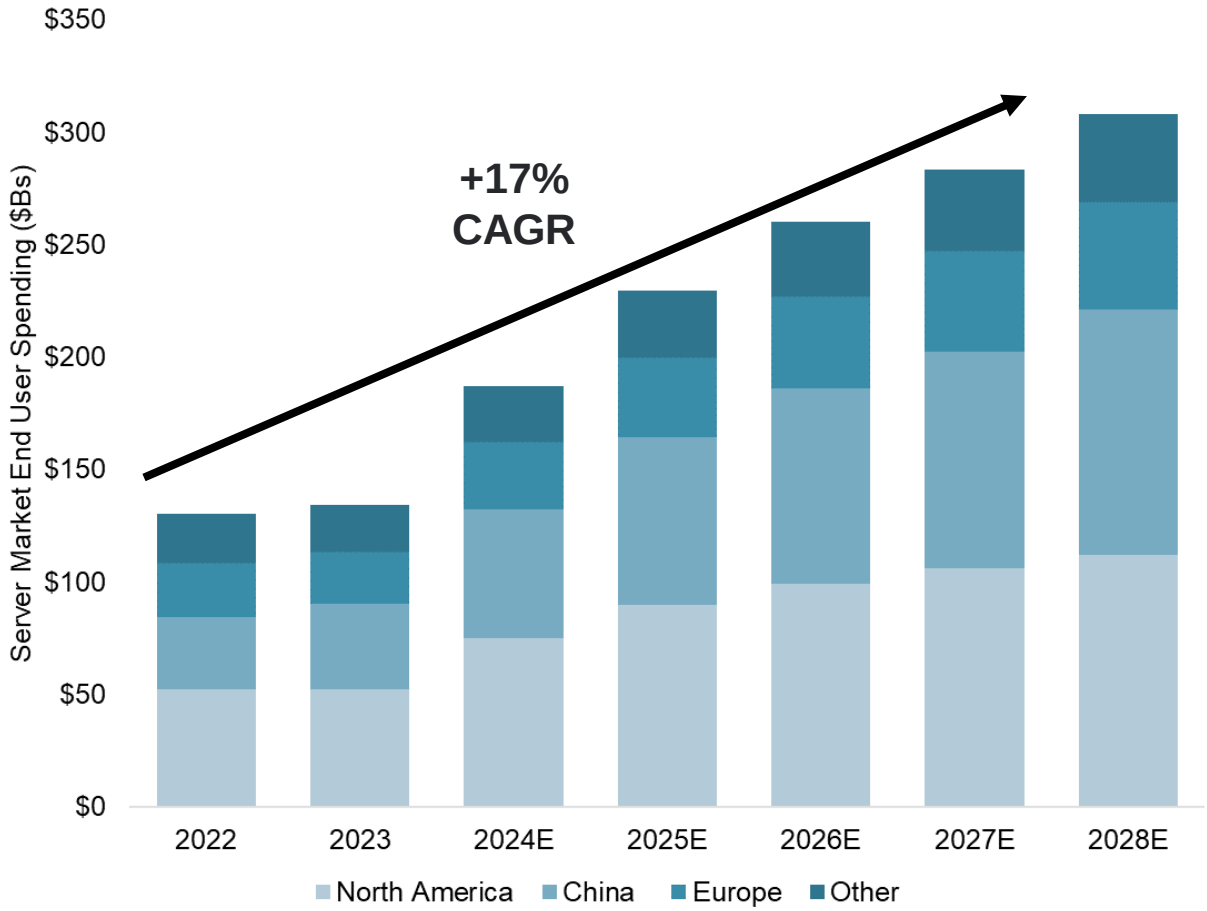

Source: Piper Sandler Research, Gartner

Storage TAM | $53B TAM by 2028 as Flash Share Gains Persist

Storage TAM Across Flash, HDD/Hybrid, HCI, and Software-Defined Storage

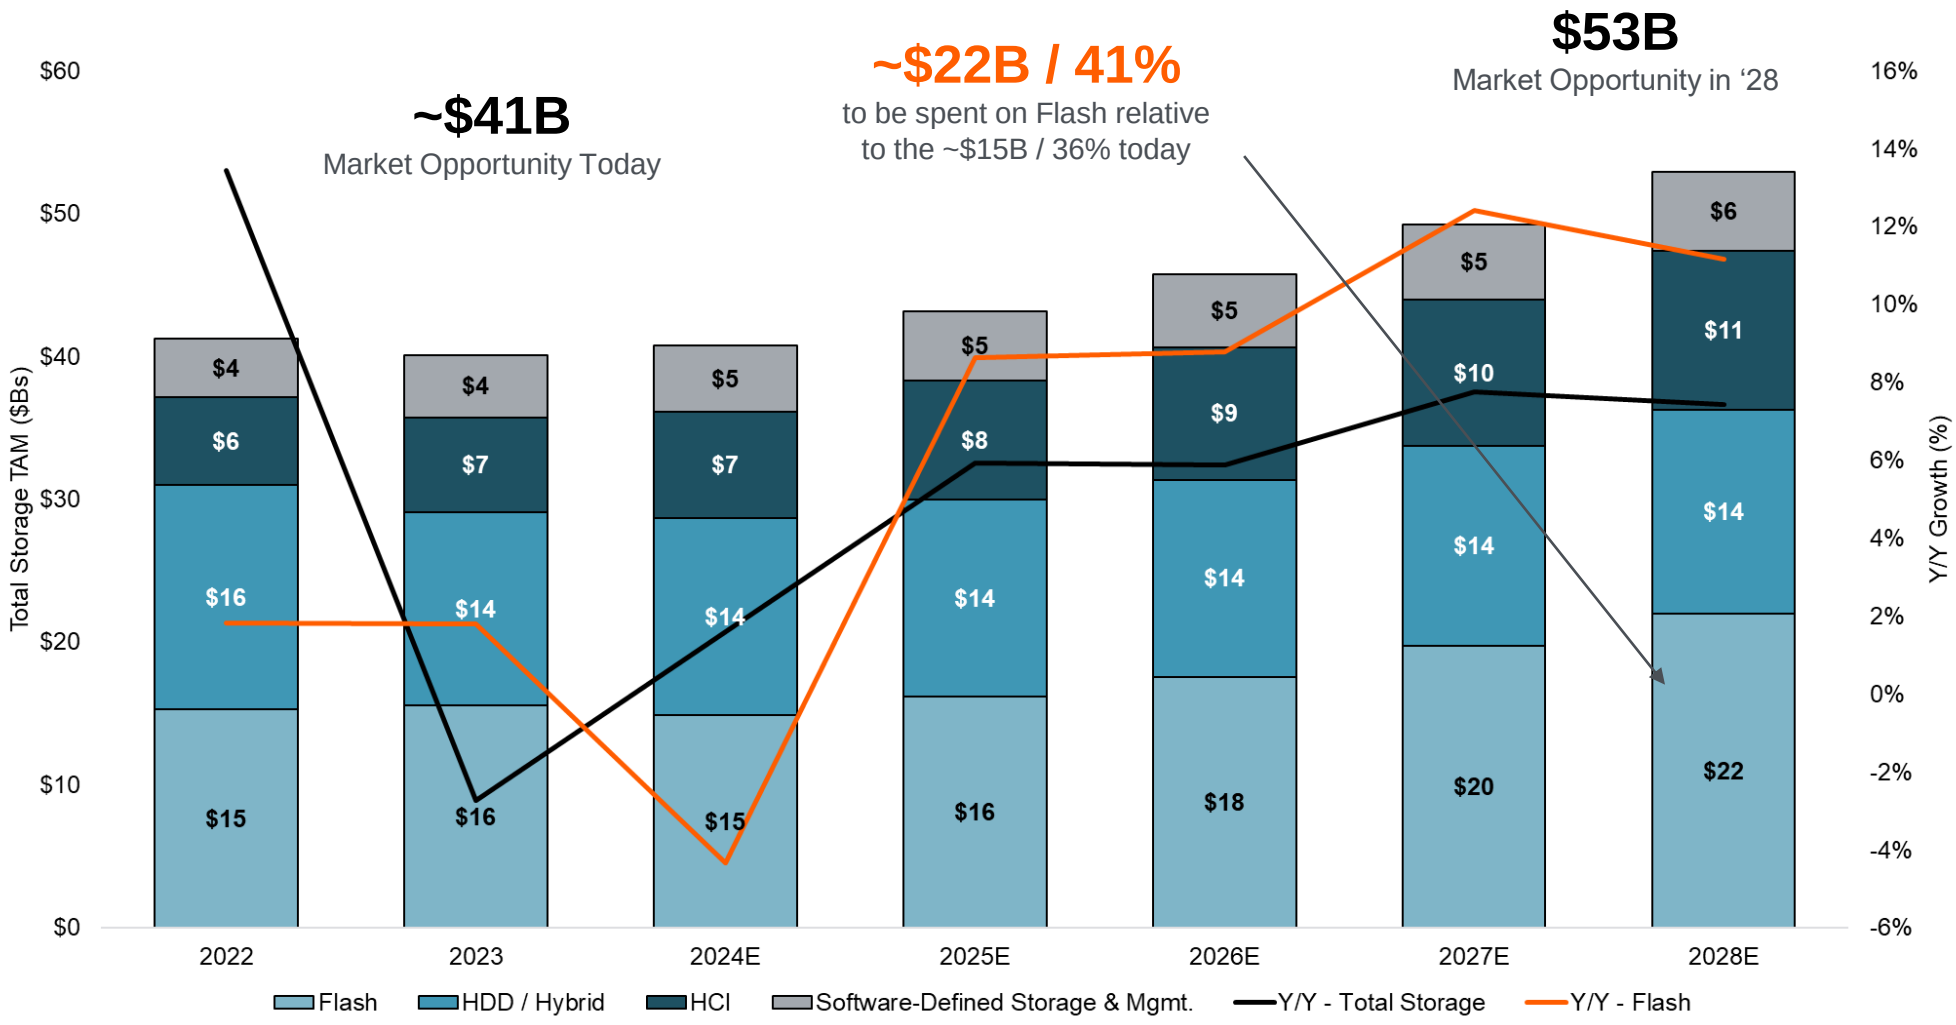

Source: Piper Sandler Research, Gartner

“DPM” TAM | +$20B Market Across Backup & Security

Backup and Recovery Still the Largest Piece of the Puzzle, Though “Shift Left” Provides New Growth Opportunity. Additional Cloud Applications Could be TAM Accretive.

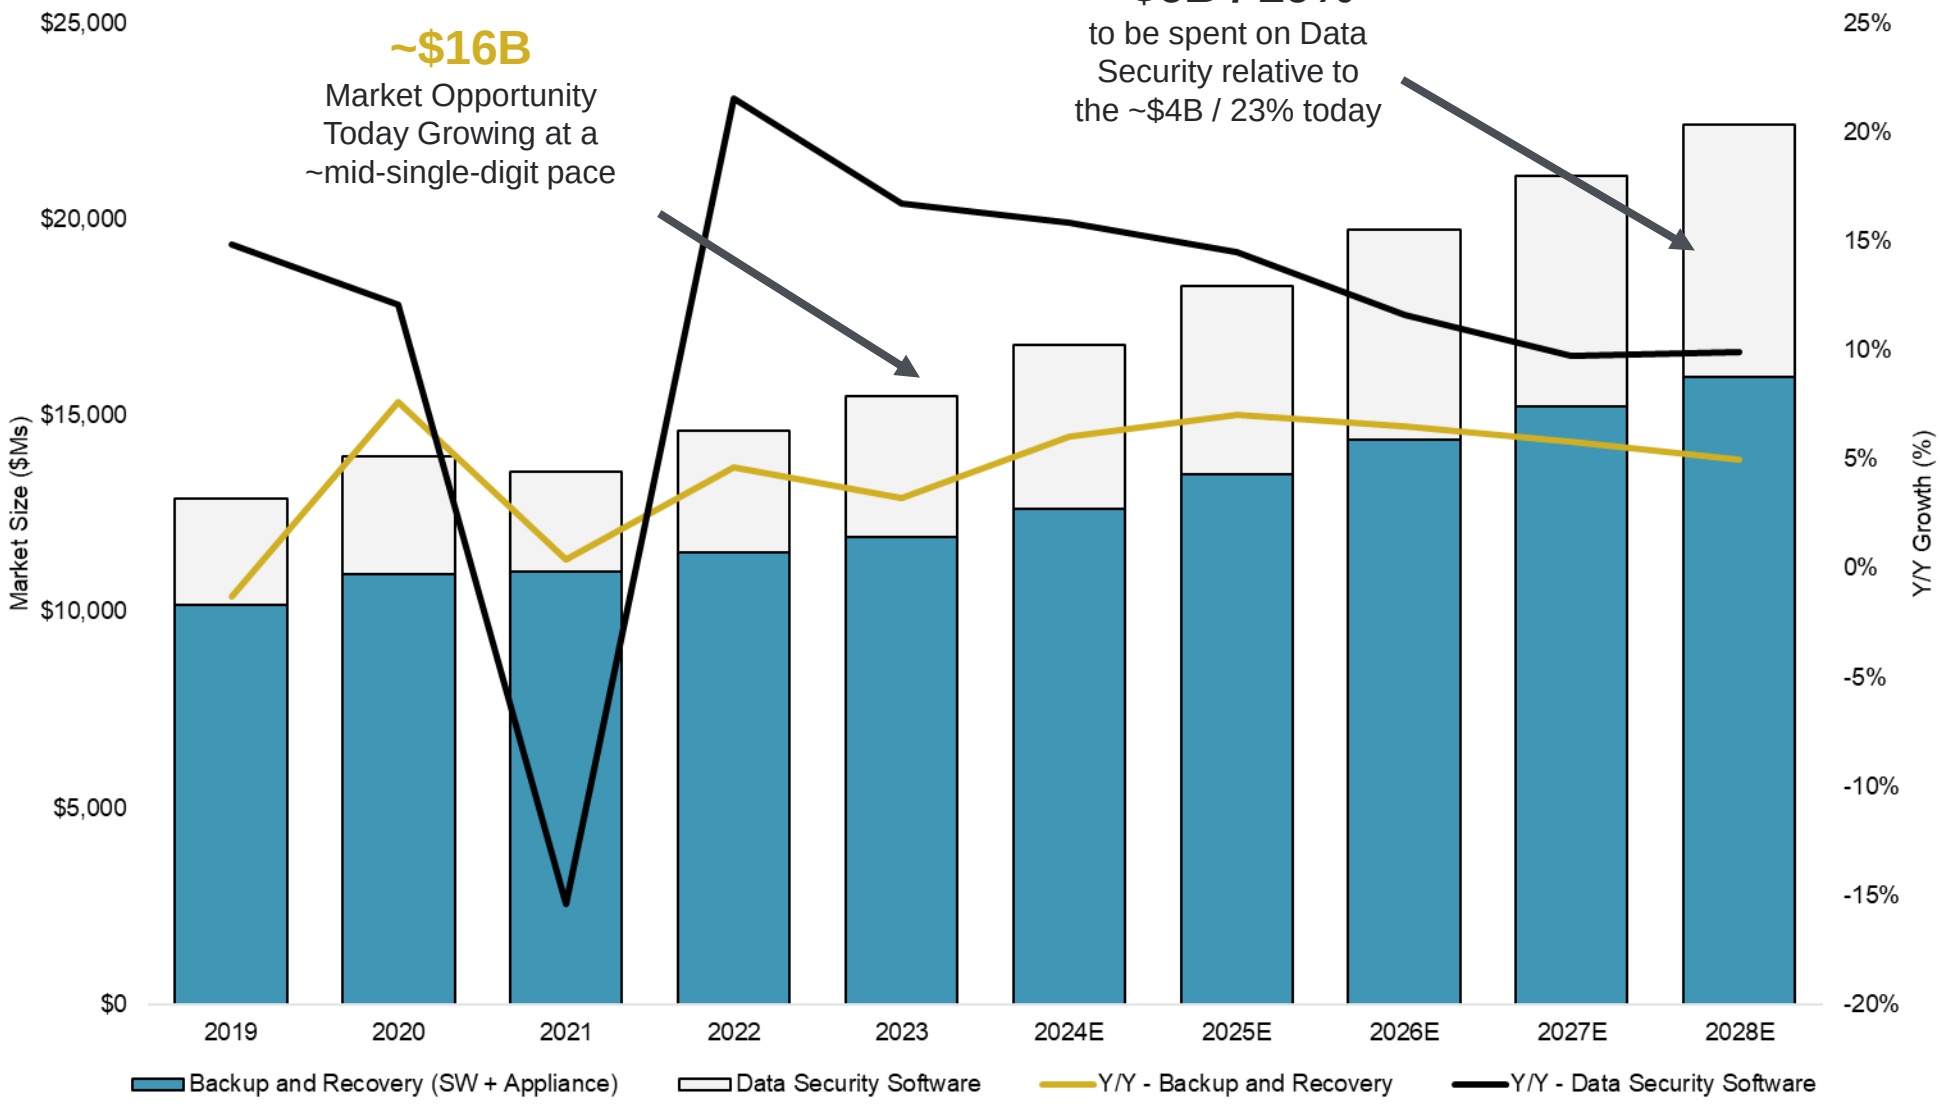

Source: Piper Sandler, Gartner

4. Compute Background & Dynamics

Net Spending Intentions for Servers Rise to Highest Levels

2H24 Piper Sandler CIO Survey Server Spending Intent

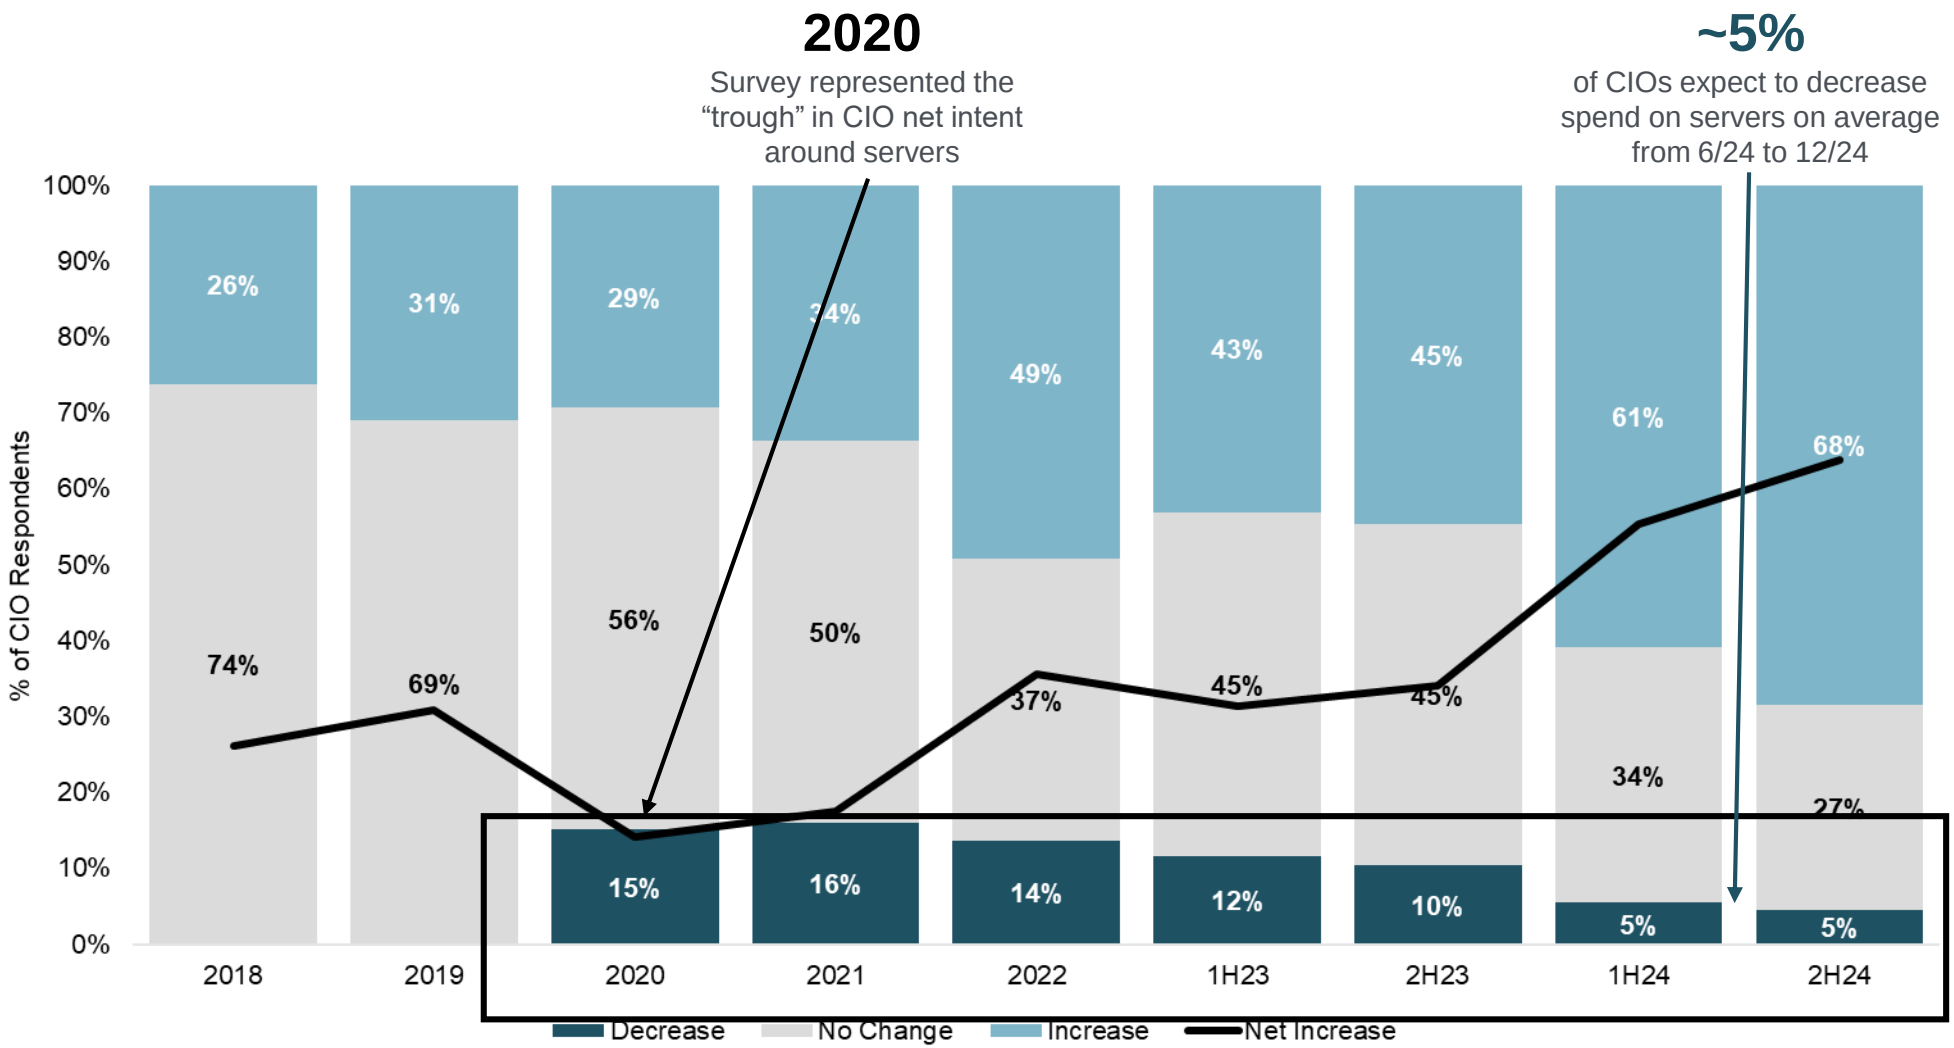

Source: Piper Sandler Research

How Does a Server Work?

A server is a system that serves resources, data, services or programs to clients over a network. Servers represent the powerhouse of the datacenter. Visualizing a hardware based server, the equipment is made up of various components including: a motherboard; CPUs; Memory; Primary Storage; Network Connection; Power Supply & GPUs. Typically, these components are aggregated utilizing an operating system (OS) software to manage and allow apps to run on the server. Servers are a general term for the basic equipment, and the specifications of the server (the mix of components) can be adjusted to fit the end need.

Cisco UCS C-Series Rack Server

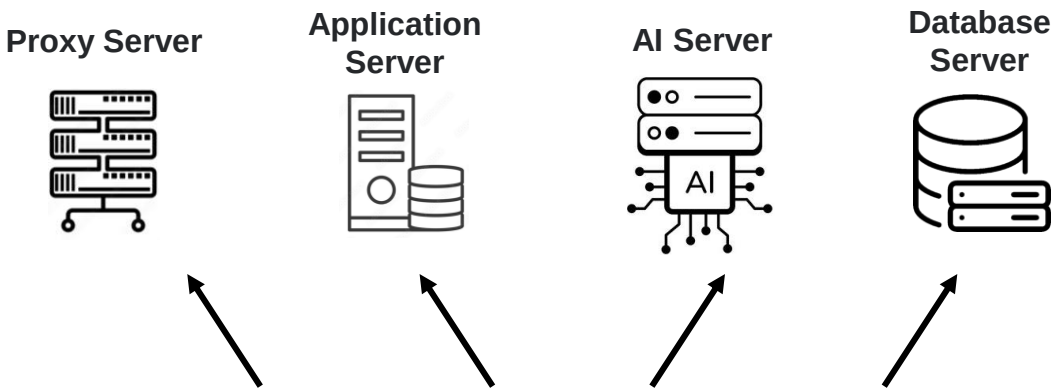

Operating System (OS)

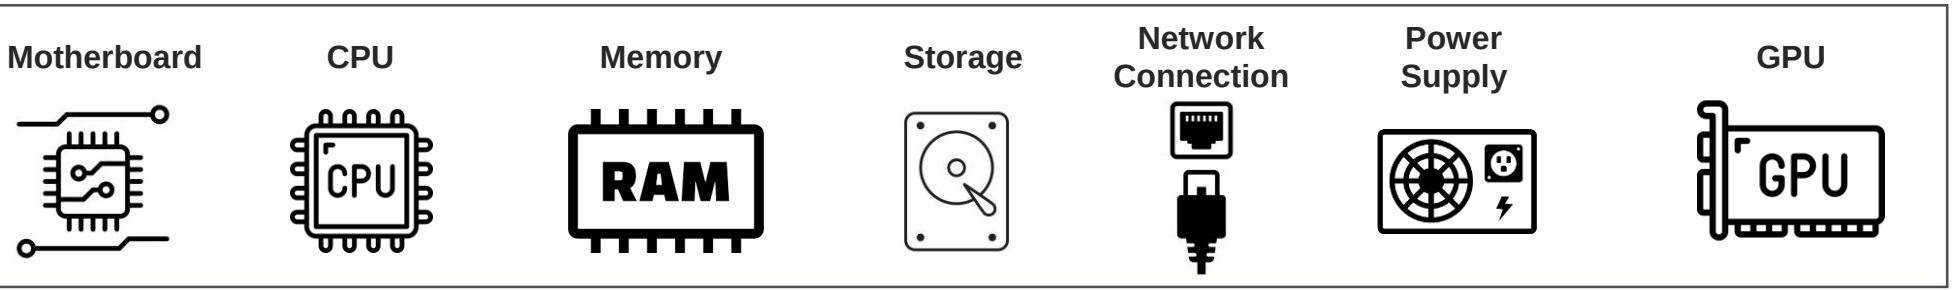

Source: Piper Sandler Research, Cisco

Types of Compute | Rack vs. Blade vs. Tower

Rack Servers

Servers that are stacked in a metal enclosure (Rack) for easy access. Rack servers have standard dimensions to fit in rack enclosures. Rack servers are smaller than standard tower servers, but require more cooling and temperature monitoring given more consolidated architecture. The majority of servers are rack servers.

Blade Servers

Blade servers tend to be smaller than rack servers, and sometimes only contain a CPU, memory and network controllers. Blade servers are used for large processing clusters, but produce more heat than rack servers. Typically Blade servers are more expensive than Tower and Rack. Blade servers are installed in bays, allowing for more servers in an area.

Tower Servers

Tower servers are stand-alone computers. They tend to look like a standard desktop, but have additional server resources installed into the machine. Very small businesses or home networks typically utilize Tower servers, but not any sort of business or datacenter with scale.

Source: Piper Sandler Research, Pure Storage, Cisco. Dell

Custom ASICs – Suppliers | How are they Deployed?

Application Specific Integrated Circuits (ASICs) provide an alternative solution to general-purpose merchant GPUs from the likes of NVDA and AMD. Currently, ASICs are primarily deployed by large CSPs and most often co-designed with compute providers such as Broadcom and Marvell. These chips are often cheaper alternatives to merchant accelerators as they consume less power and provide better TCO economics.

Where merchant GPUs are flexible and can be tailored for many different types of parallel processing workloads, ASICs are most often tailored to one specific function for a particular data center.

Example: MSVP for Video Transcoding

Meta Scalable Video Processor (MSVP) is Meta's first in-house-developed ASIC solution. It is designed for the processing needs of video on demand and live streaming workloads at Meta. MSVP is scalable and can be configured to support both the high-quality transcoding needed for video on demand as well as the low latency and faster processing times that live streaming requires.

Primary Custom XPU Designers

BROADCOM

alchip

MARVELL

Hyperscalers Deploying Custom ASICs

aws

Meta

Google

Microsoft

Source: Piper Sandler Research, META Blog

What is Accelerated Computing? Why GPUs Matter?

Accelerated computing involves the use of specialized chips to offload and bundle frequently occurring tasks using parallel processing. These types of workloads historically have bogged down CPUs and other computing units that execute tasks in a serial fashion. The most common type of accelerator is the GPU which are now being built to handle the demanding needs of AI processes in the data center such as machine learning, data analytics, simulations, and tokenization. Rather than CPUs, the GPU has now becoming the key building block to the most advanced computing systems built for AI functionalities. Importantly, workloads executed on an accelerated computing system are most often cheaper and more energy efficient, particularly at scale when adopted by CSPs. Simply put, in an accelerated workload, when a command hits the data center, the CPU starts a job and manages the flow of the workload through the compute engine. The GPU then performs the logic functionality which involves actually "computing" the workload. Finally, the job is then sent back to the CPU which then closes the workload.

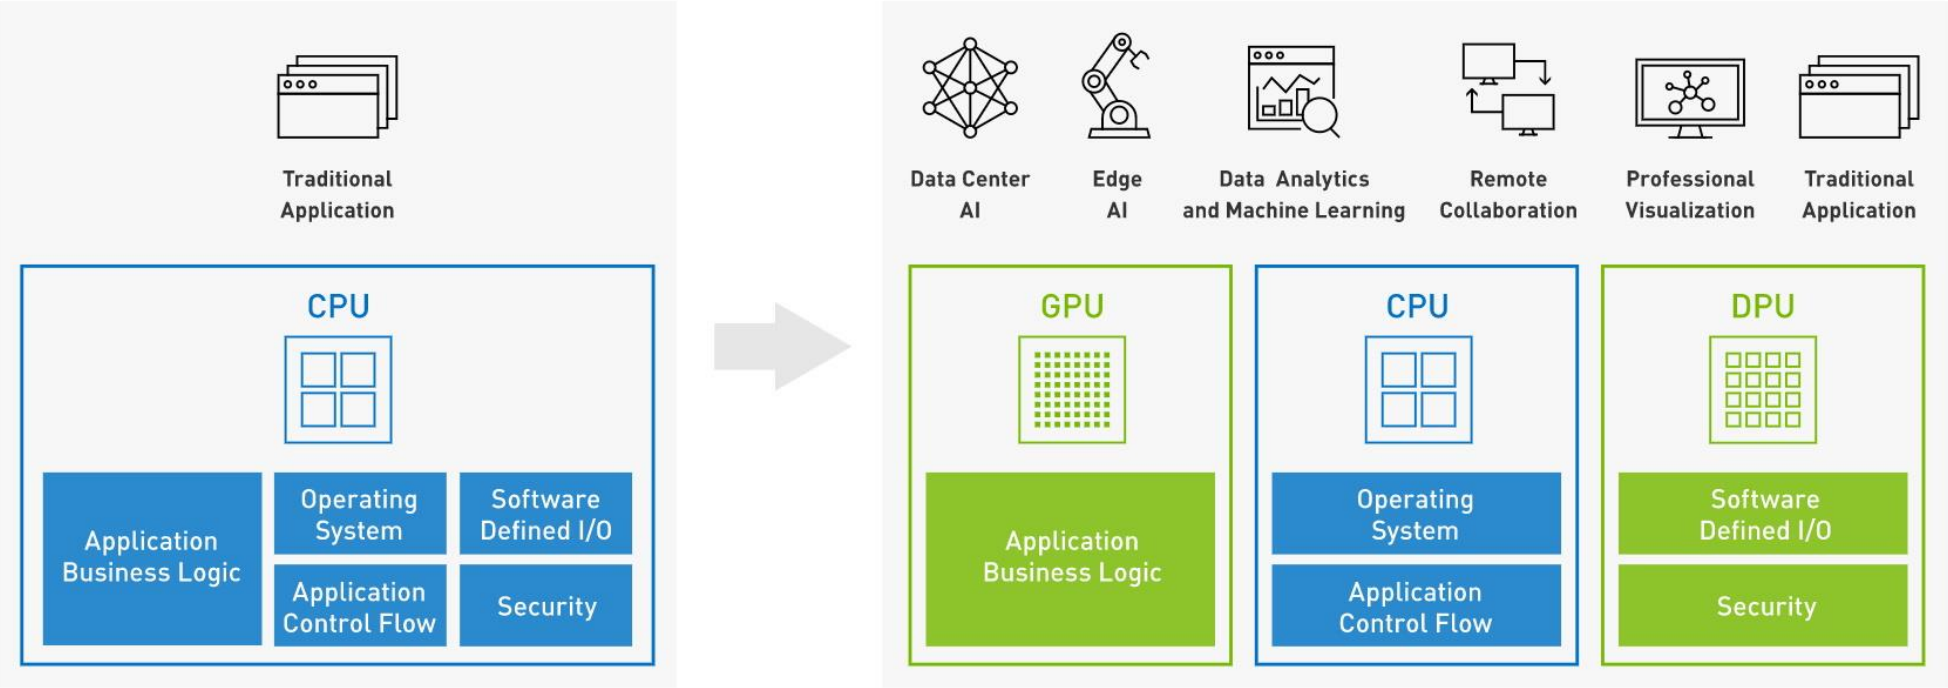

Source: Piper Sandler Research, NVIDIA

Compute / Server TAM - $500B by 2028

Per AMD, the company estimates the AI accelerator TAM to exceed $500B by 2028, representing a 60+% CAGR over that period. For reference, in 2023, the AI accelerator market totaled $45B. The current AI accelerator market is split between general merchant GPUs and custom ASICs where merchant GPUs currently dominate the overall share. Over time however, as more custom solutions are built out, we view custom taking roughly a quarter of the overall TAM in the long run. Despite the share shift, there will continue to be greater needs for more flexible general merchant compute as more AI applications are built and processes become accelerated.

Within general merchant compute, we view the majority of the share to continue to be dominated by NVDA at 90+% of this sub-segment in the long run. This is due largely to its lead in systems and rack-scale TCO for both training and inference use cases.

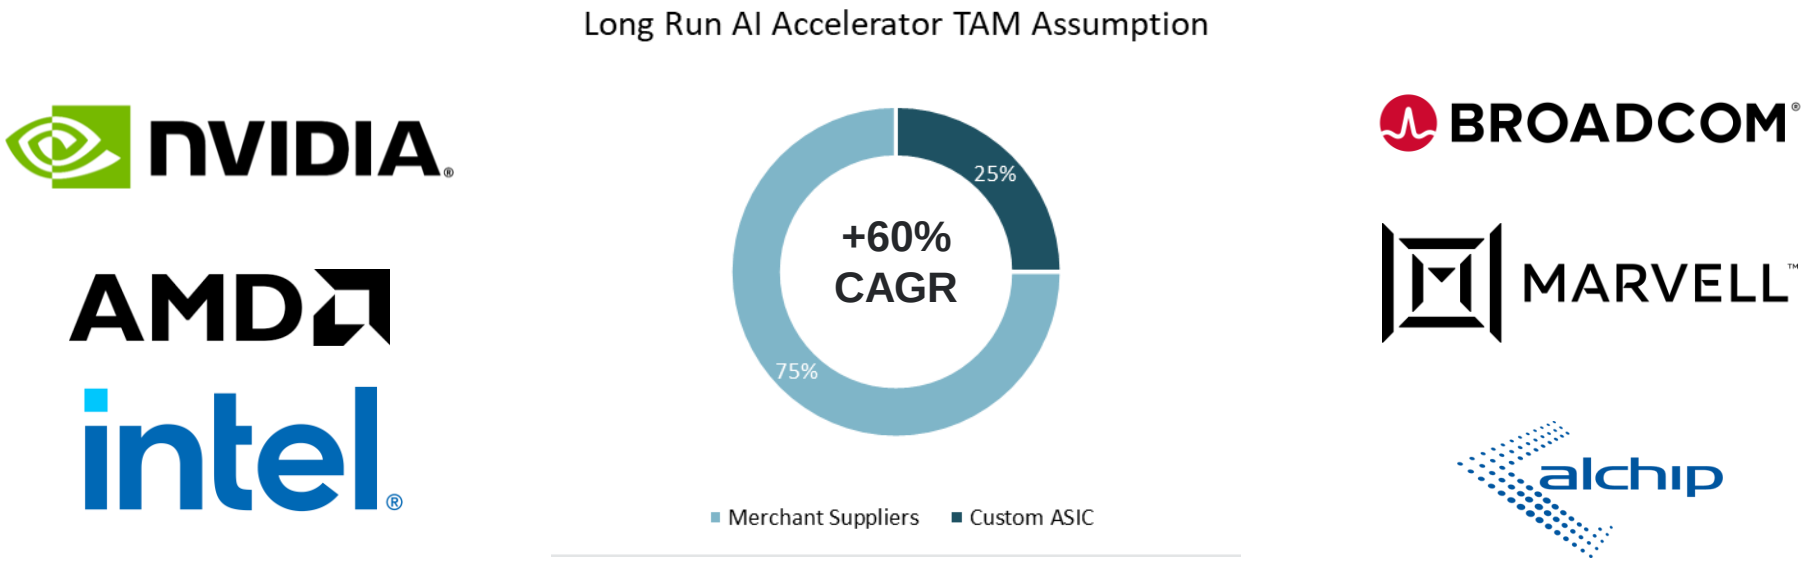

Source: Piper Sandler Research, AMD, Company Websites

5. Compute Trends & Catalysts

Themes & Trends | Edge Computing

Edge computing focuses on bringing computing as close to the source of data as possible to reduce latency and bandwidth use. Edge computing typically uses a user's computer, an IoT device, or an edge server to handle the computing workload. To enable true low-latency edge computing, vendors have to deploy edge servers distributed around the globe to reduce the distance between the server and end-device (and therefore reduce latency).

Due to the distributed location of edge servers, these devices tend to require much more rugged builds with tighter space constraints than rack or blade based servers. As latency sensitive workloads continue to grow, the demand for true distributed computing technology should continue to grow

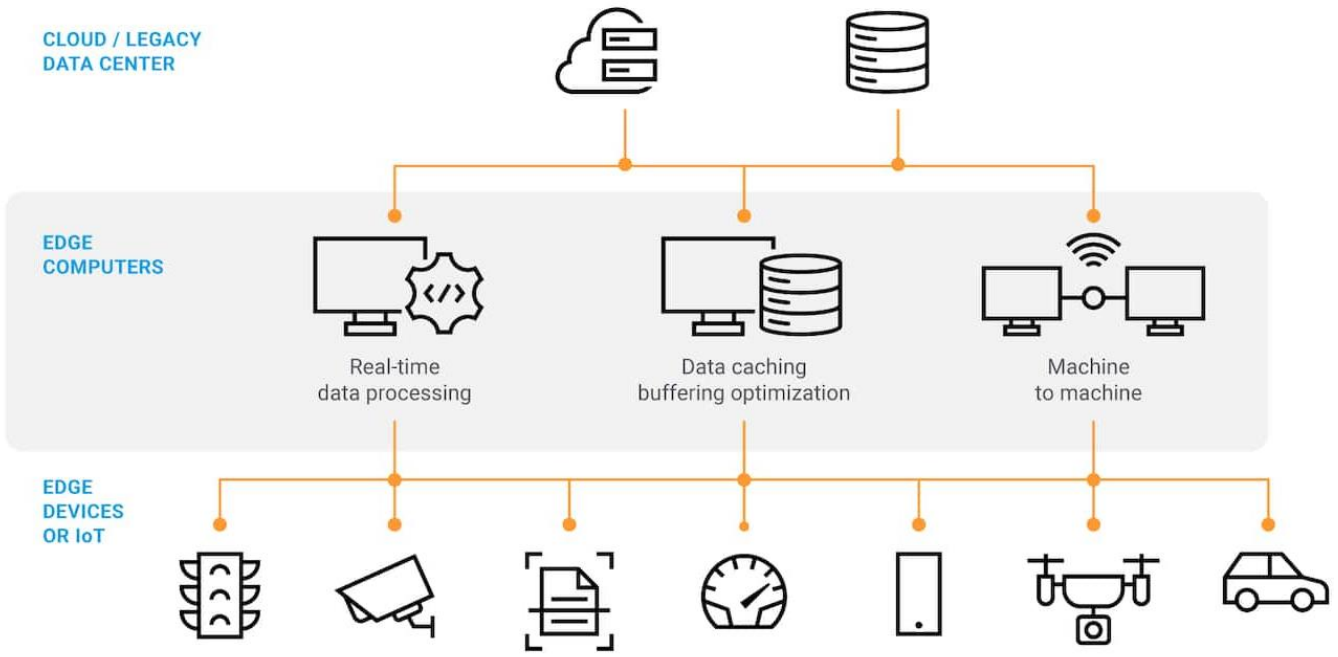

Source: Piper Sandler Research, Akamai, Dell

Nvidia Blackwell to Deliver Increased Performance and Efficiency

Nvidia's new Blackwell GPU can reduce the cost and energy consumption for trillion-parameter LLMs up to 25x over its predecessor, the H100. The Nvidia GB200 Grace Blackwell Superchip connects two Nvidia B200 Tensor Core GPUs to the Nvidia Grace CPU over a 900GB/s low-power NVLink. Integration with NVIDIA Quantum-X800 InfiniBand and Spectrum-X800 Ethernet platforms can deliver up to 800Gb/s networking speeds.

- Blackwell will ramp through CY2025 time frame with customers on an allocation process similar to the Hopper architecture.
- This is the first time in NVDA's history where the company will be selling 4 different GPU SKUs. (H200, B100, B200, GB200)
- The Blackwell is also the first GPU to be produced at scale with liquid cooling capabilities.
- The Blackwell architecture is manufactured using a custom-built 4NP TSMC process with two-reticle limit GPU dies connected with 10TB/s chip-to-chip link.
- The 2nd generation transformer engine is fueled by micro-tensor scaling support, and NVDA's advanced dynamic range algorithms integrated into TensorRT-LLM & NeMo Megatron frameworks.

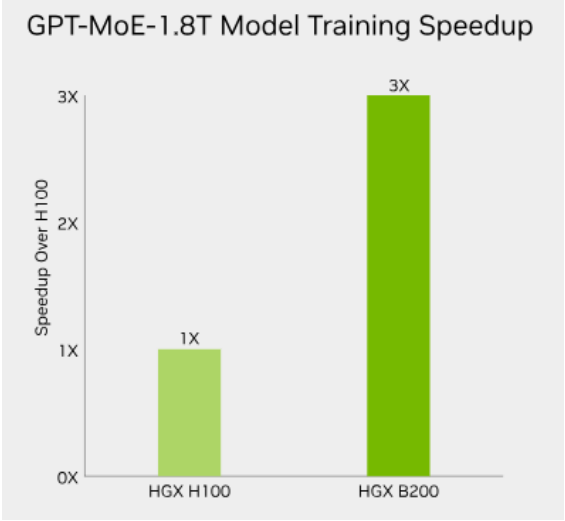

3x

Increase in GPT-MoE-1.8T Model Training Speed

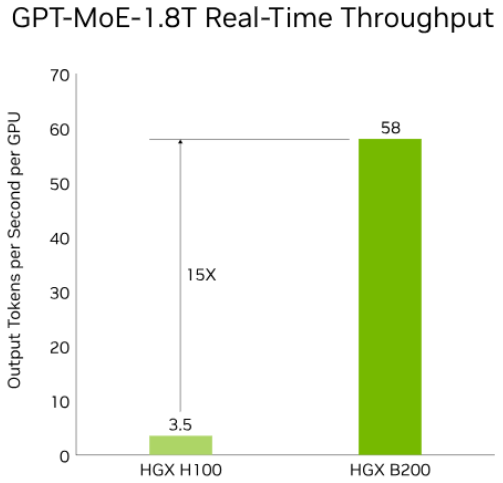

15x

Increase in GPT-MoE-1.8T real-time throughput

Source: Piper Sandler Research, Nvidia Company Materials

Legacy CPU Market Breakdown | AMD Taking Share

While GPUs have dominated the market over recent years in terms of overall compute spending, CPUs still play a valuable role in data center computing. As mentioned previously, CPUs are needed to control the flow of a workload through the hardware stack and can also be used as a more energy efficient solution in some AI inferencing use cases. Additionally, a vast majority of legacy software continues to be written to x86 architectures. The x86 server market consists of two primary players, Intel and AMD, where AMD has steadily gained market share over the previous three years. It is notable that in the most recent 3Q24, AMD's overall data center revenue (CPU+GPU) surpassed Intel's data center revenues for the quarter.

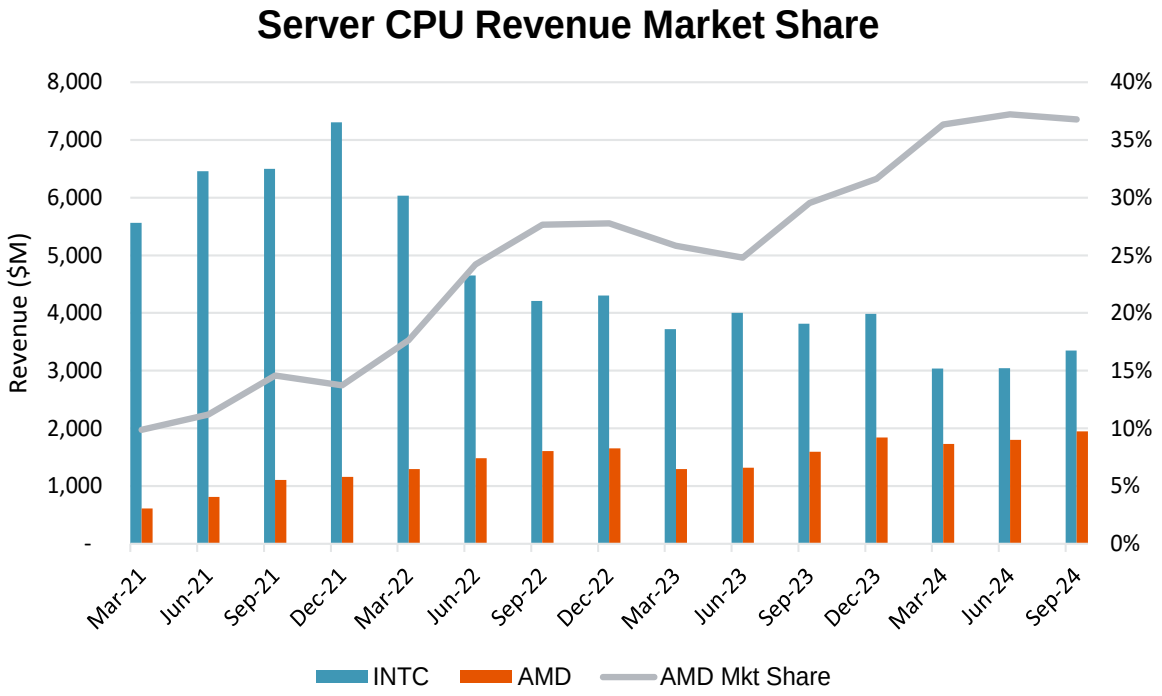

	CPU	GPU
Workload	User Applications (general processing and compatibility)	Parrallel processing analysis and graphics processing
Power	+	+++
Storage Efficiency	Low	Low
Cost	+	+++
Latency	Low Tolerance	High Tolerance
Cores	4-8	100s-1000s

Source: Piper Sandler Research, Company Filings

6. Primary Storage Background & Dynamics

How Storage Traditionally Works

Primary storage, or main storage, refers to the method of storing data that must be accessed frequently. Primary storage can include both volatile primary memory (such as RAM and DRAM) and non-volatile storage media (such as storage-class memory, hybrid storage arrays, all-flash arrays, and storage area networks (SANs)). We focus mainly on non-volatile storage in this report, which has the benefit of preserving information for long-term storage and retrieval even if power gets removed. Non-volatile storage is typically broken down further into HDDs and SDDs (Hard Disk Drive and Solid-State Drive).

Hard Disk Drive

Hard disks work by spinning the platter (disk) and the read/write head moves across the platter and writes data by sensing the electrical charge on the bits in each sector. HDDs are cheaper, but tend to be less performant and less durable than SSDs

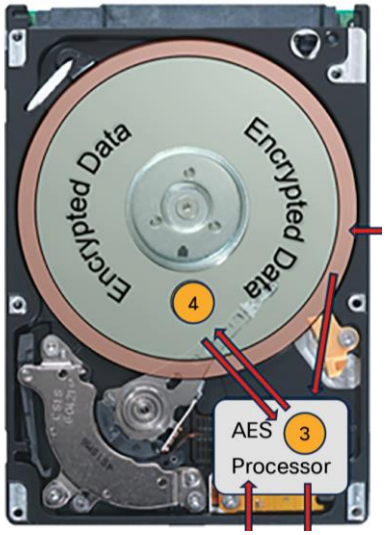

Solid-State Drive

Solid state drives work by using electronic circuits to store and retrieve data. SSDs have fewer power consumption requirements than HDDs because there are no moving components. SSDs also tend to be much higher performant, though this comes at a price with SSDs more expensive than HDD comparable drives.

Source: Piper Sandler Research, Cisco, Lenovo

Primary Storage Options

Infrastructure storage can be done across multiple paths, creating multiple key decision points before a CIO decides on one specific storage purchase. There is a gradual shift to the public cloud, but enterprises are still choosing today between on-premise / private cloud or a cloud-based storage architecture. If a CIO chooses on-premise / private cloud, then the debate turns towards going with traditional three-tier architectures (storage and compute managed separately), or with a software-defined or hyperconverged storage vendor like Nutanix, Scale, Vast, Qumulo, Weka, etc. If the CIO chooses to go with a traditional three-tier architecture, the decision tree turns to disk vs. flash. Key competitors within the all-flash segment include Pure Storage, Dell, NetApp, IBM & HPE.

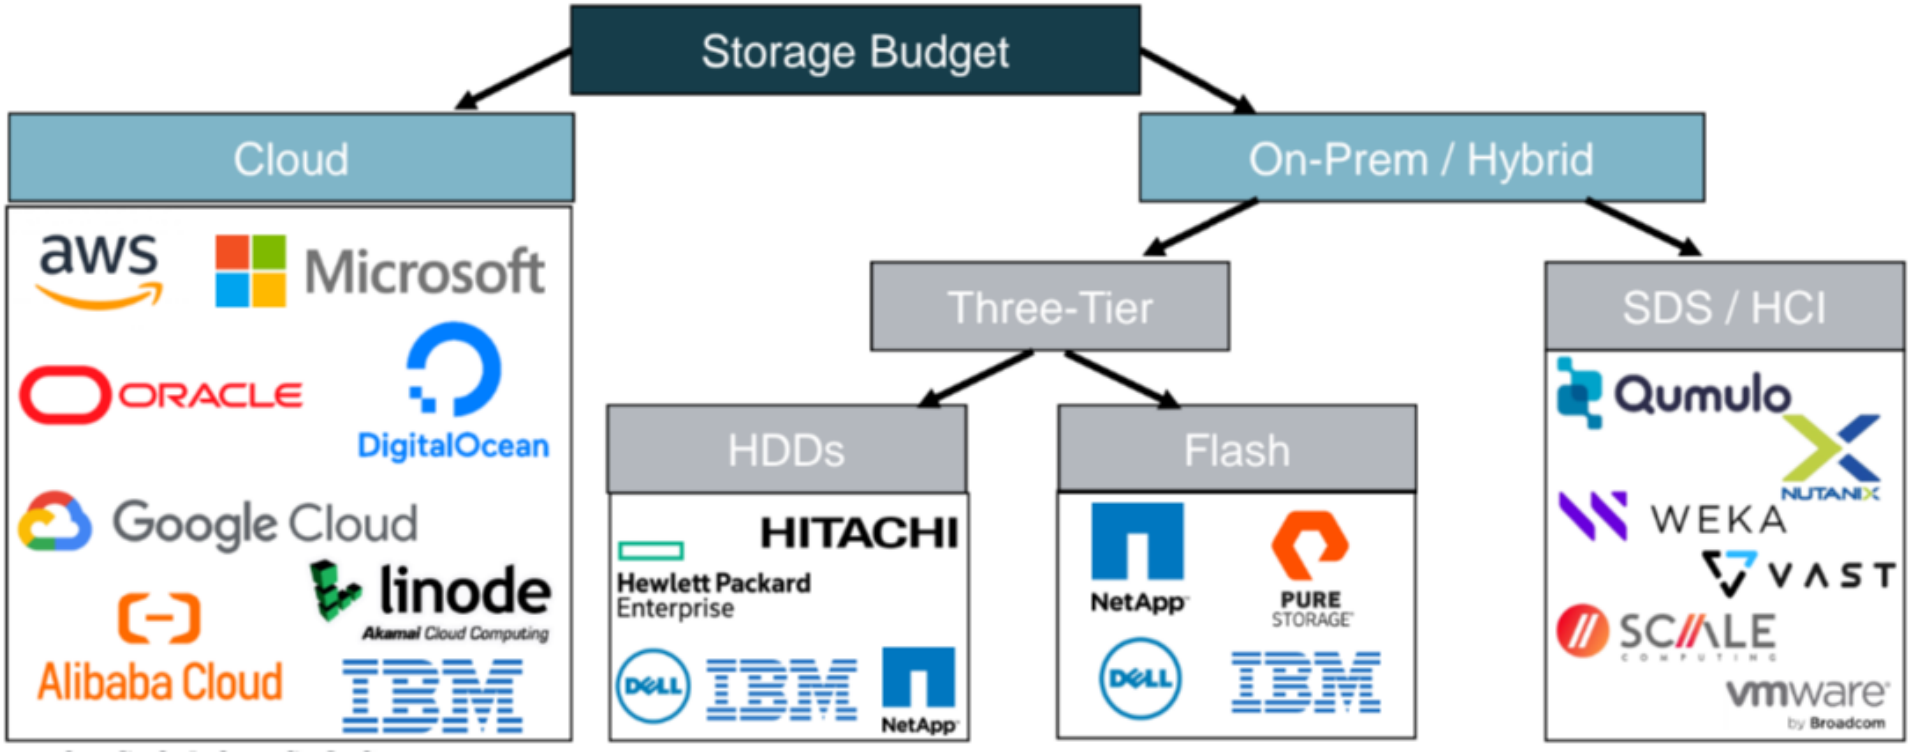

Source: Piper Sandler Research, Company Websites

Three Types of Storage | File, Block, & Object

Topic	Block	File	Object
Definition	Data broken & stored in 'blocks'	Systematically placed in folders	Organizes unstructured data
Units	Unique identifier to each block	File, which works thru a hierarchy	Individually identified objects
Advantages	Performance / Latency & Simple	Simplicity, cloud-friendly	Cost, Scale, metadata, resilience
Disadvantages	Lacks metadata, flexibility, DC-only	Cost, performance / latency, scale	Latency, reliability, changing data
Interface	Operating System	User	API / Program
Cost	$$$	$$$$	$
Use-Cases	RDMA, servers, critical apps, DBs	AI/ML, analytics, streaming, file-share, office apps	Images, PDFs, Backups, Video
Protocols	SCSI, Fibre, SATA	SMB, NFS	REST, SOAP/HTTP
Best-Suited	Frequently changing data	Shared file data / Collaboration	Static file data & cloud storage
Performance	Fastest	Middle	Slowest

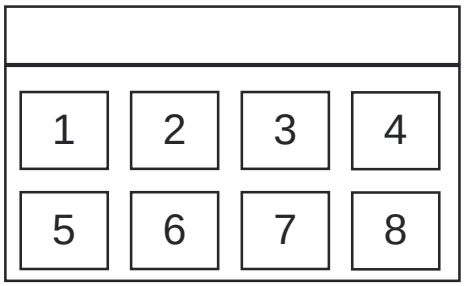

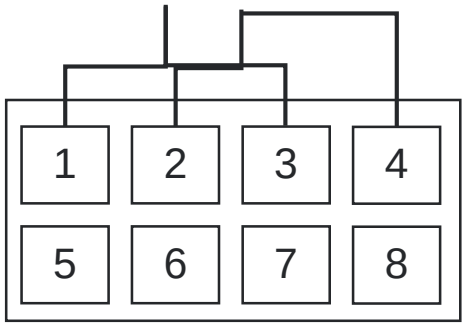

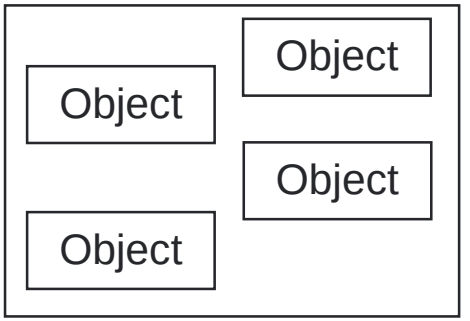

Source: Piper Sandler Research

Incumbents Losing SSA Market Share

Looking specifically at the market share of solid-state arrays (SSAs) Dell has the highest share of the market at just under 25%. Both Dell and NetApp have seen their share of the SSA space fall as competitors like Pure Storage have gained share over a multi-year period. We expect Pure to continue to gain share at the expense of "legacy" peers, helping the company grow its overall share of the Storage market.

SSA Vendor Market Share Over Time

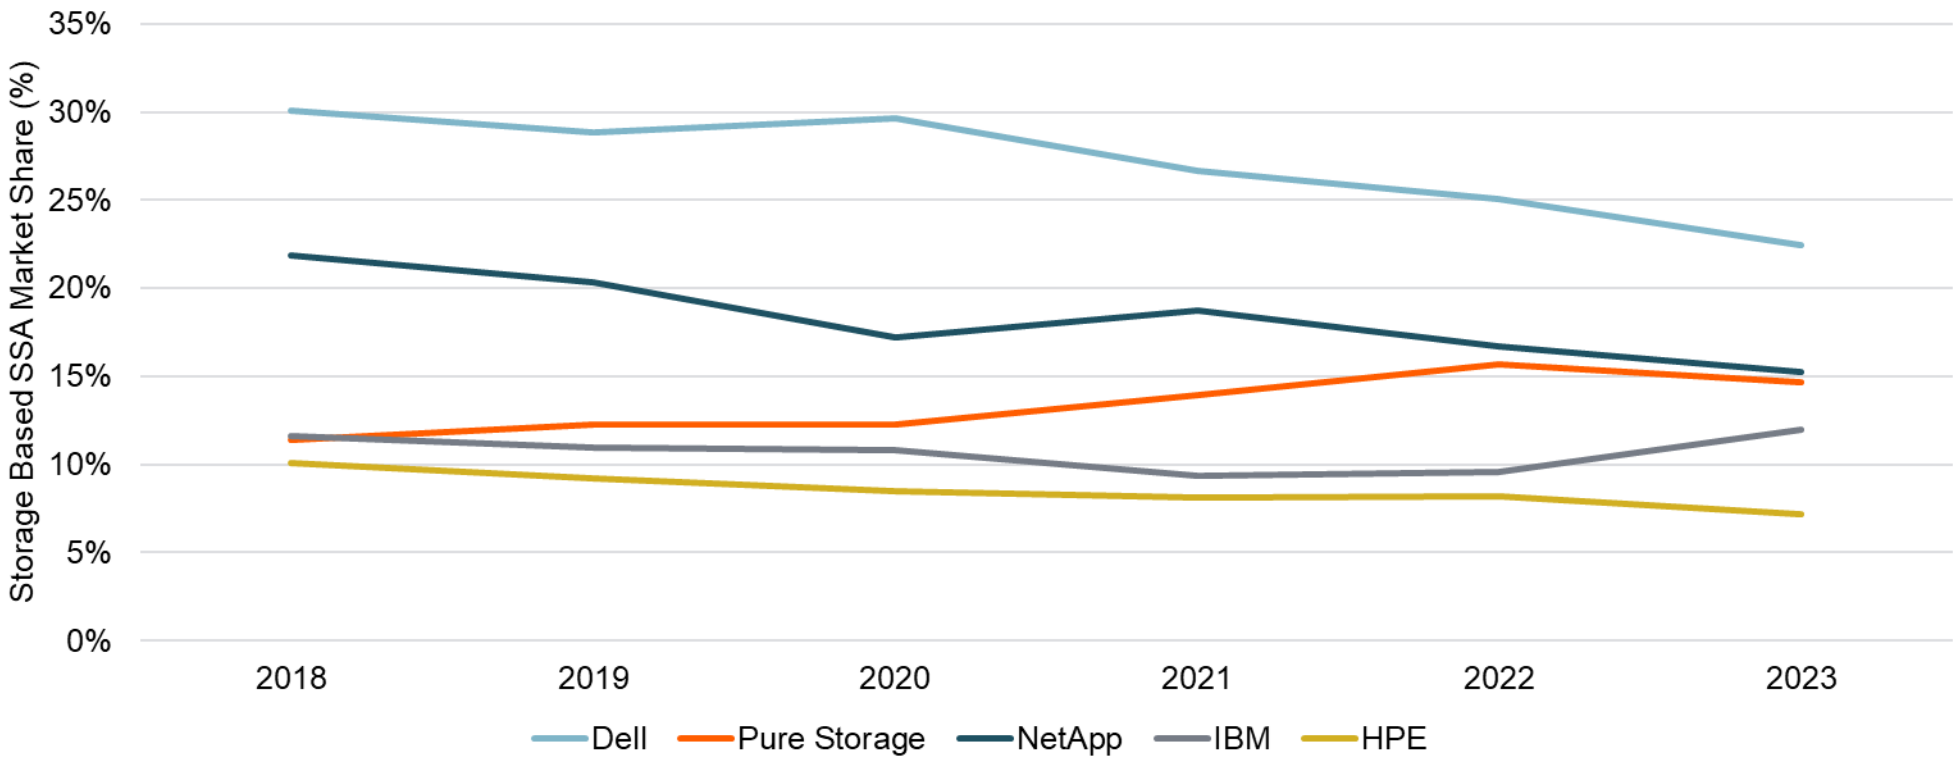

Source: Piper Sandler Research, Gartner

Net Spending Intentions Rise to Highest Levels Since 2015

2H24 Piper Sandler CIO Survey Storage Spending Intent

+42%

Net CIO intent was down slightly from peak levels in our 1H24 survey, though remains above the trailing average levels

~8%

of CIOs expect to decrease spend on storage, more inline with historical trends

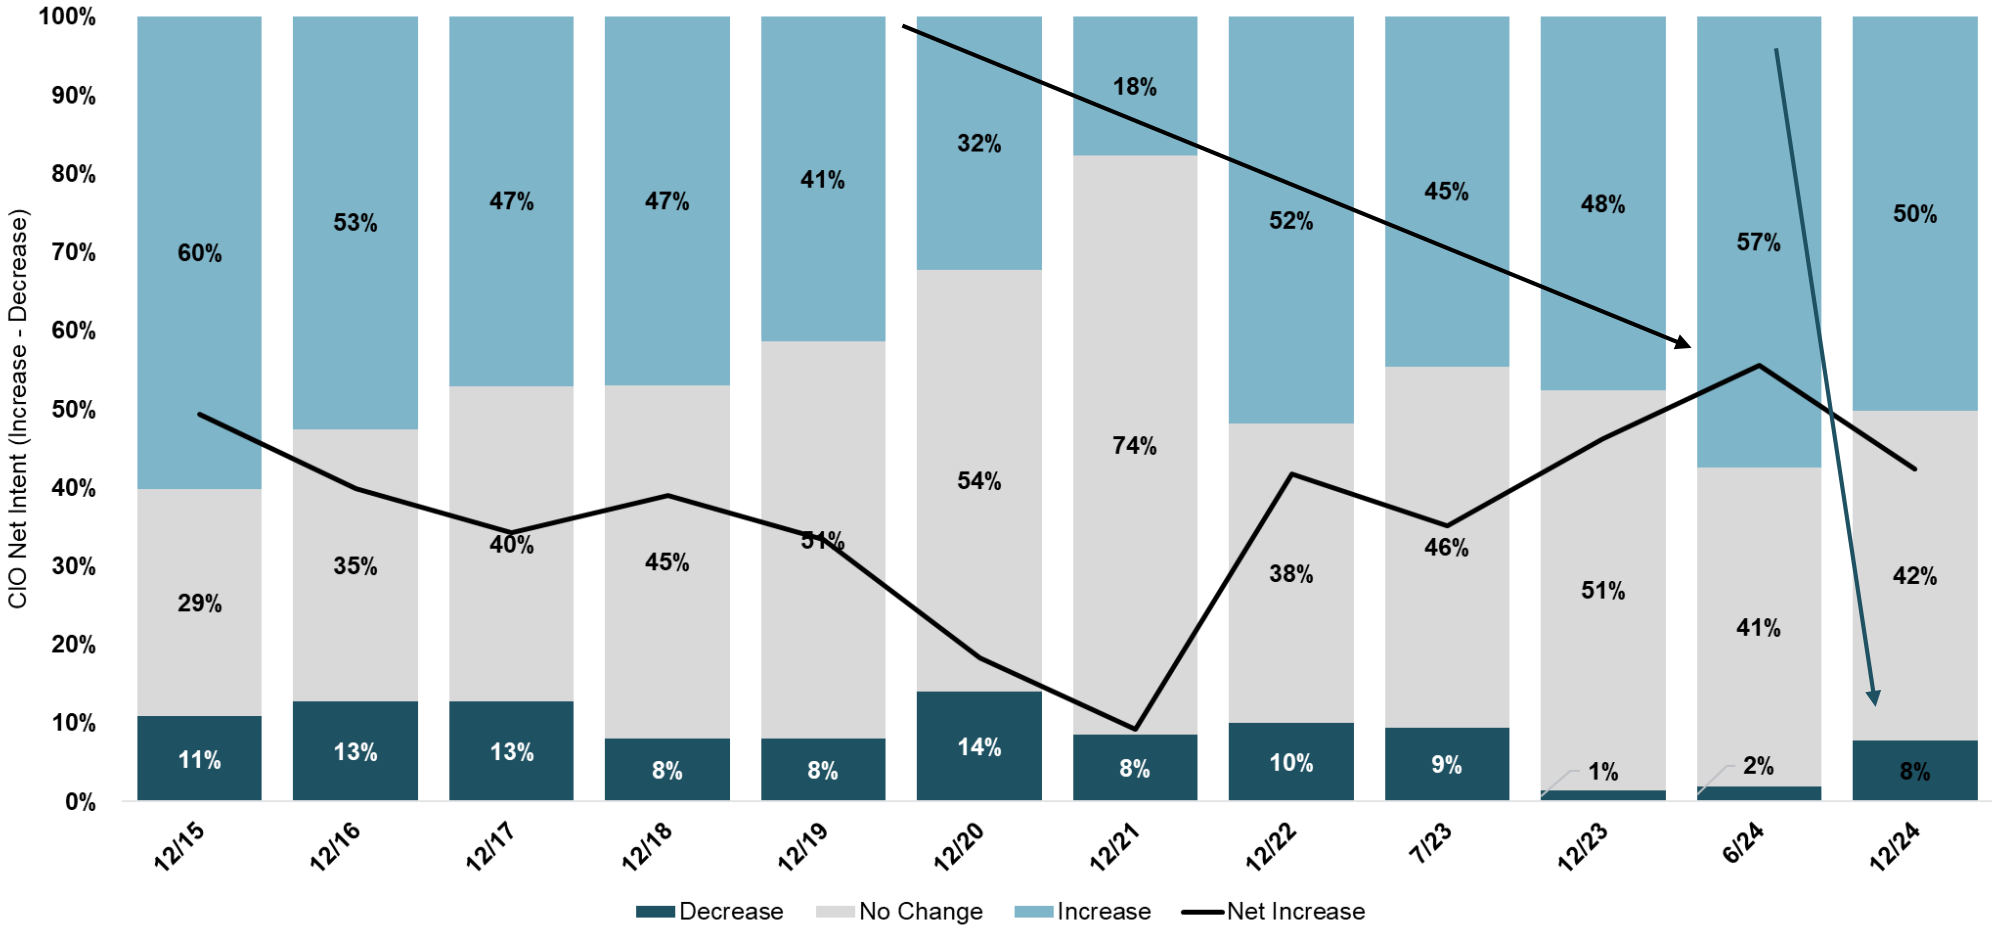

Source: Piper Sandler Research

File, Block, and Object Storage

Across File & Object Storage, Software-defined vendors (Vast, Weka, Qumulo) have the highest overall scores with Cohesity, Nutanix and Pure Storage close behind. On Block, Pure Storage is the clear leader with more “legacy” vendors like HPE Dell and NetApp as competition.

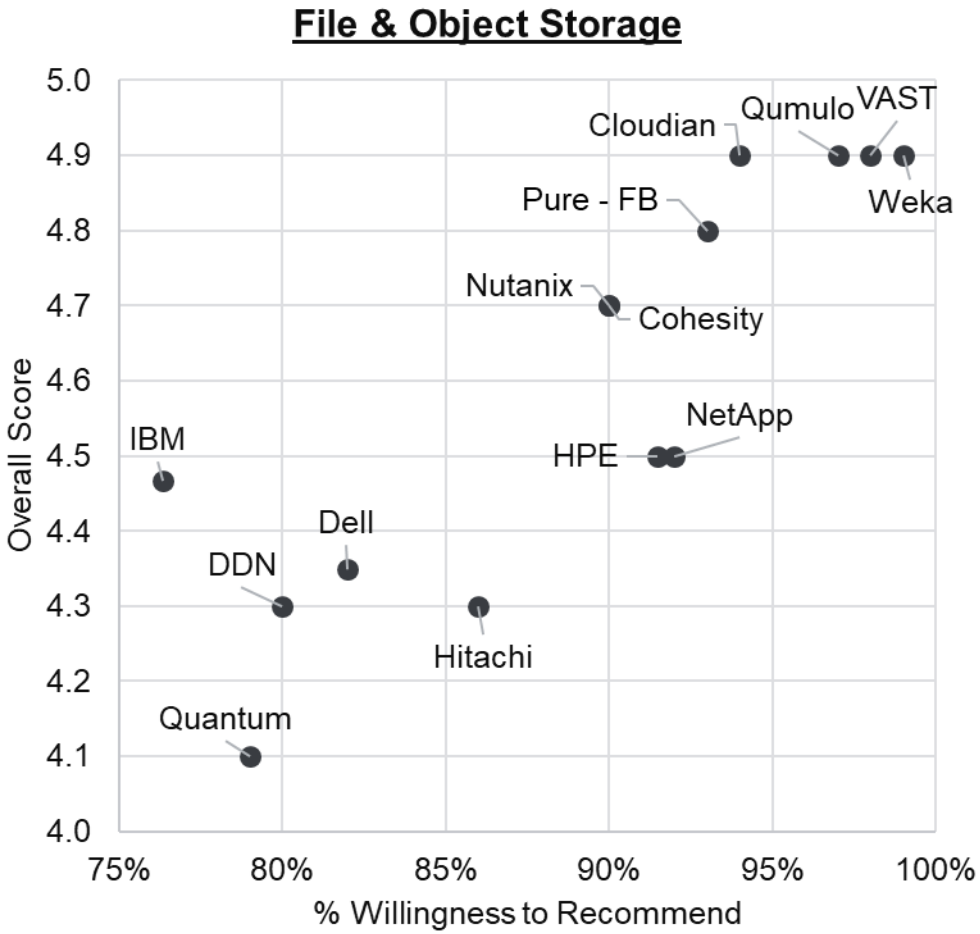

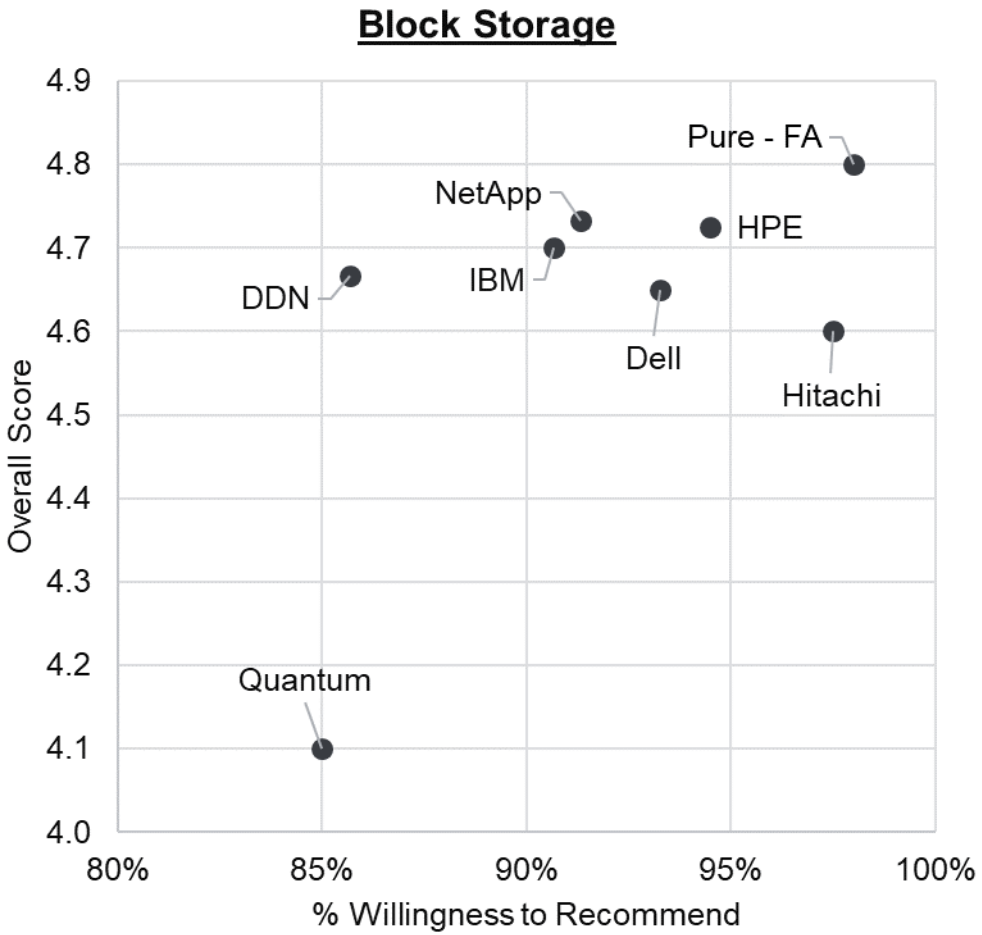

Source: Piper Sandler Research, Gartner Peer Insights

Pure Storage and Vast are the Largest Mind-Share Gainers

Within our quarterly channel survey, Pure Storage (25%), Nutanix (18%) and Dell (18%) were the respondents that partners are seeing the strongest trends within storage specifically this quarter. Pure Storage, Vast, and Dell saw notable Q/Q increases, with others remaining relatively inline. VMware did decrease Q/Q and remains in a distant 7th place.

3Q24 Piper Sandler Infrastructure Channel Survey

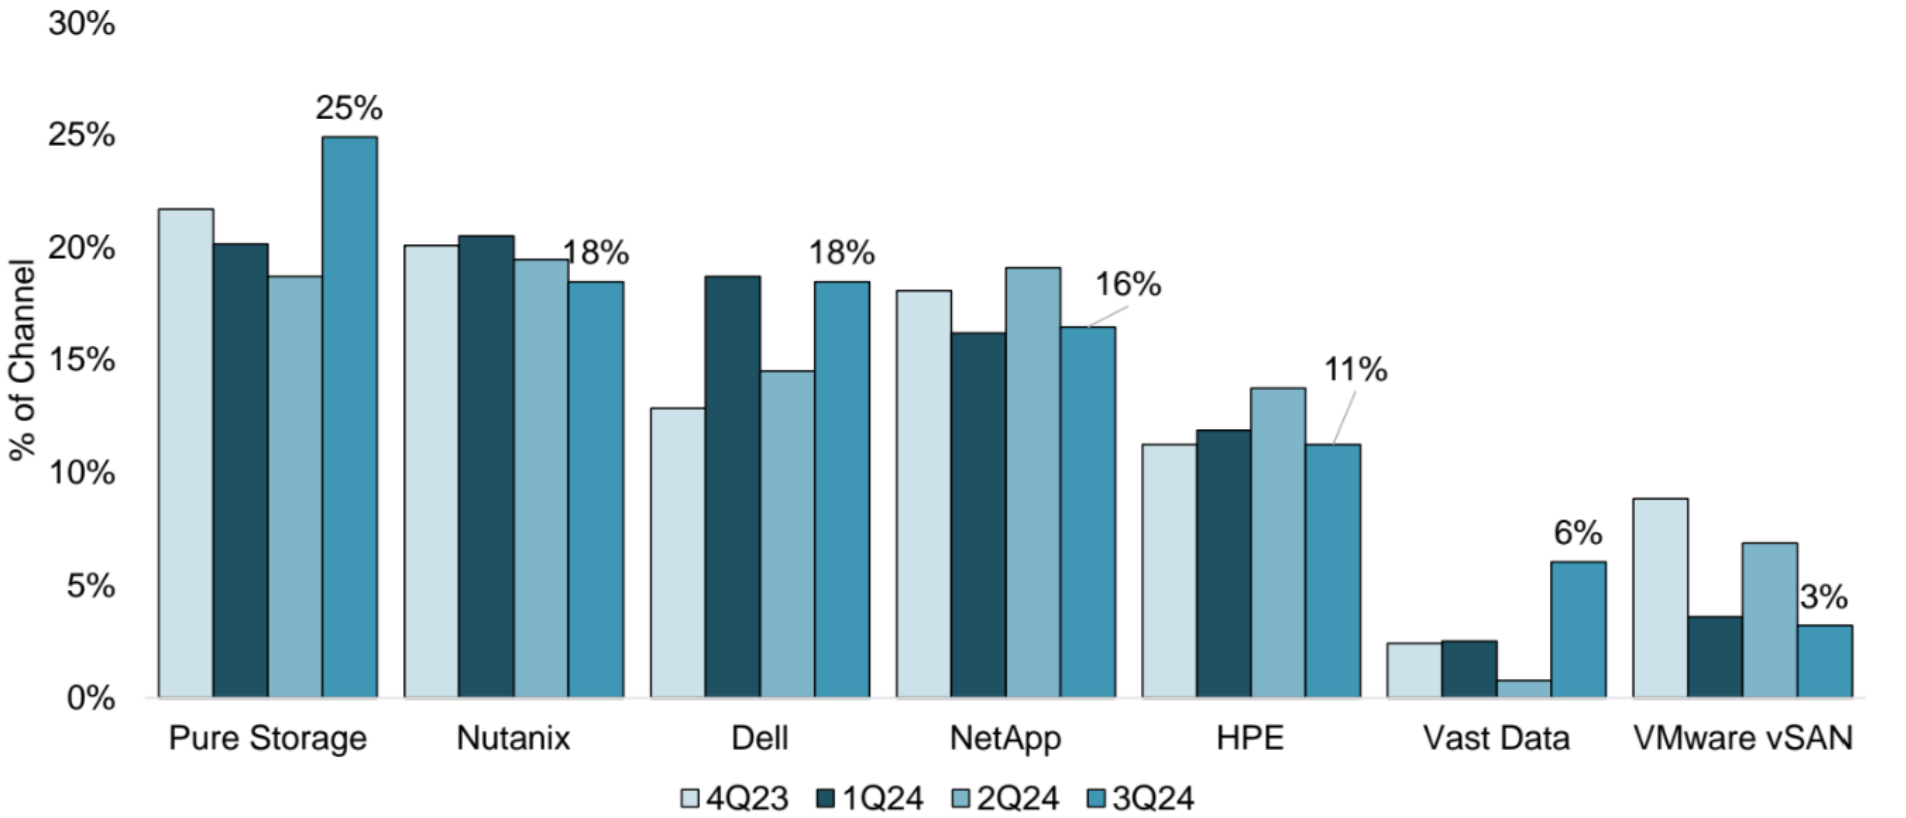

Source: Piper Sandler Research

7. Storage Trends & Catalysts

Themes & Trends | Accelerating Data Growth

Industry estimates project worldwide data growth to grow in the mid-20s through 2028 (accelerating from 2024 levels) reaching nearly 400 zettabytes by 2028. Prior to the pandemic, Cisco projected global consumer IP traffic alone would represent ~5 zettabytes by 2022, implying ~95% of data creation is coming from enterprise sources. We believe the elevated data growth would drive increased capacity needs through the next decade.

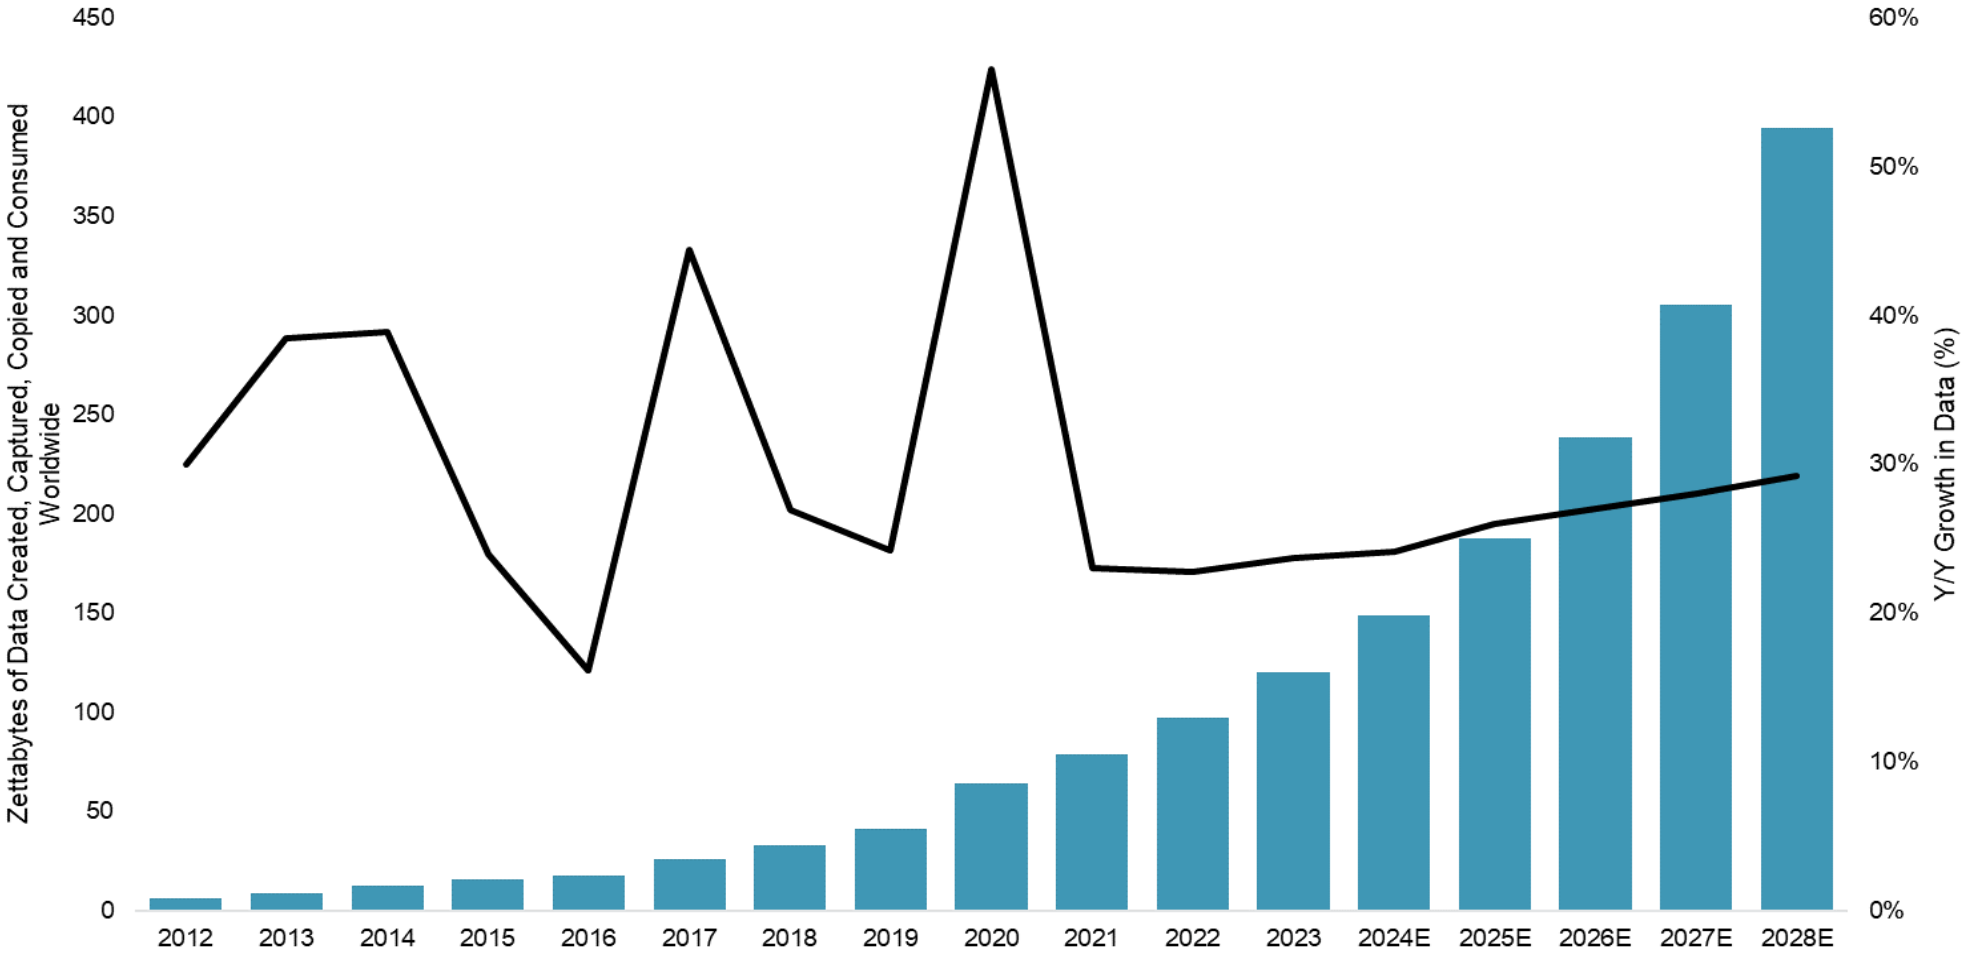

Source: Piper Sandler Research. Seagate, Gartner, Cisco

Themes & Trends | Market Shifting to Software Defined Storage

Significant growth in storage volume over the past decade has driven up overall IT budget allocation on storage. As a result, IT professionals are increasingly scrutinizing traditional operating structures, looking to reduce costs by eliminating storage infrastructure complexity, benefitting from better unit economics with industry-standard equipment, and automating storage provisioning and management.

There are a few storage datacenter architectures:

1) Direct-Attached Storage (DAS), which connects directly to the only server that can access the storage.
2) Storage Area Network (SAN), which allows multiple servers to access pooled storage.
3) Network Attached Storage (NAS) is storage connected to a network.
4) Software Defined Storage (SDS), a storage architecture that decouples storage software form hardware through virtualization, providing greater flexibility, performance, efficiency, and enabling greater levels of automation.

The "just-in-time" storage service delivery model of SDS allows datacenters to dynamically assign storage resources and capabilities as needed, instead of pre-allocating storage capacity and pre-provisioning services. SDS unlocks the converged operational model by moving to an application-centric approach, where the hypervisor manages and balances resources based on individual application needs. Moving to SDS also reduces the friction associated with upgrading and maintaining proprietary networking equipment, enabling datacenters to easily scale capacity, react to app-requirement changes, and control from a single location. As such, we expect SDS to become the standard as IT teams look to modernize and reduce friction in the storage process.

Traditional Storage vs. SDS

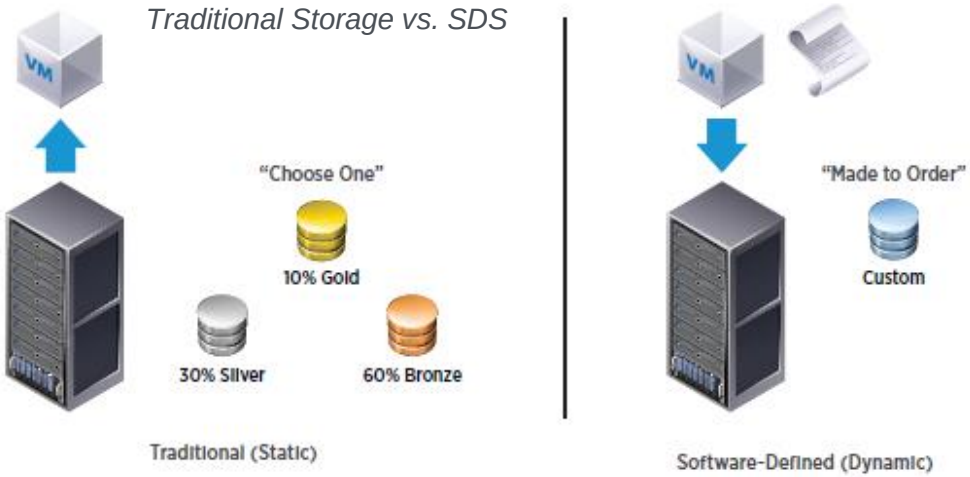

Using Windows Server 2012 to Create Software Defined Storage

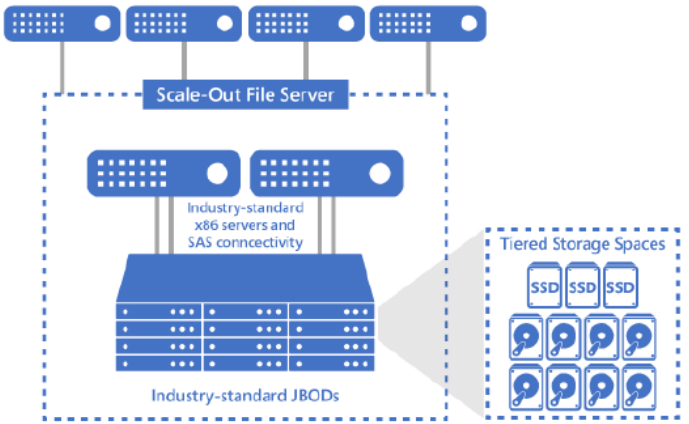

Source: Piper Sandler Research

Themes & Trends | All-Flash (SSD) is Gaining Share

Why is all-flash taking share from traditional hard-disk-drives (HDD)? While there are multiple reasons enterprises switch away from HDD disk arrays, the key factors include:

1. **Performance** – Flash can offer 10x faster throughput terms than a disk array, offering lower latency and more operations / second.
2. **Better Consumption** – Flash storage drives can consume up to 50% less power than a similar capacity HDD even with the higher throughput speeds. This in turn reduces cooling costs within Datacenters.
3. **More Reliable** – Flash drives are typically more durable, and less susceptible to catastrophic failures as traditional disk-arrays.

Because of this, Gartner believes **75% of Primary Storage will sit on all-flash arrays by 2027** (from 65% today). The downside is all-flash arrays typically come at a price premium to traditional disk arrays, though vendors like PSTG believe by utilizing cheap flash optimized with its software solutions, the overall TCO of all-flash still materially outperforms HDDs.

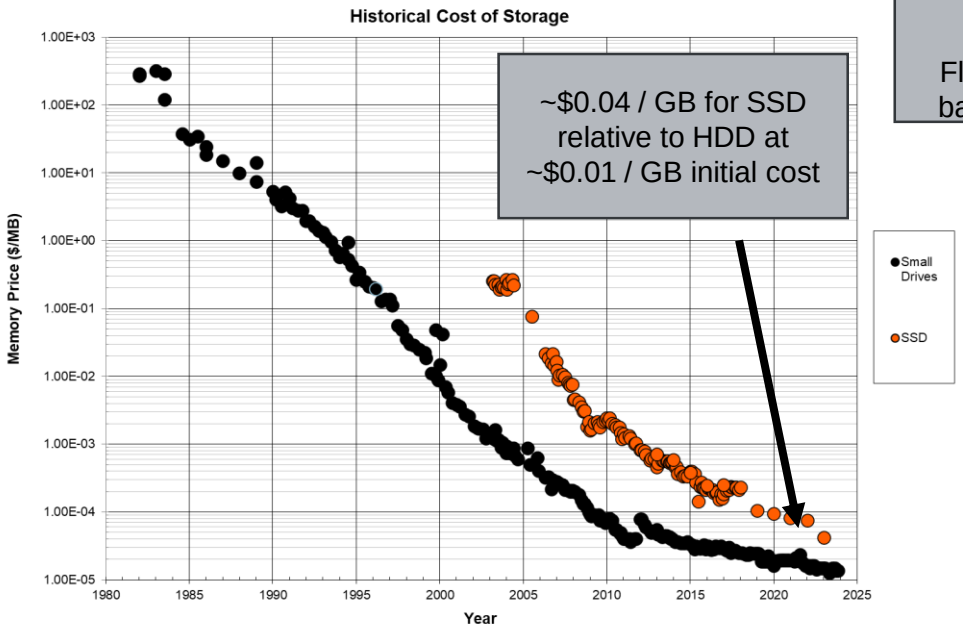

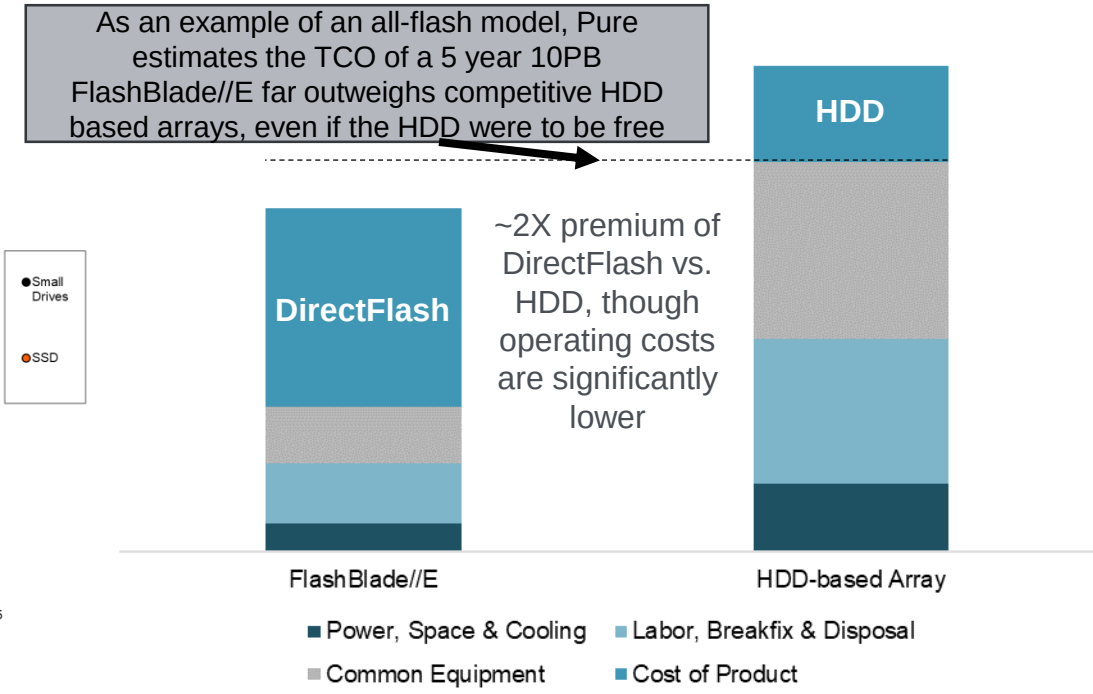

Source: Piper Sandler, Pure Storage, Gartner, Pricing plotted from https://jcmit.net/diskprice.htm

Themes & Trends | SSD Interface Improvements

The connection between storage servers or CPUs is typically called an interface, with Advanced Host Controller Interface (AHCI) commonly used as the older generation of transfer protocols. In 2008, a consortium of large IT providers created Non-Volatile Memory Express (NVMe) as a faster way to connect solid-state drives to servers. NVMe can send more commands per queue than AHCI, can communicate directly with the host CPU rather than route through a controller, and can perform parallel input/output operations for faster speeds.

Like anything, the higher performance of NVMe-based SSDs come at a higher price tag than SATA and SAS SSD interfaces, so data-heavy and ultra performant are crucial for NVMe-enabled environments. NVMe has already seen multiple enhancements including NVMe-oF (which enables memory access over network fabric such as Ethernet) and NVMe-MI as a command set for monitoring and configuring NVMe environments.

Source: Piper Sandler Research, Cisco, Seagate

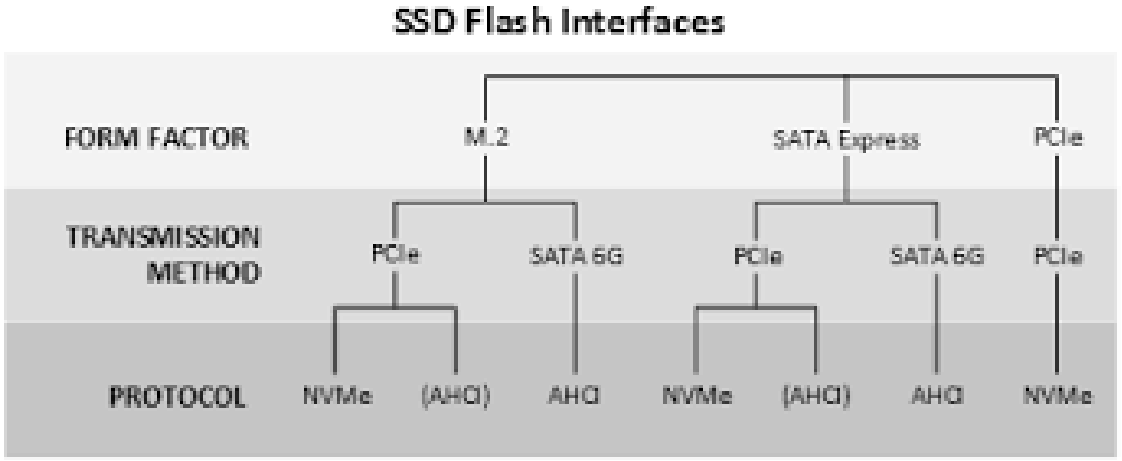

Advantages of NVMe Over Other Protocols

1. Supports up to 64k queues minimizing latency and maximizing throughput
2. NVMe connected through PCIe can provide +7.9GB/s per lane transfer speeds vs. older protocols at ~1.5GB/s per lane
3. NVMe is designed for flash memory reducing overhead vs SATA and SAS that were designed for spinning disks.

Drawbacks

1. NVMe-based PCI SSDs are more expensive and may use more power
2. NVMe is only compatible with SSD, SATA or SAS will need to be used for HDDs

Themes & Trends | AI Workloads an Accelerant for Performance

There are multiple “bottlenecks” currently hampering the data flow within artificial intelligence, with one of the most prolific being the memory bottleneck, or the fact that the GPU processor can perform operations faster than data can be moved to and from memory. Considering data is crucial across all forms of AI workloads (training, tuning &inferencing), the need for high performant memory I/O is rapidly increasing as AI workloads take off. This memory bottleneck is one of the main reasons that high performance storage, via flash, will be required for AI applications.

Data Requirements for AI Training, Fine-Tuning, and Deployment / Inferencing

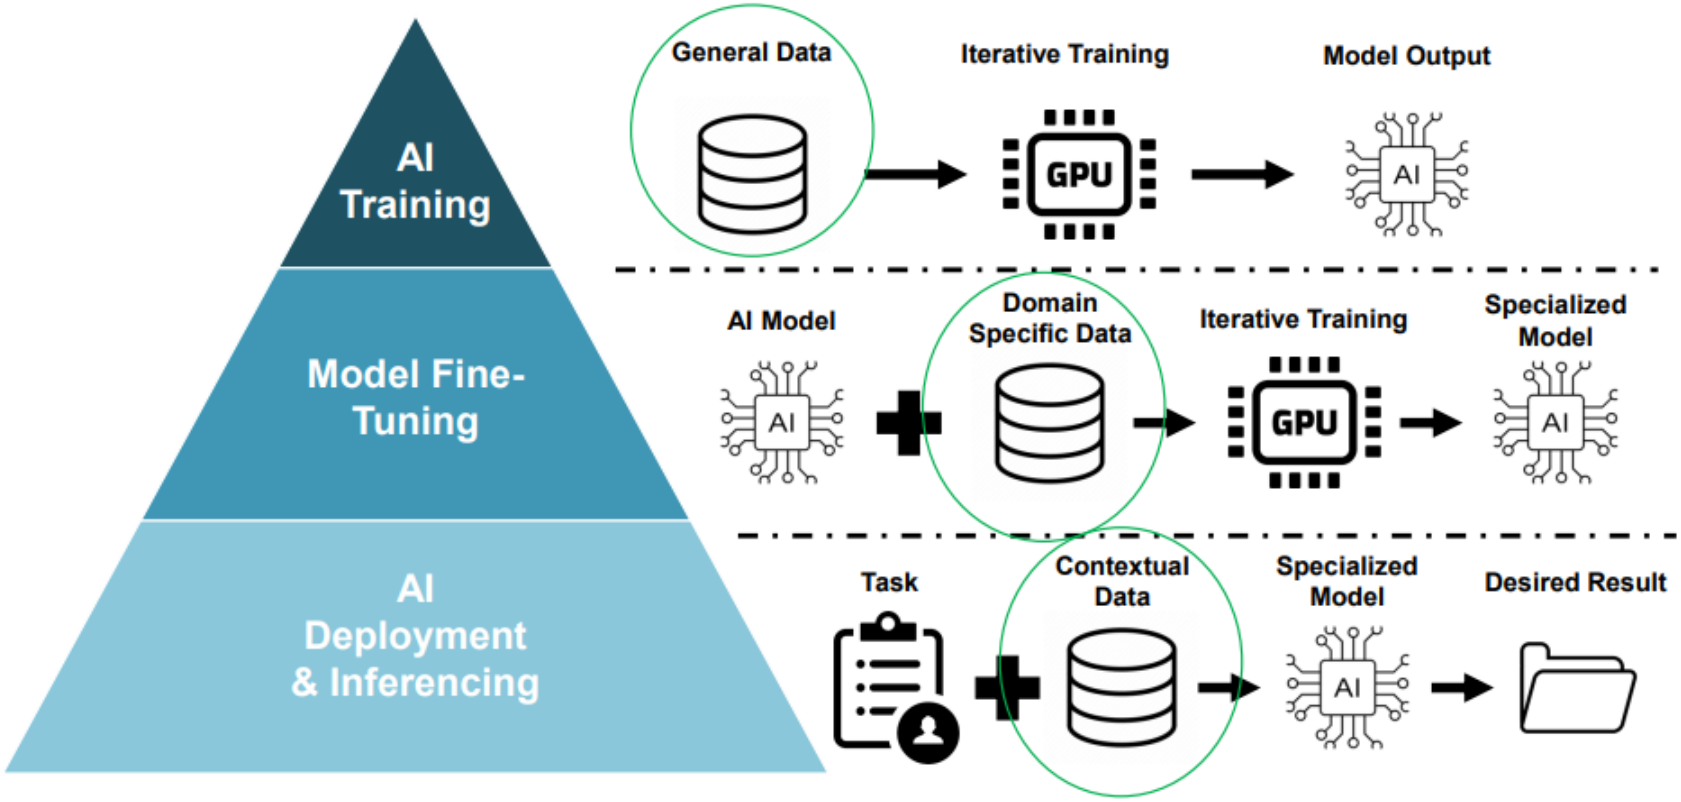

Source: Piper Sandler Research, NVIDIA, Company Reports

Themes & Trends | Hyper-Converged Infrastructure Taking Share

HCI is expected to take share of the overall storage market through 2028 as enterprise vendors rationalize hybrid environments (and HCI architectures allow for more consistent manage and control across architectures). Within our most recent partner survey, Nutanix, HPE and VMW remain the most mentioned “share gainers” in HCI, though recent CIO data points to a significant VMware replacement opportunity through 2025 and 2026

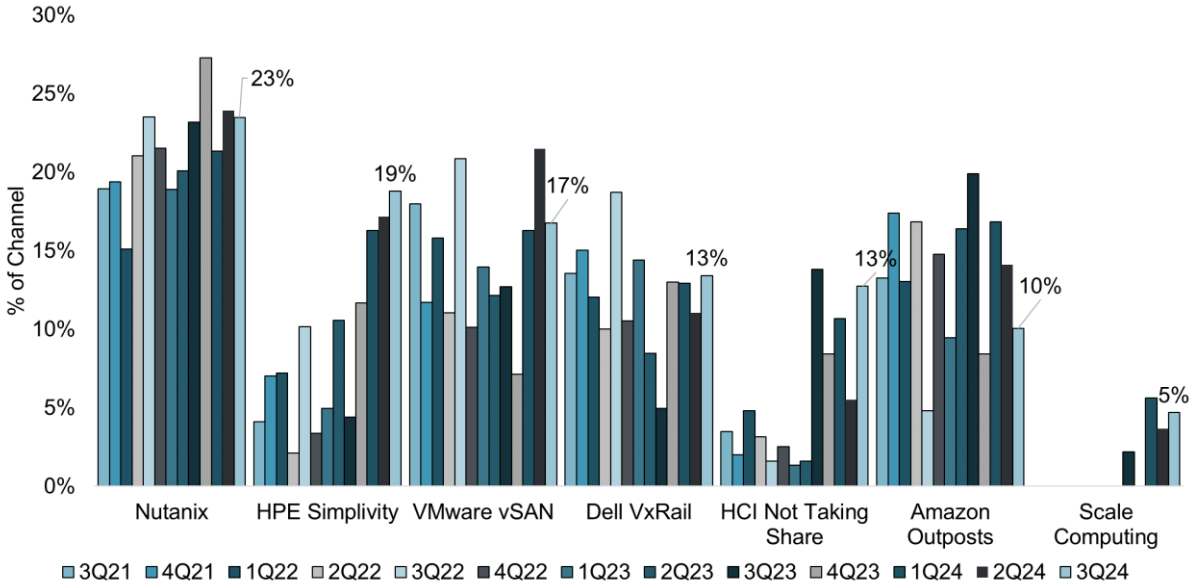

HCI to Take Share of Overall Storage Through 2028

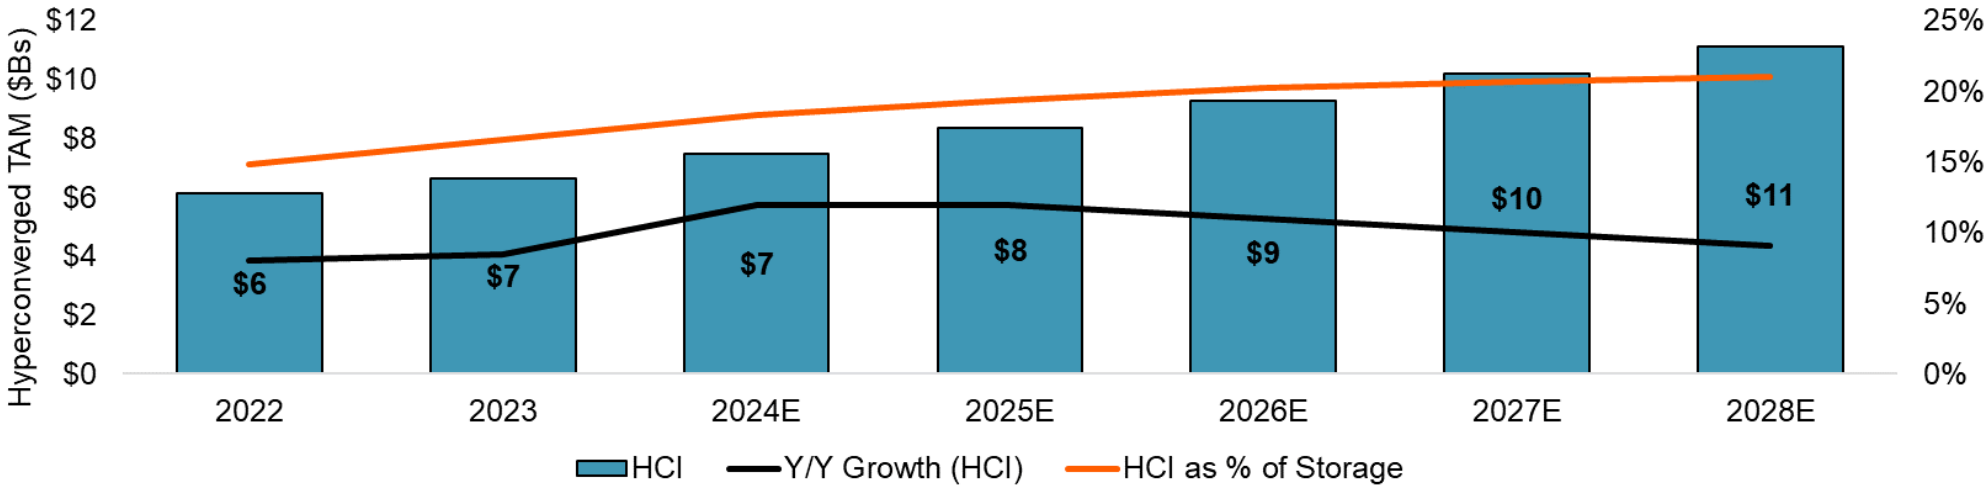

Source: Piper Sandler Research, Gartner

Themes & Trends | Storage as a Service

Gartner estimates that in 2025, **more than 40% of all on-premises IT storage admin and support costs will be replaced by managed storage as a service (STaaS) vs. less than 5% in 2020.** Massive growth in enterprise data, the flexibility of a cloud-based model, and the rise of remote work, are all contributing to the shift away from on-premise storage investments. We believe the secular shift towards STaaS should consolidate the Storage space as a whole over time, but believe Gartner may be slightly aggressive in its estimates given the push-outs in STaaS deals seen by players in the space over the last few years (ex: NetApp & Pure).

Traditional On-Premises IT Operations	Managed Storage as a Service (STaaS)
Application Support	Application Support
Change Management	Change Management
Capex/Refresh	Subscription/Renewal
Capacity Planning	Capacity On-Demand
Performance	Performance
Support Level 1 & 2	Support Level 1 & 2
Optimization	Optimization
Backup and DR	Data Protection
Config/Test/QA	Config/Test/QA

App Support is the only responsibility of the Infrastructure and Ops Teams in a STaaS model.

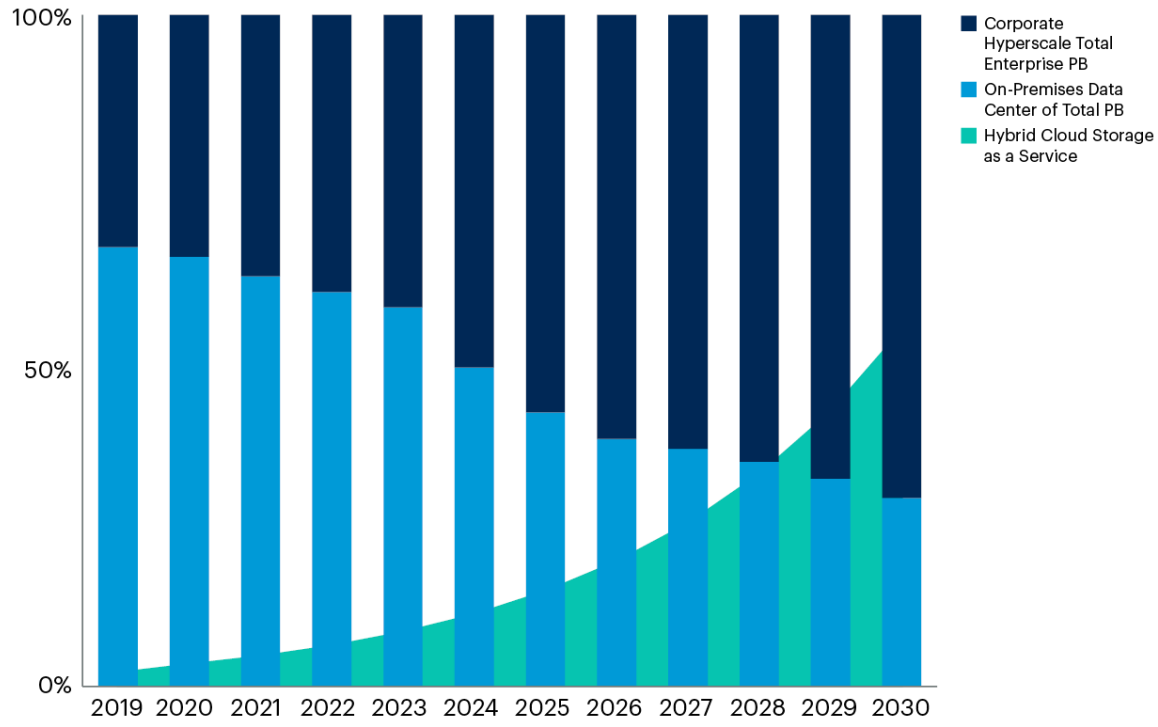

Source: Piper Sandler Research, Gartner

Themes & Trends | Container-Native Storage

A container is a unit of software that packages up code and all its dependencies for apps to run faster and more reliably across compute environments. This makes it easier to move containerized applications between environments (development, testing, production, etc.) as well as between architectures (public cloud, private cloud, on-premise).

Container-native storage refers to storage solutions designed and optimized for these containerized environments. It's typically used in big data workloads, with massive data sets that need persistent storage or high-performance computing workloads where high-throughput and low-latency is required.

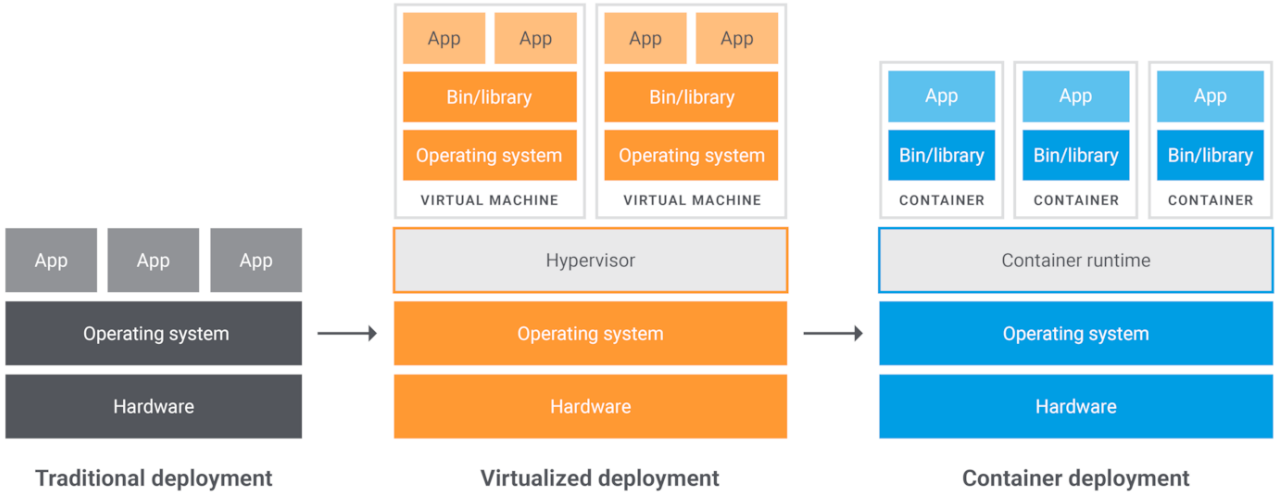

Example Container Storage Architectures

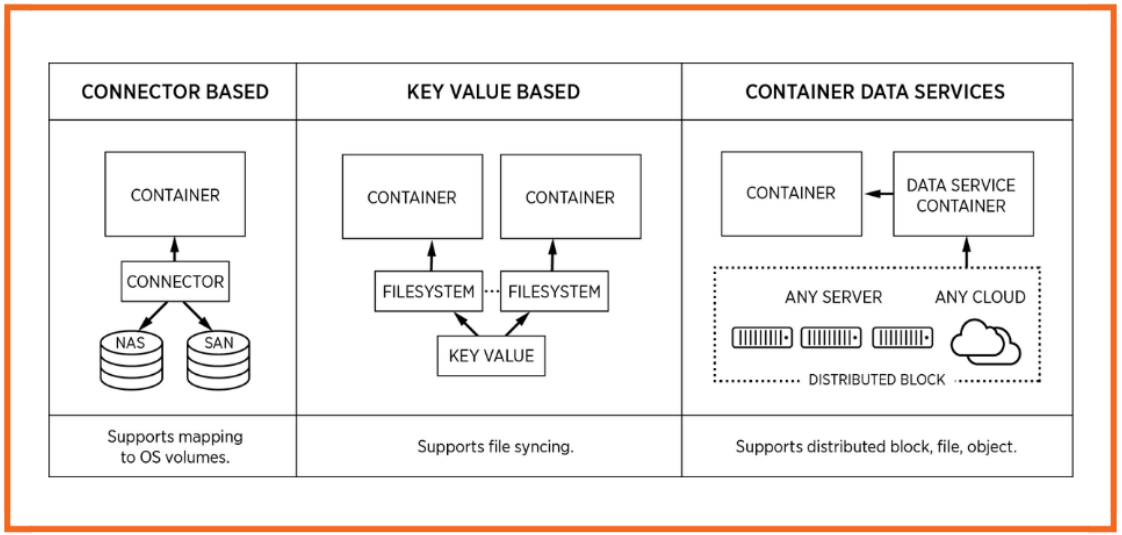

Source: Piper Sandler Research, Pure Storage, Akamai

8. Backup & Recovery Background & Dynamics

Data Management (B&R) Through the Ages

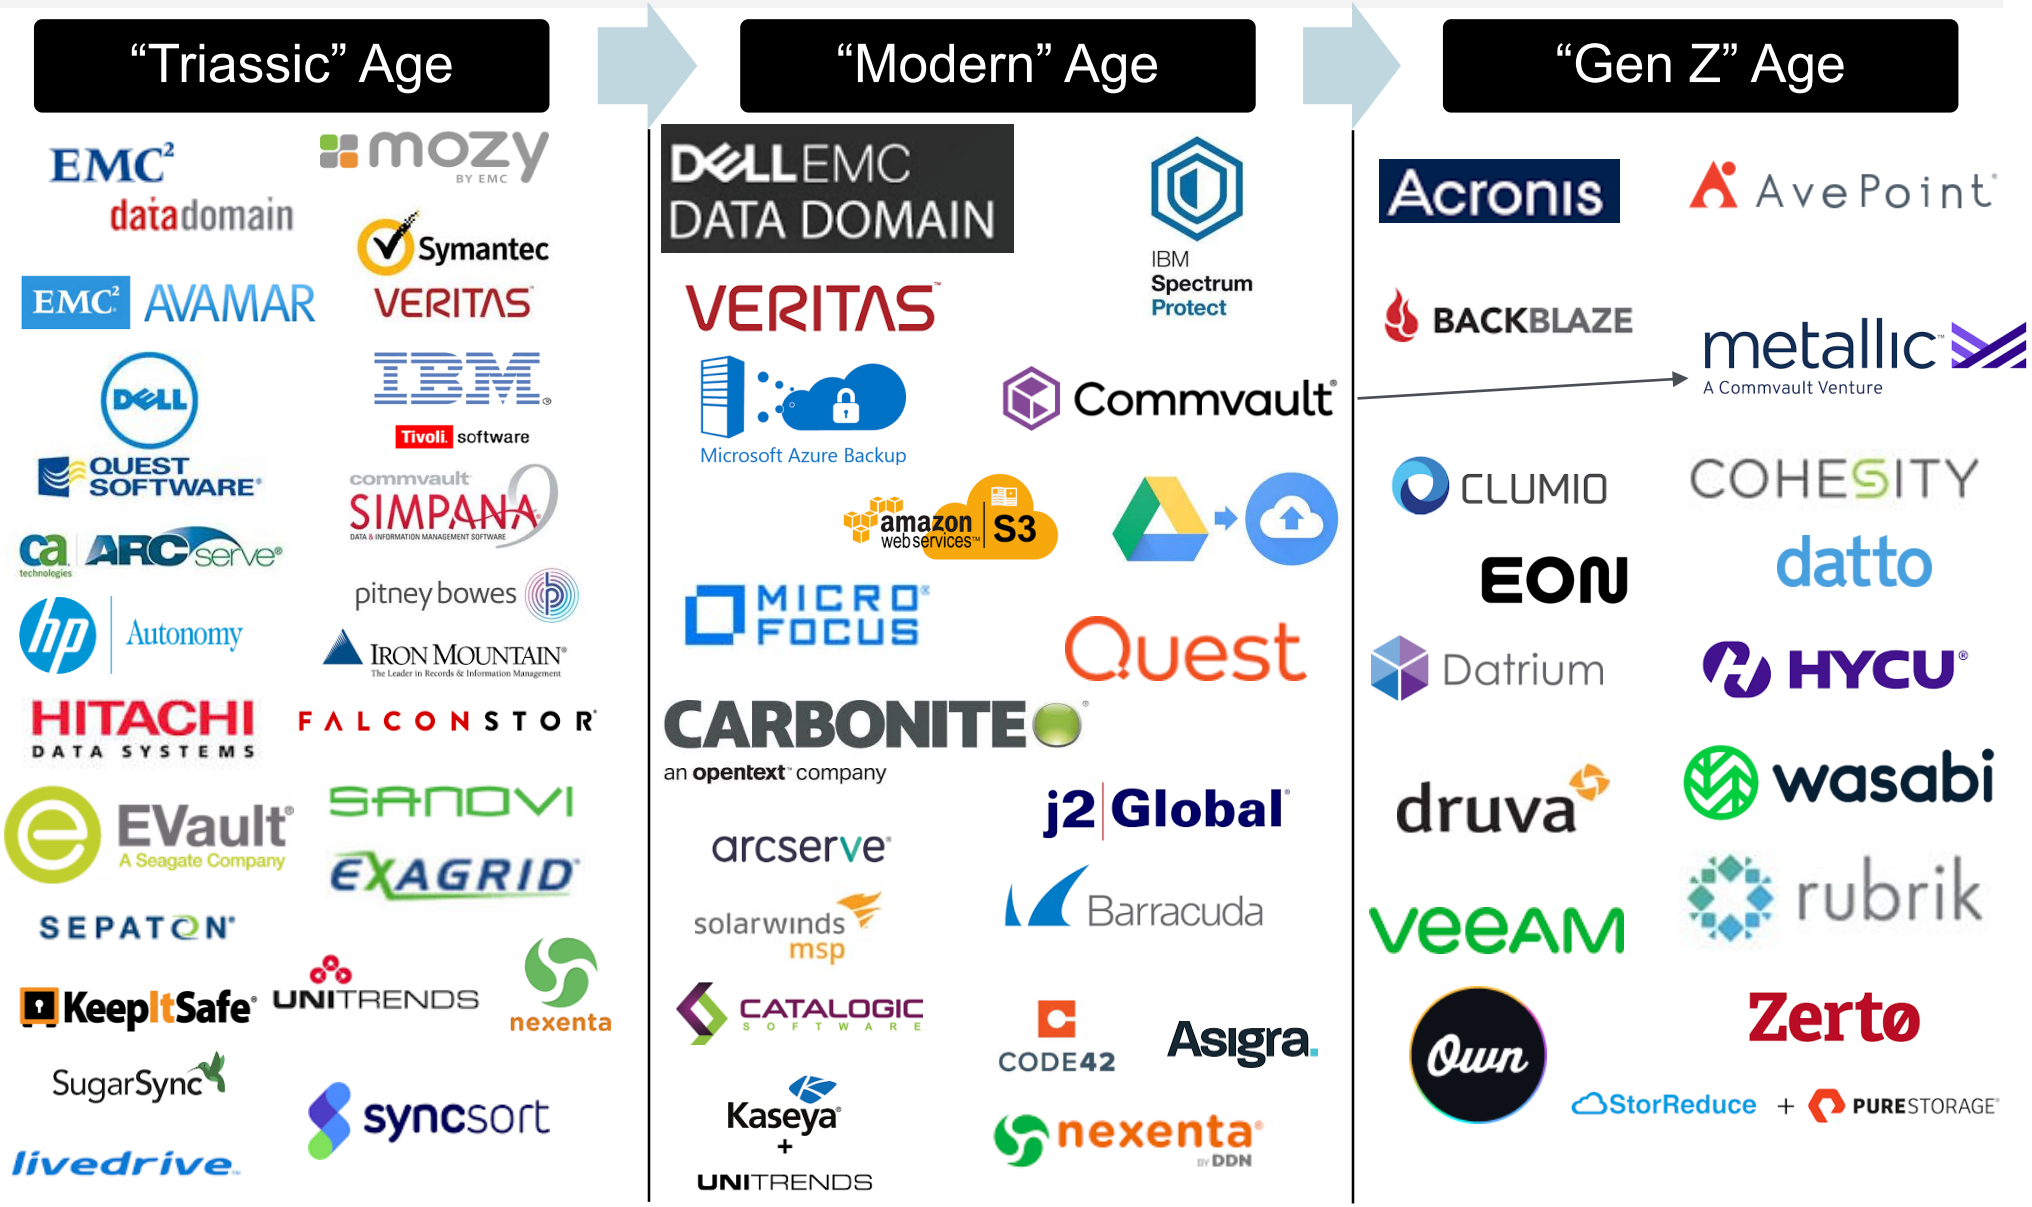

Source: Piper Sandler, Company Websites

"Gen-Z" Vendors Have the Right Architectures but What's Different?

Differentiation between "next gen players" is the #1 question we get across the space. There are various features and functionalities each vendor is "known for" in the space, but for the most part each of the up-and-coming vendors check the same boxes across cloud capabilities. Increasingly, go-to-market is the differentiating factor for why a vendor wins over others, as "getting shots on goal" is more difficult for legacy vendors.

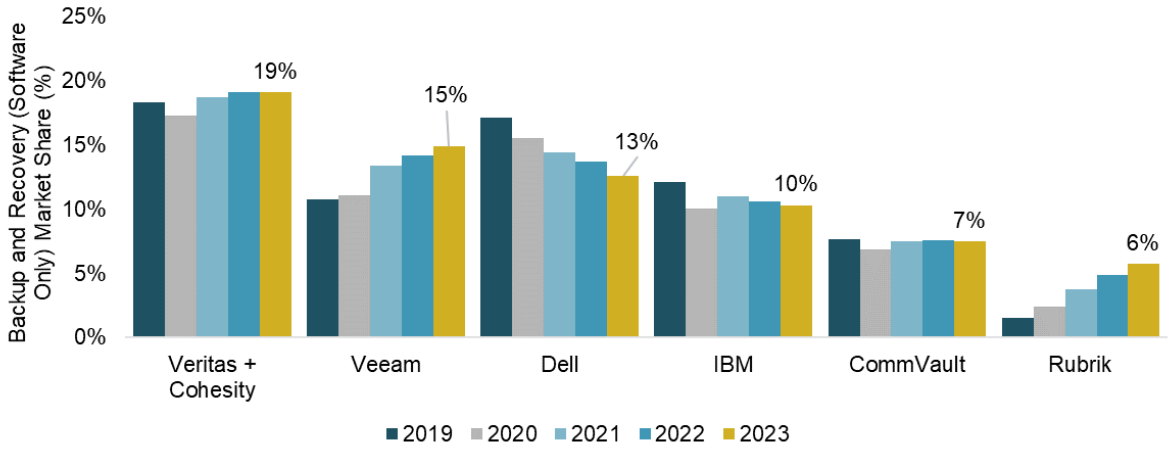

	Rubrik	Cohesity	Veritas (NB)	Commvault	Dell	Druva	HPE	IBM	Veeam
Customer Review	4.6	4.8	4.7	4.4	4.7	4.8	4.5	4.5	4.6
Product Capabilities	4.6	4.8	4.8	4.7	4.6	4.7	4.8	4.0	4.7
Implementation	4.8	4.9	4.6	4.1	5.0	4.9	4.0	4.0	4.7
Integration	4.9	4.9	4.6	4.6	5.0	4.8	4.5	4.4	4.7
Ransomware	4.7	4.6	4.7	4.5	4.4	4.7	4.3	4.6	4.4
Scalability	4.7	4.6	4.9	4.9	4.9	4.8	4.7	4.3	4.6
Simplicity	4.7	4.7	4.7	4.4	4.7	4.8	4.6	4.5	4.6
User Interface	4.7	4.7	4.7	4.2	4.5	4.7	4.7	4.2	4.6
% Recommend	91%	93%	92%	93%	93%	93%	85%	69%	99%
"Known for"	Security	Convergence & Capabilities	Scalability	Flexibility (incl. Metallic)	Hardware-Approach	True BaaS (via AWS)	Multiple Products	Reliability	Simplicity, VMs, & Leadership
Zero Trust Architecture	X	X	\	X	X	X	X	\	X
Single Control / 1 UI	X	X	\	X		X			X
Fully Extensible APIs	X	X	\	X	\	X	\	\	X
Immutable	X	X	X	X	X	X	X	X	X
Air-gapping	X	X	X	X	X	X	X	X	X

Source: Piper Sandler, Company Websites

Data Resilience Seeing Increasing Conversation / Prioritization

Given the interest in the space, we recently asked channel partners if they are seeing a material increase in customer priorities / conversations due to the outages at vendors like Microsoft and CrowdStrike.

51% of respondents noted an increase in customer priorities / conversations regarding data resilience **while only 3% note**d a decreasing priority (resulting in a +48% "net"). Following multiple ransomware / vendor outages, regulation implementations like DORA, recent IPOs, and vendor consolidation, we believe Data Resilience will remain in the limelight through 2025

How has the prioritization of data resilience and recovery (particularly for backup & recovery) following the CrowdStrike & Microsoft outages impacted data resilience spending intentions in 2H24 relative to prior quarters?

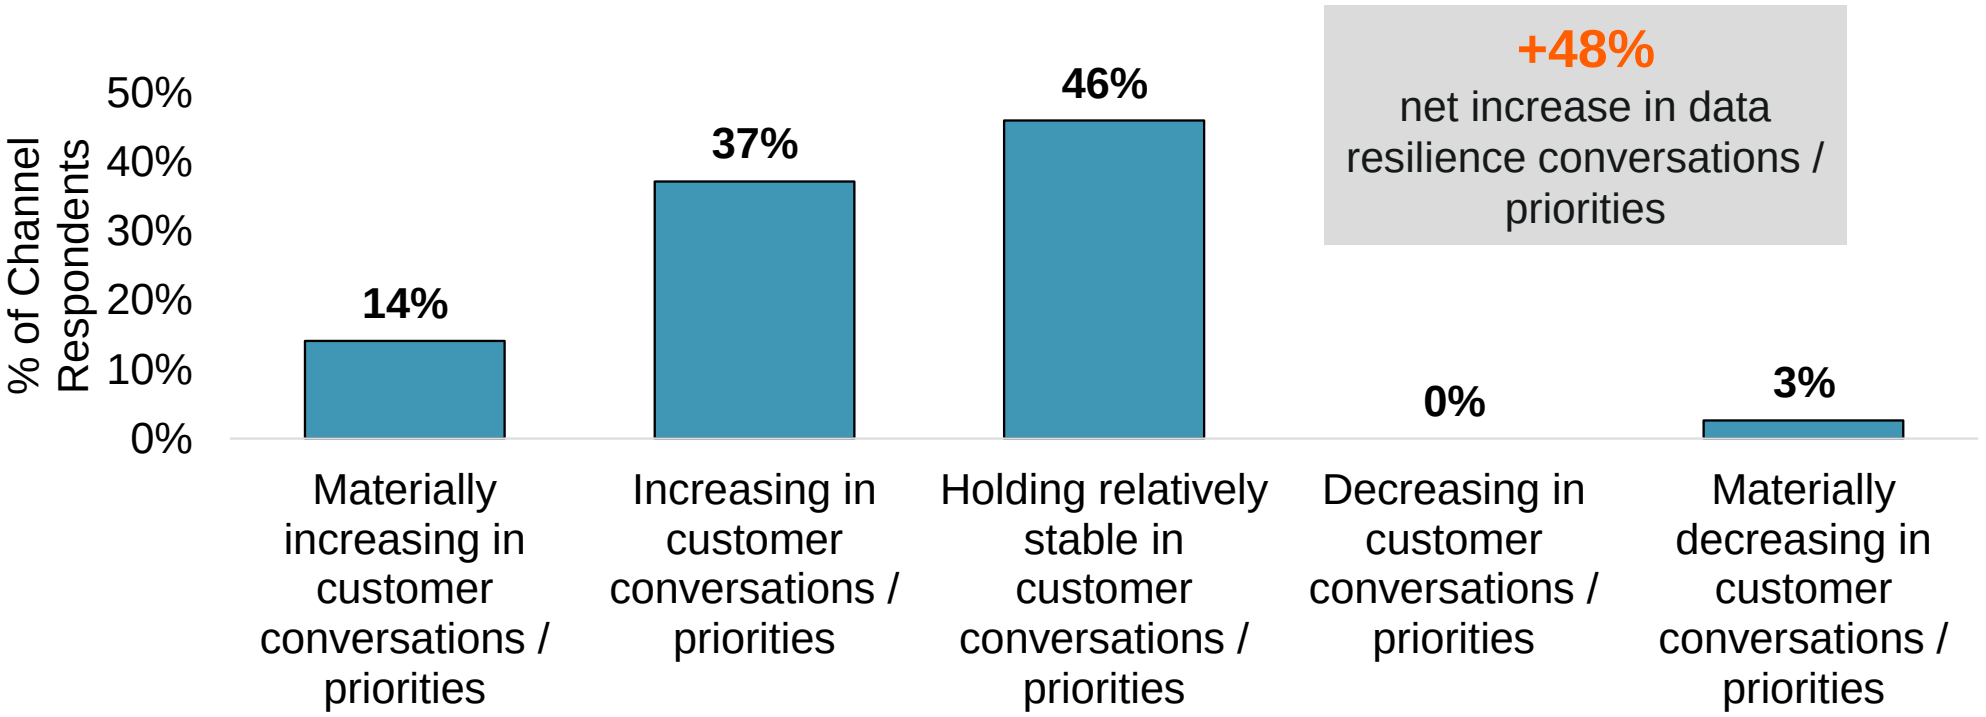

Source: Piper Sandler Research

Net Spending Intentions Backup & Recovery

2H24 Piper Sandler CIO Survey Storage Spending Intent

+67%

Net CIO intent is at all time highs for the backup and recovery space with intent bottoming two years ago

~4%

of CIOs expect to decrease spend on backup and recovery, more inline with historical trends

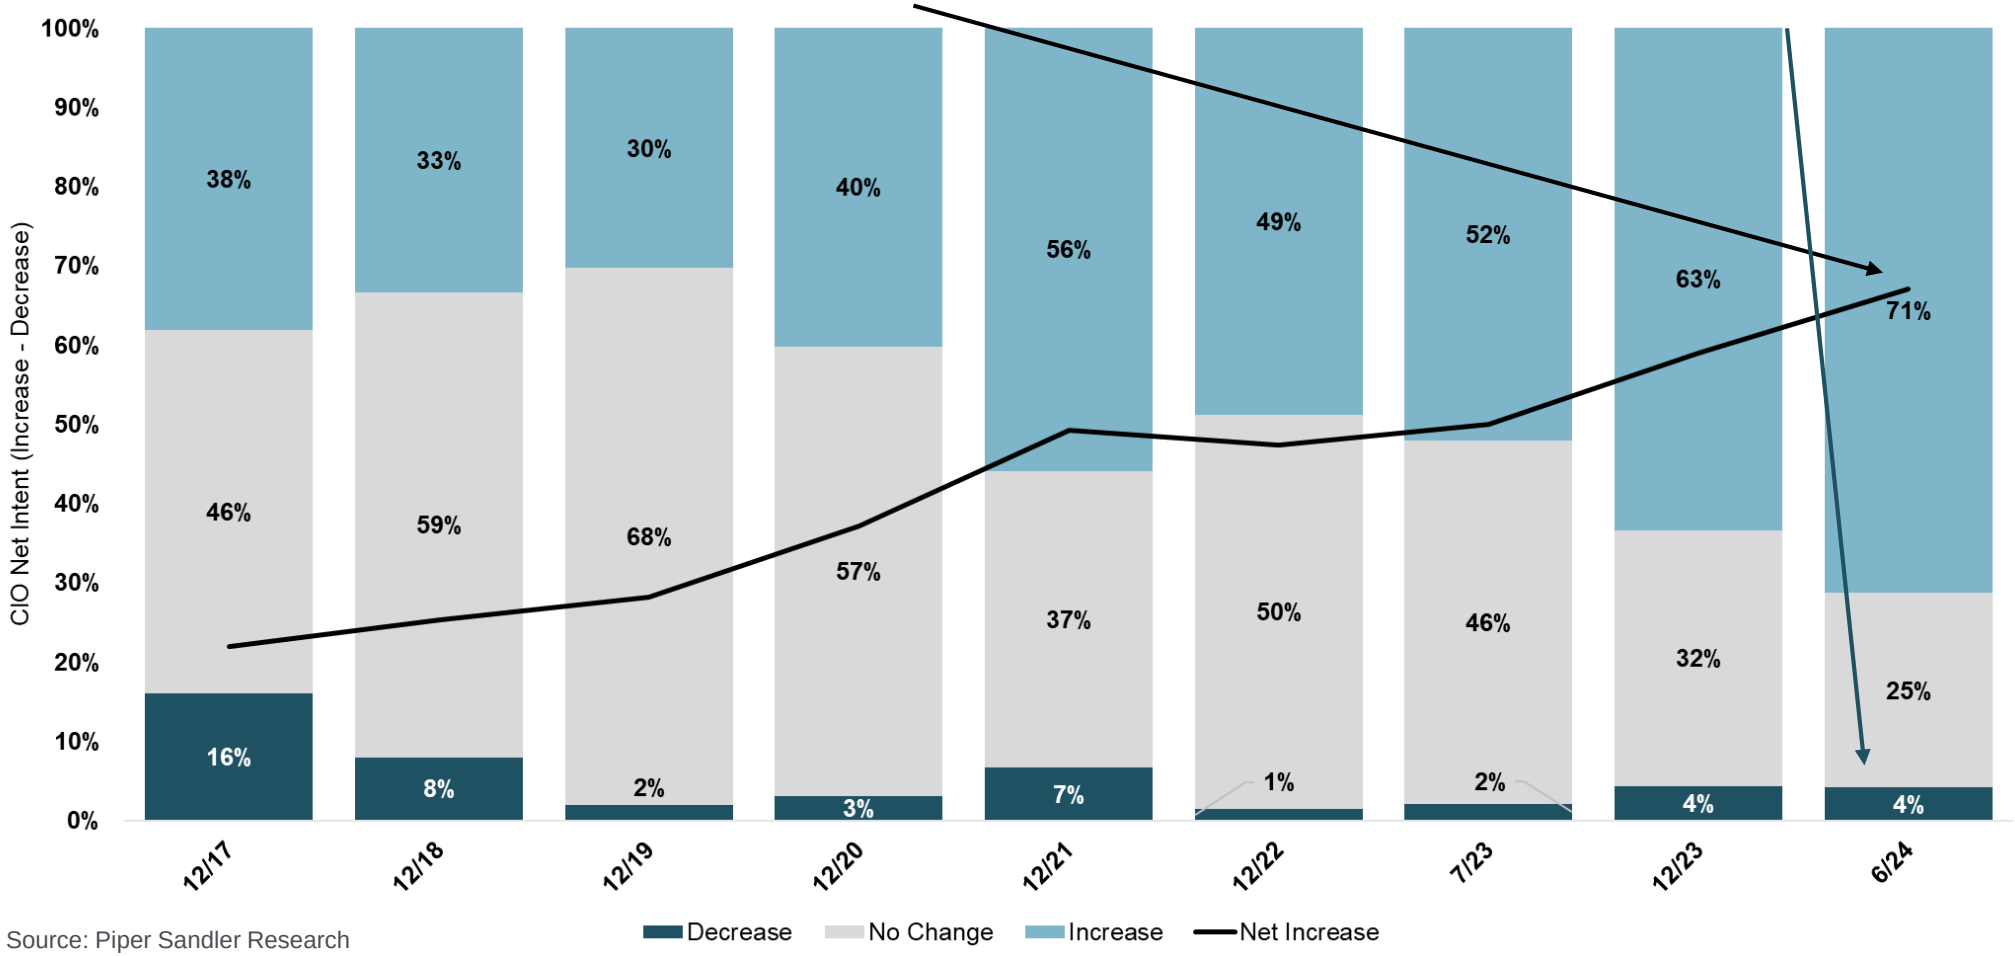

Source: Piper Sandler Research

9. Backup & Recovery Trends & Catalysts

Data Protection & Management (DPM) | Market Drivers & Dynamics

Scale-Out vs. Scale-Up: The market has largely shifted from the traditional scale-up architectures (ex: Arcserve, Dell-EMC, Veritas) towards scale-out architectures (ex: Cohesity, Commvault, Rubrik, Veeam). The mature and well-understood scale-up architecture relies on an organization having to size and configure its solution appropriately over the next few years, making it better suited for more predictable and minimal data growth environments. Scale-out (or "hyper-converged") architectures are comprised of nodes that are logically grouped together that enables an enterprise to purchase more nodes as needed. Scale-out is best-suited for unexpected or rapid-growth environments (a feature necessary in the cloud world).

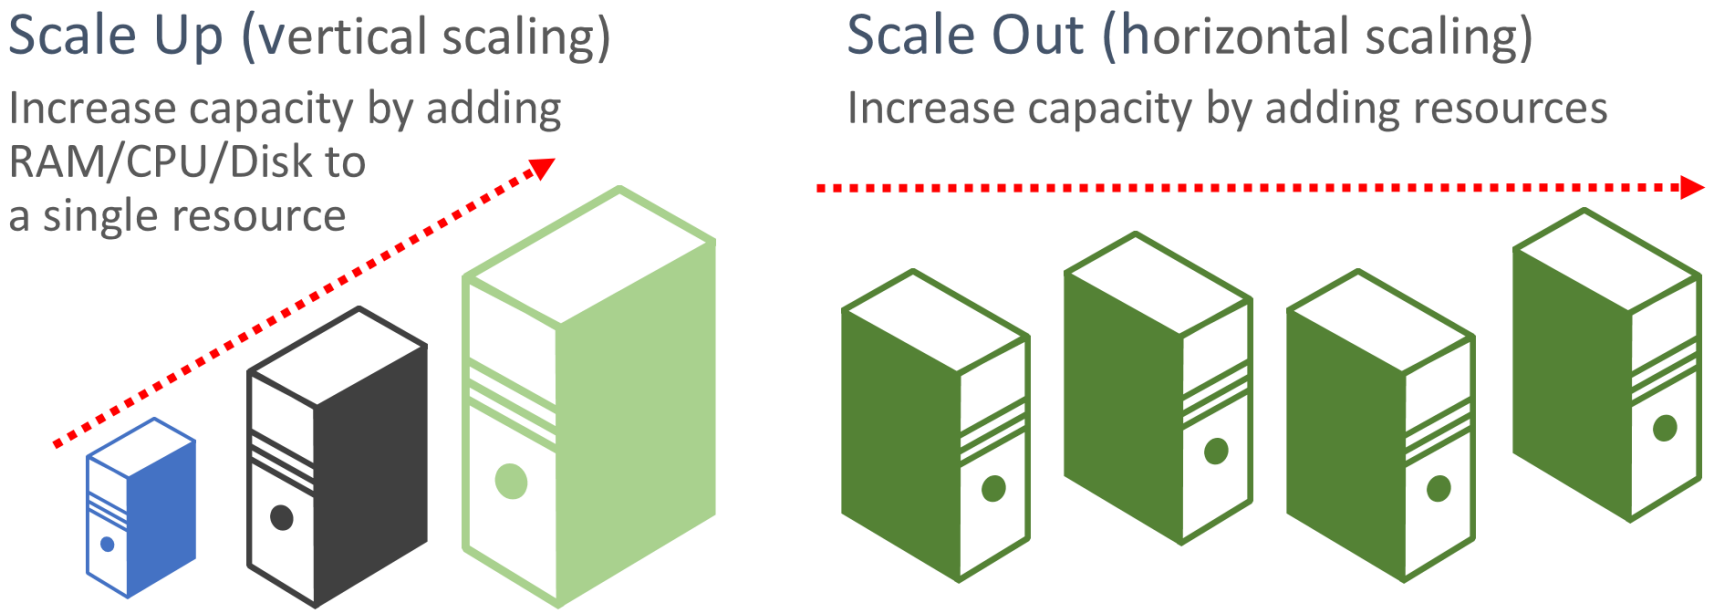

Source: Piper Sandler Research, Microsoft, F5, Inc, Gartner

DPM | Market Drivers & Dynamics

Multi-Cloud Drives Federated Needs: Many organizations will operate a hybrid environment over the next decade, which means organizations will need to back up endpoints, systems, applications, and directories / across all these components. This necessitates a "federated" solution across enterprises' environments and support for multi-cloud. Data is already spread across locations, with the "lion's share" still sitting in traditional & private cloud datacenters according to F5. The spread across multiple hyperscalers creates complexity with going with a CSP backup offering.

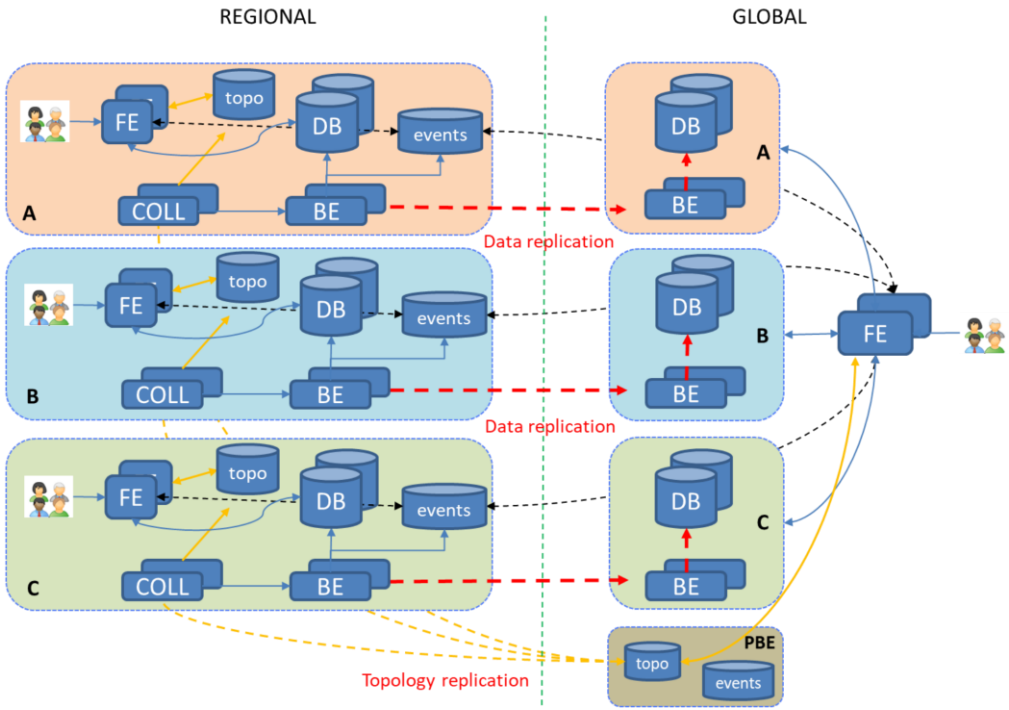

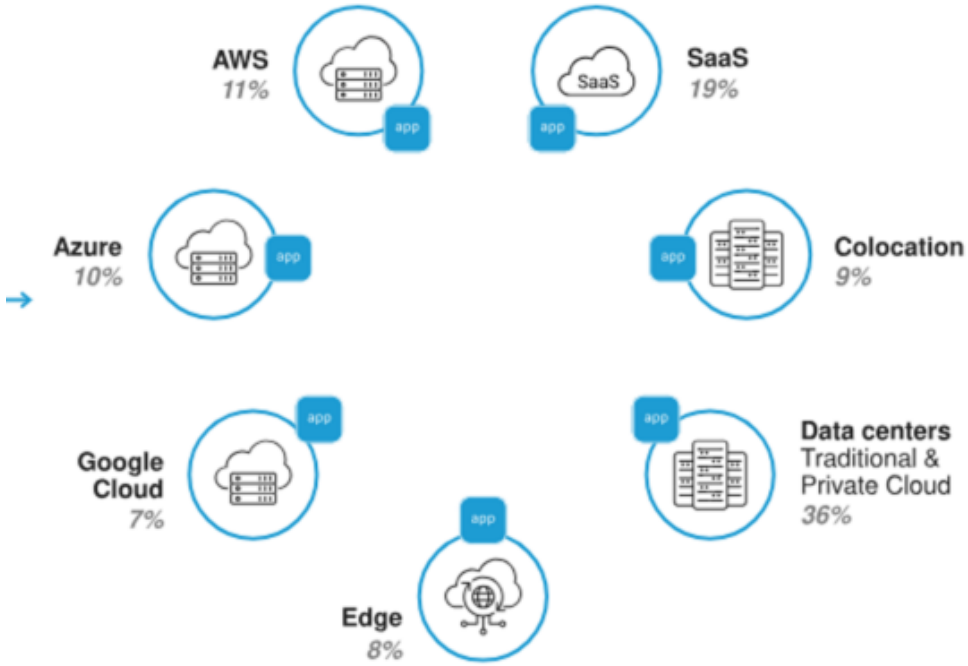

Source: Piper Sandler, Microsoft, F5, Inc, Gartner

DPM | Market Drivers & Dynamics

Accelerating Cloud-Application Usage: Enterprises are increasingly using more cloud-based applications. To avoid vendor lock-in, and have longer retention periods and better security and compliance postures, enterprises should leverage third-party data-management vendors. Most of the large cloud applications like O365 (the most in-demand app for backup) and Salesforce expose their APIs to enable data management. O365 is one of the fastest growing applications but only approximately 15% of users today are backed-up, creating a ~$2B opportunity today that could expand to $5B over time alone at a measly $2/m/u price point. Other applications like Salesforce, ServiceNow, Workday, and other Microsoft apps (Teams, Sharepoint, etc.) could provide even further market expansion and necessary SaaS backup. Rubrik's State of Data Security research report found SaaS data grew the fastest, while overall data grew 42% over the prior 18 months.

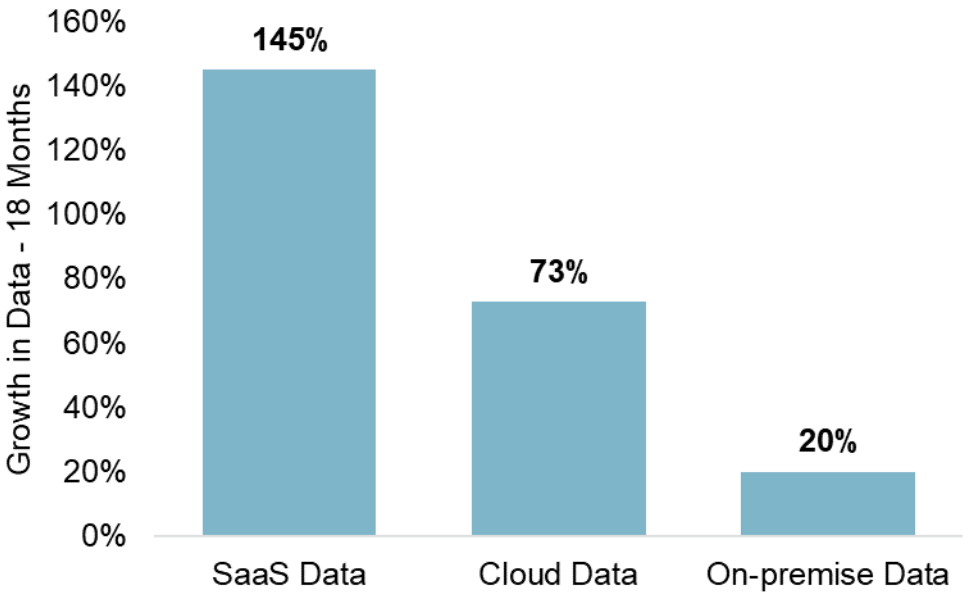

Microsoft O365 Opportunity

~60M or 15%
of O365 users are being backed up as of mid-2024 vs. 10% last year

→

~$5B
opportunity alone at $2/m/u price point and 50% penetration

Source: Piper Sandler, F5, Inc, Gartner, Microsoft, Rubrik

DPM | Market Drivers & Dynamics

Keep An Eye on the Regulation: GDPR and related regulations like the California Consumer Privacy Act, the Digital Operational Resilience Act (DORA) and India's Personal Data Protection Bill, promote the usage of data management solutions for data protection use-cases (ex: The Right to be Forgotten). With different states and nations having different regulations, in our view, data management vendors will increasingly be relied upon to ensure enterprises are meeting different requirements across the globe.

Commvault's GDPR Strategy

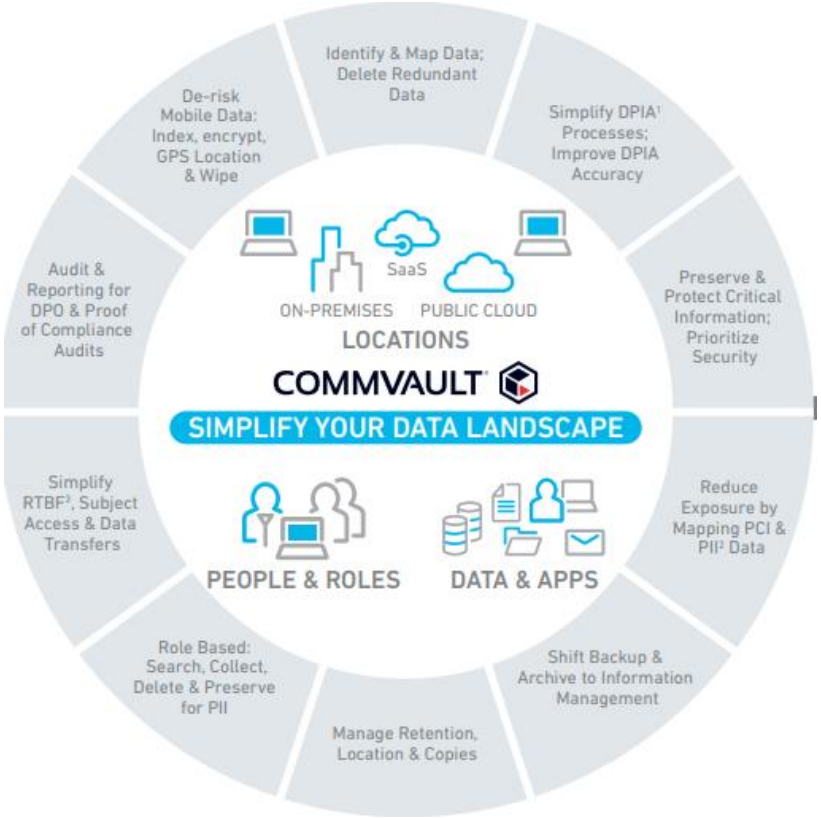

5 Pillars of DORA

- Risk Management
- Incident Management
- Resilience Testing
- Management of Third-Party Risk
- Information Sharing

Source: Piper Sandler Research, Commvault, Gartner, Cohesity

DPM | Market Drivers & Dynamics

Data Discovery: Data discovery is the process of detecting patterns and outliers within existing data by navigating stored data and applying advanced analytics. The traditional backup and recovery space has been slowly shifting towards a more "Data Centric" pitch, with add-on offerings like Data Analytics, Data Discovery / eDiscovery, and Data Visibility (including for metadata) aimed, at helping customers get more use out of the existing data "at rest".

Competitive Disruptions: The space has seen major changes and consolidation over the last decade, with the largest players having critical changes. EMC was acquired by Dell, which now has >10 different systems across its data management suite. Veritas was first sold by Symantec, and just recently combined with Cohesity. Most of the "Modern Period" data management vendors are either large conglomerates, bundlers of solutions, or are owned by a form of private-equity or private-equity like entities, with acquisitions in the space accelerating in recent years.

Date	Target / Merger	Acquirer / Merger	EV	Rev	EV / Rev
Sep-24	Clumio	Commvault	NA	NA	NA
Sep-24	Own	Salesforce	$1,900	NA	NA
Apr-24	Appranix	Commvault	NA	NA	NA
Feb-24	Veritas*	Cohesity*	$7,000	$1,600	4.4x
Aug-23	Laminar	Rubrik	$225	$105	2.1x
Aug-22	Micro Focus International	Open Text	$6,000	$2,727	2.2x
Apr-22	Datto	Kaseya	$6,200	723	8.6x
Apr-22	Barracuda Networks (TB)	KKR	$4,000	$500	8.0x
Nov-21	Quest Software	Clearlake Capital	$5,400	$858	6.3x
Jul-21	Zerto	HPE	$374	$130	2.9x
Mar-21	StorageCraft*	Arcserve*	NA	NA	NA
Dec-20	Actifio	Google	$1,300	$221	5.9x
Nov-20	sfApex	Druva	NA	NA	NA
Sep-20	Portworx	Pure Storage	$370	NA	NA
Jul-20	Datrium	VMware	$137	NA	NA
Mar-20	Xmedius	OpenText	$75	$40	1.9x
Jan-20	Veeam Software	Insight Software	$5,000	$1,000	5.0x
Nov-19	Carbonite	OpenText	$1,420	$450	3.2x
Aug-18	StorReduce	Pure Storage	$27	NA	NA
Feb-18	Mozy	Carbonite	$146	NA	NA
Nov-17	Barracuda Networks	Thoma Bravo	$1,272	$353	3.6x
Jan-16	Veritas	Carlyle Group	$7,400	$2,528	2.9x
Oct-15	EMC*	Dell*	$64,671	$24,750	2.6x
Jul-14	Arcserve	Marlin Equity Partners	NA	NA	NA
Oct-07	Voyence	EMC	$100	$10	10.0x
Nov-06	Avamar	EMC Corporation	$165	NA	NA
Dec-04	Veritas*	Symantec	$13,500	$2,042	6.6x
Feb-96	Tivoli	IBM	$743	$50	14.9x
Avg			**$5,706**	**$2,240**	**5.4x**
Avg Over the Last 5 Years			**$2,885**	**$759**	**4.6x**

Source: Piper Sandler, Gartner, FactSet

DPM | Market Drivers & Dynamics

AI & Backup: Artificial Intelligence is going to significantly change the Data Management & Protection markets. Attackers can leverage AI for attacks, meaning organizations need to be utilizing AI & ML to combat new techniques rather than relying on human intervention. Additionally with a proper solution in place that leverages AI, LLMs can better utilize the data stored by backup solutions. Backups leveraging AI will automatically back up data based on priorities, assess the environment to ensure processes are working correctly, compare data to regulatory requirements, check data quality faster and more precisely, track unusual activity, and leverage AI assistants to remove manual processes.

Ransomware: Ransomware attacks have become rampant, with bad actors even running "Ransomware-as-a-Service". Attackers aren't just targeting a single victim, but entire supply chains, typically attacking unpatched systems and VPNs, and entering via phishing campaigns. High profile ransomware attacks such as NotPetya, the Colonial Pipeline attack, & WannaCry have occurred in recent years. Chainanalysis estimates >$1B of ransomware payments in 2023, Verizon estimates ransomware attacks were involved in 24% of all breaches in 2023, and Veeam found 85% of organizations suffered at least 1 ransomware attack over the last year. The debate of course is whether to pay the ransom and risk facing further extortion, as customers may not know if their backups are clean of the malware that could lead to further ransomware attacks. This makes having an immutable backup paramount to one's security posture. Organizations are now required to disclose when breaches have occurred, putting further emphasis on security solutions (including data security).

Back up data regularly on external media or immutable storage. Refrain from relying exclusively on cloud storage.
– Palo Alto Networks' Unit42 on safeguarding against Ransomware

By augmenting human expertise and strengthening defense measures, AI-driven cybersecurity solutions provide a robust shield against a vast array of attacks.
– Check Point on preventing ransomware

Verizon: Ransomware in Breaches

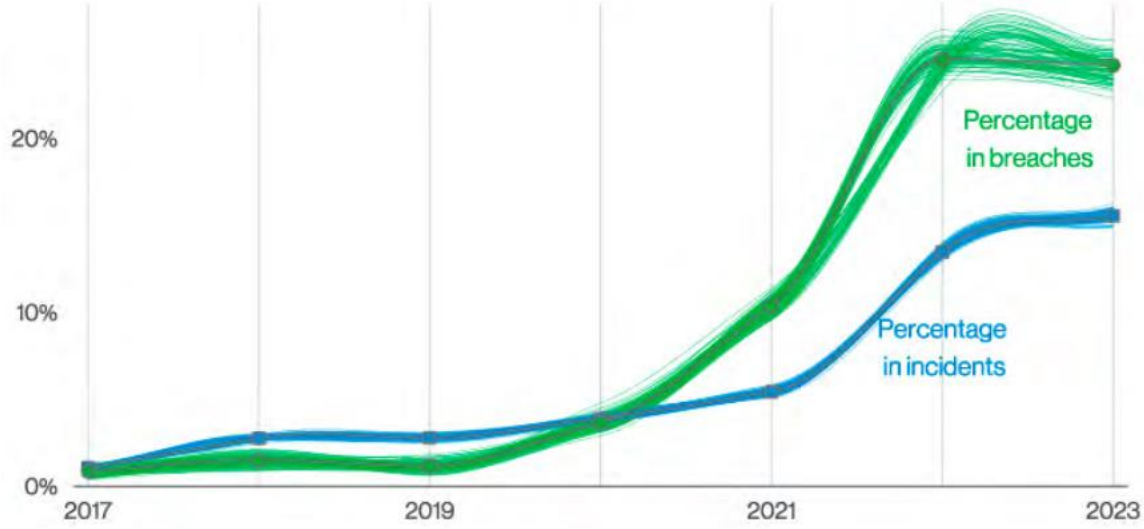

Source: Piper Sandler, Verizon 2023 Data Breach Investigations Report, Palo Alto & Check Point (Cov'd by R. Owens)

DPM | Market Drivers & Dynamics

Management PLUS Protection: Backup and Recovery vendors have increasingly talked about "shifting left", essentially representing a move **away** from the traditional backup and recovery offering **towards** a Data Security offering. This is being more recently classified as the Data Security Posture Management Space (DSPM). The goal of DSPM is to focus on the overall data security posture of the organization, including on-premises and cloud environments, by helping customers identify and assess risks, monitor controls, and plan incident response. Historically, backup vendors have played a role in data security by protecting **data at rest,** particularly with ransomware as described above. However, backup/recovery providers now sit in a unique opportunity to protect **data in motion by moving into the DSPM** market. By shifting left, backup vendors can attack a much faster growing space, though begin to compete with larger competitors. For example, Microsoft launched Purview for data governance, private player Wiz has shown momentum, Rubrik acquired Laminar, CrowdStrike acquired Flow, and Palo Alto Networks acquired Dig Security in DSPM, all while traditional DLP / data security vendors shift towards this next generation.

DSPM provides organizations with an approach to protecting cloud data by ensuring sensitive and regulated data have the correct security posture, **regardless of where the data resides or is moved to**.

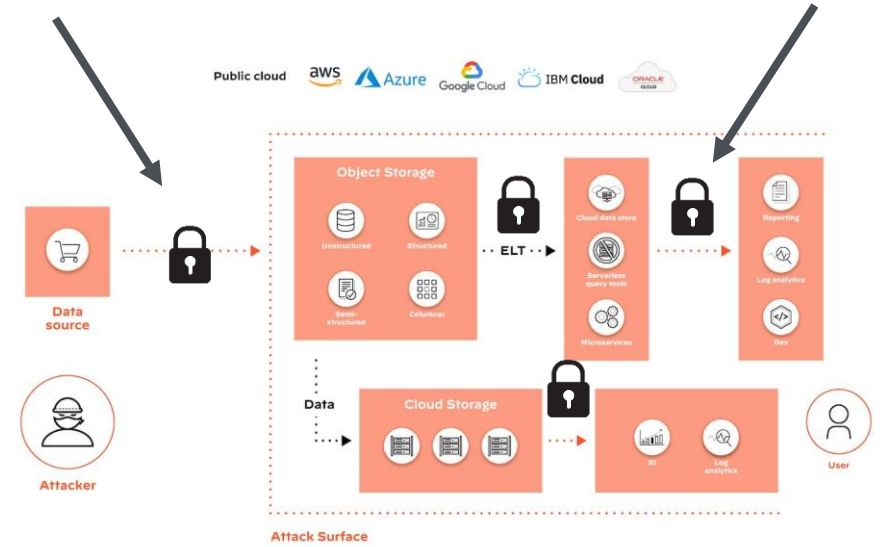

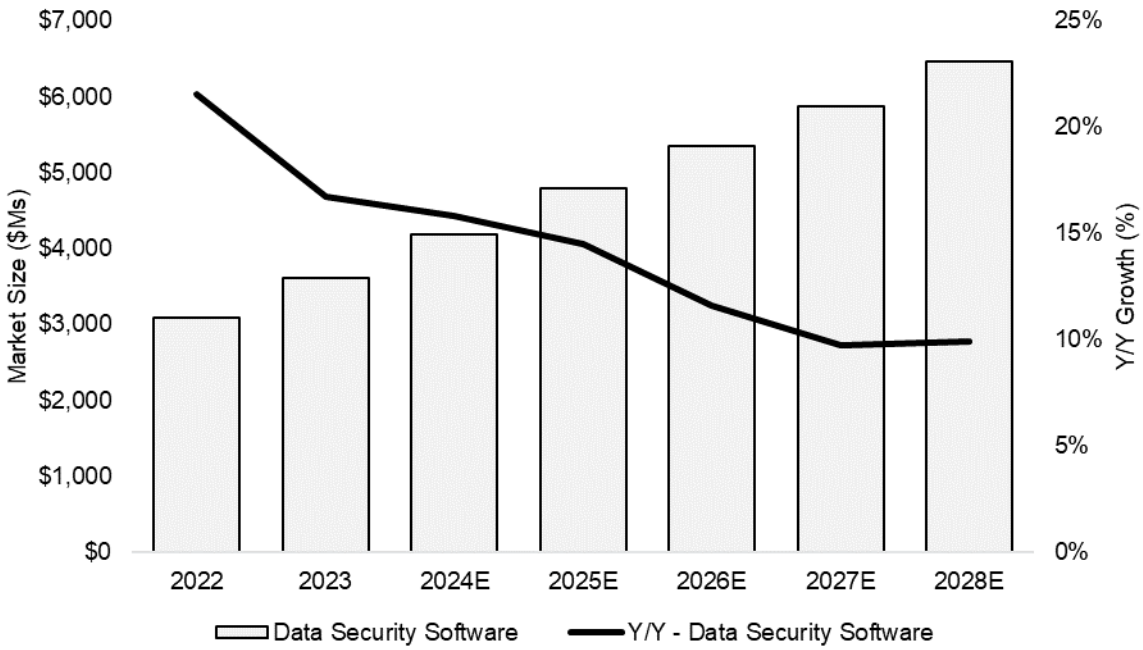

Source: Piper Sandler, Gartner, Palo Alto

10. Public & Private Beneficiaries

Public Vendors – Servers / Chips / Silicon

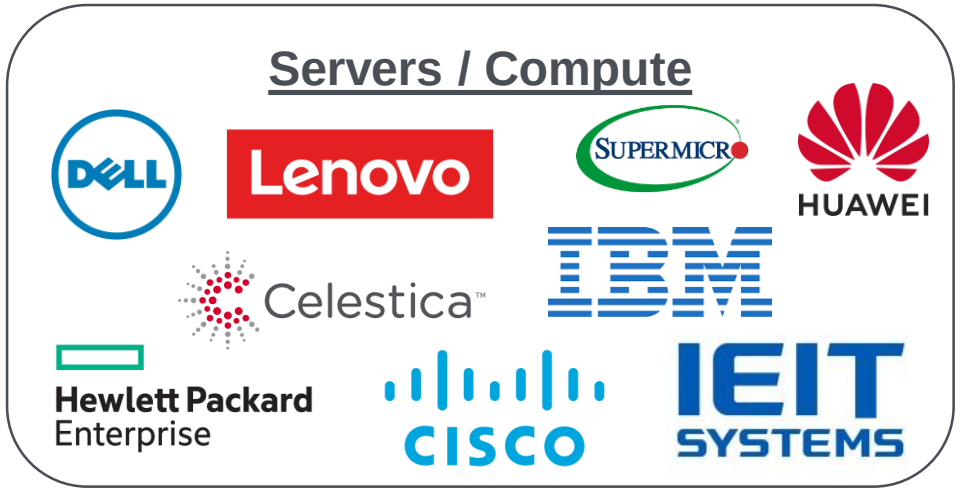

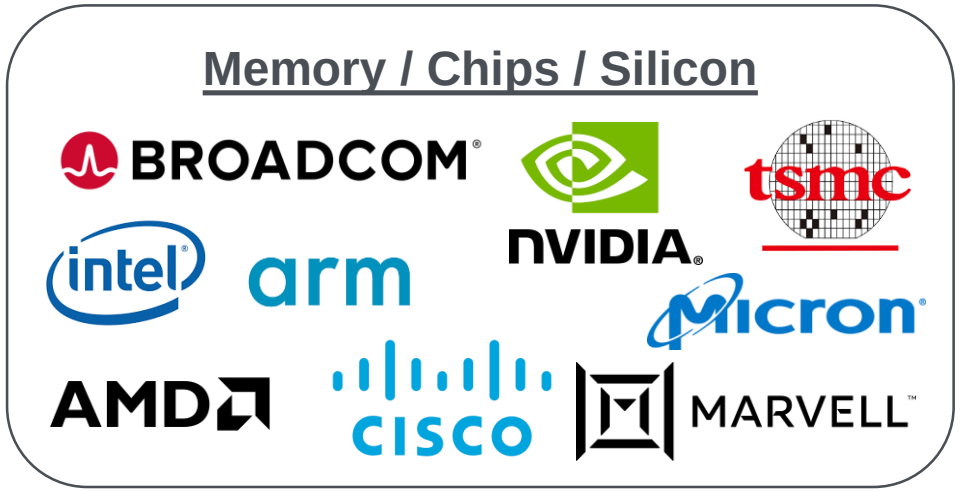

AVGO – OW, $250 PT

Covering Analyst: Harsh Kumar

Investment Thesis

We view AVGO as the best positioned custom ASIC provider currently and primary beneficiary of a move to custom silicon given increased power efficiency requirements coupled with higher data rates. The company currently cites three customers for its ASIC product, however we feel AVGO has the potential to increase this customer count, reducing key customer risk in this business.

MRVL – OW, $120 PT

Covering Analyst: Harsh Kumar

Investment Thesis

The expanded agreement to collaborate with AWS has allowed MRVL to strengthen its positioning within the broader custom ASIC market. The agreement also calls for collaboration on multiple fronts in which MRVL will support AWS with additional data center solutions beyond its ASIC chip. MRVL also expects AI revenues to track ahead of the original $1.5B and $2.5B company estimates for this year and next year, respectively.

Source: Piper Sandler Research, Company Materials & Websites

Public Vendors – Servers / Chips / Silicon

Investment Theses

NVDA – OW, $175 PT

Covering Analyst: Harsh Kumar

We currently view NVDA as our top large-cap pick in our coverage universe given the company's dominant positioning in AI accelerators and the upcoming launch of the Blackwell architecture. Our viewpoint is rooted in the belief that the overall TAM for AI accelerators to rise in 2025 and we see NVDA well positioned to capture most of the incremental TAM increases while ceding only a small bit to its merchant chip competition.

AMD – OW, $180 PT

Covering Analyst: Harsh Kumar

With each new architecture launch, AMD continues to position itself as a key second source to the accelerator market. Importantly, AMD's chips are cheaper on a relative basis to the competition and provide customers with a TCO advantage. Within the upcoming two years, we believe the acquisition of ZT Systems will aid in closing the gap in overall system performance between the two merchant GPU suppliers.

CSCO – N, $57 PT

Covering Analyst: James Fish

While Cisco is mainly known for its networking segment within the datacenter landscape, the company's acquisitions and investments have led to a strategy centered around silicon and optics (ex: Acacia) via "SiliconOne". There is some debate the SiliconOne wins cannibalize Cisco's systems-base, but it does allow for Cisco to win new or retain footprint that would have gone potentially to an ODM. We still view Enterprise networking spending as Cisco's biggest driver, but the SiliconOne strategy is a strong approach to increasing Cloud & AI exposure.

Source: Piper Sandler Research, Company Websites

Public Vendors – Primary Storage

PSTG – OW, $76 PT

Covering Analyst: James Fish

Investment Thesis

Following Pure's win of a "Top 4" hyperscaler, we upgraded the name to Overweight based on: 1) Hyperscaler win creating a catalyst for upside, with opportunity for others to follow; 2) Pure's differentiation in all-flash, which is gaining share of the storage market; and 3) we see further estimate upside and room for multiple expansion, while the narrative around the hyperscaler win gives long-only investors something to play for outside of the stalling STaaS transition.

NTNX – OW, $83 PT

Covering Analyst: James Fish

Investment Thesis

Nutanix has been our top-pick as the company has multiple catalysts. With the subscription transition inflecting higher, hyperconverged infrastructure gaining share of the broader storage market, a favorable competitive landscape that gives Nutanix a direct benefit via Cisco and a indirect catalyst via VMware, an improved go-to-market motion, new opportunity around standalone hypervisor attach, & FCF upside, we see room for estimate and multiple upside.

Source: Piper Sandler, Company Websites

Public Vendors – Backup & Recovery

RBRK – OW, $72 PT

Covering Analyst: James Fish

Investment Thesis

We recently initiated coverage, with a thesis based upon: 1) Attractive competitive landscape with ~50% of the Data Protection & Management market held by legacy players; 2) Ransomware, data growth, A.I. and regulations will continue to drive market growth, with Rubrik differentiation around go-to-market & security; 3) SaaS transition is largely complete and should help total sustainable growth; 4) Between sales, cost-to-book, and other leverage, we see higher FCF potential than current estimates project; 5) Mix shift towards cloud should help stabilize overall NRR.

CVLT – N, $154 PT

Covering Analyst: James Fish

Investment Thesis

Commvault competes in an attractive space that has helped the stock significantly outperform in 2024. We expect Commvault to continue to benefit from cyber resilience, ransomware, and regulatory tailwinds, while acquisitions help hold up reported ARR growth. The product portfolio is now competitive in our view and should be considered “Gen-Z” (including with Metallic), but we worry about the transition’s completion, gaining net-new total customers, and organic growth realization vs. current valuation.

Source: Piper Sandler, Company Websites

Private Vendors – Chips / Silicon

Vendors in Focus

Cerebras Systems

Platform Overview & Differentiation: Cerebras Systems develops accelerated computing chips for Generative AI inference and training, and HPC applications. Their solutions include the WSE-3 (wafer-scale engine) which boasts 900K compute cores, 44G of on-chip memory, and 21 PB of fabric bandwidth for cluster-scale performance on a single chip. The WSE-3 powers the Cerebras CS-3 AI supercomputer, a system less than one rack in size capable of training LLMs with 24 trillion parameters in a single local memory space, without partitioning or refactoring. Cerebras Inference delivers high throughput at low cost, with the ability to support hundreds of concurrent users.

Groq

Platform Overview & Differentiation: Groq is a hardware and software platform for AI inference. The Groq LPU is a purpose-built compute solution for AI that leverages a compiler defined architecture (kernel-less) to deliver high-performance and scalability. Groq's on-prem system, the GroqRack, features eight compute and one redundant Groq Node server, each containing eight Groq Cards with compute and memory co-located on the chip to eliminate resource bottlenecks. Groq Cloud is a platform for AI applications, with tools for real-time streaming, transforming natural language into actionable API calls, and connecting to external sources. The platform has grown to 480K developers since launching in Feb 2024.

Source: Piper Sandler, Company Websites

Private Vendors – Primary Storage & Data Platforms

Source: Piper Sandler, Company Websites

Vendors in Focus

Vast | $9.1B Valuation

Platform Overview & Differentiation: The Vast Data Platform is designed for deep learning in a scale-up architecture to enable machine learning from datasets by leveraging the DataEngine at the core across DataStore, DataBase, and DataSpace. It is a software-defined and globally distributed platform that is leveraged by large brands like Pixar, Zoom, NASA, CoreWeave, Lambda, and ServiceNow. Differentiation comes via its enterprise grade features and functionality, particularly around unstructured data that eliminates data-tiering needs, alongside the higher performance (which leverages all-flash) and scalability relative to traditional storage players. The combination of a single, scalable offering across storage, database, and compute is unique, in our view.

Weka | $1.6B Valuation

Platform Overview & Differentiation: Weka offers a file-based software data platform built to deliver 10x performance and scale relative to traditional storage peers, and is tailored for cloud and AI based workloads. Weka runs in the cloud and on-premise, and is hardware agnostic, allowing end-customers like Samsung, Wharton, and Contextual.ai more flexibility than traditional storage deployments running proprietary hardware. Weka operates its own network and hypervisor that runs across cloud and on-premise environments. By doing this, Weka can optimize flash media both on device and on the object storage metadata and lower storage costs by 65%. Weka is bringing more enterprise-grade features and functionalities to market, and integrating with partners such as NVIDIA, Cisco, Super Micro, AWS, Google, & more.

Private Vendors – Backup & Recovery

Source: Piper Sandler, Company Websites

Vendors in Focus

Cohesity | $7B Valuation

Platform Overview & Differentiation: The platform consolidates secondary data and apps, including backup, files, objects, and analytics in a single, centrally managed API-driven platform. The software is based on a scale-out file system / hyperconverged architecture and offers on-prem, hybrid and SaaS workload protection from a single platform. Cohesity differentiates via its Automated platform leveraging AI, multi-cloud, API-first approach, enterprise functionality, neutrality between OEMs, Cohesity Marketplace, & strategic relationships. Cohesity recently completed its acquisition of the larger Veritas, giving Cohesity a strong "hunting ground" with a proven go-to-market organization together. The combined entity has +12K customers (including 85% of the F100), +$1.7B revenue, $1.5B+ ARR, and +100 EBs of data protected.

Veeam | $15B Valuation

Platform Overview & Differentiation: The Veeam Data Platform is a complete data management platform across backup, replication, monitoring and analytics, with an agentless approach that leverages multiple hypervisors. Veeam can back up and manage the data no matter where the customer's data resides, particularly with the new Veeam Data Cloud (VDC) offering. Veeam differentiates via its agentless approach with built-in advanced replication and instant + automated machine recovery, very intuitive GUI / ease of use, fast deployment and storage agnosticism, large installed base (+550K customers), ease of deployment (especially remote branches), VM-backup capabilities, availability across all 3 major IaaS, & large partner ecosystem that has helped Veeam generate $1.7B ARR.

IMPORTANT RESEARCH DISCLOSURES

Notes: The boxes on the Rating and Price Target History chart above indicate the date of the fundamental Equity Research Note, the rating and the price target. Each box represents a date on which an analyst made a change to a rating or price target, except for the first box, which may only represent the first Note written during the past three years.
Legend:
I: Initiating Coverage
R: Resuming Coverage
T: Transferring Coverage
D: Discontinuing Coverage
S: Suspending Coverage
OW: Overweight
N: Neutral
UW: Underweight
NA: Not Available
UR: Under Review

Distribution of Ratings/IB Services
Piper Sandler

			IB Serv./Past 12 Mos.	
Rating	**Count**	**Percent**	**Count**	**Percent**
BUY [OW]	**539**	**58.33**	**111**	**20.59**
HOLD [N]	**353**	**38.20**	**38**	**10.76**
SELL [UW]	**32**	**3.46**	**2**	**6.25**

Note: Distribution of Ratings/IB Services shows the number of companies currently covered by fundamental equity research in each rating category from which Piper Sandler and its affiliates received compensation for investment banking services within the past 12 months. FINRA rules require disclosure of which ratings most closely correspond with "buy," "hold," and "sell" recommendations. Piper Sandler ratings are not the equivalent of buy, hold or sell, but instead represent recommended relative weightings. Nevertheless, Overweight corresponds most closely with buy, Neutral with hold and Underweight with sell. See Stock Rating definitions below.

Analyst Certification

The analysts James E. Fish and Harsh V. Kumar, primarily responsible for the preparation of this research report, attest to the following:

The views expressed in this report accurately reflect our personal views about the subject company and the subject security. In addition, no part of our compensation was, is, or will be directly or indirectly related to the specific recommendations or views contained in this report.

Piper Sandler research analysts receive compensation that is based, in part, on overall firm revenues, which include investment banking revenues.

Time of dissemination: 18 December 2024 03:34EST.

Research Disclosures

All statements in this report attributable to Gartner represent Piper Sandler's interpretation of data, research opinion or viewpoints published as part of a syndicated subscription service by Gartner, Inc., and have not been reviewed by Gartner. Each Gartner publication speaks as of its original publication date (and not as of the date of this report). The opinions expressed in Gartner publications are not representations of fact, and are subject to change without notice.

Affiliate disclosures: Piper Sandler is the trade name and registered trademark under which the corporate and investment banking products and services of Piper Sandler Companies and its subsidiaries Piper Sandler & Co. and Piper Sandler Ltd. are marketed. This report has been prepared by Piper Sandler & Co. and/or its affiliate Piper Sandler Ltd. Piper Sandler & Co. is regulated by FINRA, NYSE and the United States Securities and Exchange Commission, and its headquarters are located at 800 Nicollet Mall, Minneapolis, MN 55402. Piper Sandler Ltd. is authorised and regulated by the Financial Conduct Authority, and is located at 6th Floor, 2 Gresham Street, London EC2V 7AD. This report is distributed in the European Union by Piper Sandler Europe SAS, a French *"société par actions simplifiée"* incorporated in France under the number RCS Paris 932 083 769 whose registered office is 10, place Vendôme, 75001, Paris, France. Disclosures in this section and in the Other Important Information section referencing Piper Sandler include all affiliated entities unless otherwise specified.

Rating Definitions

Stock Ratings: Piper Sandler fundamental research ratings are indicators of expected total return (price appreciation plus dividend) within the next 12 months. At times analysts may specify a different investment horizon or may include additional investment time horizons for specific stocks. Stock performance is measured relative to the group of stocks covered by each analyst. Lists of the stocks covered by each are available at www.pipersandler.com/researchdisclosures. Stock ratings and/or stock coverage may be suspended from time to time in the event that there is no active analyst opinion or analyst coverage, but the opinion or coverage is expected to resume. Research reports and ratings should not be relied upon as individual investment advice. As always, an investor's decision to buy or sell a security must depend on individual circumstances, including existing holdings, time horizons and risk tolerance. Piper Sandler sales and trading personnel may provide written or oral commentary, trade ideas, or other information about a particular stock to clients or internal trading desks reflecting different opinions than those expressed by the research analyst. Piper Sandler will buy and sell securities of the subject company on a principal basis, regardless of whether it is a registered market maker in those securities. In addition, Piper Sandler offers technical research products that are based on different methodologies, may contradict the opinions contained in fundamental research reports, and could impact the price of the subject security. Recommendations based on technical analysis are intended for the professional trader, while fundamental opinions are typically suited for the longer-term institutional investor.

Overweight (OW): Anticipated to outperform relative to the median of the group of stocks covered by the analyst.

Neutral (N): Anticipated to perform in line relative to the median of the group of stocks covered by the analyst.

Underweight (UW): Anticipated to underperform relative to the median of the group of stocks covered by the analyst.

Other Important Information

The material regarding the subject company is based on data obtained from sources we deem to be reliable; it is not guaranteed as to accuracy and does not purport to be complete. This report is solely for informational purposes and is not intended to be used as the primary basis of investment decisions. Piper Sandler has not assessed the suitability of the subject company for any person. Because of individual client requirements, it is not, and it should not be construed as, advice designed to meet the particular investment needs of any investor. This report is not an offer or the solicitation of an offer to sell or buy any security. Unless otherwise noted, the price of a security mentioned in this report is the most recently available closing market price. Piper Sandler does not maintain a predetermined schedule for publication of research and will not necessarily update this report. Piper Sandler policy generally prohibits research analysts from sending draft research reports to subject companies; however, it should be presumed that the fundamental equity analyst(s) who authored this report has had discussions with the subject company to ensure factual accuracy prior to publication, and has had assistance from the company in conducting diligence, including visits to company sites and meetings with company management and other representatives.

Notice to customers: This material is not directed to, or intended for distribution to or use by, any person or entity if Piper Sandler is prohibited or restricted by any legislation or regulation in any jurisdiction from making it available to such person or entity. Customers in any of the jurisdictions where Piper Sandler and its affiliates do business who wish to effect a transaction in the securities discussed in this report should contact their local Piper Sandler representative, or as otherwise noted below. **Canada:** Piper Sandler & Co. is not registered as a dealer in Canada and relies on the "international dealer exemption" set forth in National Instrument 31-103 – Registration Requirements, Exemptions and Ongoing Registrant Obligations of the Canadian Securities Administrators. This research report has not been prepared in accordance with the disclosure requirements of Dealer Member Rule 3400 – Research Restrictions and Disclosure Requirements of the Investment Industry Regulatory Organization of Canada. **United Kingdom:** This material is for the use of intended recipients only and only for distribution to professional and institutional investors, i.e., persons who are authorised persons or exempted persons within the meaning of the Financial Services and Markets Act 2000 of the United Kingdom, or persons who have been categorised by Piper Sandler Ltd. as professional clients under the rules of the Financial Conduct Authority. **European Union:** This material is intended only for institutional investors (professional clients and eligible counterparties within the meaning of the Markets in Financial Instruments Directive (Directive 2014/65/EU of the European Parliament and of the Council of 15 May 2014)). Piper Sandler Europe SAS is not licensed as an investment services provider. It carries out the activity of a financial analyst not associated with an investment services provider as described in Articles 327-1 et seq. of the General Regulation of the *Autorité des marchés financiers* (“RG AMF”). Piper Sandler Europe SAS complies, in particular, with Articles 327-2 to 327-6 of the RG AMF as well as the relevant provisions of the Regulation (EU) No 596/2014 (the Market Abuse Regulation or “MAR”) and the Commission Delegated Regulation n°2016/958 applicable to the production or dissemination of “investment recommendations” (as defined under MAR). Piper Sandler Europe SAS does not provide investment services. **United States:** This report is distributed in the United States by Piper Sandler & Co., member SIPC, FINRA and NYSE, Inc., which accepts responsibility for its contents. The securities described in this report may not have been registered under the U.S. Securities Act of 1933 and, in such case, may not be offered or sold in the United States or to U.S. persons unless they have been so registered, or an exemption from the registration requirements is available.