



Flash Memory Summit

Hybrid Cloud and AI through the lens of 3D-NAND technology

Ranjana Godse, Jung Yoon & Andy Walls
IBM Corporation

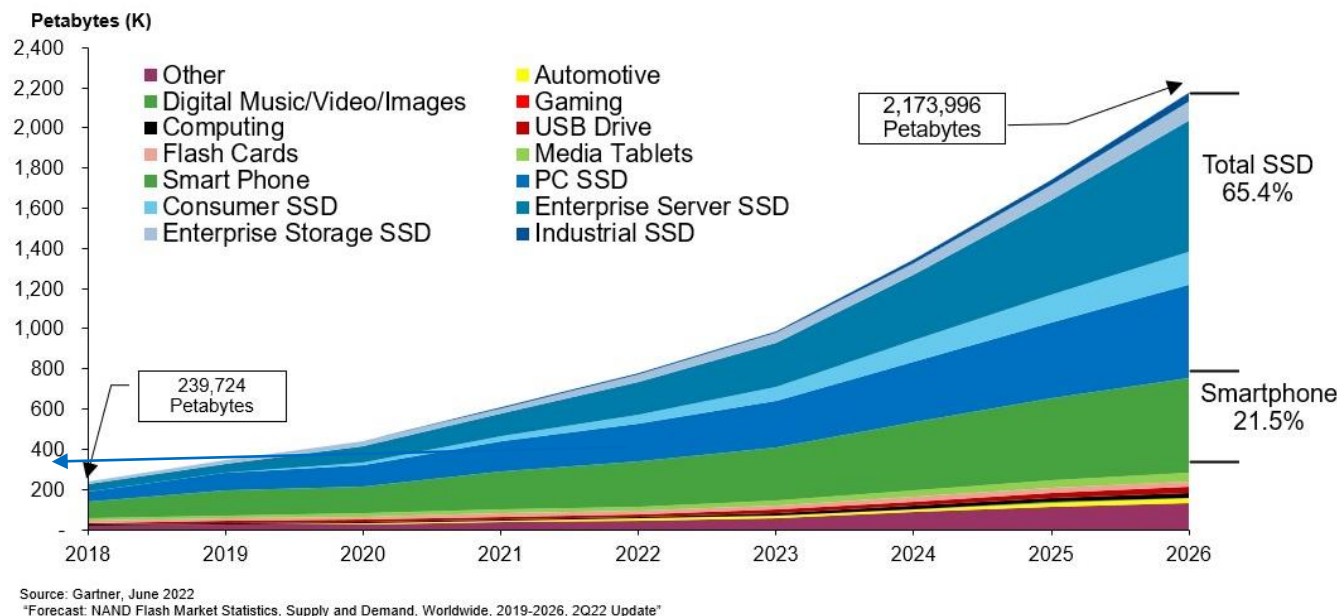
Agenda



Flash Memory Summit

- 3D NAND flash demand drivers
- Flash enabling IT infrastructure
- Flash Scaling and Reliability Outlook
- Block, File and Object Storage
- Flash for Hybrid Cloud & AI

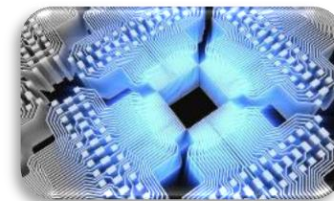
Data explosion and NAND Demand drivers:



- COVID-19, IT infrastructure, 5G, EV & autonomous driving, Gaming etc. driving semiconductor technology innovations at an unprecedented rate
- Flash and SSD demand driven by ever increasing nature of structured and unstructured data
- SSD remains a vital demand driver for Flash
- Enterprise Server SSD segment expected to grow 4x by 2026
- Complex AI algos are driving needs for more Memory and Storage to process, analyze and store ever-growing data

Flash driving IT infrastructure transformation

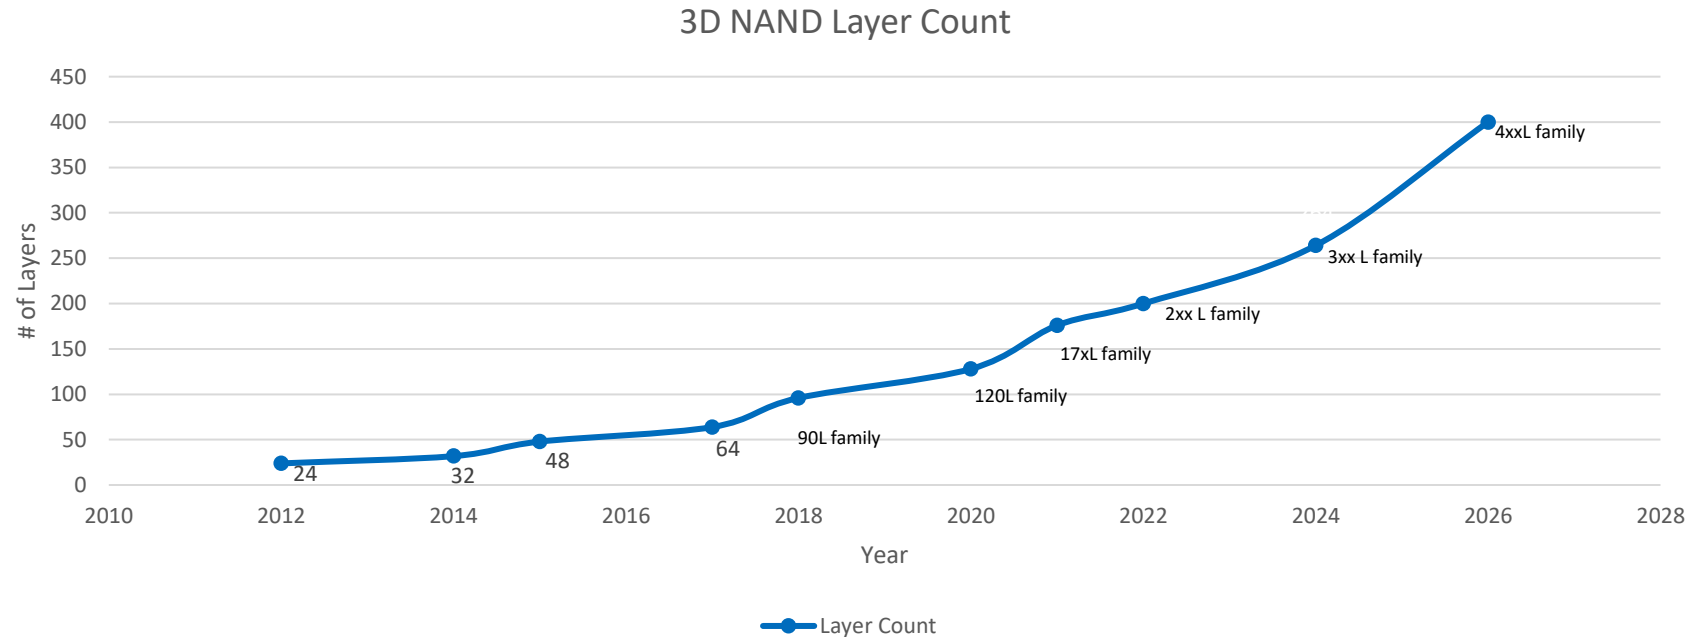
- Use Cases (AI/ML, Big Data analytics, Database applications) => enabled by Flash
- As Public Cloud IaaS adoption grows, flash-SSD driven performance storage services essential
- AI is the fastest growing application in Data Centers – “More Data”, “FASTER”
- AI/ML applications require vast amounts of Low Latency, High IOPS, High Thruput Flash Storage at Low TCO
- 3D-NAND scaling, TLC>QLC and Computational Storage are key enabling technologies
- Cloud vs Edge processing, a critical business decision => both supported by Flash Technology development



No Flash

All Flash

3D NAND Layer Count Growth



- Growing #of layers enabling 3D NAND manufacturers to drive lower \$/GB
- Channel etch AR challenges driving need for Multi Stacking – all suppliers transitioning to 2 high Multi-stacking process at >120 Layer 3D-NAND generation,
- Currently suppliers are transitioning to ~ 17x Layer in 2H22
- Manufacturing and Performance challenges comes with growing #of layers

Flash Scaling outlook

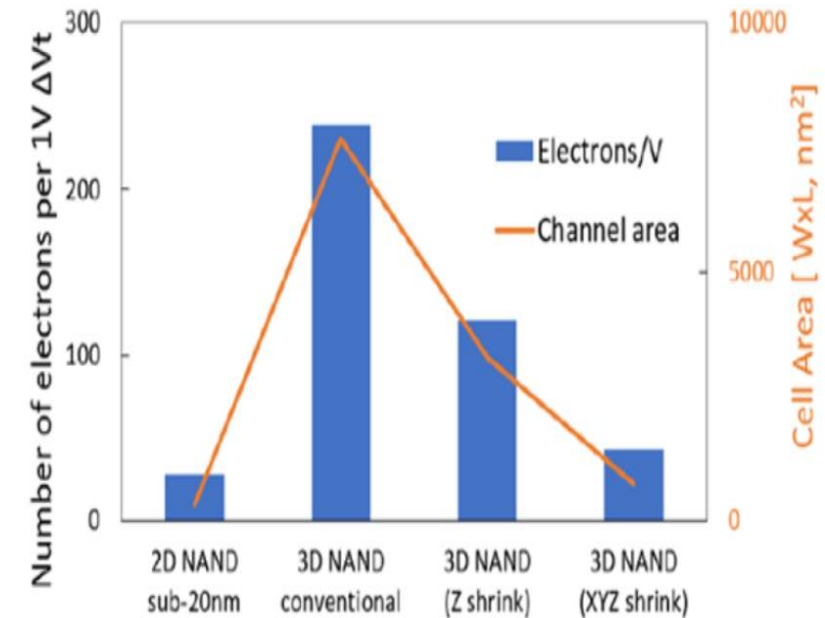


Flash Memory Summit

- **Higher #of layers challenges:**

- High Aspect Ratio (HAR) etch is associated with insufficient thin film step coverage causing poor retention on lower WL
- P/E speed of lower WL is faster due to higher coupling ratio because of smaller channel hole CD
- Higher Z height can create issues like wafer warpage
- With layer count growing 100+L, the stack height reaching > 10 μm . 16DP pkg with 300+L expected to take half of the wafer thickness (~350 μm) posing

- **Vertical (Z directional) scaling can pose challenges of storing fewer electrons/cell** - small number of electrons (charge loss/gain) can result in dramatic effects on V_t



3D NAND \$/GB Cost modelling



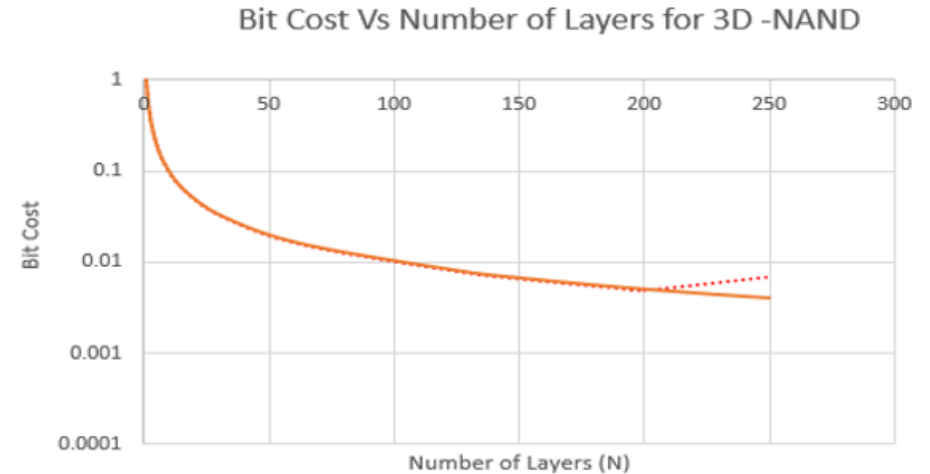
Flash Memory Summit

$$\text{3D NAND die cost} = (1/N * \alpha)$$

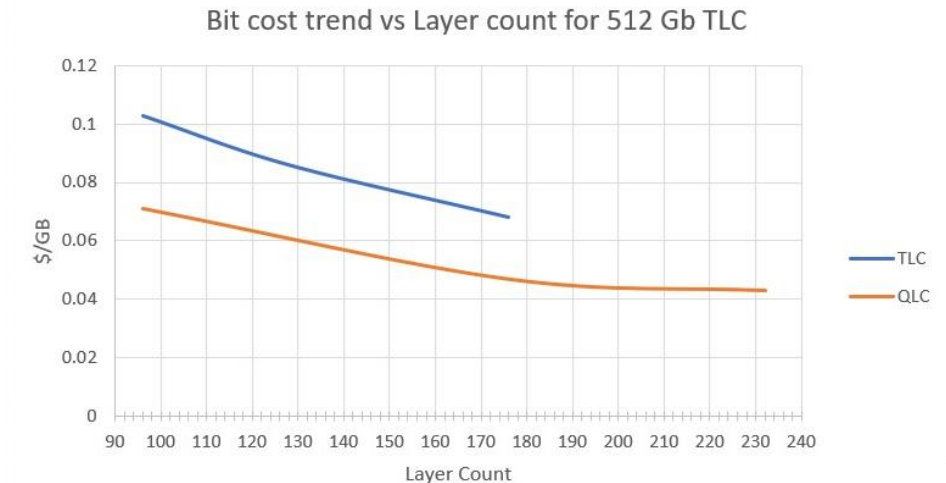
Where, N = #of layers for 3D NAND generation
 α = Process complexity factor

$$\frac{\$}{GB} = \frac{\text{Wafer cost}}{DPW * \text{Yield} * GB/Die} + \text{Package cost/GB}$$

- Flash bit cost is majorly driven by layer count growth. In the cost model, we try to correlate flash cost vs growing #of layers for 3D NAND.
- We estimate that the flash cost is inversely proportional to #of layers (N) till ~400 layers.
- For 400+L, flash manufacturer will need to invest in more sophisticated etch equipment as the AR increases, posing challenges for vertical channel etch.
- Process complexity factor α will be of significance for >400+L gen



Ref : 3D-NAND scaling & 3D-SCM – Implications to Enterprise Storage. FMS 2017



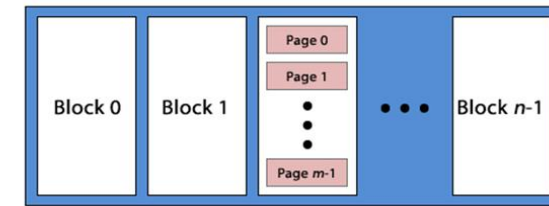
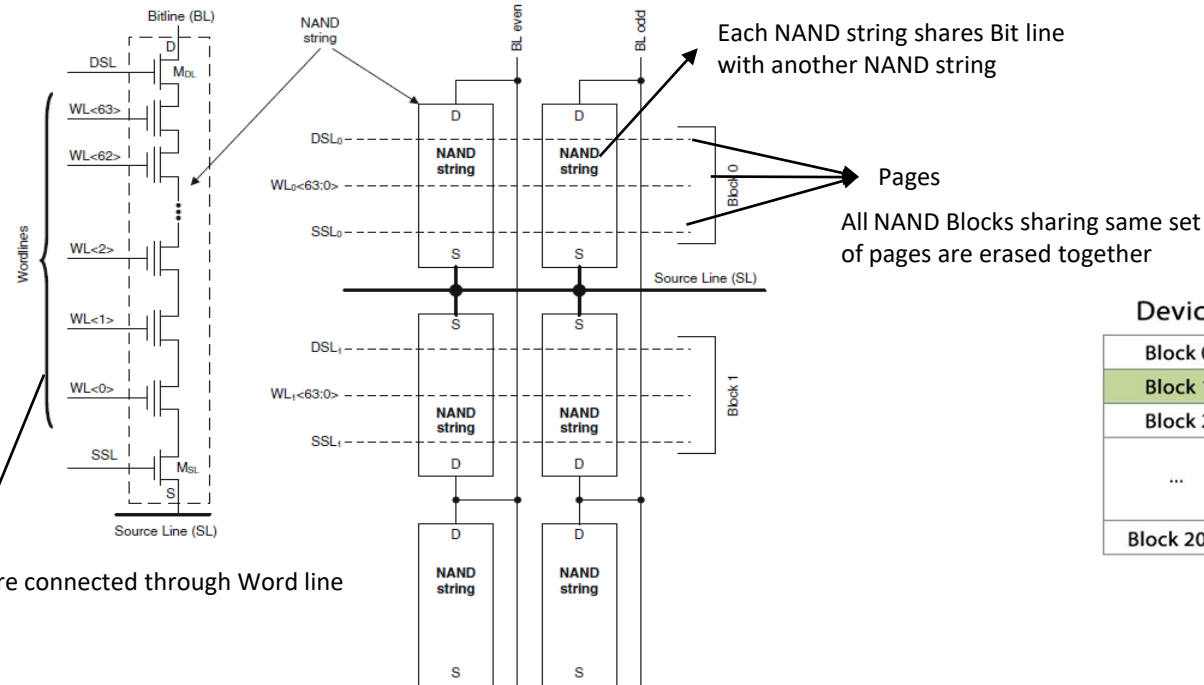
3D NAND Architecture



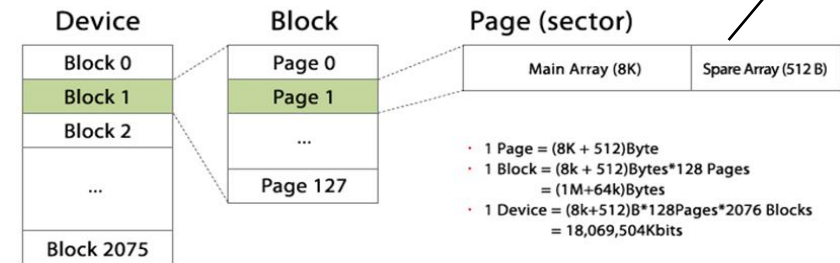
Flash Memory Summit



Control gates are connected through Word line



ECC and Firmware Meta data



- Flash is block based device where each block is made up of several pages (WL)
- Each NAND string contains several pages and shares Bit line with another NAND string
- NAND is erased on a block level

3D NAND Scaling - Key Factors for Hybrid Cloud & AI

1. Density

- Low latency access to vast amount of data with TLC, QLC multi-Terabit density
- 3D-Flash is ideal for Cloud – data center power, cooling, footprint advantage
- Stacked package >16DP, TSV driving industry's highest Density/mm³ package, but may face technology limits in Si wafer thinning with 3D-NAND scaling, and Bus bottlenecks limiting throughput

2. Flash Architecture

- Increasing number of Pages/Block, increasing Block size, driving large amount of data in FTL with Block retirement scheme
- Cell current for sense margin decrease with 3D-NAND layer count driving increased Block size
- Page retirement data structure is complex. Driving complexity in garbage collection and increase in Over Provisioning

3. Cost

- 3D-NAND scaling via layer count increase
- Memory Array vs peripheral circuit efficiency – Die size vs added process steps
- TLC vs QLC cost factors

3D NAND Scaling - Key Factors for Hybrid Cloud & AI



Flash Memory Summit

4. Performance

- Real time AI & workloads require ingestion of large amount of data at high rate and very high throughput
- As datasets used for training ML/DL are growing, Flash with its low latency and high throughput is optimal for AI storage
- Flash enables high IOPS with low read latency, lower power and smaller footprint compared to HDD

5. Reliability

- PE Cycling endurance, Vth window, 3D-NAND process controls, consistency & uniformity
- Data Retention, Complex 3D-NAND charge loss mechanisms
- AI driving read intensive workloads – understanding of read disturb mechanisms & process improvements

6. Power

- Program & Read energy/Byte scaling vs 3D-NAND generation critical
- AI workload & utilization, Cloud system power linkage with 3D-NAND Vcc, Icc, tProg/tRead key parameters

7. Resiliency

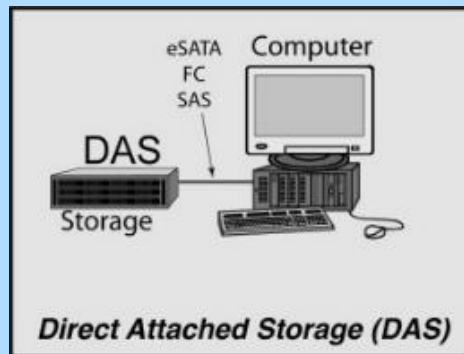
- Snapshot & Restore is an important feature for data back and recovery



DAS vs SAN vs NAS

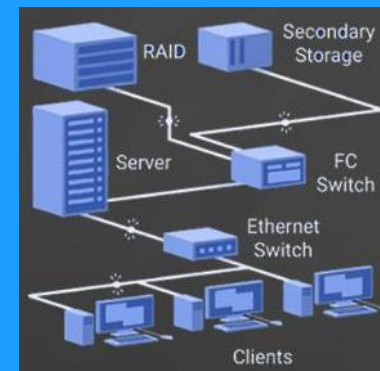
Direct Attached Storage (DAS)

- Anything that is attached to your computer/server
- Cheaper
- Not scalable
- Not sharable



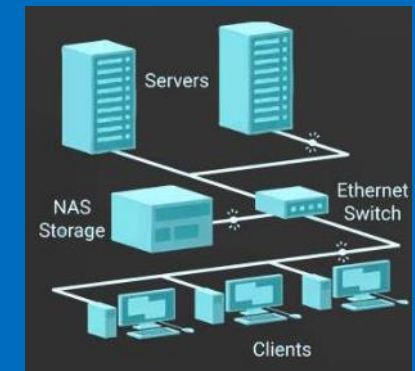
Storage Area Network (SAN)

- Bunch of hard drives (Block based storage) that are connected via the Fiber Network
- High Performance
- Scalable & Shareable
- Resilient and highly available
- Great for many read/write operation



Network Attached Storage (NAS)

- Files are transferred over network (Wi-Fi)
- Scalable
- Data security
- RAID feature for back up

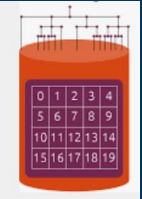


Types of Cloud Storage



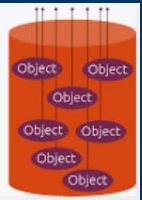
Block

- Fixed sized blocks of data shared over Fiber Channel
- msec - μ sec latency (Flash)



File

- Similar to traditional file storage system
- Latency vs Storage trade off



Object

- Horizontal hierarchy
- Persistent key/value store
- Latency tolerant

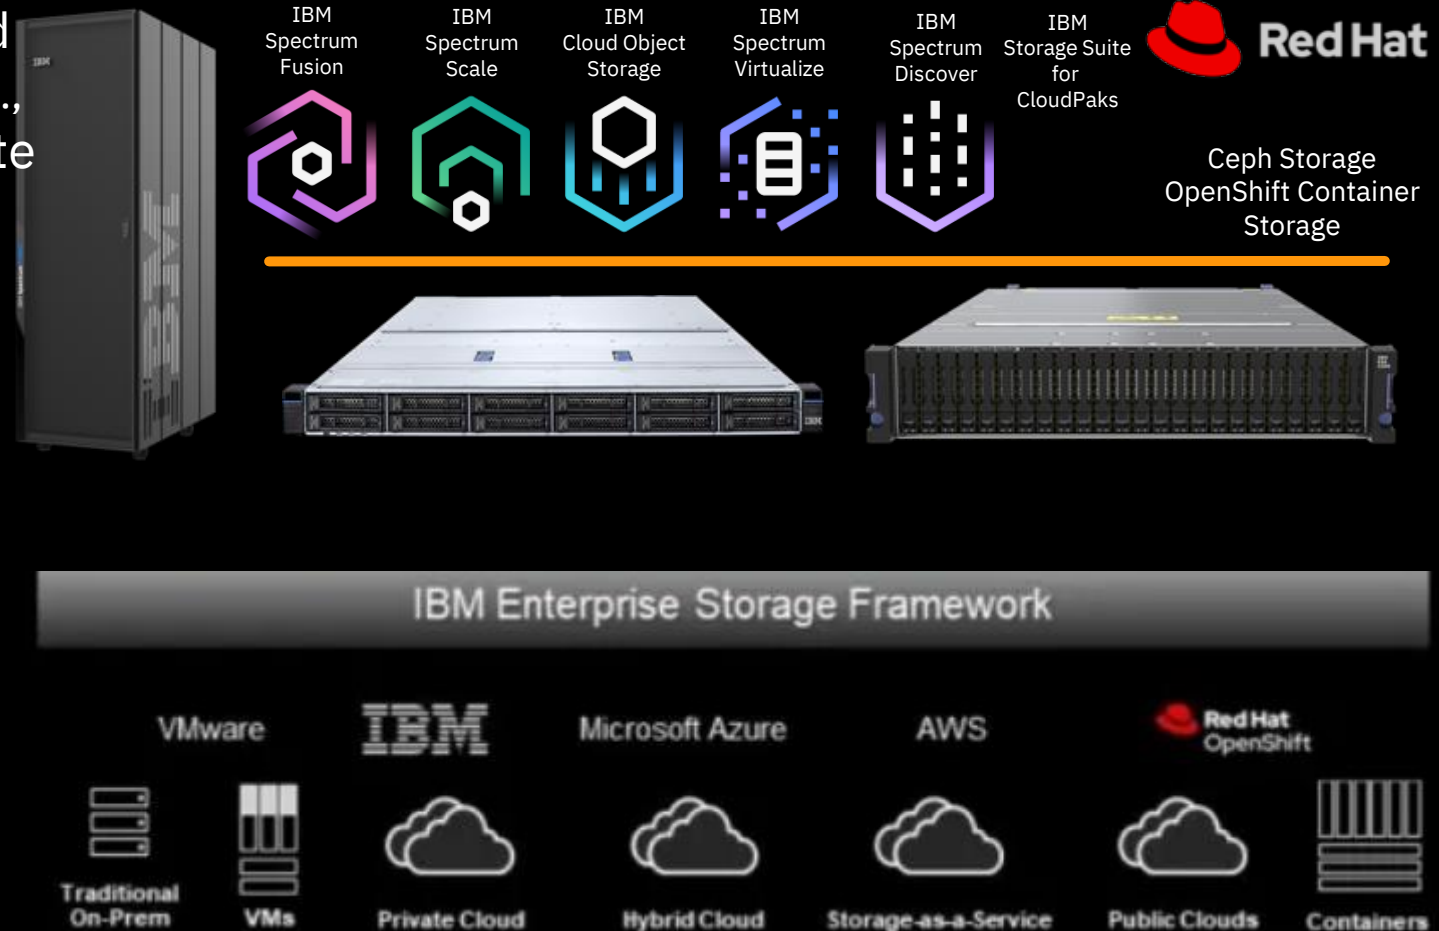
	Block	File	Object
Interface	Direct attached or SAN	User	Program (API)
Performance	Highest	Medium	Lowest
Cost	\$\$-\$\$\$	\$\$-\$\$\$\$	\$
Scalability	Low	Medium	High
Use case	Performance intensive storage - Structured Database	Document sharing, Web content, Media and Entertainment	Unstructured data - Image, PDF, Video

SSD Storage for Hybrid Cloud

- Flash being optimized for workload requiring frequent read/write operations, high IOPS makes a perfect choice for Block Storage
- Containerization, where multiple containers are orchestrated in an enterprise environment, benefits from the speed of block storage and the native ability for a single host to mount multiple blocks
- A rapidly growing portion of SAN deployments leverages all-flash storage to gain its high performance, consistent low latency, and lower total cost when compared to spinning disk
- Developments in NVMe, NVMe-oF and PCIe can accelerate adoption of Flash based storage for Hybrid Cloud

IBM Storage for Hybrid Cloud

IBM Storage provides seamless hybrid cloud and container integration for multi-modal storage (i.e., block, file, or object) with full support for all compute architectures (x86, IBM POWER 10, or IBM Z)



Flash, NVMe and FC advancements are enabling Hybrid cloud storage

- 3D-NAND scaling will continue to drive exponential flash market growth focused on AI and Cloud workloads
- Layer scaling and architectural improvements will continue to drive \$/GB reduction
- 3D-NAND Flash can provide a perfect trade of between Scale, Performance and Price for Cloud and AI application over HDD

When you interact with IBM, this serves as your authorization to Flash Memory Summit or its vendor to provide your contact information to IBM in order for IBM to follow up on your interaction.

IBM's use of your contact information is governed by the IBM Privacy Policy.

