SAP HANA Solutions on Flash

based on Cascade Lake Processors

Saroj Mohapatra
SAP SME
August 2019

Business Architecture SAP Business Application Revolution

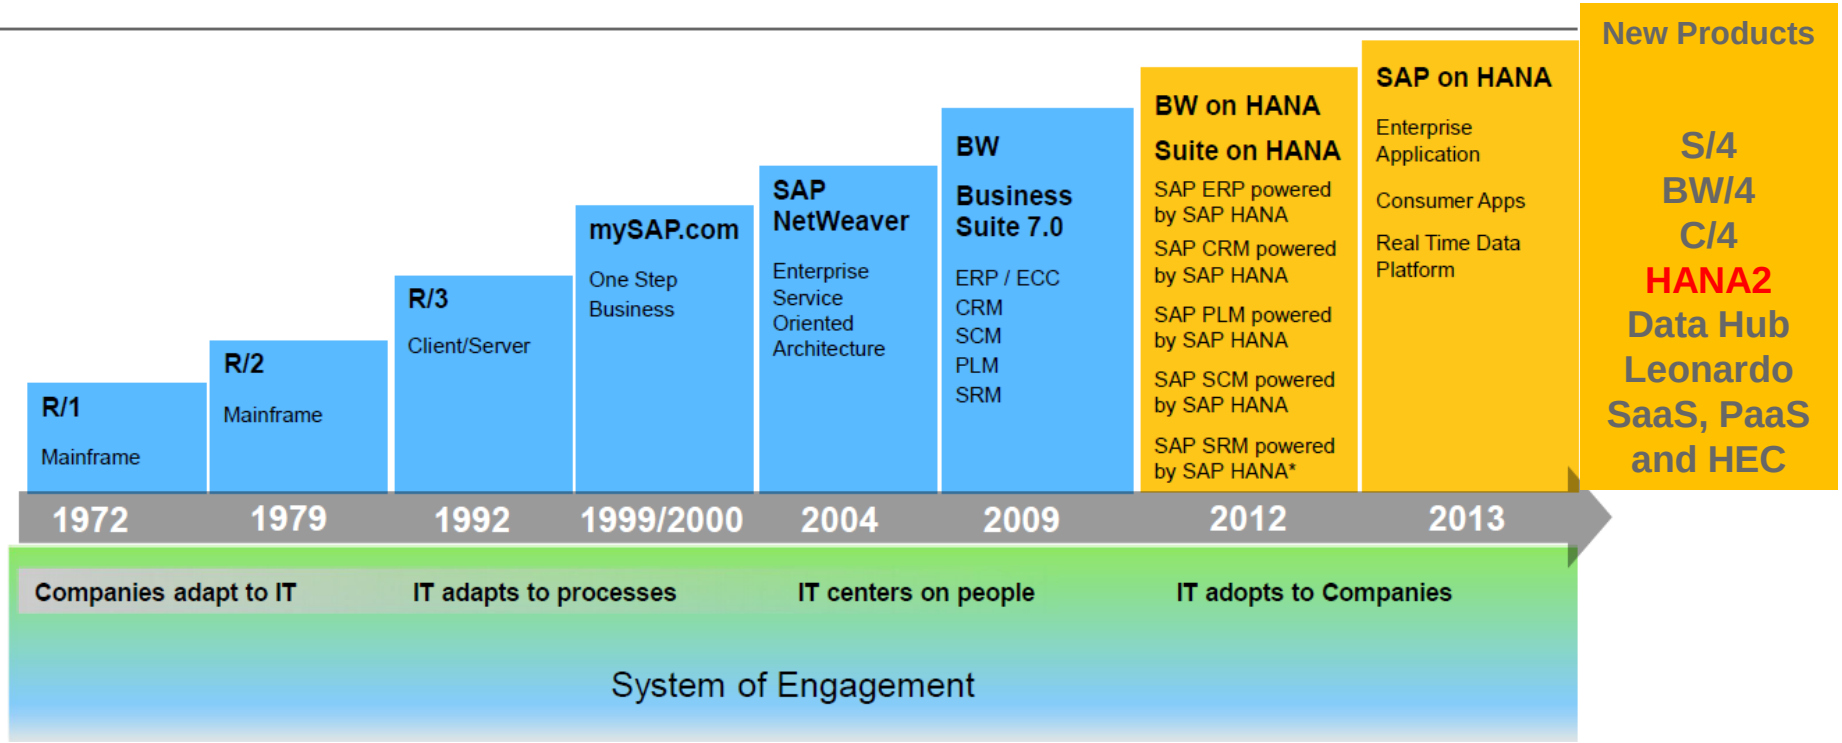

* All SAP customers needs to move to HANA by 2025

SAP Customer Concerns Before Optane DC

Do I have to go to a bigger server to accommodate my HANA memory foot print?

* HANA Memory to socket ratio rule with DDR4 memory

Can I use 2 socket or 4 socket server and reduce my TCO.

* Avoiding 8 socket and bigger servers

I can move some of my data to warm tier but that project is going to take some time and I don't have fund for it.

* Using solution as Dynamic Tiering (Warm tier) to reduce in memory database size

Intel Processors with Optane DC

Cisco M1/M2 Servers (EOLed)	Cisco M3 Servers (EOLed)	Cisco M4 Servers	Cisco M5 Servers
Thurley Platform	Romley Platform	Grantley Platform	Purley Platform
Nehalem 45nm New Microarchitecture	Westmere 32nm New Processor Technology	Sandy Bridge 32nm New Microarchitecture	Ivy Bridge 22nm New Processor Technology	Haswell 22nm New Microarchitecture	Broadwell 14nm New Processor Technology	Skylake 14nm New Microarchitecture	**Cascade Lake** 14nm New Processor Technology
DDR3 Memory Up to	DDR3 Memory Up to 1866 MHz	DDR4 Memory 1, 2 DPC= 2400 MHz, 3DPC = up to 1866 MHz	DDR4 Memory 2933 MHz 2DPC 3D Xpoint Ready
PCIe 3.0: 36 lanes per CPU	PCIe 3.0: 40 lanes per CPU	PCIe 3.0: 40 lanes per CPU	PCIe: 48 lanes per CPU

Raw Component Performance Comparison

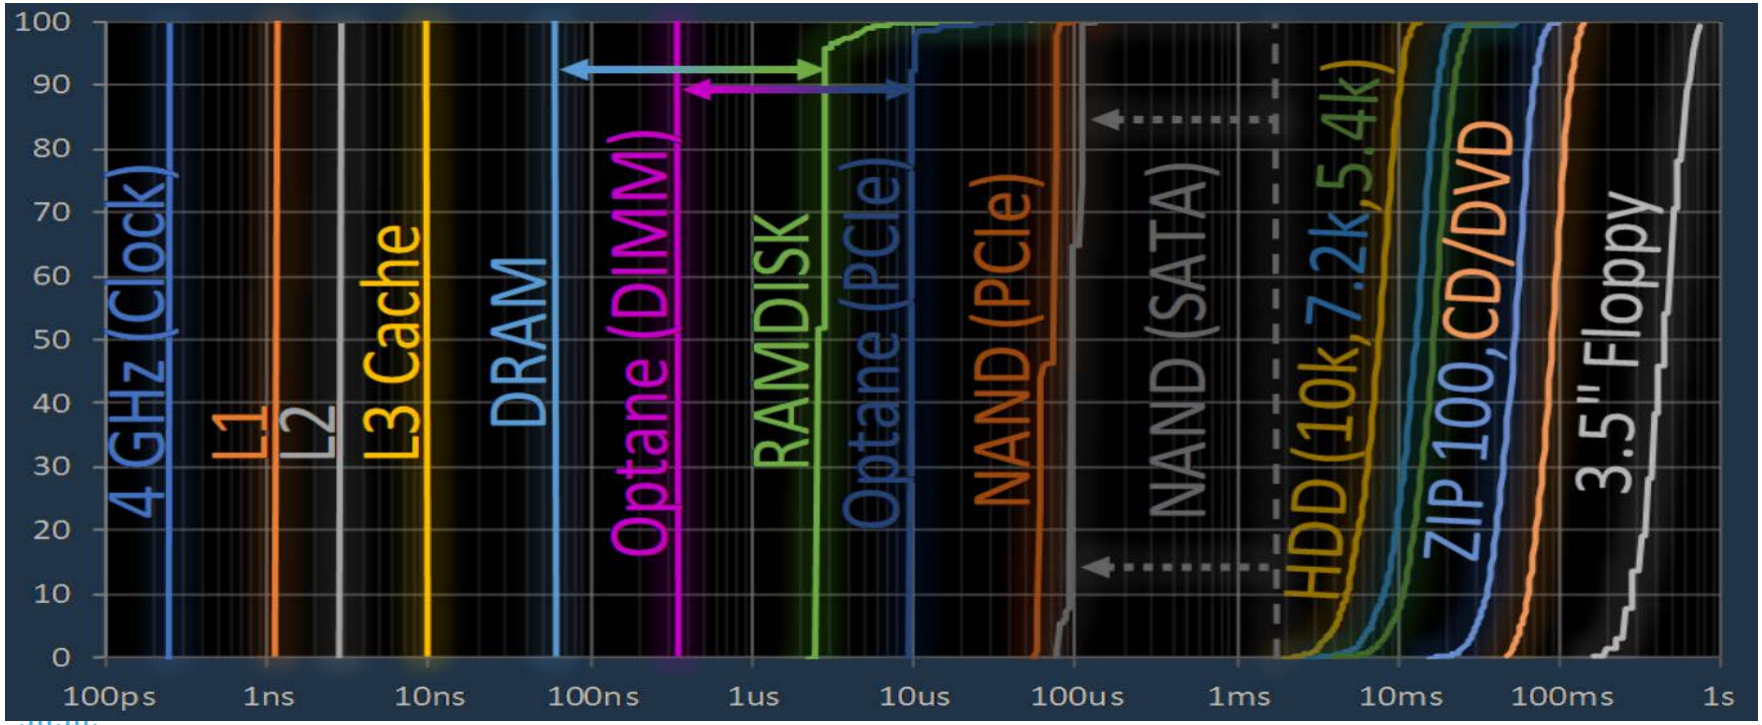

REIMAGINING THE DATA CENTER MEMORY AND STORAGE HIERARCHY

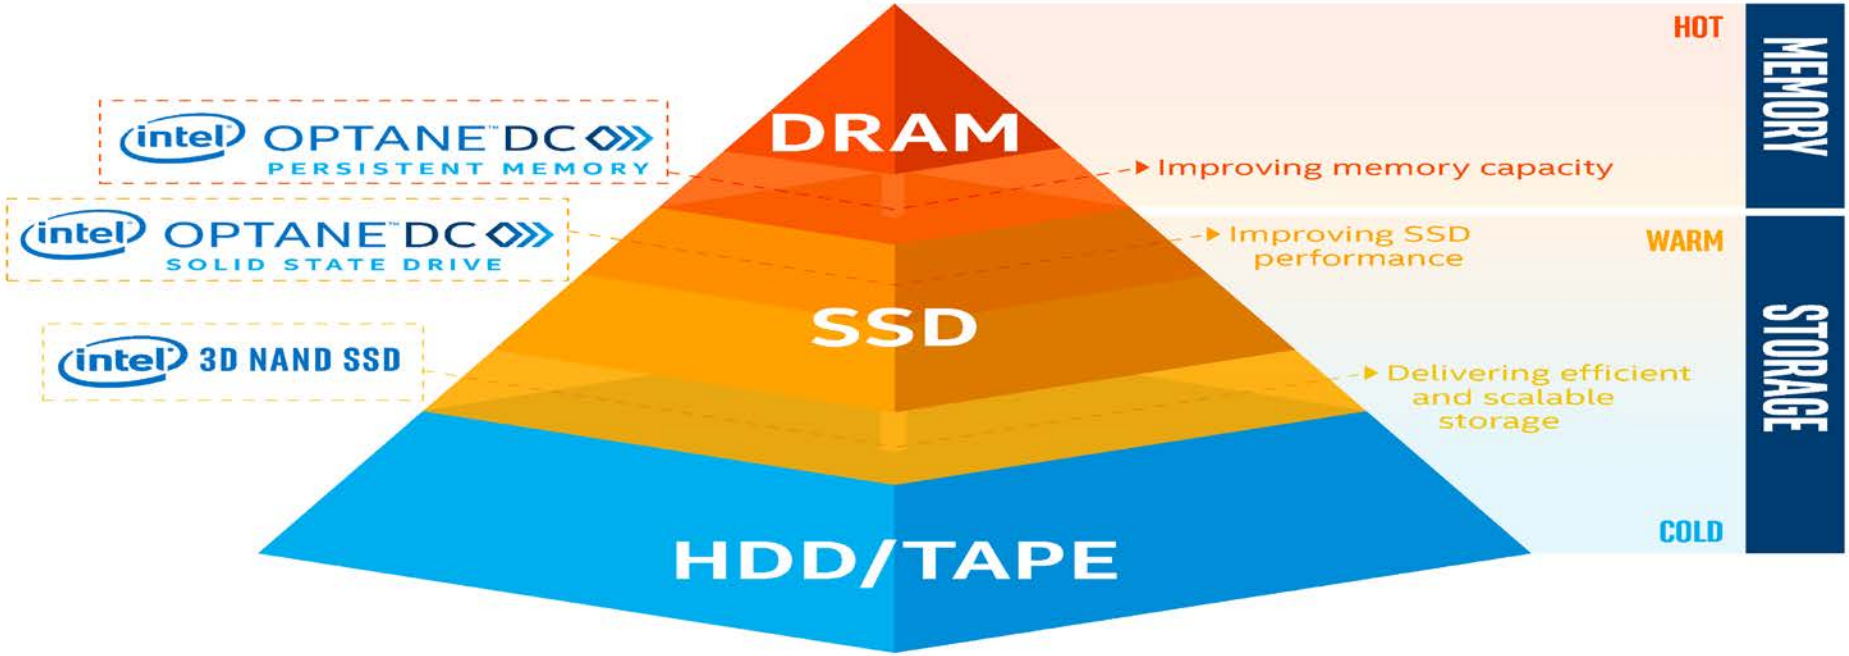

Two Functional Modes – App Direct and Memory

App Direct Mode: Both DCPMM and DRAM capacity counts towards to CPU Tiering
Memory Mode: DRAM as cache. Only DCPMM capacity counts towards to CPU Tiering
Mix Mode: DRAM as cache. Only DCPMM capacity counts towards to CPU Tiering

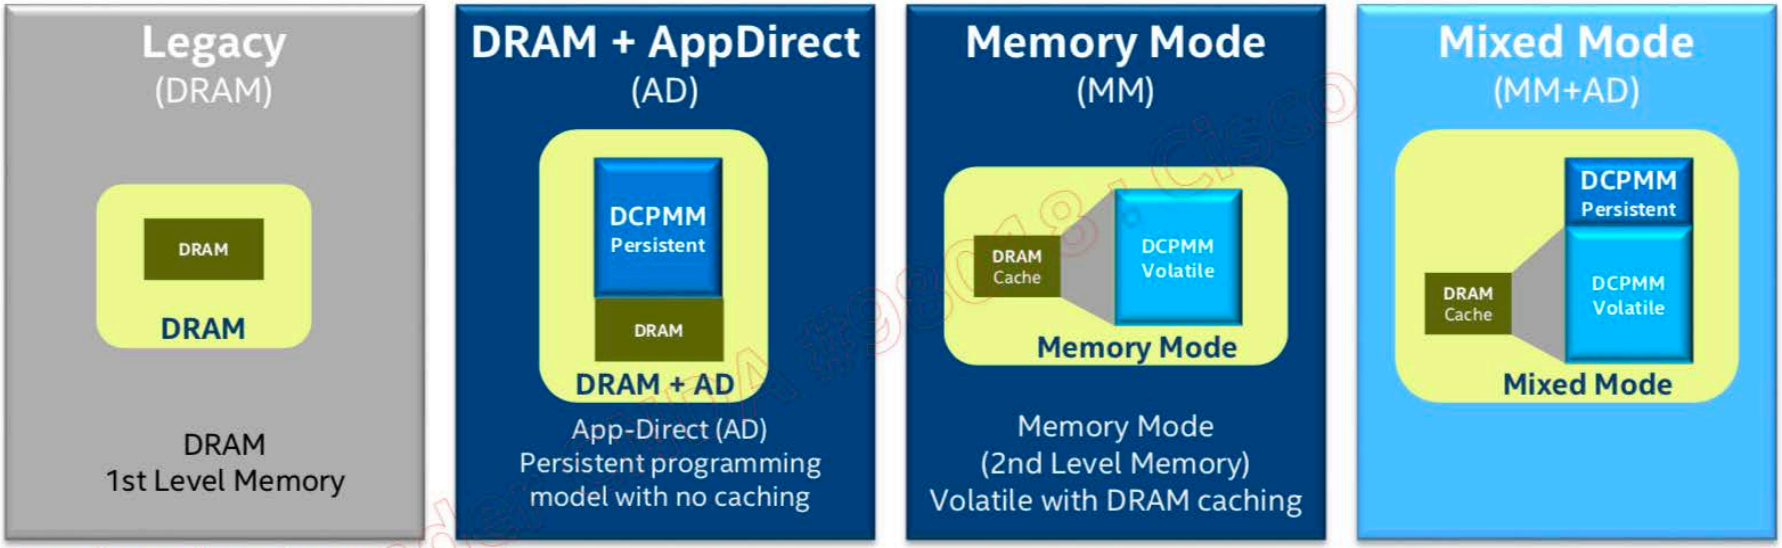

Example: For Illustration Purposes Only

2S DRAM Capacity	768GB (24X32)	384GB (12x32)	384GB (12x32)	384GB (12x32)
2S DCPMM DIMM Capacity	NA	1.5TB (12x128)	1.5TB (12x128)	1.5TB (12x128)
What OS Sees	768GB (24x32)	384GB (12x32) + 1.5TB (12x128)	1.5TB (12x128)	1.28TB (10x128) + 256GB (2x128)
Persistent	NA	1.5TB (12x128)	NA	256GB (2x128, if 1/6th cap is AD)

Skylake and Cascade Lake Balanced DDR4 Memory Configurations

Skylake Memory Configuration

	Controller #	0						1						Capacity (GB) per CPU	Capacity (GB) with #CPUs			
	Channel #	0		1		2		0		1		2						
	DIMM slot #	1	2	3	4	5	6	7	8	9	10	11	12		2	4	6	8
1DPC	16GB DIMM	16		16		16		16		16		16		96	192	384	576	768
2DPC	16GB DIMM	16	16	16	16	16	16	16	16	16	16	16	16	192	384	768	1152	1536
2DPC	*16/32G mix*	*32*	*16*	*32*	*16*	*32*	*16*	*32*	*16*	*32*	*16*	*32*	*16*	288	576	1152	1728	2304
1DPC	32GB DIMM	32		32		32		32		32		32		192	384	768	1152	1536
2DPC	32GB DIMM	32	32	32	32	32	32	32	32	32	32	32	32	384	768	1536	2304	3072
1DPC	64GB DIMM	64		64		64		64		64		64		384	768	1536	2304	3072
2DPC	*32/64G mix*	*64*	*32*	*64*	*32*	*64*	*32*	*64*	*32*	*64*	*32*	*64*	*32*	576	1152	2304	3456	4608
2DPC	64GB DIMM	64	64	64	64	64	64	64	64	64	64	64	64	768	1536	3072	4608	6144
1DPC	128GB DIMM	128		128		128		128		128		128		768	1536	3072	4608	6144
2DPC	*64/128G mix*	*128*	*64*	*128*	*64*	*128*	*64*	*128*	*64*	*128*	*64*	*128*	*64*	*1152*	2304	4608	6912	9216
2DPC	128GB DIMM	128	128	128	128	128	128	128	128	128	128	128	128	1536	3072	6144	9216	12288

Notes:

1. Only use Cisco UCS validated DIMMs
2. All DIMMs must be same technology, i.e. LRDIMM, RDIMM, SD3, etc
3. Only adjacent sized DIMMs can be used, e.g. 32G/64G, 64G/128G. Cannot use 16G/64G DIMMs
4. Must have either 6 or 12 DIMMs per CPU, with every DIMM channel filled
5. Observe Intel requirements, including DIMM ranking, etc.

Cascade Lake memory config

CLX Memory Configuration

Controller #		0						1						CPU Type	Capacity (GB) per CPU	Capacity (GB) with #CPUs				DDR/AEP ratio 1:x
Channel #		0		1		2		0		1		2								
DIMM slot #		1	2	3	4	5	6	7	8	9	10	11	12			2	4	6	8	
2DPC	128GB AEP + 32GB DRAM	128	32	128	32	128	32	128	32	128	32	128	32	base	960	***1920***	***3840***	***5760***	***7680***	4
2DPC	128GB AEP + 64GB DRAM	128	64	128	64	128	64	128	64	128	64	128	64	M	1152	***2304***	***4608***	***6912***	***9216***	2
2DPC	128GB AEP + 128GB DRAM	128	128	128	128	128	128	128	128	128	128	128	128	M	1536	***3072***	***6144***	***9216***	***12288***	1
2DPC	256GB AEP + 64GB DRAM	256	64	256	64	256	64	256	64	256	64	256	64	M	1920	***3840***	***7680***	***11520***	***15360***	4
2DPC	128GB AEP + 256GB DRAM	128	256	128	256	128	256	128	256	128	256	128	256	L	2304	***4608***	***9216***	***13824***	***18432***	0,5
2DPC	256GB AEP + 128GB DRAM	256	128	256	128	256	128	256	128	256	128	256	128	L	2304	***4608***	***9216***	***13824***	***18432***	2
2DPC	256GB AEP + 256GB DRAM	256	256	256	256	256	256	256	256	256	256	256	256	L	3072	***6144***	***12288***	***18432***	***24576***	1
2DPC	512GB AEP + 128GB DRAM	512	128	512	128	512	128	512	128	512	128	512	128	L	3840	***7680***	***15360***	***23040***	30720	4
2DPC	512GB AEP + 256GB DRAM	512	256	512	256	512	256	512	256	512	256	512	256	L	4608	***9216***	***18432***	27648	36864	2

´= Supported configurations

´= not supported

What's New in SAP HANA 2.0 SPS03: Persistence

- Non-volatile Memory
  - HANA 2.0 SPS 03 supports persistent memory, e.g., based on Intels Optane DC memory technology.
    - DIMM form-factor, replacing DRAM.
  - Column Store Main is placed on persistent memory
- Affects more than 95% of persisted data in most SAP HANA systems. About 95percent of data in main and it is 10–20x compressed
  - Significantly **lower restart times**. Column store does not have to be loaded at startup.
  - **Lower TCO** than DRAM.
- No changes to the persistence layer.

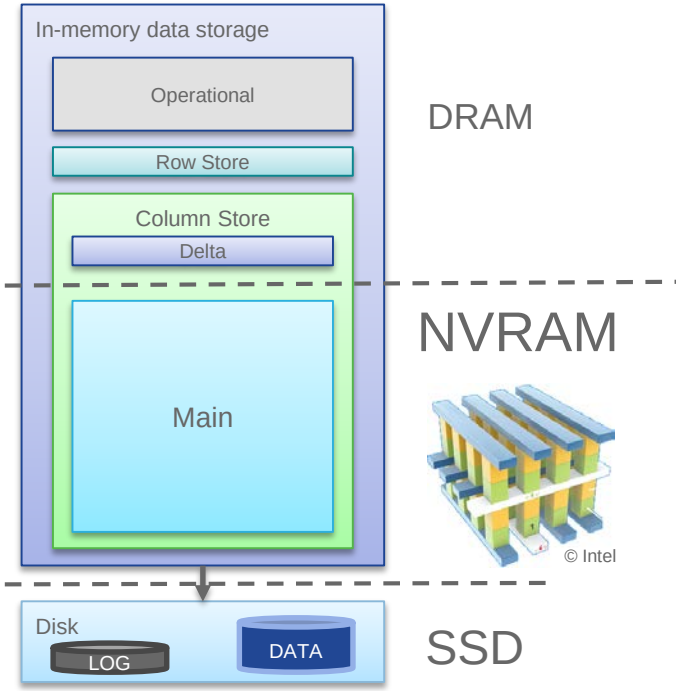

Memory Configuration

Example: 2 sockets, largest DRAM and largest PMEM configurations

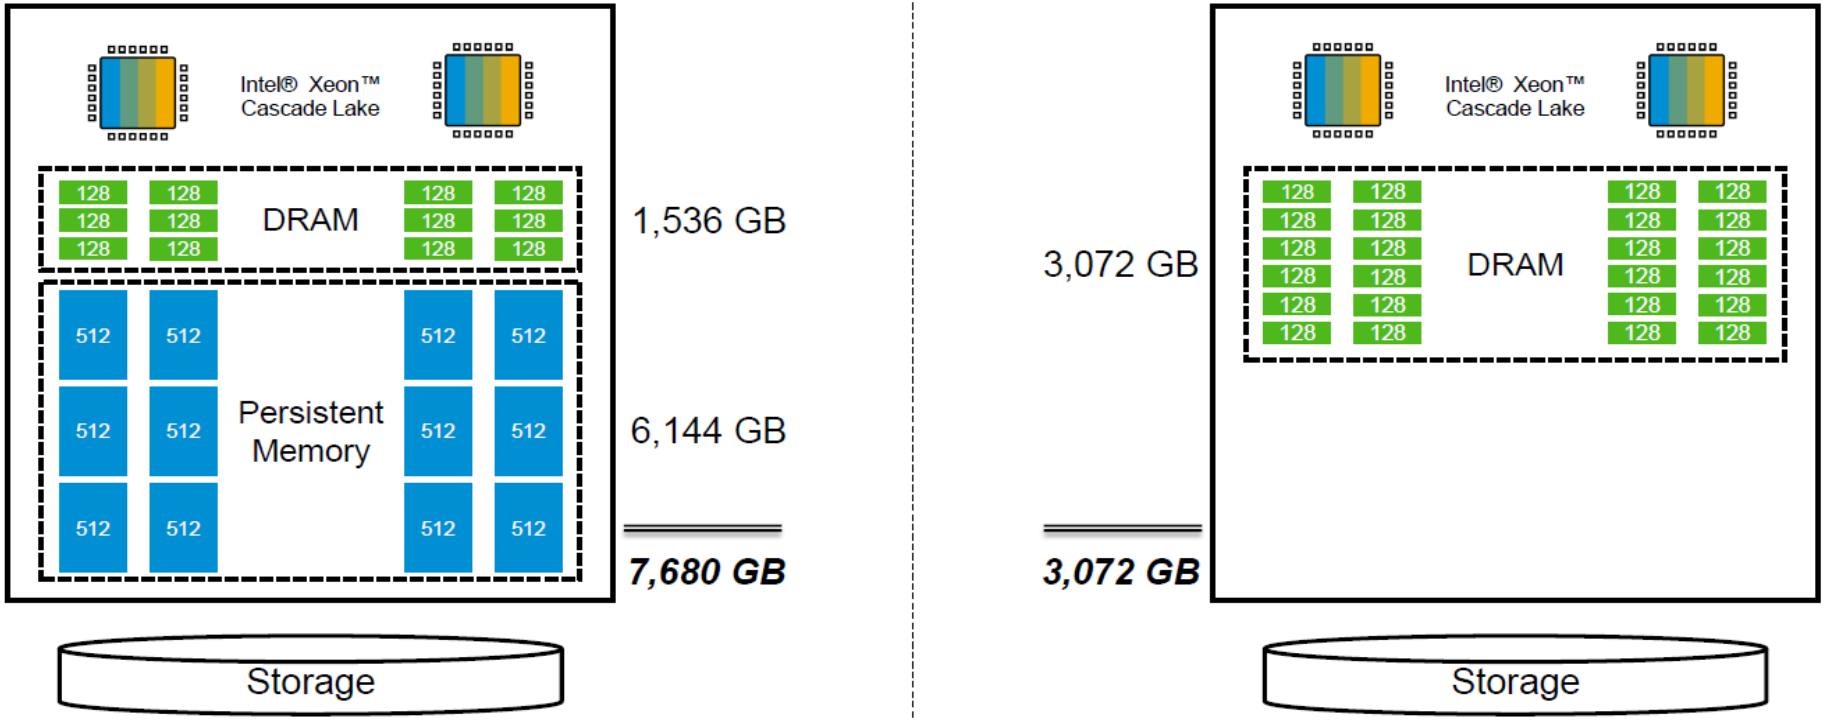

Actual configuration and ratios between DRAM, PMEM and CPU depend on application sizing.

What's New in SAP HANA 2.0 SPS03: Persistence

Non-Volatile Memory

intel OPTANE™ DC PERSISTENT MEMORY

BIG

Enables 4S systems with >12 TB memory capacity

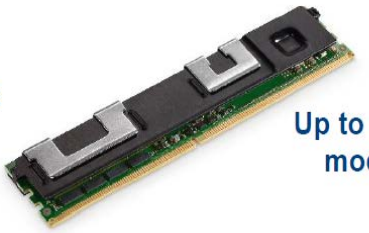

Up to 512 GB modules

- Column Store Main is placed on persistent memory
- Affects more than 95% of persisted data in most SAP HANA systems
- About 95percent of data in main and it is 10–20x compressed
- Significantly lower restart times. Column store does not have to be loaded at startup.

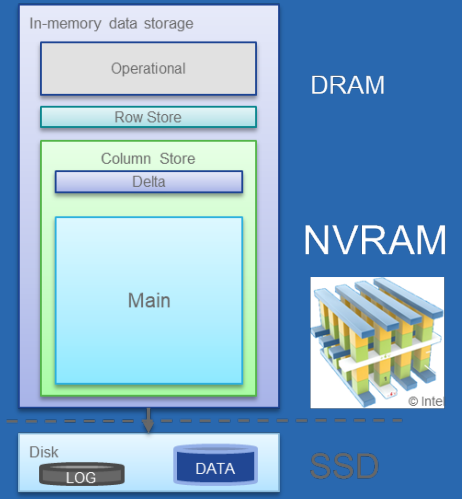

Faster starts help ensure business continuity and service-level agreements

4-MINUTE

average SAP HANA® 2.0 SPS 3.0 index server start time using Intel® Optane™ DC persistent memory—a **12.5x decrease** compared to DRAM alone[11]

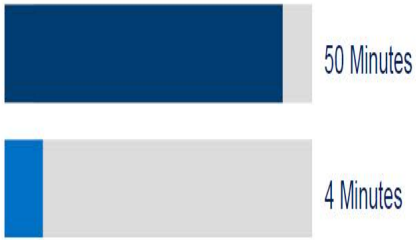

Use of Intel® Optane™ DC with Extension Node

Extension node now uses Intel® Optane™ DC instead of disk.

Benefit

- No disk loading/unloading to RAM.
- Expected significant speed benefit when accessing tiered/cold data.
- Backup and restore is exactly the same. Failover node can support the extension node.

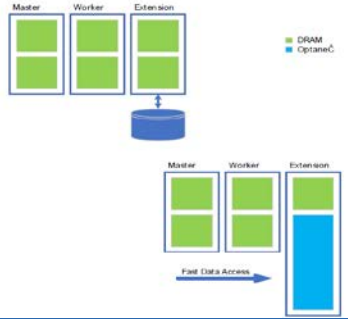

Intel Optane DC PMEM Configuration Rules

- A sizing exercise is required to properly configure the server(s).
- All system board memory slots must be fully populated (no half loads)
- Must be equal number of DDR & PMEM per CPU socket
- **SAP HANA only supports DDR/PMEM capacity ratios of 2:1,1:1, 1:2, and 1:4.**
- App Direct is only mode supported
- No half load CPU systems – It has to be fully populated

How can SAP HANA benefit from OptaneDC

Lower Platform TCO

Restarts up to 12.5x faster

Increased memory capacity


CISCO