NANTERO

Memory Class Storage

Bill Gervasi
Principal Systems Architect
bilge@Nantero.com

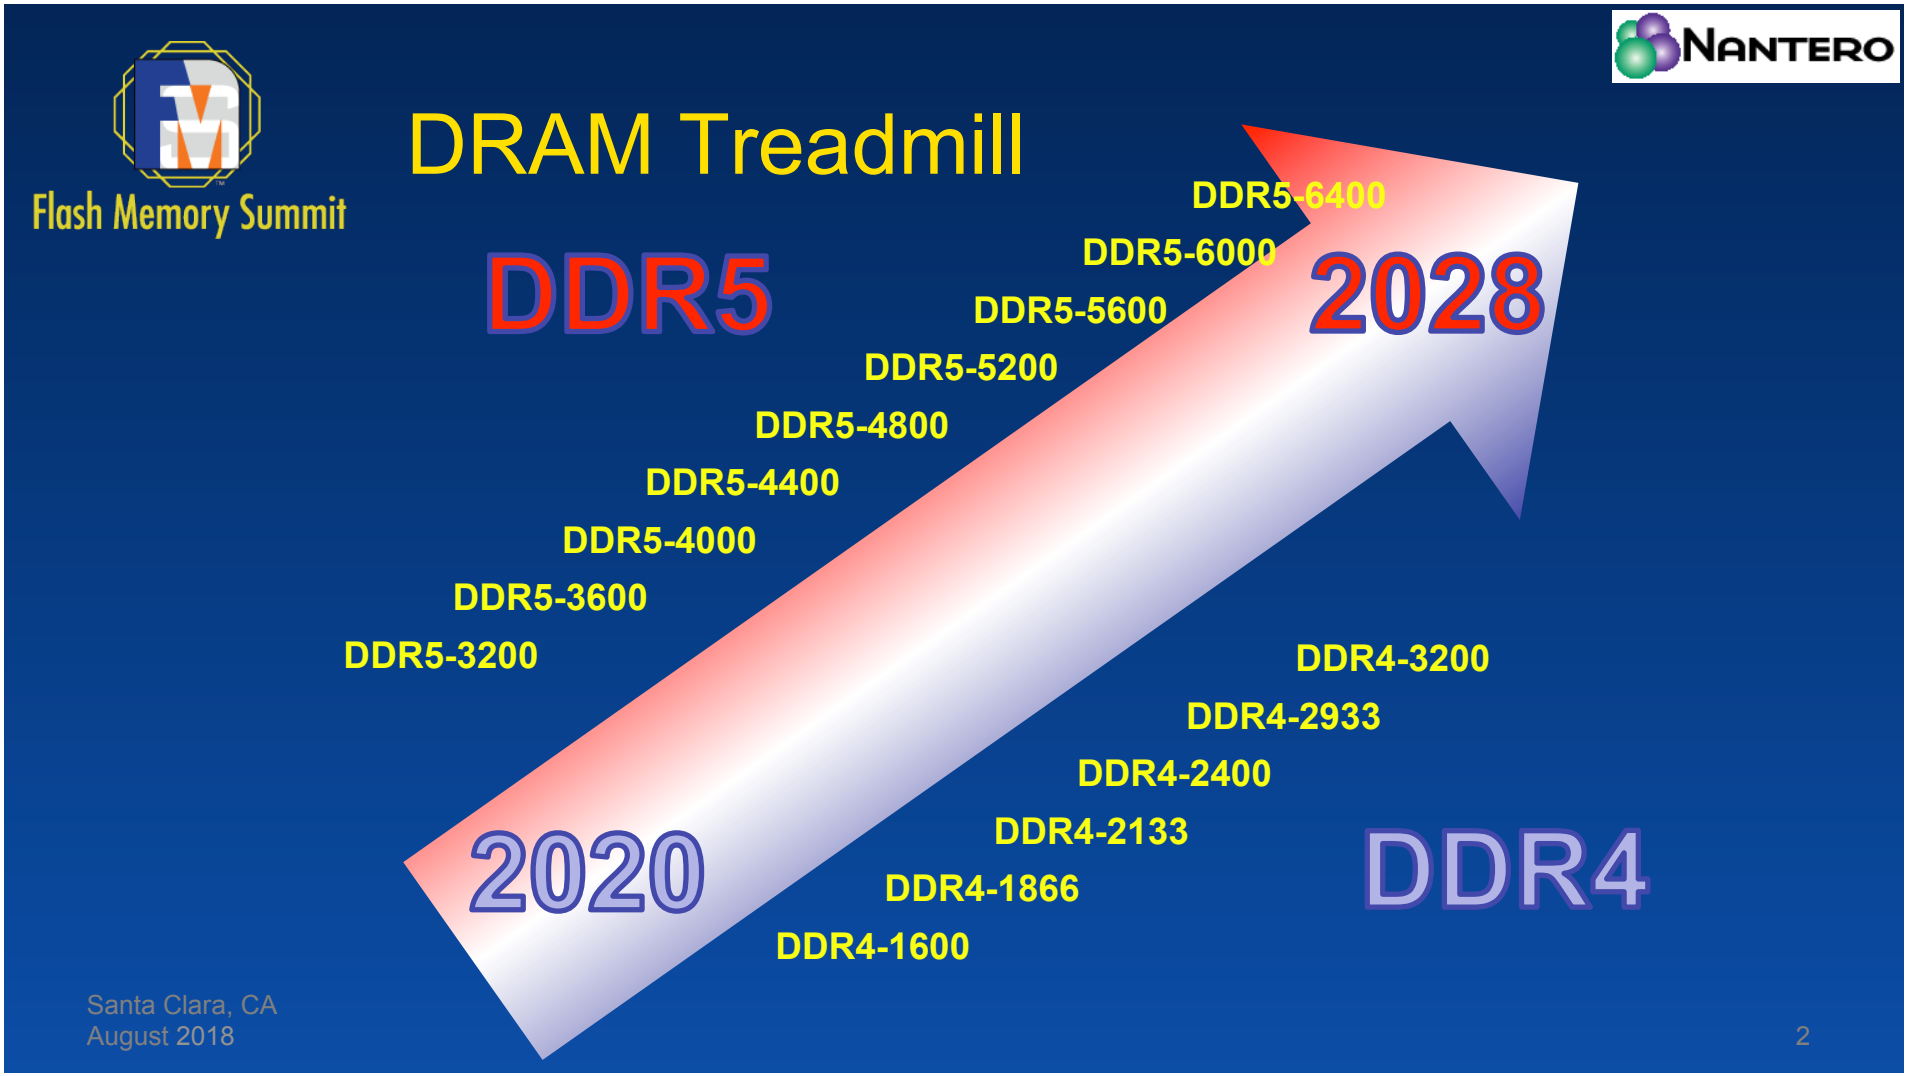
Flash Memory Summit
NANTERO
DRAM Treadmill
DDR5
2028
DDR5-6400
DDR5-6000
DDR5-5600
DDR5-5200
DDR5-4800
DDR5-4400
DDR5-4000
DDR5-3600
DDR5-3200
DDR4-3200
DDR4-2933
DDR4-2400
DDR4-2133
DDR4-1866
DDR4-1600
2020
DDR4
Santa Clara, CA
August 2018
2

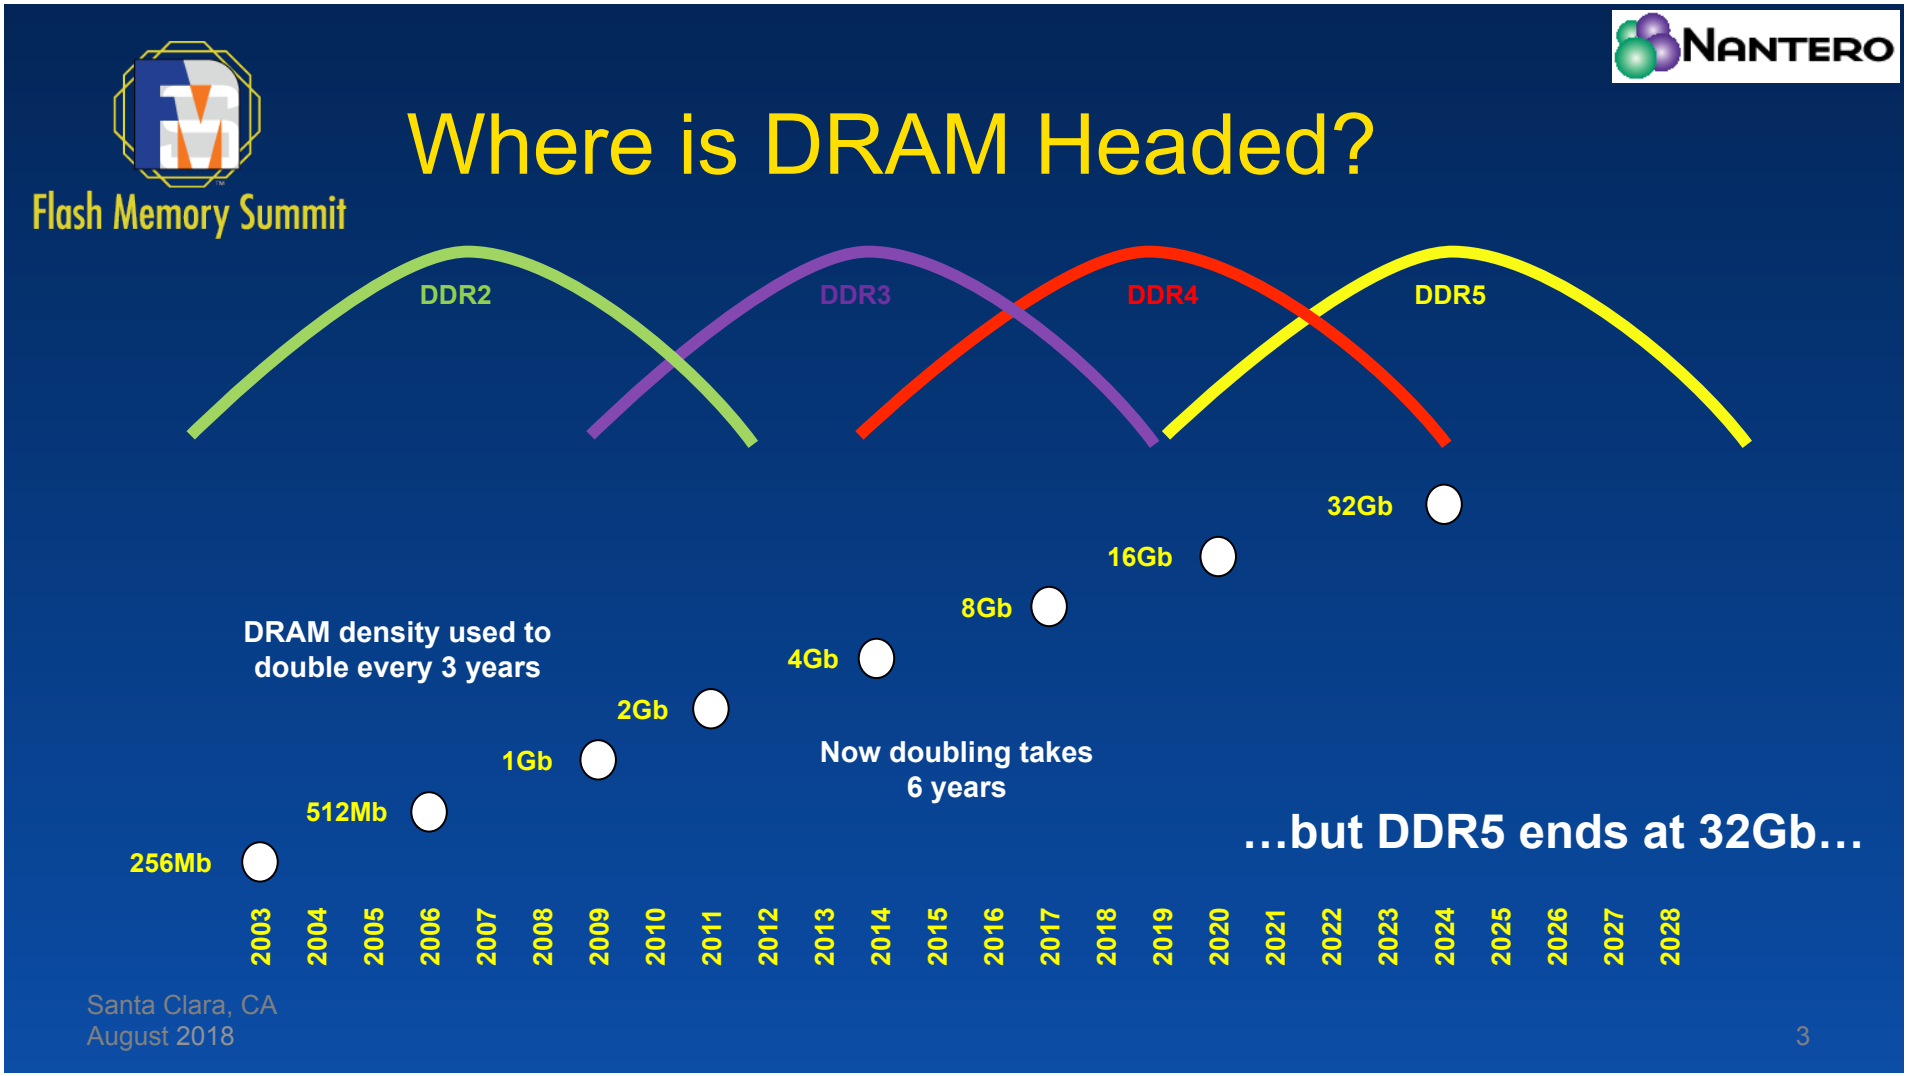
Flash Memory Summit
NANTERO
Where is DRAM Headed?
DDR2
DDR3
DDR4
DDR5
32Gb
16Gb
8Gb
4Gb
2Gb
1Gb
512Mb
256Mb
DRAM density used to double every 3 years
Now doubling takes 6 years
…but DDR5 ends at 32Gb…
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
Santa Clara, CA
August 2018
3

NANTERO

DDR5 Ends at 32Gb??? Really???

But the DDR5 data sheet shows a 64Gb chip!

However, to get 64Gb, you have to give up 3DS…

32 Gb Addressing Table

Configuration		8 Gb x4	4 Gb x8	2 Gb x16
Bank Address	# of Bank Groups	8	8	4
	BG Address	BG0~BG2	BG0~BG2	BG0~BG1
	Bank Address in a BG	BA0~BA1	BA0~BA1	BA0~BA1
Row Address		R0~R16	R0~R16	R0~R16
Column Address		C0~C10	C0~C9	C0~C9
Page size		1KB	1KB	2KB

64 Gb Addressing Table

Configuration		16 Gb x4	8 Gb x8	4 Gb x16
Bank Address	# of Bank Groups	8	8	4
	BG Address	BG0~BG2	BG0~BG2	BG0~BG1
	Bank Address in a BG			
Row Address		R0~R17	R0~R17	R0~R17
Page size		1KB	1KB	2KB

Function	Abbreviation	CS	CA Pins													
			CA0	CA1	CA2	CA3	CA4	CA5	CA6	CA7	CA8	CA9	CA10	CA11	CA12	CA13
Activate	ACT	L	L	L	R0	R1	R2	R3	BA0	BA1	BG0	BG1	BG2	CID0	CID1	DDPID
		H	R4	R5	R6	R7	R8	R9	R10	R11	R12	R13	R14	R15	R16	CID3/ R17
Write	WR	L	H	L	H	H	L	BL*=L	BA0	BA1	BG0	BG1	BG2	CID0	CID1	CID2/ DDPID
		H	V	C3	C4	C5	C6	C7	C8	C9	C10	V	AP=L	WR_ Partial=L	V	CID3
Read	RD	L	H	L	H	H	H	BL*=L	BA0	BA1	BG0	BG1	BG2	CID0	CID1	CID2/ DDPID
		H	C2	C3	C4	C5	C6	C7	C8	C9	C10	V	AP=L	V	V	CID3

Flash Memory Summit

NANTERO

Remind Me What 3DS Is…

Chip IDs (CIDs) select a DRAM within a stack

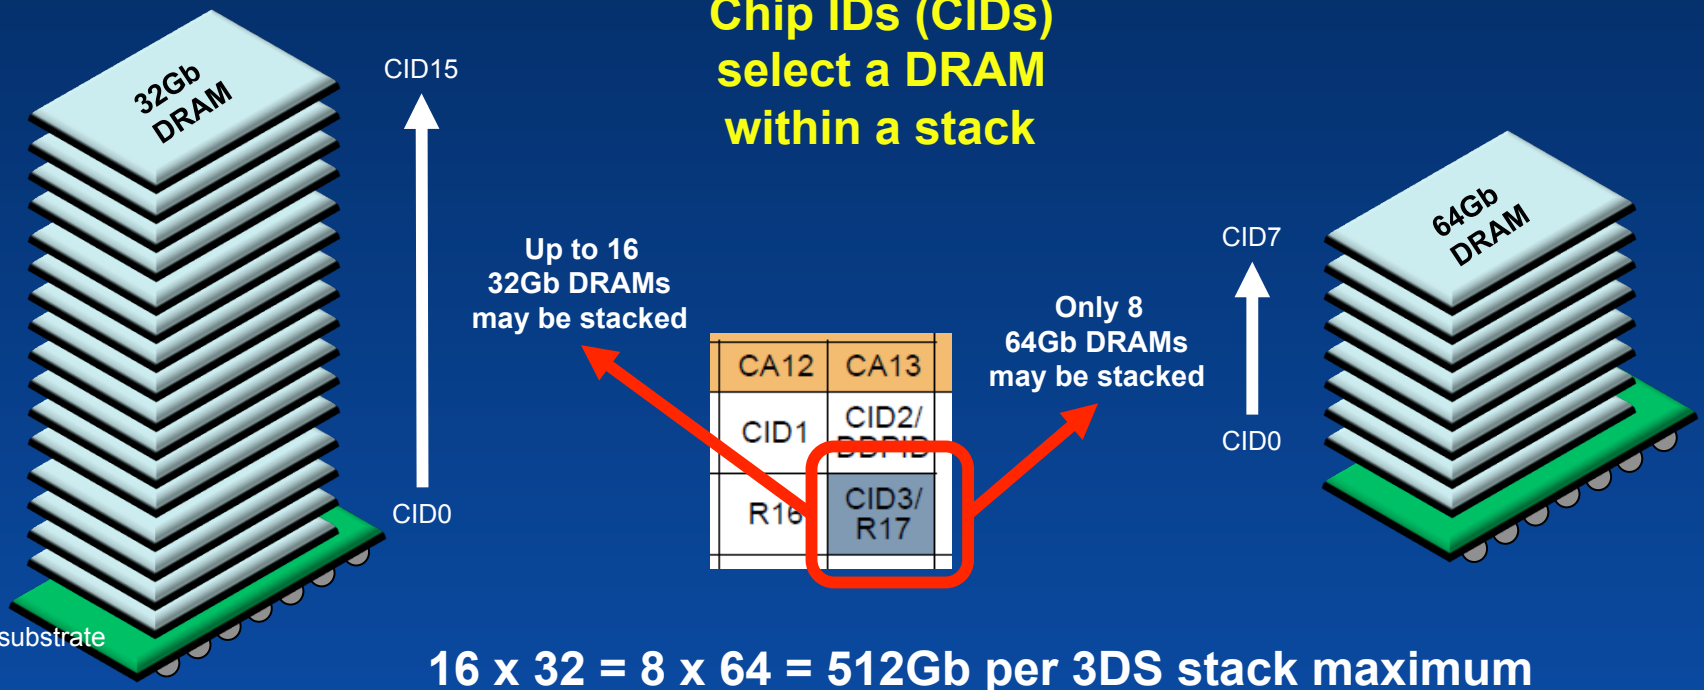

16 x 32 = 8 x 64 = 512Gb per 3DS stack maximum

The Universe of Storage Class Memories (Persistent Memories)

Flash Memory Summit

NANTERO

Many technologies coming online to fill the gap between DRAM and Flash

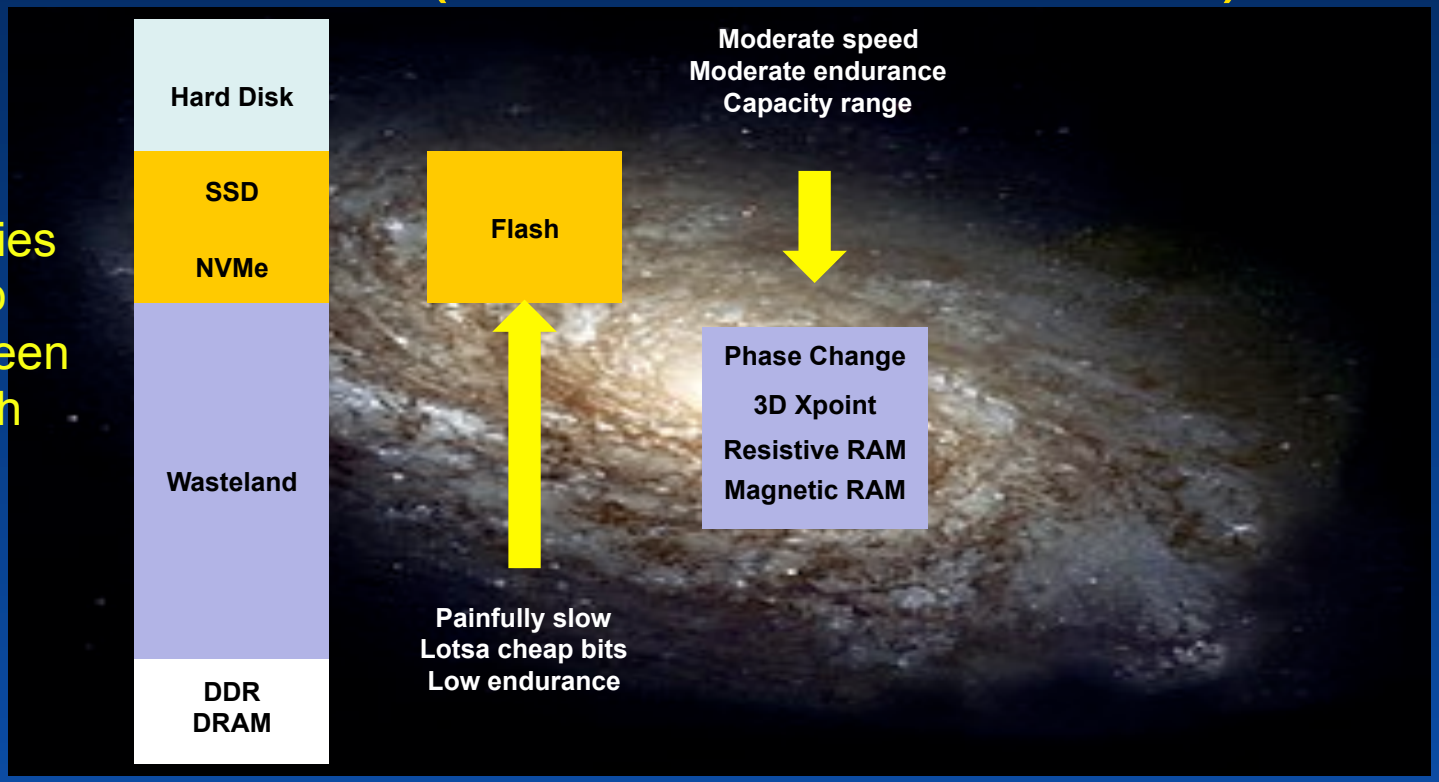

NANTERO

Overview: Storage Class Memories

	DRAM	MRAM	ReRAM	PCM / 3DXpoint	FeRAM	Flash
Non-volatility	No	Yes	Yes	Yes	Yes	Yes
Endurance	10^{15}	10^{9}	10^{6}	10^{6}	10^{14}	10^{3}
Read Time	10 ns	50 ns	200 ns	100 ns	50 ns	50M ns
Write Time	10 ns	1000 ns	1000 ns	1000 ns	50 ns	25M ns
	Memory		**Storage Class Memory**			**Storage**

Chasm

Changing "SCM" to "Persistent Memory" doesn't help…
In fact, it may make it less clear…

NANTERO

The NVDIMM-P Protocol

NVDIMM-P protocol invented to handle memory with low endurance

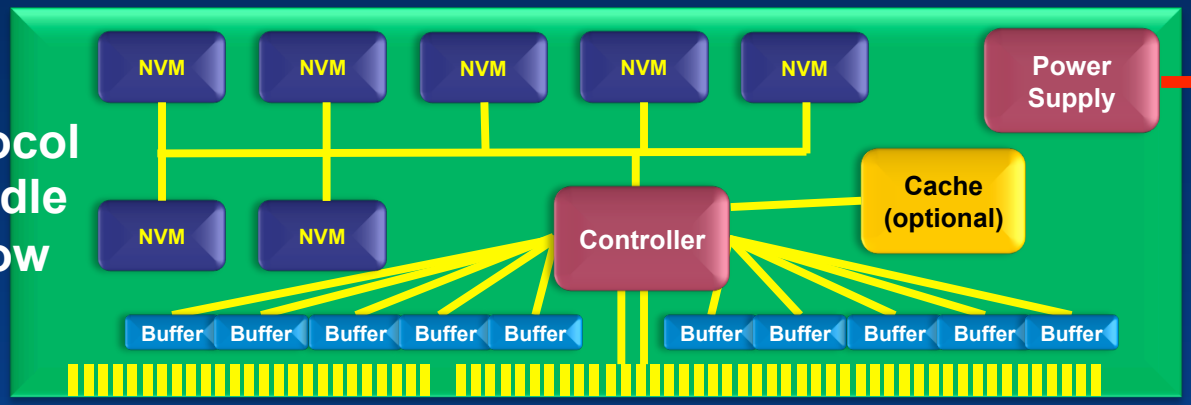

Optional external energy source

Non-deterministic credit based system allows time for "clean-up"

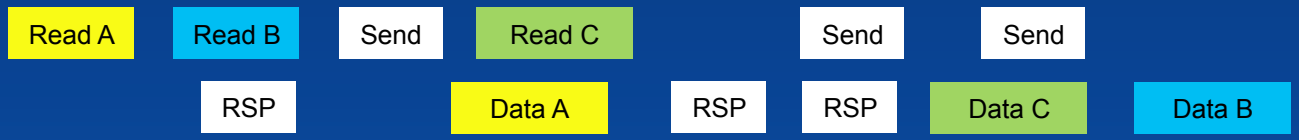

Out-of-order data returned with ID

Nantero NRAM™

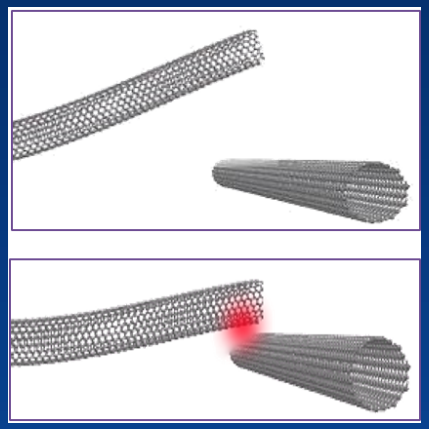

Carbon nanotubes connected or split using electrostatic force

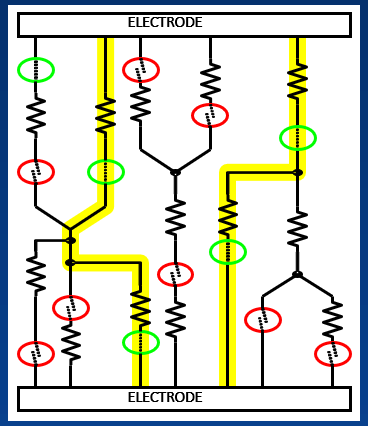

Stochastic array of tubes forms a resistive matrix

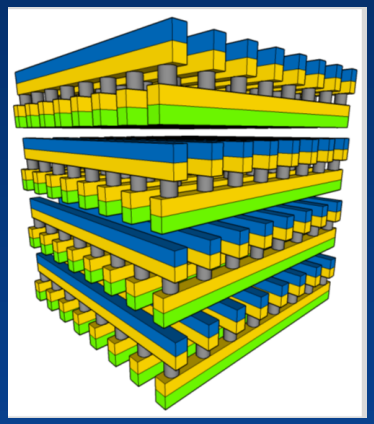

Stackable into a crosspoint memory array

For details, attend my talk at Hot Chips, September 2018

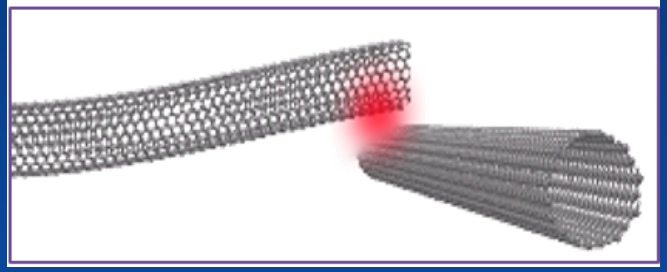

Unlimited Write Endurance

NRAM carbon nanotube bonds are atomic

Once made, they stay connected until disrupted

Insensitive to heat, radiation, etc.

No CNT failure has been detected

Unlimited write endurance expected

Flash Memory Summit

NANTERO

DDR4 NRAM

What can you do with a 10 ns core speed?

Not only perform as a DDR4 SDRAM

With non-volatility

And add the ECC on-the-fly function of DDR4E

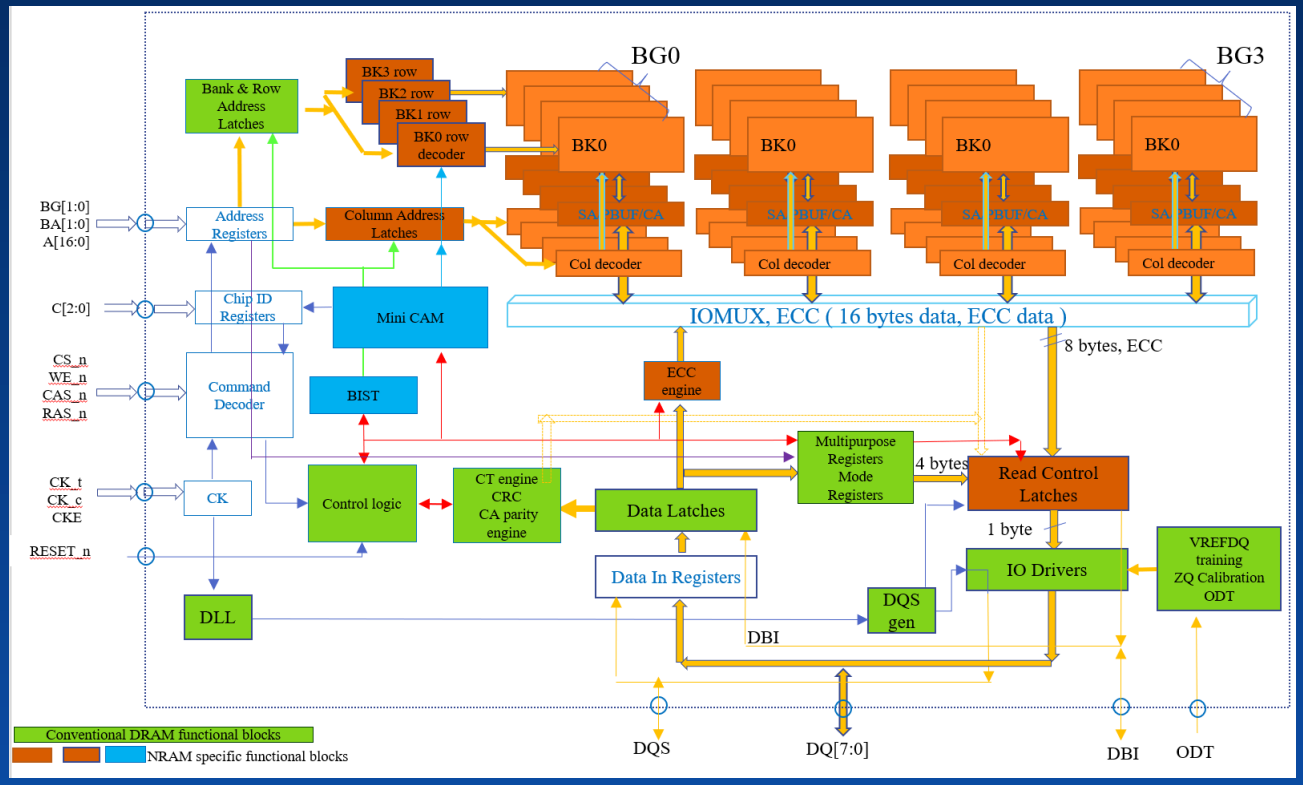

NANTERO

RDIMM = NVDIMM

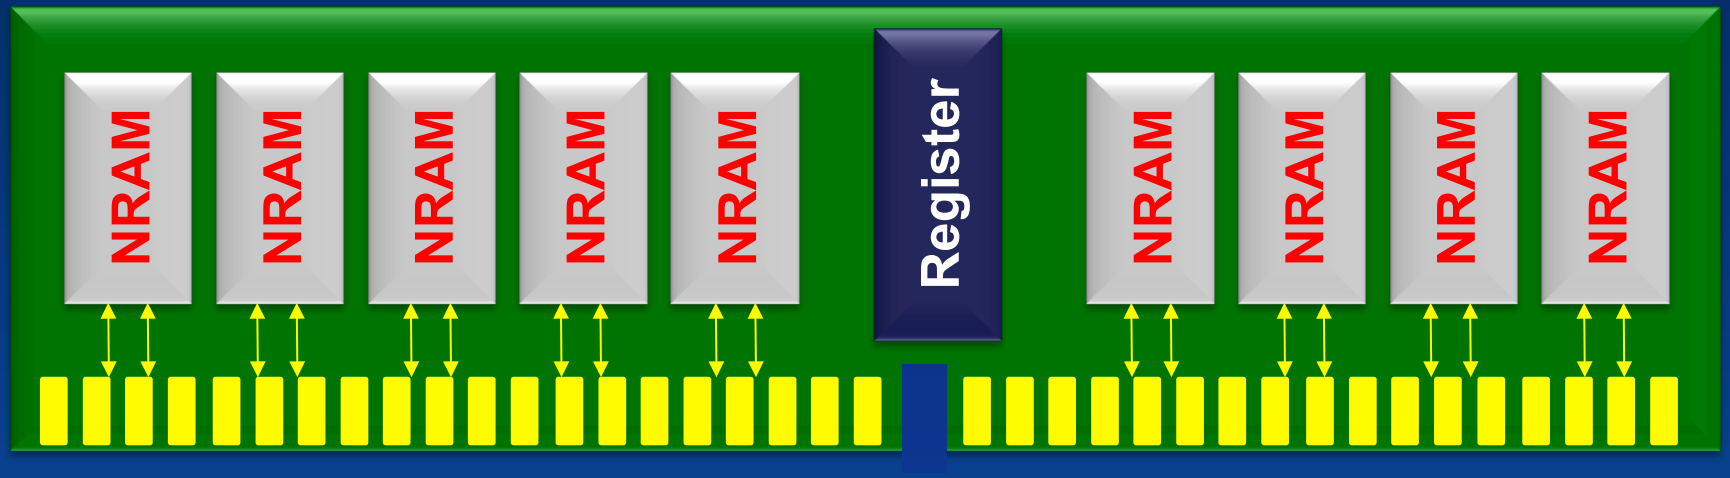

With inherent non-volatility, an RDIMM is an NVDIMM

No need for backup & restore procedure nor for battery backup

NANTERO

Introducing Memory Class Storage

Industry primed for a DRAM replacement technology

Hard Disk

SSD

NVMe

Wasteland

DDR DRAM

Flash

Painfully slow
Lotsa cheap bits
Low endurance

Moderate speed
Moderate endurance
Capacity range

Phase Change

3D Xpoint

Resistive RAM

Magnetic RAM

> DRAM performance
= DRAM endurance
> DRAM capacity
< DRAM price

DDR NRAM

NANTERO

Memory Class Storage vs SCM

	DRAM	NRAM	MRAM	ReRAM	PCM / 3DXpoint	FeRAM	Flash
Non-volatility	No	Yes	Yes	Yes	Yes	Yes	Yes
Endurance	10^{15}	10^{15}	10^{9}	10^{6}	10^{6}	10^{14}	10^{3}
Read Time	10 ns	10 ns	50 ns	200 ns	100 ns	50 ns	50M ns
Write Time	10 ns	10 ns	1000 ns	1000 ns	1000 ns	50 ns	25M ns
	Memory & Memory Class Storage		**Storage Class Memory**				**Storage**

…aka "Persistent Memory"

When DRAM Ends…

NANTERO

DRAM

End of the line
32Gb mono
512Gb 3DS

8Gb→16Gb

DDR4→5
3DS 8H → 16H

16Gb→32Gb

Flash Memory Summit

NANTERO

…NRAM just begins…

DRAM

End of the line
32Gb mono
512Gb 3DS

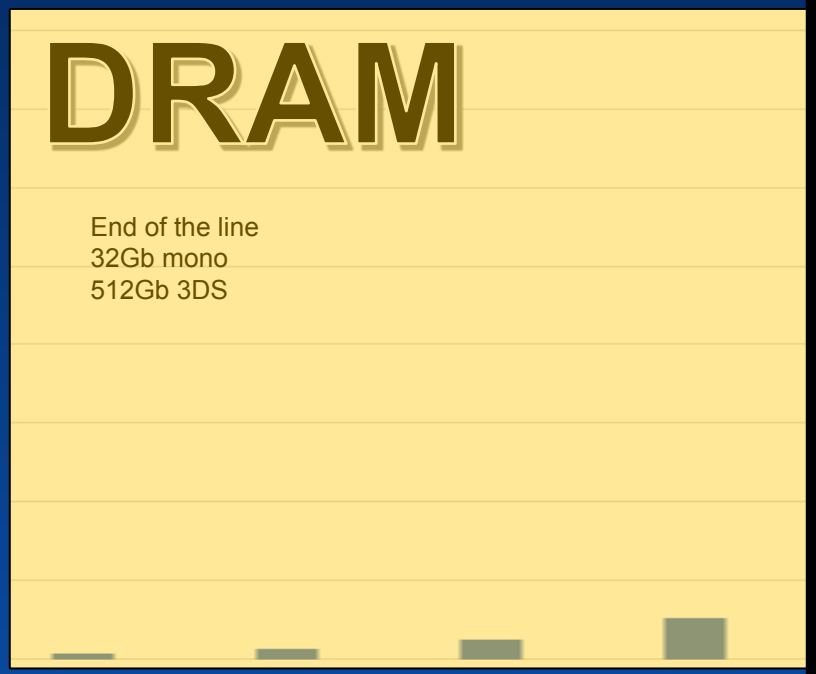

NRAM

Roadmap continues
512Gb mono
8Tb 3DS

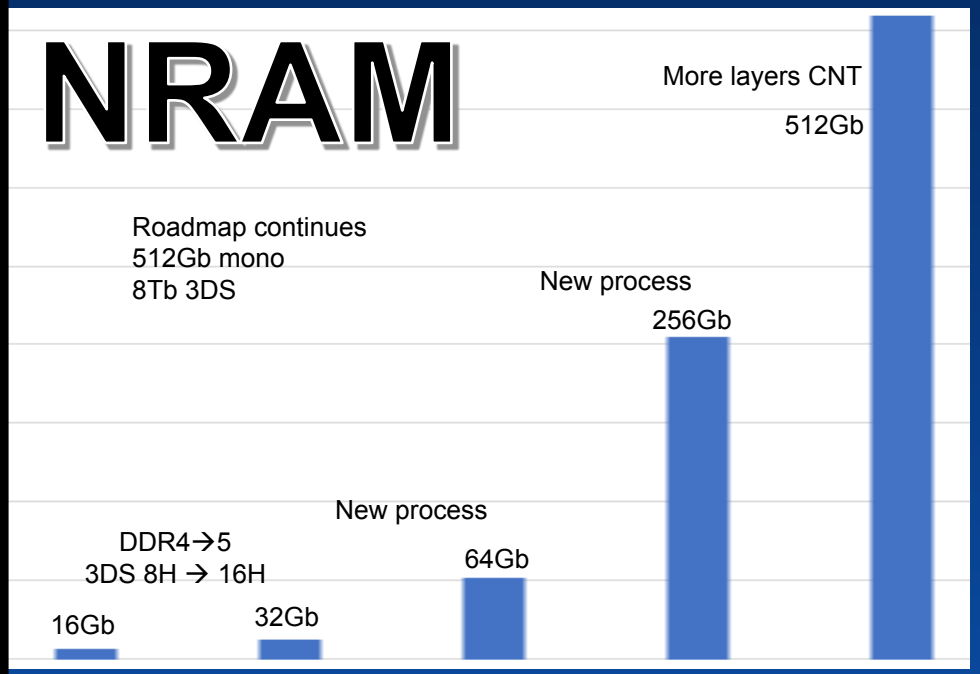

Flash Memory Summit

NANTERO

Density Roadmap

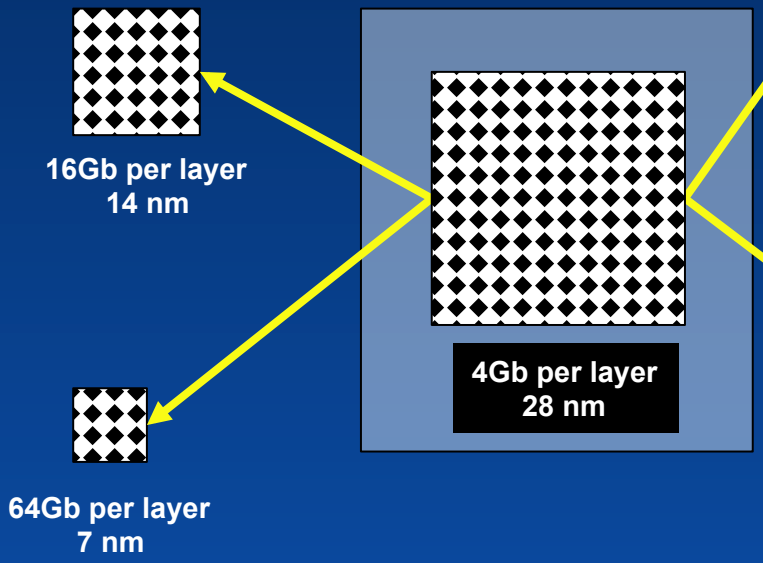

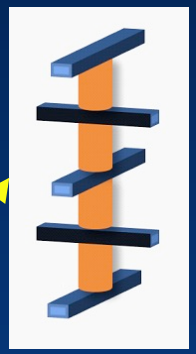

4 layers of CNT
4X density

16Gb @ 28 nm
64Gb@ 14 nm
256Gb @ 7 nm

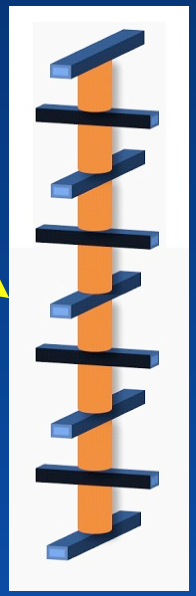

8 layers of CNT
8X density

32Gb@ 28 nm
128Gb@ 14 nm
512Gb @ 7 nm

Will DDR5 Let Us Deliver?

NANTERO

Existing

Function	Abbrevia-tion	CS	CA Pins													
			CA0	CA1	CA2	CA3	CA4	CA5	CA6	CA7	CA8	CA9	CA10	CA11	CA12	CA13
Activate	ACT	L	L	L	R0	R1	R2	R3	BA0	BA1	BG0	BG1	BG2	CID0	CID1	CID2/ DDPID
		H	R4	R5	R6	R7	R8	R9	R10	R11	R12	R13	R14	R15	R16	CID3/ R17

Row Extension	REXT	L	H	H	H	L	L	R17	R18	R19	R20	R21	R22	R23	V	V

New Command

Row Extension

	A+ROW	A+ROW	A+ROW		B+ROW	B+ROW		C+ROW		A+ROW	
REXT A	ACT	ACT	ACT	REXT B	ACT	ACT	REXT C	ACT	REXT A	ACT	

Proposal before the JEDEC Committee

NANTERO

Extended Addressing Protocol

64 Gb Addressing Table, Option 2				
Configuration		16 Gb x4	8 Gb x8	4 Gb x16
Bank Address	Bank Group Address	BG0~BG2	BG0~BG2	BG0~BG1
	Bank Address in a BG	BA0~BA1	BA0~BA1	BA0~BA1
	# BG / # Banks per BG / # Banks	8 / 4 / 32	8 / 4 / 32	4 / 4 / 16
Extended Row Address		R17	R17	R17
Row Address		R0~R16	R0~R16	R0~R16
Column Address		C0~C10	C0~C9	C0~C9
Page Size		1 KB	1 KB	2 KB
Chip IDs / Maximum 3DS Stack Height		CID3~0 / 16H	CID3~0 / 16H	CID3~0 / 16H

1 Tb Addressing Table				
Configuration		256 Gb x4	128 Gb x8	64 Gb x16
Bank Address	Bank Group Address	BG0~BG2	BG0~BG2	BG0~BG1
	Bank Address in a BG	BA0~BA1	BA0~BA1	BA0~BA1
	# BG / # Banks per BG / # Banks	8 / 4 / 32	8 / 4 / 32	4 / 4 / 16
Extended Row Address		R17~R21	R17~R21	R17~R21
Row Address		R0~R16	R0~R16	R0~R16
Column Address		C0~C10	C0~C9	C0~C9
Page Size		1 KB	1 KB	2 KB
Chip IDs / Maximum 3DS Stack Height		CID3~0 / 16H	CID3~0 / 16H	CID3~0 / 16H

Extends addressing to 4Tb (512 GB) per die or 64Tb (8 TB) per 3DS stack

NANTERO

The Post-DRAM Age is Inevitable

Deep Learning

Self Driving

Data Mining

Artificial Intelligence

In-memory Computing

Memory demand is only headed north

What are you going to do after 32Gb?

NANTERO

Summary

DRAM is going away after 32Gb

“Storage Class Memory” fills a gap between DRAM and Flash

NVDIMM-P protocol allows SCM to share a DRAM channel

NRAM defines a new “Memory Class Storage” category

MCS = Unlimited write endurance and full DRAM speed

NRAM process roadmap extends beyond DRAM by > 16X

Questions?

Bill Gervasi
Principal Systems Architect
bilge@Nantero.com