

Lenovo Update

Benchmarking DAOS in Lenovo's HPC Innovation Center



Lenovo

19-Nov-2021

| Sigrun Eggerling, Lenovo DAOS PoC Lead (seggerling@lenovo.com)

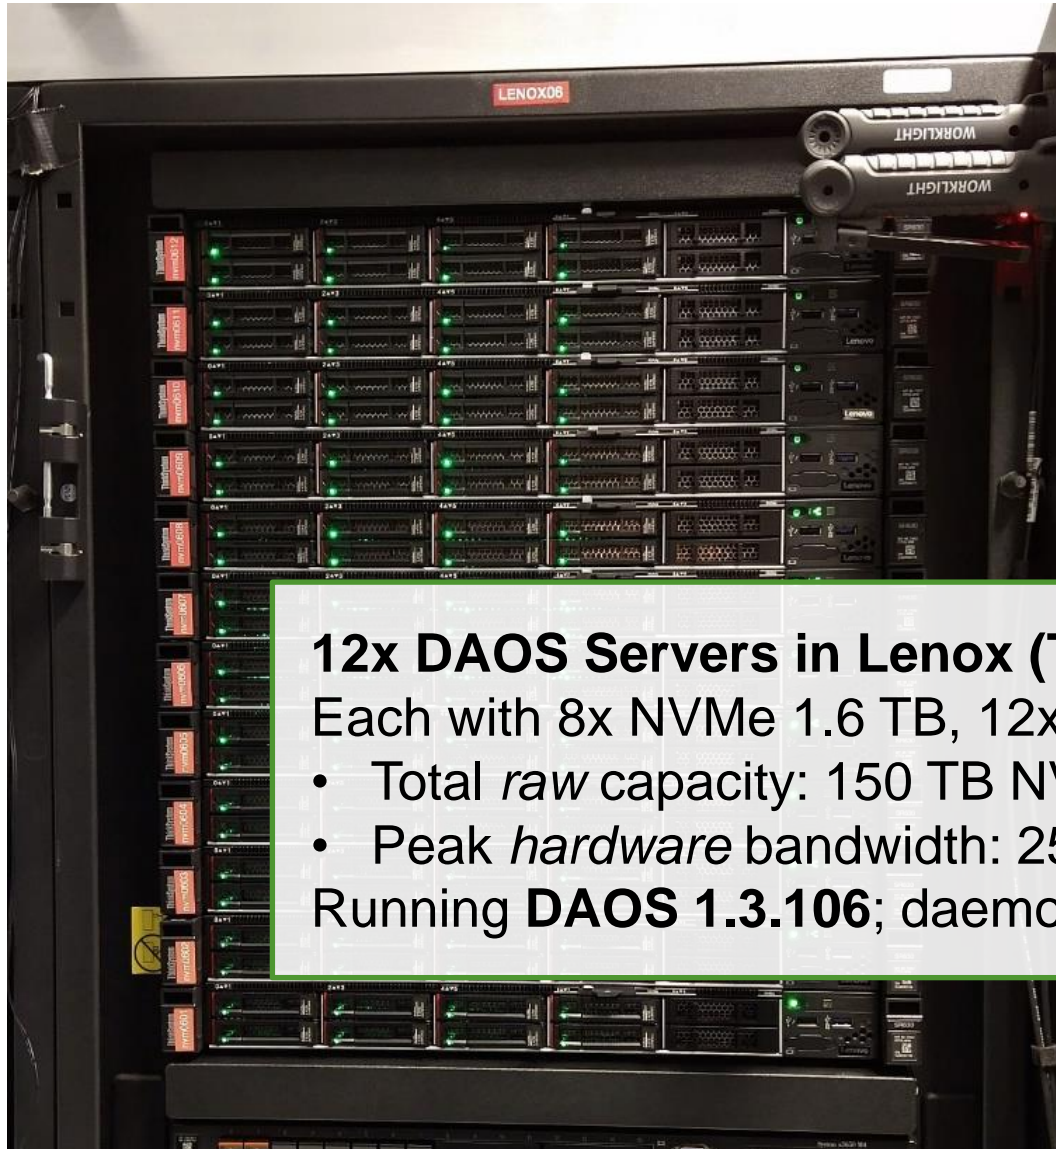
| Steve Eiland, Lenovo HPC Storage Product Management (seiland@lenovo.com)

Lenovo activities around DAOS

- **DAOS benchmarking facility as part of the HPC / AI Innovation Center**
 - Cascade Lake / Icelake based servers running DAOS
 - Integrated into LENOX compute infrastructure
 - Close collaboration with Intel since 2019
 - IO500 benchmark published June 2021



DAOS Servers in Lenovo's HPC Innovation Center Stuttgart



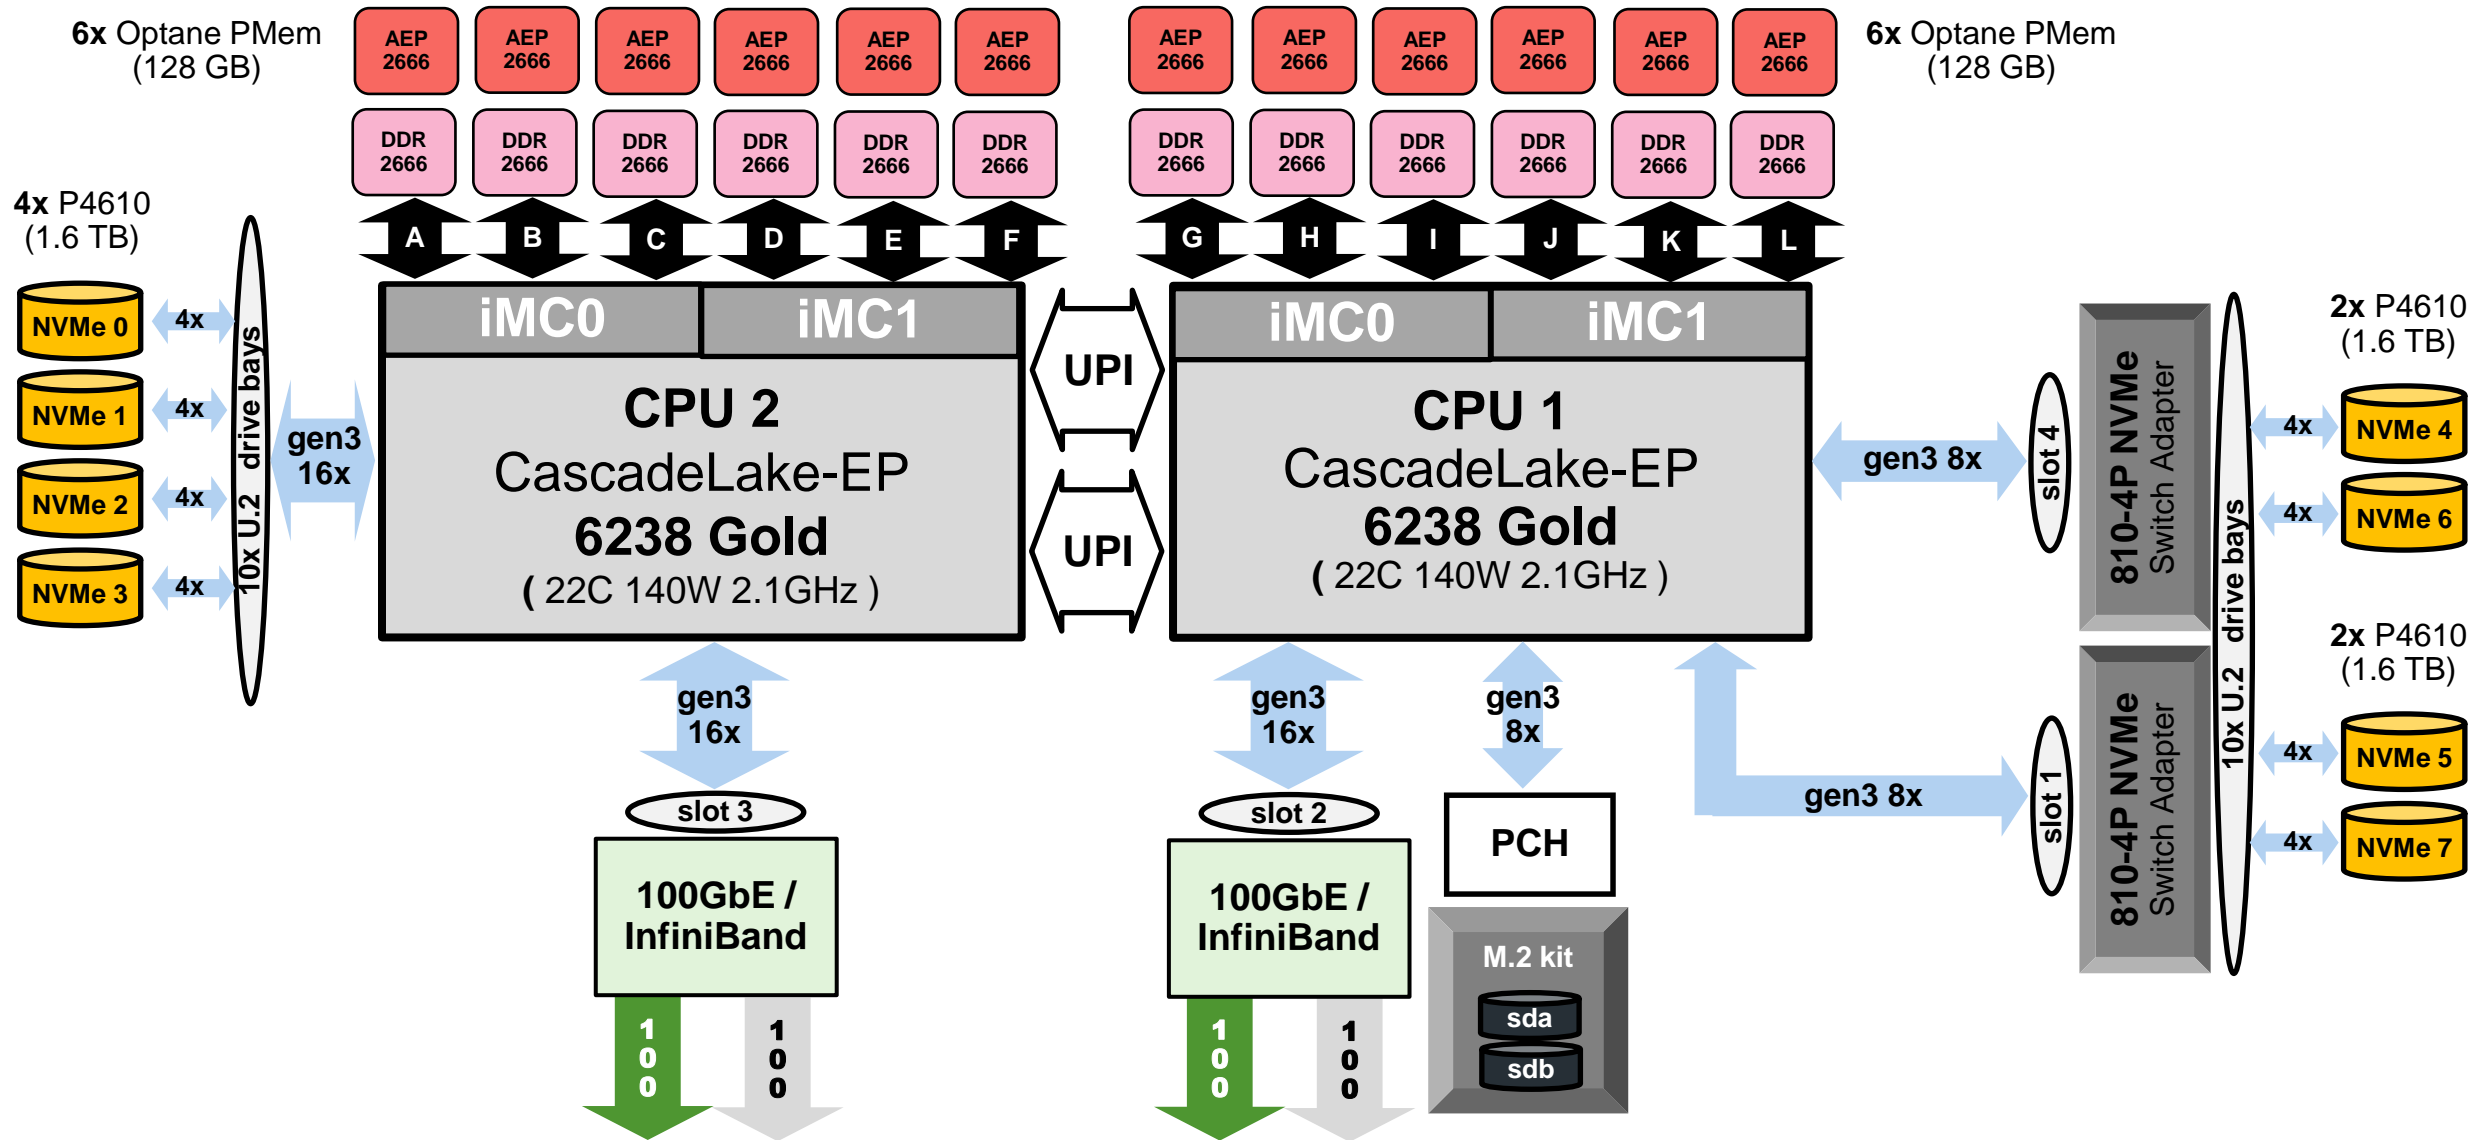
12x DAOS Servers in Lenox (ThinkSystem SR630, CLX)

Each with 8x NVMe 1.6 TB, 12x Optane PMem 128 GB, 2x EDR

- Total *raw* capacity: 150 TB NVMe + 18 TB PMem
- Peak *hardware* bandwidth: 250 GB/s read (180 GB/s raw write)

Running **DAOS 1.3.106**; daemons managed with `systemd`

DAOS CLX Server Architecture: Lenovo ThinkSystem SR630



IO500 SC21 Full List (Lenovo DAOS @ 12 SR630 is #7)


IO500 SC21 List										
<input type="button" value="IO500"/> <input type="button" value="10 Node"/> <input type="button" value="Full"/> <input type="button" value="Historical"/> <input type="button" value="Customize"/> <input type="button" value="Download"/>										
# ↑	INFORMATION									
	BOF	INSTITUTION	SYSTEM	STORAGE VENDOR	FILE SYSTEM TYPE	CLIENT NODES	TOTAL CLIENT PROC.	SCORE ↑	IO500 BW (GIB/S)	MD (KIOP/S)
1	ISC21	Pengcheng Laboratory	Pengcheng Cloudbrain-II on Atlas 900	Pengcheng	MadFS	512	36,864	36,850.40	3,421.62	396,872.82
2	SC21	Huawei HPDA Lab	Athena	Huawei	OceanFS	10	1,720	2,395.03	314.56	18,235.71
3	SC21	Olympus Lab	OceanStor Pacific	Huawei	OceanFS	10	1,720	2,298.69	317.07	16,664.88
4	SC21	Huawei Cloud		PDSL	Flashfs	15	1,560	2,016.70	109.82	37,034.00
5	ISC21	Intel	Endeavour	Intel	DAOS	10	1,440	1,859.56	398.77	8,671.65
6	ISC20	Intel	Wolf	Intel	DAOS	52	1,664	1,792.98	371.67	8,649.57
7	ISC21	Lenovo	Lenovo-Lenox	Lenovo	DAOS	36	3,456	988.99	176.37	5,545.61
8	SC21	BPFS Lab	Kongming		BPFS	10	800	972.60	96.26	9,827.09
9	SC19	WekaIO	WekaIO on AWS	WekaIO	WekaIO Matrix	345	8,625	938.95	174.74	5,045.33
10	ISC20	TACC	Frontera	Intel	DAOS	60	1,440	763.80	78.31	7,449.56

<https://io500.org/> (2021-11-18)



2021 Lenovo. All rights reserved.

DAOS Servers in Lenovo's HPC Innovation Center Stuttgart



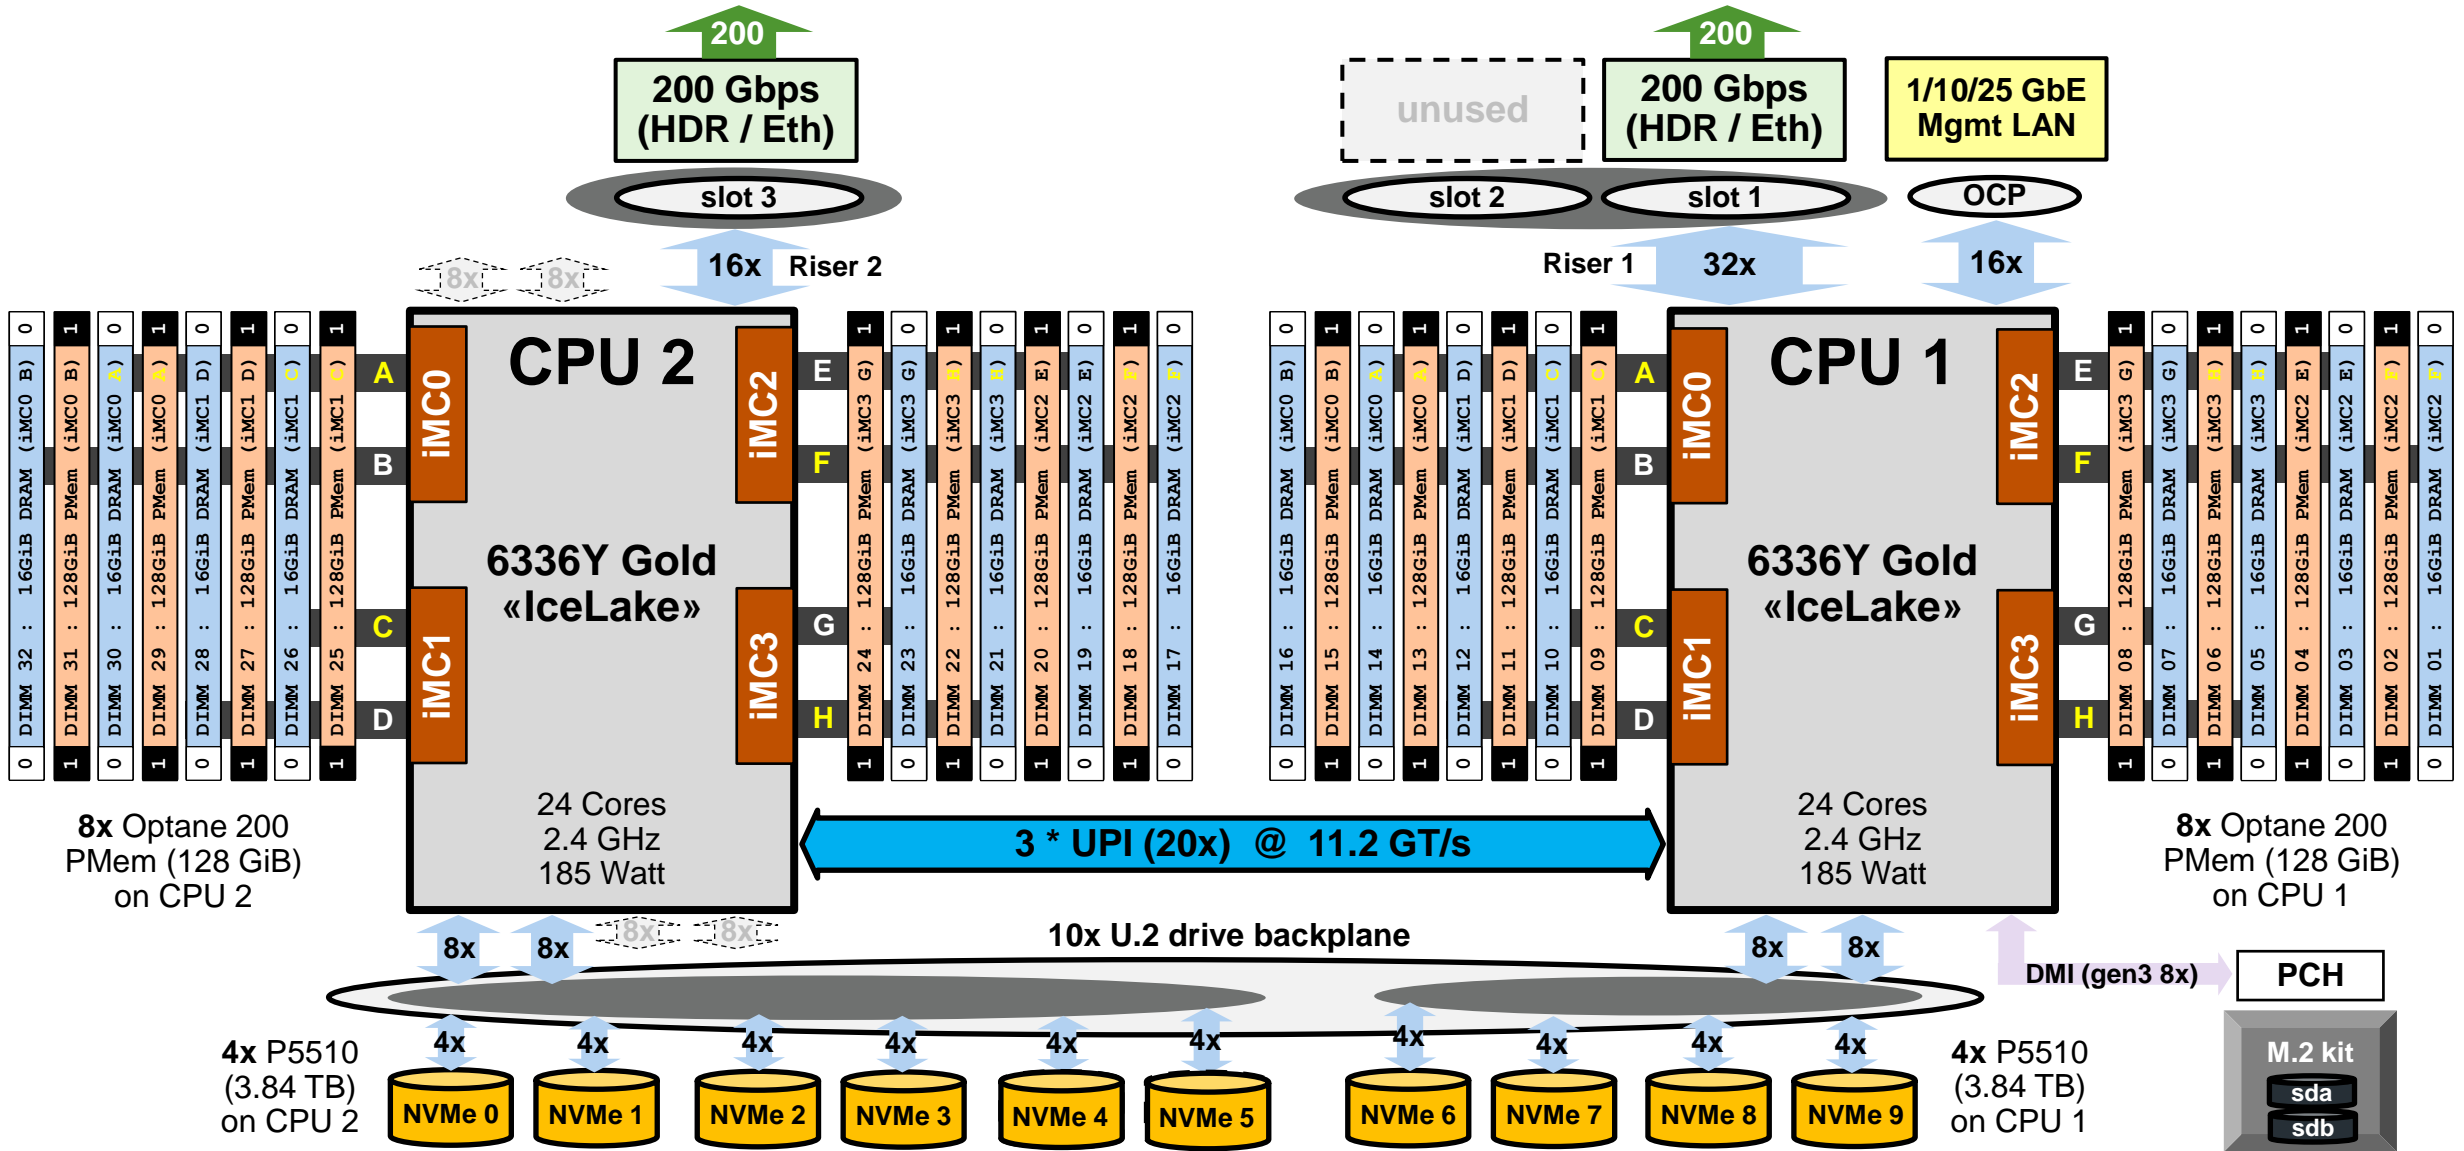
6x DAOS SR630V2 Servers in Lenox (ICX, with PCIe gen4)

Each with 10x NVMe 3.84 TB, 16x Optane 200 PMem 128 GB, 2x HDR

- Total *raw* capacity: 220 TB NVMe + 12 TB PMem
- Peak *hardware* bandwidth: 280 GB/s read (210 GB/s raw write)

Running **DAOS 1.3.106**; daemons managed with systemd

ThinkSystem SR630 V2 for DAOS: 8x U.2 with 16x PMem



Lenovo activities around DAOS

- **DAOS benchmarking facility as part of the HPC / AI Innovation Center**
 - Cascade Lake / Icelake based servers running DAOS
 - Integrated into LENOX compute infrastructure
 - Close collaboration with Intel since 2019
 - IO500 benchmark published June 2021
- **PoC's and projects**
 - Working with customers & prospects on DAOS use cases
 - Customer deployments



Proof of Concept & customers

- Proof of Concepts / Validation activities



Leibniz Supercomputing Centre
of the Bavarian Academy of Sciences and Humanities



- Customer activities



https://www.lrz.de/presse/ereignisse/2020-10-10_ExtremeScalingWorkshop_DE/

SuperMUC-NG already offers enormous computing power, but will now be upgraded for more diverse tasks with this expansion: Some of the new technology is currently being tested in the LRZ test environment BEAST (Bavarian, Energy, Architecture, Software Testbed) to better understand its capability in a future large-scale HPC system. To ensure that phase 2 of SuperMUC-NG continues to operate as energy-efficient as possible, the 240 Intel compute nodes are integrated into Lenovo's SD650-I v3 platform, which is directly cooled with warm water, and connected to the DAOS storage system via a high-speed network. Its capacity is 1 petabyte of data storage, but more importantly, this technology enables fast throughput of large data volumes. This system architecture is particularly well-suited then for highly scalable, compute- and data-intensive workloads and artificial intelligence applications.

https://www.lrz.de/presse/ereignisse/2021-05-04-SuperMUC-NG-Phase-2_ENG/

Lenovo





2021 Lenovo. All rights reserved.

Lenovo activities around DAOS

- **DAOS benchmarking facility as part of the HPC / AI Innovation Center**
 - Cascade Lake / Icelake based servers running DAOS
 - Integrated into LENOX compute infrastructure
 - Close collaboration with Intel since 2019
 - IO500 benchmark published June 2021
- **PoC's and projects**
 - Working with customers & prospects on DAOS use cases
 - Early customer deployments
- **DAOS – Next steps**
 - DAOS Solutions Design Guide
 - Installation automation
 - Monitoring
 - Support offering



DAOS @ Lenovo – Next steps



**DAOS on Lenovo SR630,
SR630v2**
(LenovoPress LP1398, LP1421)

<https://tinyurl.com/ub54t5jw>

- Demonstrate feasibility
- Gather experience in test/demo environment
- Verify performance improvements for individual use-case

- Reduce deployment time
- Ensure identical setups across installations
- Deploy a tested solution stack
- Reduce time to PoC

- Gathering data about inventory, utilization and performance of the DAOS systems
- Collect DAOS specific details
- System health data

- Lenovo to cover Level 1 and Level 2
- Customer contact
- Ticket handling
- Problem determination, log analysis, troubleshooting
- Integration and configuration management
- Escalation to Level 3 (Intel) if needed

thanks.

Lenovo