



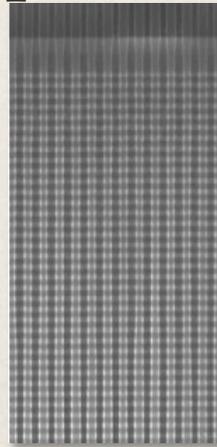
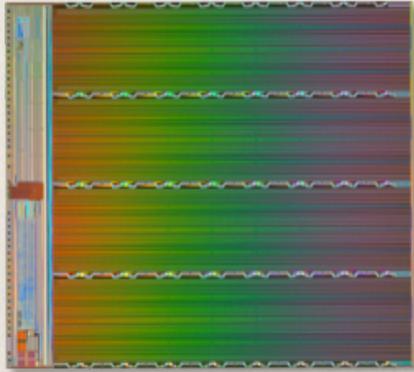
# Developments in Digital Data Storage Persistent Memory and Binary Encoding Implications on Archives

September 2016

Rob Peglar, SVP & CTO, Symbolic IO

@peglarr

# 3D NAND Comparison with Planar NAND Scaling



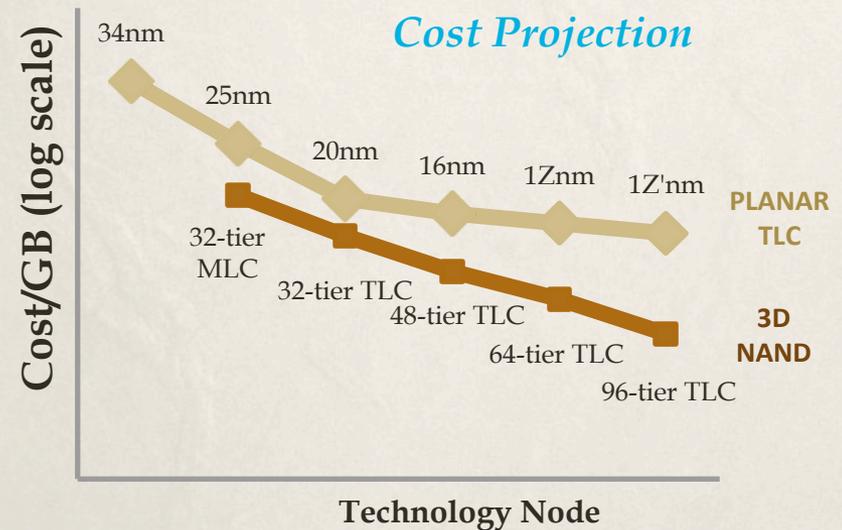
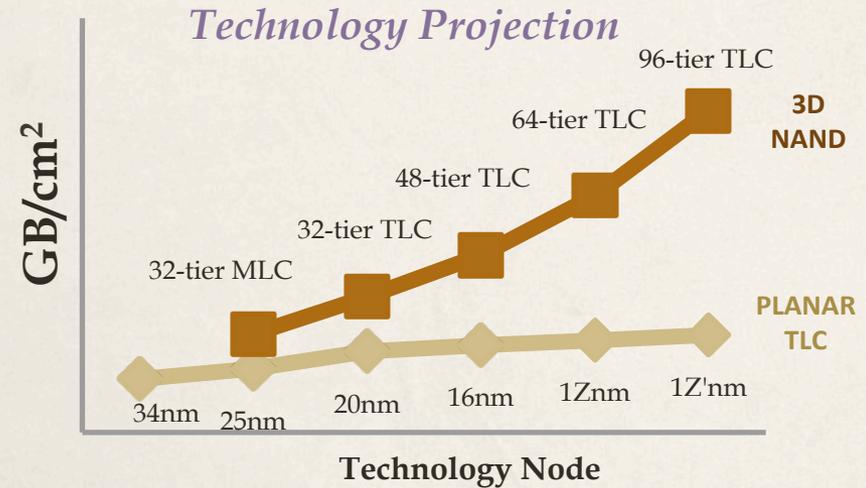
## \* Planar NAND scaling

- Planar can be scaled below 16nm, but performance and cost are not competitive with 3D NAND
- Industry emphasizing 3D NAND advantages

## \* 3D NAND scaling

- 3D NAND cost improvement over planar expands with subsequent nodes
- 3D NAND cell architecture enables significant performance improvement relative to planar technology
- Today's single-die (~1 cm<sup>2</sup>) density is 384Gb/cm<sup>2</sup> (soon to be 768)
- Near-future single-die density planned for 1Tb/cm<sup>2</sup> TLC
- All four NAND manufacturers working on QLC designs > 1Tb/cm<sup>2</sup>

September 7, 2016



# Driving Persistent Memory Adoption



Source: TCO Analysis, Wikibon 2015

	2015	2016	2017	2018
Highest Available Capacity HDD	10TB	11TB	12TB	12.5TB
Highest Available Capacity SSD	8TB	25.6TB	33TB	43TB

SSD Capacity / HDD Capacity 0.8x 2.3x 2.75x **3.4x**

Source: Gartner 2015

## Compelling 4-Yr TCO SSD vs. HDD

- \$/GB tipping point in 2016 with SSD capacity 2.3x HDD capacity, growing to 3.4x by 2018
- TCO benefit of SSDs will drive rapid adoption of the solid state-dominant data center over next 4-5 years
  - Lower power/cooling footprint
    - e.g. idle power 20mW
  - Higher reliability
    - AFR < 0.1%
  - Greater density
  - Reduce surrounding infra

# What is Possible - Near-Future SSD Target Platform using TLC/QLC Media - Very High Capacity Archive

- \* PCIe Gen 3/4, 32TB+, 2.5"-15mm form factor, U.2 interface (NVMe)
  - \* 1.536 PB raw in 2U
  - \* Other media form factors possible - modules, 3.5" (60 TB+), etc.
- \* Usable with intelligent encoding → 16PB in 2U
- \* See US Patent 9,304,703 (Ignomirello)
  - \* Computational Defined Storage, adaptive bit markers, secure encoding
  - \* Approaches Shannon limit for channel capacity
- \* Endurance: 500 sequential fills over life (16 PB TBW)
- \* Shelf Retention: 3 months @ 40C, 12 months @ 25C
- \* 200 Gb/s (20 GB/s) throughput currently (8x25GbE)
- \* Enterprise Data Integrity
  - \* Drive-level RAIN, host-level EC, cluster-wide EC, e.g. Ceph
- \* Performance: x4 saturation on reads (3.2 GB/s); 1-3 GB/s writes based on power
- \* Workload: 90%/10% Read/Write +/- 5%
  - \* Writes are 32K-4MB IO sizes, smaller host access routed to NVM buffer/namespace
  - \* Reads are 4K-4MB IO sizes



# Tape Analysis

## Tape Worksheet for <redacted>

### Tape Archive Parameters

Typical Working MB/sec/lib	210.00
GB/hr/lib	756.00
TB/day/lib	18.14
PB/month/lib	0.55
Total data in library (TB)	10000
Total data in library (PB)	10.00
TB/day ingest max per drive	3.63
Carts/day written (Teject=0)	0.73
Time to fill one cart (hr)	33.07
Total library cost	\$1,800,000
Total drive cost	\$360,000
Total cart cost	\$932,000
Add-on h/w annual cost	\$372,800
Annual data added (TB)	2000
Annual data added (PB)	2.00
Daily data added (TB)	5.48

### Variables

Price of library	\$450,000
Price of drive	\$20,000
Price of lib maint (annual)	\$0
Price of drive maint (annual)	\$2,400
Drive lifetime (years)	7
Cart fill level (% of raw)	100%
Library lifetime (years)	25
Price of lib h/w maint (yr)	\$33,750
Price of lib s/w maint (yr)	\$30,000
Price of FC switch(es)	\$50,000
Price of FC HBA	\$1,000
Price of FC switch maint (yr)	\$5,000
Price of FC HBA maint (yr)	\$0
First copy drives added (yr)	0
First copy carts added (yr)	800
Second copy carts added (yr)	800
Price of HSM server	\$50,000
Price of cartridge	\$233
Price of HSM s/w	\$100,000
Annual Maint of HSM s/w	\$10,000
Instances of HSM (# servers)/site	2
First copy weekly ingest (TB)	38.46
Second copy weekly ingest (TB)	38.46
Second copy drives added (yr)	0
Price of SAN storage for HSM	\$500,000
SAN capacity for HSM (TB)	100
Maint for SAN storage (\$/yr)	\$70,000
Maint of HSM server h/w (yr)	\$5,000
Cost of power (cents/KW/hr)	8

### Constants

Library make/model	<redacted>
First copy library count	2
Drive make/model	<redacted>
Second copy library count	2
Library max slots	38115
Library software	<redacted>
Drive-to-drive copy	yes - unused
RAIT	not used
Current drive utilization (%)	30%
Cartridge make/model	<redacted>
Cartridge capacity (raw TB)	5
Drive throughput (MB/sec)	120
Current #of FC switches	2
Current # of FC HBAs	16
First copy cartridge count	1000
Second copy cartridge count	1000
First copy drive count/lib	5
Second copy drive count/lib	4
Drive speed (MB/sec)	140
Watts/library	750
Cooling/library (BTU)	2448
Watts/FC switch	800
Cooling/FC switch (BTU)	2700
Watts/SAN array	4500
Cooling/SAN array (BTU)	14000
Watts/HSM server	1500
Cooling/HSM server (BTU)	5000
Watts/tape drive (active)	27
Cooling/tape drive (active)	80
BTUs/watt (constant)	3.412
Number of sites	2

# Tape TCO

Tape Archive Costs		Notes			
Day 0 Cost	<b>\$2,996,000</b>	lib, drives, carts, SAN, HSM	Day 0 Library Cost	\$1,800,000	
Delta annual h/w addons	\$372,800	drives & carts	Day 0 Drive Cost	\$360,000	
Delta annual maint	\$0	maint on drive adds	Day 0 Cart Cost	\$932,000	
Year 1 Maint	\$508,200	lib h/w, lib s/w, drives, SAN, HSM	Day 0 SAN Cost	\$1,104,000	array, switches & HBAs
Year 2 Maint	\$508,200		Day 0 HSM Cost	\$600,000	h/w & s/w
Year 3 Maint	\$508,200				
Year 4 Maint	\$508,200		Annual Library Maint	\$255,000	h/w & s/w
Year 5 Maint	\$508,200		Annual Drive Maint	\$43,200	
FTE annual cost (admins)	\$125,000	1 admin(s) @ 125K/yr	Annual FC Fabric Maint	\$10,000	
Watts consumed (kW-hr)	23.86	includes cooling	Annual HSM server Maint	\$60,000	includes cool
Watts consumed (annual)	209,004		Annual SAN Maint	\$140,000	
5 years of h/w addons	\$1,864,000	drives & carts			
5 years of maint	\$2,541,000				
5 years of admins	\$625,000				
5 years of power	\$83,602				
5 years total OpEx	<b>\$5,322,630</b>				
Total spend over 5 years	<b>\$8,318,630</b>	Day 0 cost plus 5 yrs opex	Watts consumed by libraries (kW-hr)	5.87	
Total data end of year 5 (PB)	<b>18.00</b>		Watts consumed by drives (kW-hr)	0.27	
Cost/PB/year	<b>\$462,146</b>	Two sites	Watts consumed by FC switches (kW-hr)	3.18	
			Watts consumed by SAN array (kW-hr)	8.60	
			Watts consumed by HSM servers (kW-hr)	5.93	

# Scale-out Magnetic Disk TCO

Type of Node	Scale-out NAS	Notes		Notes
# of drives/node	60	3.5" SATA	Year 1 total cost	\$1,904,548 includes nodes, power, cooling, FTE
Drive Technology (TB)	10		Year 2	\$470,034
Raw capacity/node	600		Year 3	\$478,377
Storage Efficiency (%)	84%	16+3 EC	Year 4	\$486,719
Usable capacity/node (TB)	504		Year 5	\$495,062
List price/TB (raw)	\$500.00			
Discount off list (%)	70%		5-year total	<b>\$3,834,740</b>
Buy price/node	<b>\$90,000</b>	Incl. maint + s/w lic	Cost/PB/year	<b>\$213,041</b> One site
Buy price/TB (usable)	<b>\$178.57</b>		Cost/PB/year	<b>\$426,082</b> Two sites
FTE annual cost	\$62,500	0.5 FTE @ \$125K		
Watts/node (nominal)	1500	not including cooling		
Nodes required year 0	20			
Node adds per year	3.968			
kW-hr consumed year 1	525600	Includes cooling	Year 1 cost of power	\$42,048
kW-hr consumed year 2	629886		Year 2 cost of power	\$50,391
kW-hr consumed year 3	734171		Year 3 cost of power	\$58,734
kW-hr consumed year 4	838457		Year 4 cost of power	\$67,077
kW-hr consumed year 5	942743		Year 5 cost of power	\$75,419
			5-year total cost of power	<b>\$293,669</b>

# 3D NAND TCO

Type of Node	SSD/SDS	Notes		Notes
# of drives/node	48	2.5" NVMe	Year 0 total cost	\$1,683,149 includes nodes, power, FTE
Drive Technology (TB)	32		Year 1	\$391,906
Raw capacity/node	1536		Year 2	\$393,464
Storage Efficiency (%)	93.75%	15+1 EC	Year 3	\$395,021
Usable capacity/node (TB)	1440		Year 4	\$396,578
List price/TB (raw)	\$500.00			
Discount off list (%)	70%		5-year total	<b>\$3,260,118</b>
Buy price/node	<b>\$230,400</b>	Incl. maint + s/w lic	Cost/PB/year	<b>\$181,118</b> One site
Buy price/TB (usable)	<b>\$160.00</b>		Cost/PB/year	<b>\$362,235</b> Two sites
FTE annual cost	\$62,500	0.5 FTE @ \$125K		
Watts/node (nominal)	800	not including cooling		
Nodes required year 0	7			
Node adds per year	1.389			
kW-hr consumed year 1	98112	Includes cooling	Year 0 cost of power	\$7,849
kW-hr consumed year 2	117579		Year 1 cost of power	\$9,406
kW-hr consumed year 3	137045		Year 2 cost of power	\$10,964
kW-hr consumed year 4	156512		Year 3 cost of power	\$12,521
kW-hr consumed year 5	175979		Year 4 cost of power	\$14,078
			5-year total cost of power	<b>\$54,818</b>



# Developments in Digital Data Storage Persistent Memory and Binary Encoding Implications on Archives

September 2016

Rob Peglar, SVP & CTO, Symbolic IO

@peglarr