Cooling Environments

Deploying Immersion Native ITE in AI Neocloud: An OCP Collaboration Story

Immersion Project

Amy Short, PhD *Executive Officer – Advanced Data Center Technologies, Denvr Dataworks*

Austin Hipes *Chief Technologist, VP of Engineering, UNICOM Engineering*

OCP BRONZE Member

OCP SILVER Member

Speaker Introductions & Tech Deep Dive QR



Amy Short, PhD

Exec. Officer – Advanced Data Center Technologies
Denvr Dataworks



Austin Hipes

Chief Technologist & VP, Engineering
UNICOM Engineering

"After Hours" Technical Deep Dive

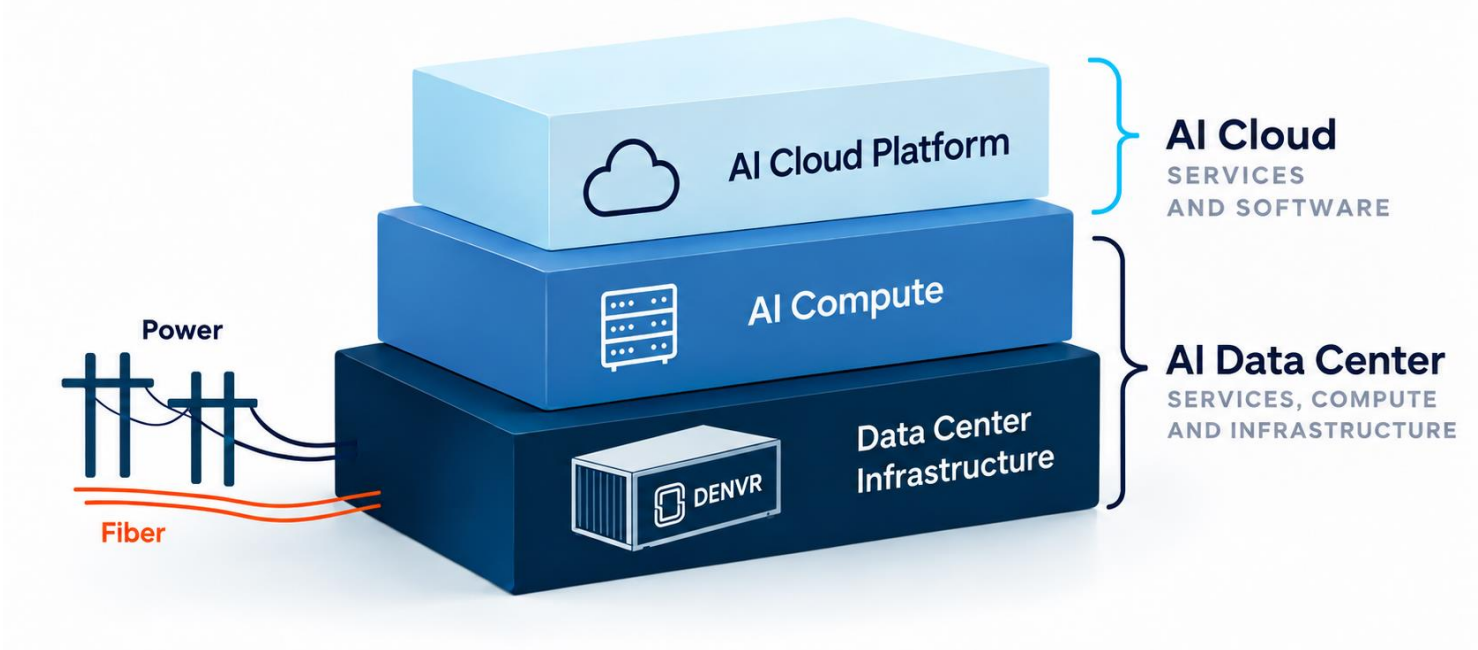
Download our Global Summit slides, whitepapers, and podcast episodes

OCP GLOBAL SUMMIT | 2026

Challenges in AI Deployments

Technical & Commercial Challenges

- Power & thermal constraints
- Water & sustainability needs
- Deployment speed
- Warranty and long-term reliability



Why We're Here Today

- Few production examples for high density GPU platforms
- To prove immersion is scalable & performant – we have not yet seen a ceiling
- Solves a real COMMERCIAL challenge – it is not a lab experiment
- Clear demand for differentiated AI infrastructure

Denvr Approach to Infrastructure at Scale

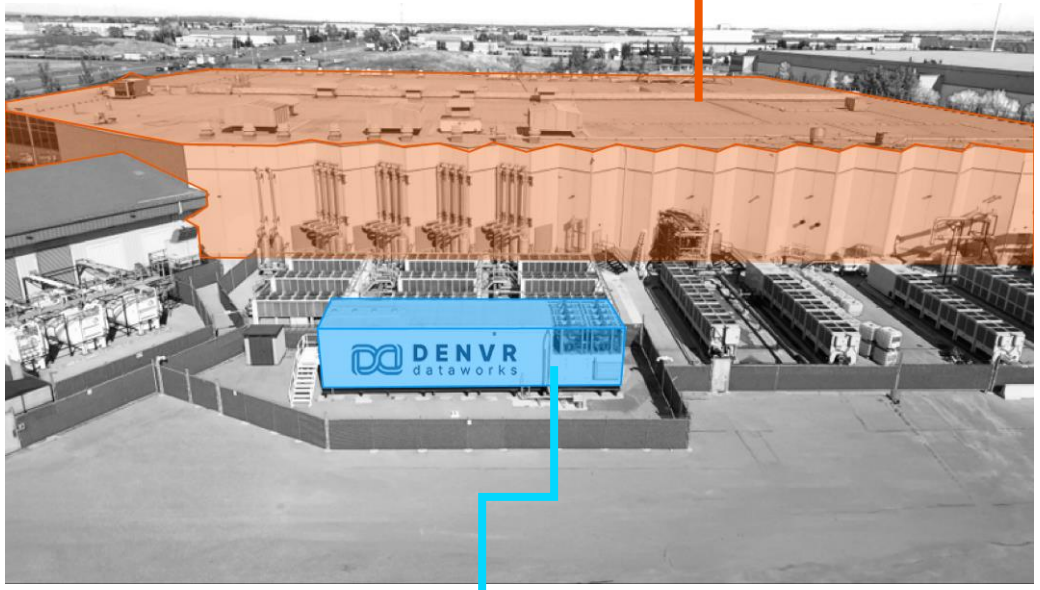
- **Modular** Standardized and repeatable
- **High Density** 128 GPUs in 900 sq ft
- **Sustainable** Zero-water, PUE <1.04
- **Built for rapid scale** 6-12 months to deploy



3OU H200 SXM servers deployed 14 per rack at Denvr

Air Cooled Co-Lo

1 MW per 6400 sq ft



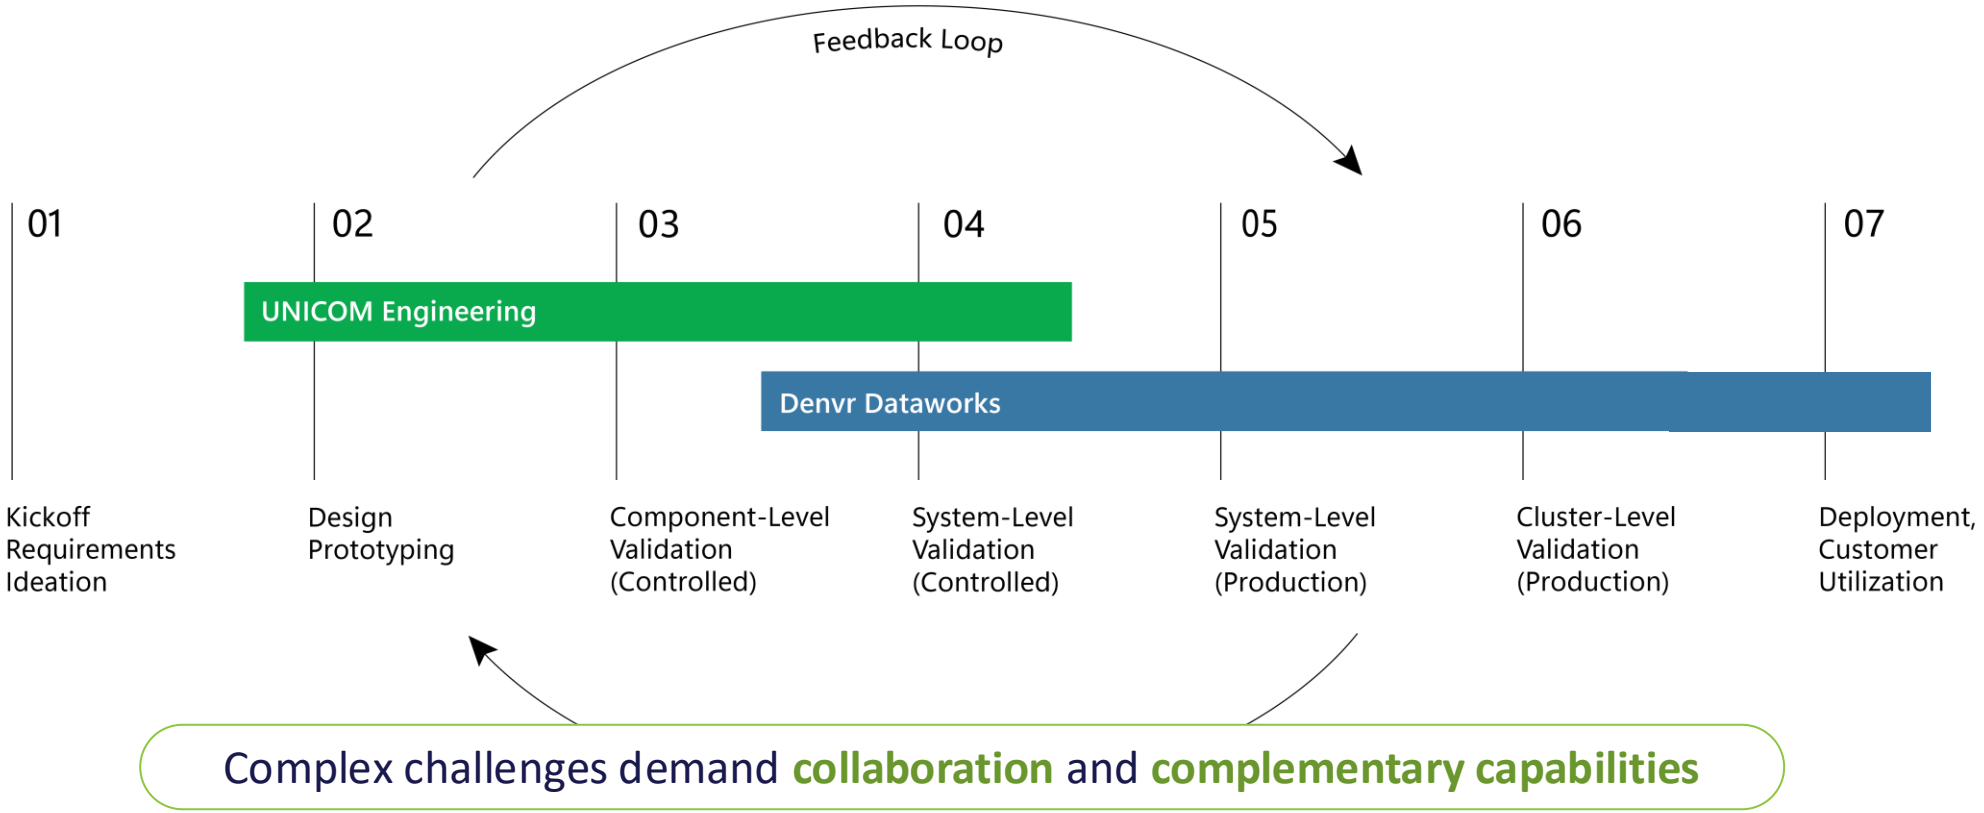
Denvr Dataworks Modular Data Center

1 MW per 600 sq ft

Collaborative AI Server Development Program

UNICOM Engineering Component-level and server-level validation

Denvr Dataworks Server-level and cluster-level validation; commissioning and deployment



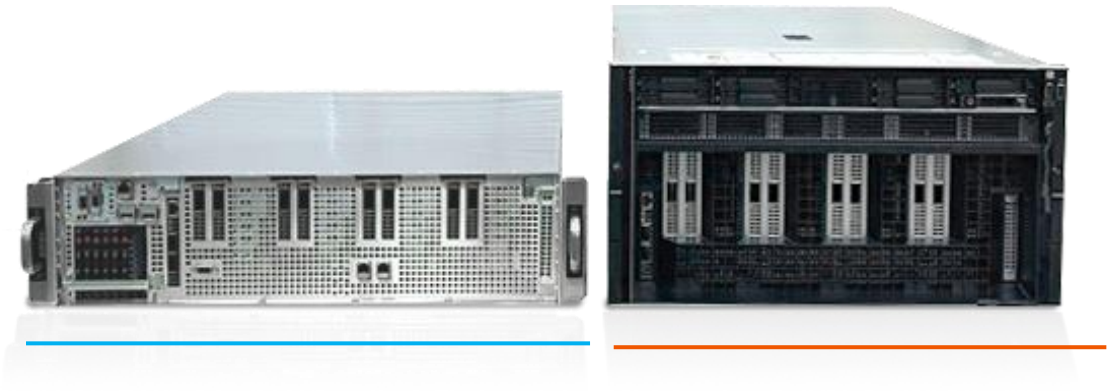
Complex challenges demand **collaboration** and **complementary capabilities**

Immersion-Optimized SXM Server Design

2x server size reduction

4x servers per rack (14x H200s)

~800 W per server power savings



Denvr Immersion Optimized Server Design

H200 SXM server

Air Cooled XE9680 Server Design

H200 SXM server

An OCP-Aligned foundation establishes a standardized approach to address power, thermal, density, and deployment challenges

Rethinking Server Design for AI Use Cases

Warranty Coverage Based on 2.5 Years of Reliability Testing

New Physical Design

- Reduced physical footprint by 2x (6U → 3OU form factor)
- Based on the Dell XE9680 Platform for H100 and H200 GPUs

Power Architecture Updates

- 102 kW power savings per 128 server scalable unit
- Inspired by OCP Open Rack designs
- Adaptable for ORV2.2 and ORV3

Thermal Redesign

- Immersion-ready copper heat sinks
- Evaluated in system temperatures up to 60 °C

Software & Firmware Optimization

- Updated firmware
- Co-development with Dell



Fourteen H200 SXM servers with 3OU chassis supported in a single Denvr production environment tank.

Thermal Performance

Power, Frequency, and Thermals are **Predictable and Stable**

Thermals, not power, limit density in air-cooled racks

- Immersion increases usable compute per square foot

Stable thermals improve reliability and warranty confidence

- Reduced CPU/GPU variability
- Validated via NVQual and production benchmarking

NVQual Thermal Test Result

Test Application	TDP
Test Result	**Pass**

Test Result (Immersion)

Parameter	GPU<5d.00.0>	GPU<3b.00.0>	GPU<db.00.0>	GPU<19.00.0>	GPU<4c.00.0>	GPI<bb.00.0>	GPU<cb.00.0>	GPU<9b.00.0>
GPU Tj avg: max. measured (max. spec)	**65.0 (85)**	**70.2 (85)**	**71.2 (85)**	**70.4 (85)**	**75.0 (85)**	**70.2 (85)**	**64.5 (85)**	**70.5 (85)**
Memory Tj: max. measured (max. spec)	**70.0 (95)**	**74.0 (95)**	**72.0 (95)**	**69.0 (95)**	**75.0 (95)**	**76.0 (95)**	**67.0 (95)**	**67.0 (95)**
Power: avg measured in Watts (Limit)	**685.038 (700)**	**685.612 (700)**	**685.598 (700)**	**684.909 (700)**	**685.326 (700)**	**685.608 (700)**	**685.987 (700)**	**684.132 (700)**
Tlimit: min. measured (min. spec)	**19.19 (0)**	**6.69 (0)**	**4.00 (0)**	**7.69 (0)**	**4.92 (0)**	**13.16 (0)**	**22.09 (0)**	**8.72 (0)**

Full data set available on Denvr-UNICOM Engineering "After Hours" landing page linked at end of presentation

Signal Integrity

Signal Integrity must be designed and validated for Immersion

- Not assumed from air-cooled architectures

Data-driven Validation

- NVQual testing confirms signal integrity
- Benchmarking shows stable, sustained performance

NVQual PCIe Qualification Test Result

Test Result PASS

GPU	7] [db.00.0
Error Code	0
Loop Counts	3

Depth 0

Link Speed	32 Gbps (PCIe Gen 5.0)
Link Width	16
Total Elapsed Time (min)	40.755567

Loop 1

BER Local Correctable (allowed)	0.000000 (32.000000)
BER Local CRCErrors (allowed)	0.000000 (32.000000)
BER Local FailedL0sExits (allowed)	0.000000 (32.000000)
BER Local Fatal (allowed)	0.000000 (0.000000)
BER Local LineErrors (allowed)	0.000000 (32.000000)
BER Local NAKs Received (allowed)	0.000000 (32.000000)
BER Local NAKs Sent (allowed)	0.000000 (32.000000)
BER Local Non-Fatal (allowed)	0.000000 (0.000000)

Loop 2

BER Local Correctable (allowed)	0.000000 (32.000000)
BER Local CRCErrors (allowed)	0.000000 (32.000000)
BER Local FailedL0sExits (allowed)	0.000000 (32.000000)
BER Local Fatal (allowed)	0.000000 (0.000000)
BER Local LineErrors (allowed)	0.000000 (32.000000)
BER Local NAKs Received (allowed)	0.000000 (32.000000)
BER Local NAKs Sent (allowed)	0.000000 (32.000000)
BER Local Non-Fatal (allowed)	0.000000 (0.000000)

Loop 3

BER Local Correctable (allowed)	0.000000 (32.000000)
BER Local CRCErrors (allowed)	0.000000 (32.000000)
BER Local FailedL0sExits (allowed)	0.000000 (32.000000)
BER Local Fatal (allowed)	0.000000 (0.000000)
BER Local LineErrors (allowed)	0.000000 (32.000000)
BER Local NAKs Received (allowed)	0.000000 (32.000000)
BER Local NAKs Sent (allowed)	0.000000 (32.000000)
BER Local Non-Fatal (allowed)	0.000000 (0.000000)

Signal Integrity **maintained in immersion**

	Non-zero entries within threshold	Entries exceeding threshold
	total entries [result per entry]	*total entries [result per entry]*
Host Correctable (3 entries)	0	0
Host Fatal (3 entries)	0	0
Host Non-Fatal (3 entries)	0	0

Depth 1

	Non-zero entries within threshold	Entries exceeding threshold
	total entries [result per entry]	*total entries [result per entry]*
Host Correctable (24051 entries)	0	0
Host Fatal (24051 entries)	0	0
Host Non-Fatal (24051 entries)	0	0
Local Correctable (12027 entries)	0	0
Local Fatal (12027 entries)	0	0
Local Non-Fatal (12027 entries)	0	0

Depth

	Non-zero entries within threshold	Entries exceeding threshold
	total entries [result per entry]	*total entries [result per entry]*
Host Correctable (12024 entries)	0	0
Host Fatal (12024 entries)	0	0
Host Non-Fatal (12024 entries)	0	0
Local Correctable (6012 entries)	0	0
Local Fatal (6012 entries)	0	0
Local Non-Fatal (6012 entries)	0	0

NVQual NVLink Bandwidth Test Results

Source GPU	Source Link	Remote GPU	Remote Link	Input BW (%) Min	Input BW (%) Max	Output BW (%) Min	Output BW (%) Max
0	0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17	2	0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17	99.4	99.7	99.3	99.6
5	0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17	6	0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17	99.3	99.5	99.3	99.5
1	0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17	7	0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17	99.4	99.4	99.3	99.4
3	0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17	4	0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17	99.4	99.5	99.3	99.5

Test Loop #1		Test Loop #2		Test Loop #3	
Input	Output	Input	Output	Input	Output
99.4	99.3	99.6	99.6	99.7	99.6
99.3	99.3	99.5	99.5	99.5	99.4
99.4	99.4	99.4	99.3	99.4	99.4
99.4	99.3	99.4	99.3	99.5	99.5

Full data set available on Denvr-UNICOM Engineering "After Hours" landing page linked at end of presentation

Demonstrated In Real AI Workloads

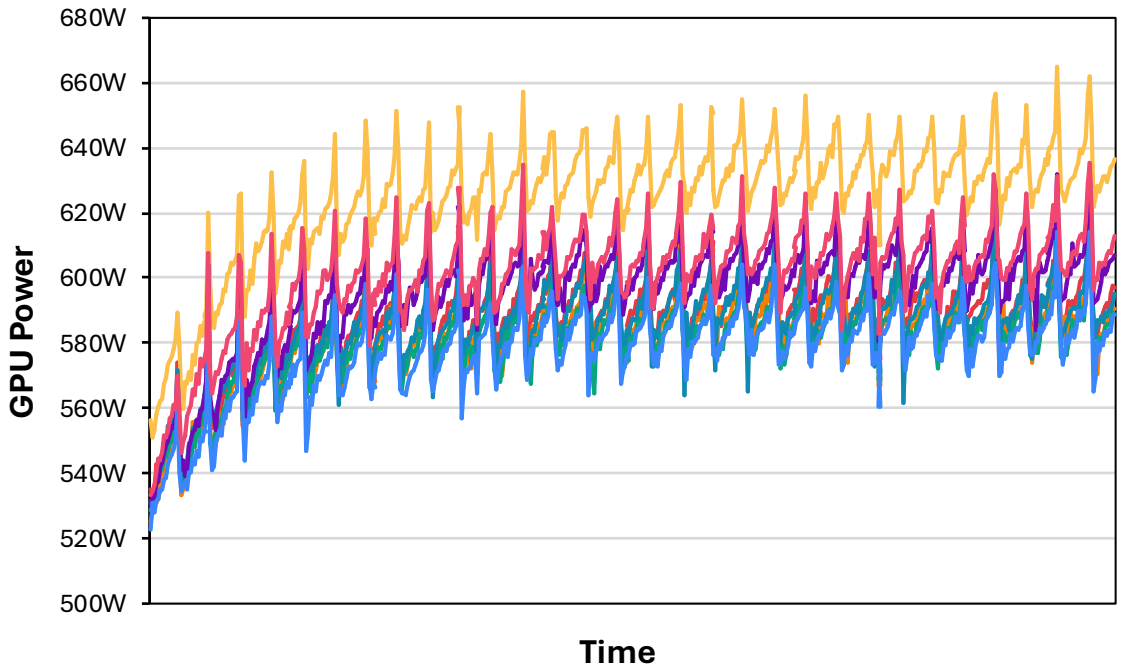
Hardware: 8x GPU H200-SXM5-141GB Denvr MSC1

GPU TDP: 700 W

Test Type: Continuously batched LLM inferencing of Llama-3.1-Instruct 70B

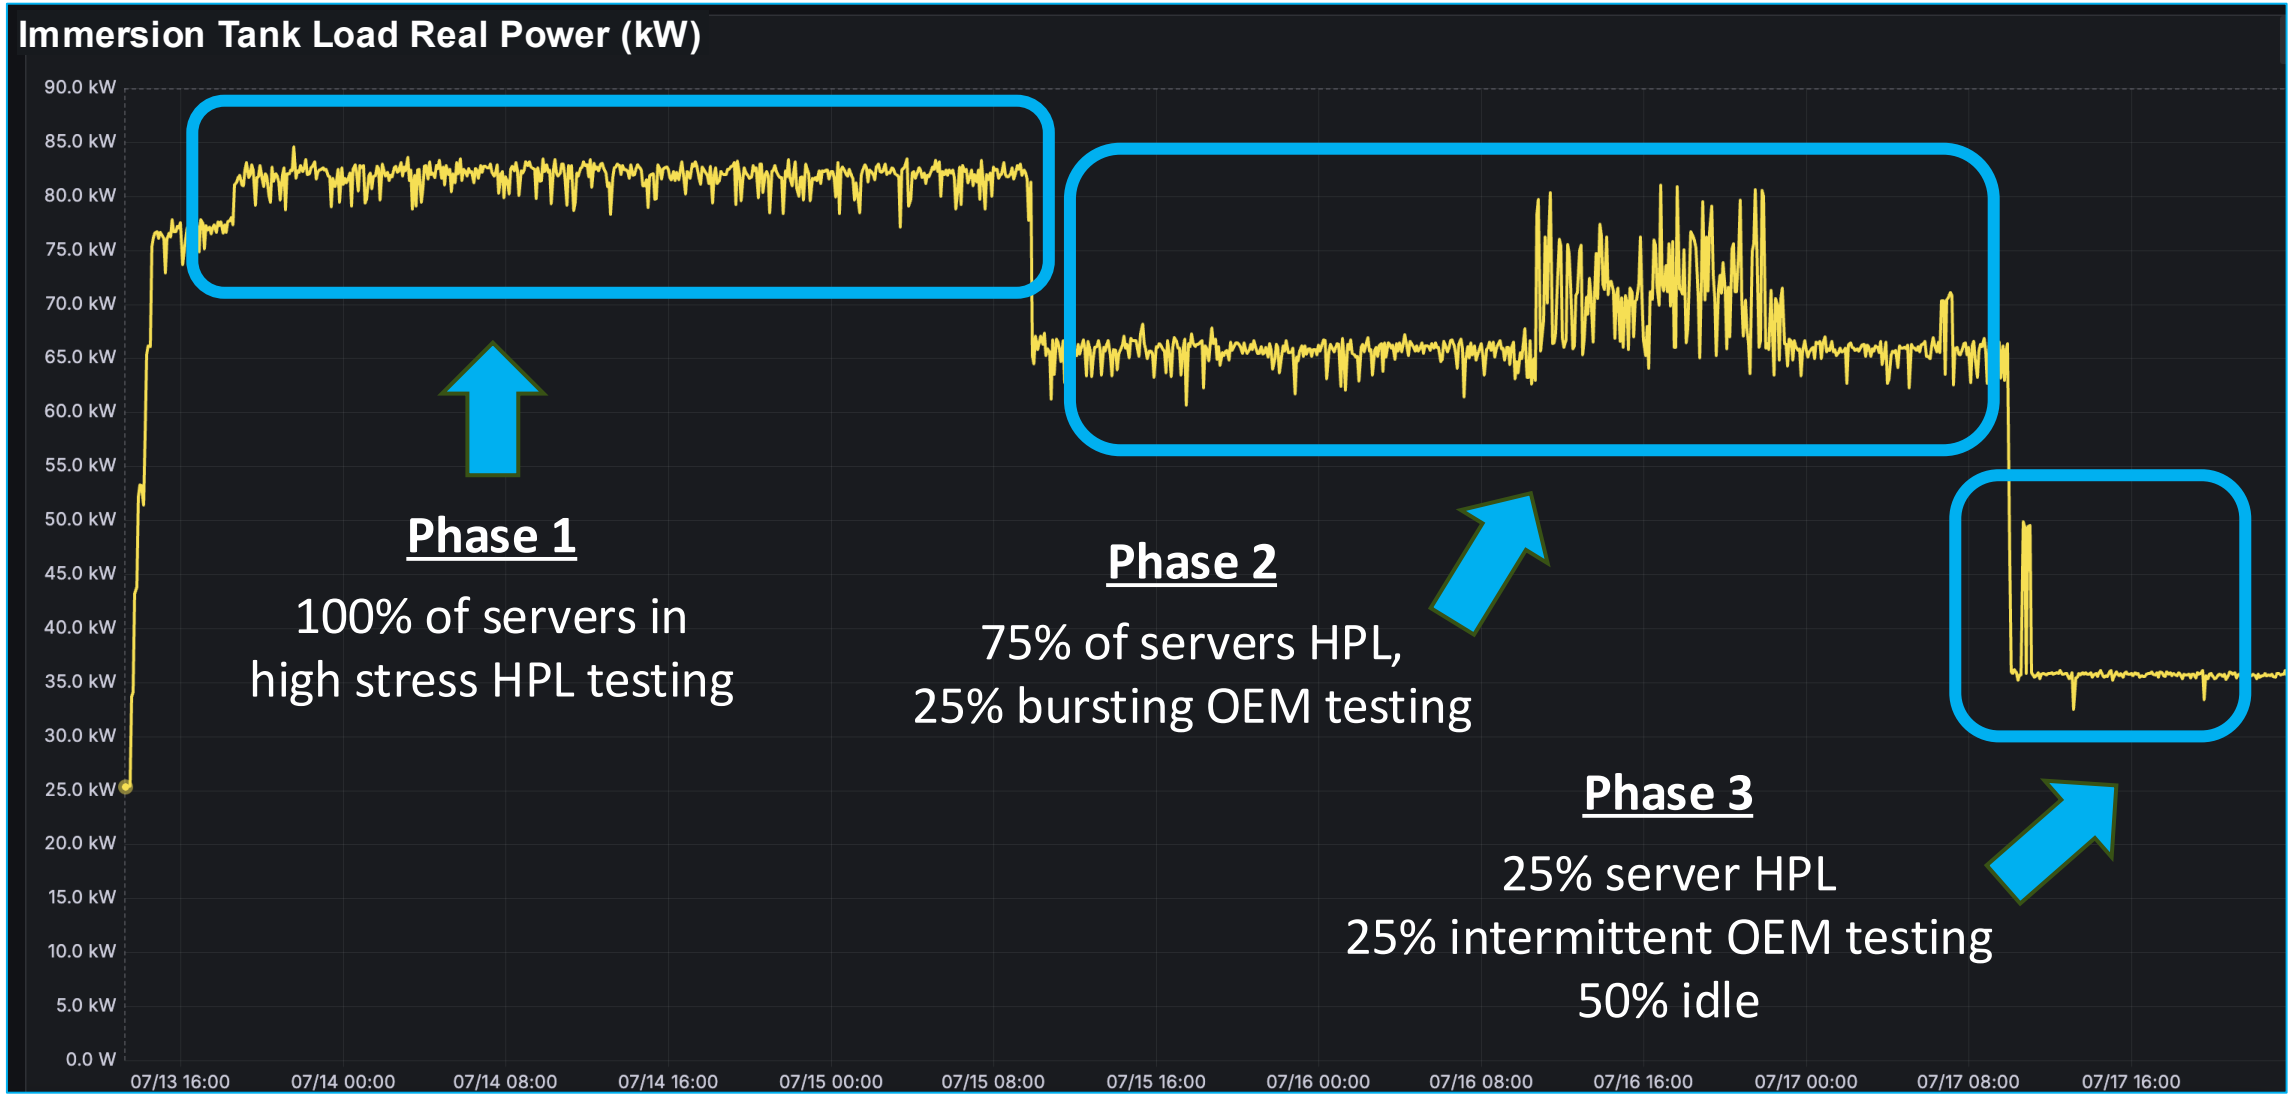
Details:

- Tensor Parallelism across all 8X H200 accelerators
- Context length per batched sequence: 4096 toks (3072tok/prompt, 1024tok/response)
- Maximum of 150 concurrent requests
- OTPS (output toks per second): 8068
- RPS (responses per second): 7.9
- P99 per-response latency (1024 toks): 19.82s



Test data by Denvr; Full data set available on Denvr-UNICOM Engineering "After Hours" landing page linked at end of presentation

Denvr Infra Response to Multi-Day AI Stress Test



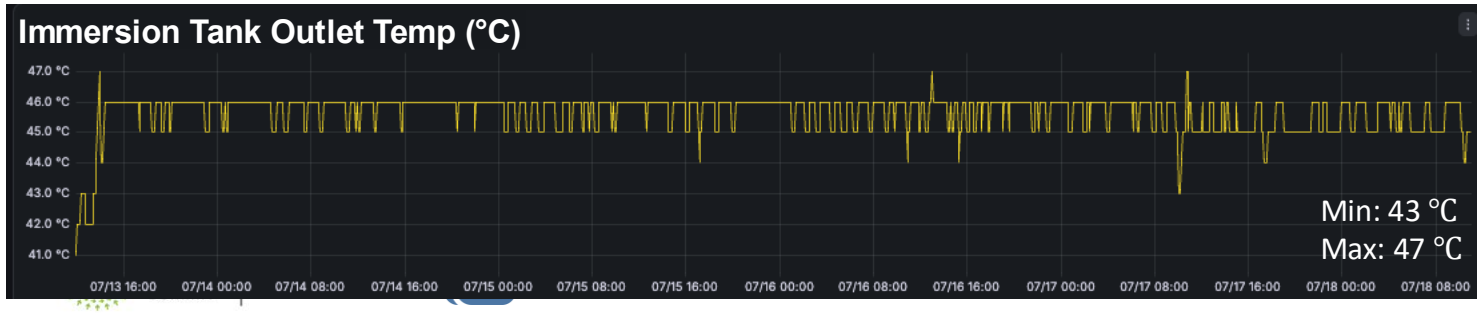
Representative example from Denvr H200 SXM cluster commissioning

Denvr Infra Response to Immersion H200 SXM Stress Test: Reliability, Responsiveness & Scalability









Power Draw

- Measured IT Load responds as expected.
- ~800 W per server power savings vs. air cooled parent configuration.

Flow Rate

- Increases as expected in response to IT load.
- Stabilizes quickly in each test phase.

Mild Temperatures, Longer Lifetime Servers

- Peak system temp increases only 6°C at and quickly stabilizes at ~45 °C.
- Rack inlet/outlet ΔT ~10°C maintained as targeted.
- FMEA testing demonstrated 15 min-1 hr response time possible due to mild env.

Representative example from Denvr H200 SXM cluster commissioning

DENVR

Design Updates For Blackwell Generation GPUs & Beyond

Immersion server architecture validated for sustained 1,000W+ class GPU operation, PCIe/NVLink performance, and scalable thermal management requires additional enhancements.

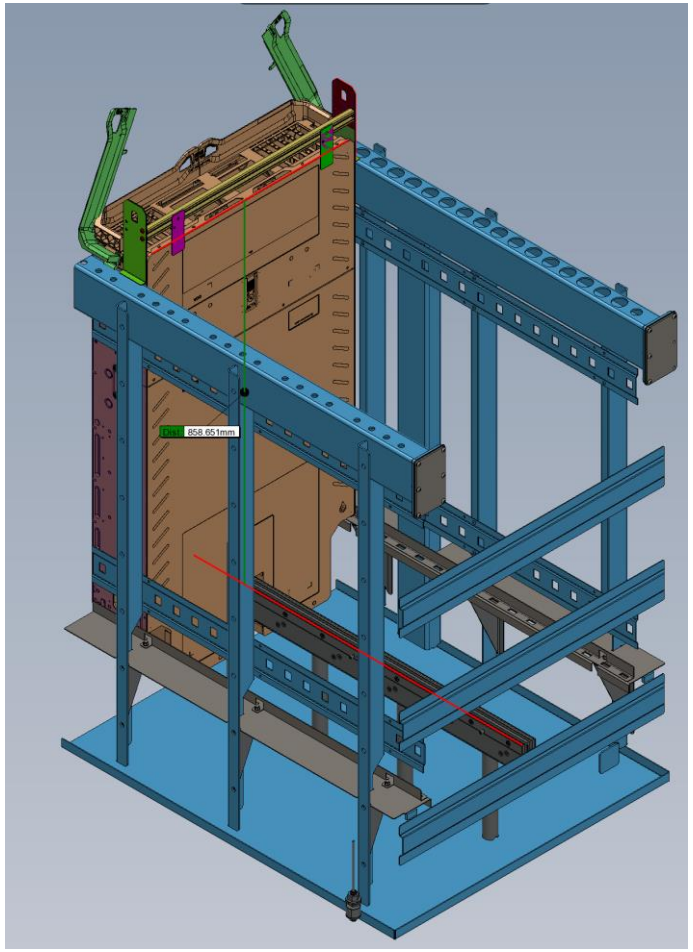
1,000 W+
Target power per GPU

Targeted Flow

3000W+ Per OU

Design for the next thermal envelope

Immersion-optimized heatsinks, direct flow delivery, redundancy, and server-managed controls create a path beyond today's Blackwell validation point.



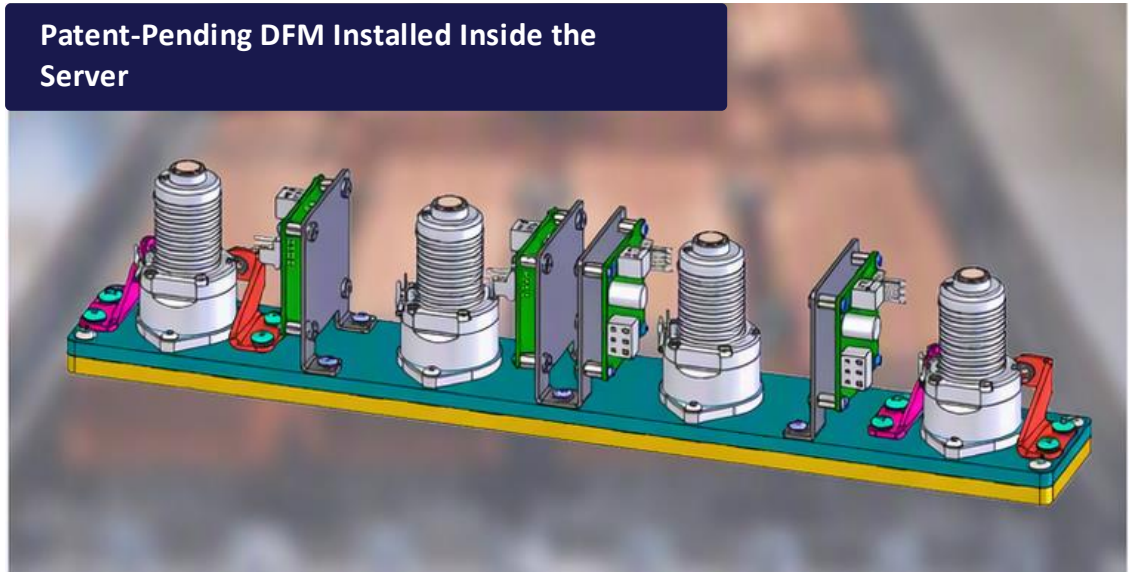
Physical Modifications Enhance Immersion Cooling for 1000W+ GPUs

Server-level thermal hardware, redundant flow delivery, and closed-loop management work together as one cooling architecture

Immersion-Optimized Heatsinks Installed



Patent-Pending DFM Installed Inside the Server



Cold plates replaced
Immersion-optimized copper heatsinks increase wetted surface area for direct heat transfer.

Direct Flow Manifold
Integrated, low-profile DFM hardware delivers coolant to critical components.

N+1 redundancy
Redundant flow paths maintain enhanced cooling capability.

Server-managed control
System management adjusts flow to individual server thermal needs.

Energy-aware operation
Flow and pumping power are optimized together. Cooling power consumption ranges between 12.5W – 75W total.

Immersion optimized design: optimized heatsinks transfer heat to the fluid, while patent-pending DFM hardware and server system management deliver redundant, demand-based flow with lower unnecessary pumping power.

Blackwell SXM Sustains Full Power Under Load

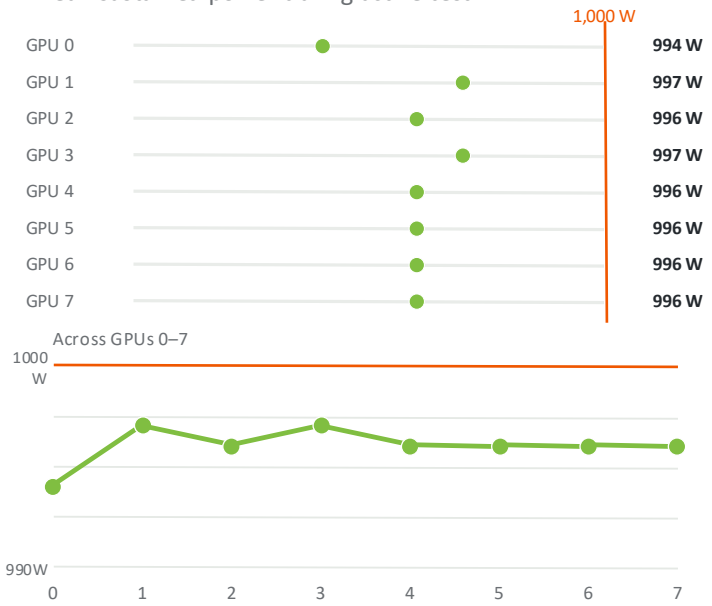
8 SXM GPUs • 1,000 W target per GPU • 600-second DCGM targeted-power test • 35°C Dielectric Fluid Inlet

996 W	0.4%	77°C	83°C	8 / 8
Mean GPU power	Below 1,000 W target	Peak GPU core	Core thermal limit	DCGM pass

Power and clock remain tightly grouped while core temperatures stay below the stated limit.

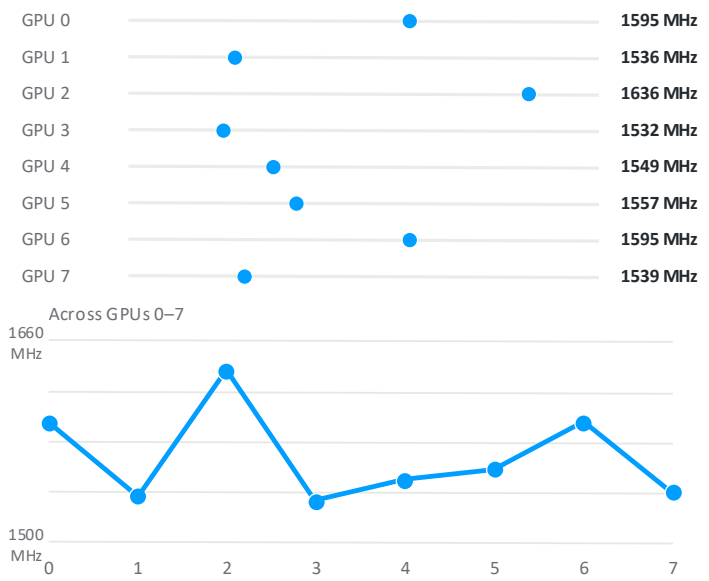
Power

Mean sustained power during active test



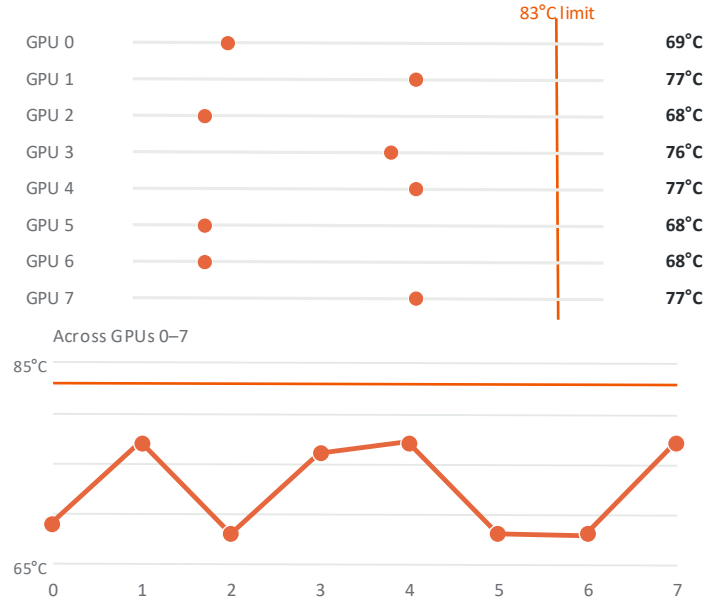
Clock

Average GPU clock during active test



GPU Core Temperature

Maximum measured Tcore



Takeaway: The 8-GPU system sustained 99.6% of target power while peak core temperature remained 6°C below the stated limit.

Delivers Comparable Full Power Performance in Immersion & DLC

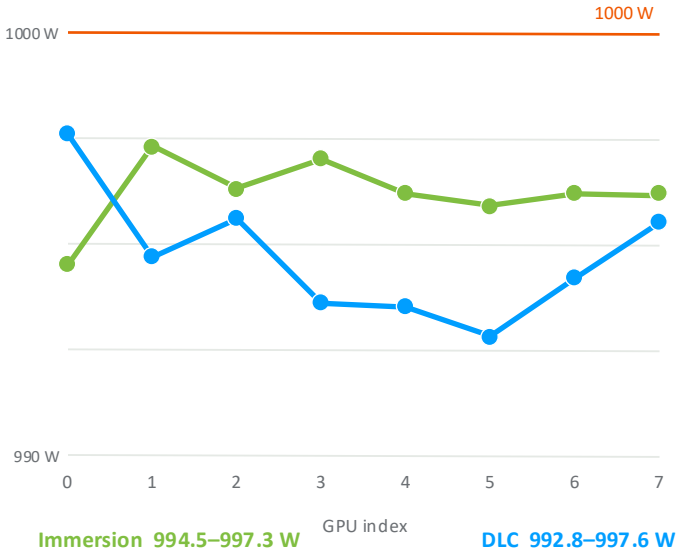
Same platform • 8 × SXM GPUs • 1,000 W target per GPU • 600-second DCGM targeted-power tests

Immersion, 35°C DLC, 30°C PG25

	Immersion, 35°C	DLC, 30°C PG25
Mean power	996 W	995 W
Mean clock	1,567 MHz	1,594 MHz
Peak Tcore	77°C	64°C
DCGM result	8 / 8 PASS	8 / 8 PASS

Power

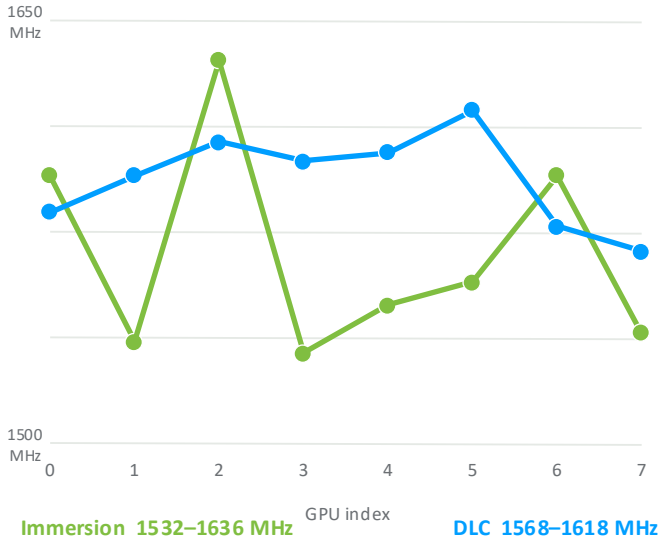
Mean sustained power by GPU



Immersion 994.5–997.3 W DLC 992.8–997.6 W

Clock

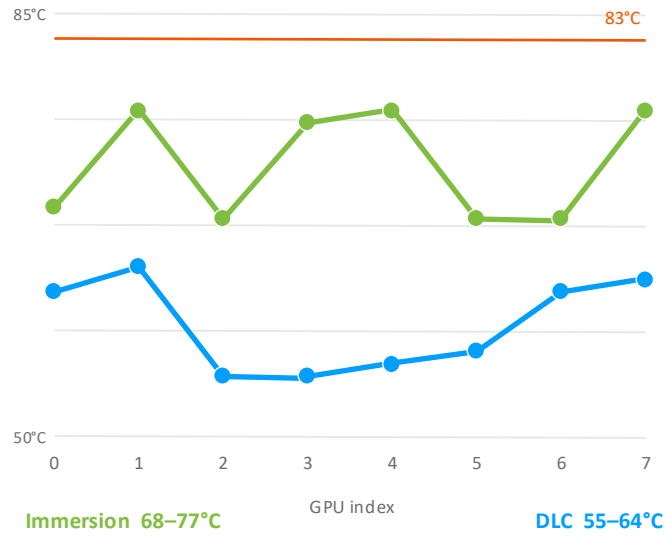
Average GPU clock during load



Immersion 1532–1636 MHz DLC 1568–1618 MHz

GPU Core Temperature

Maximum measured Tcore

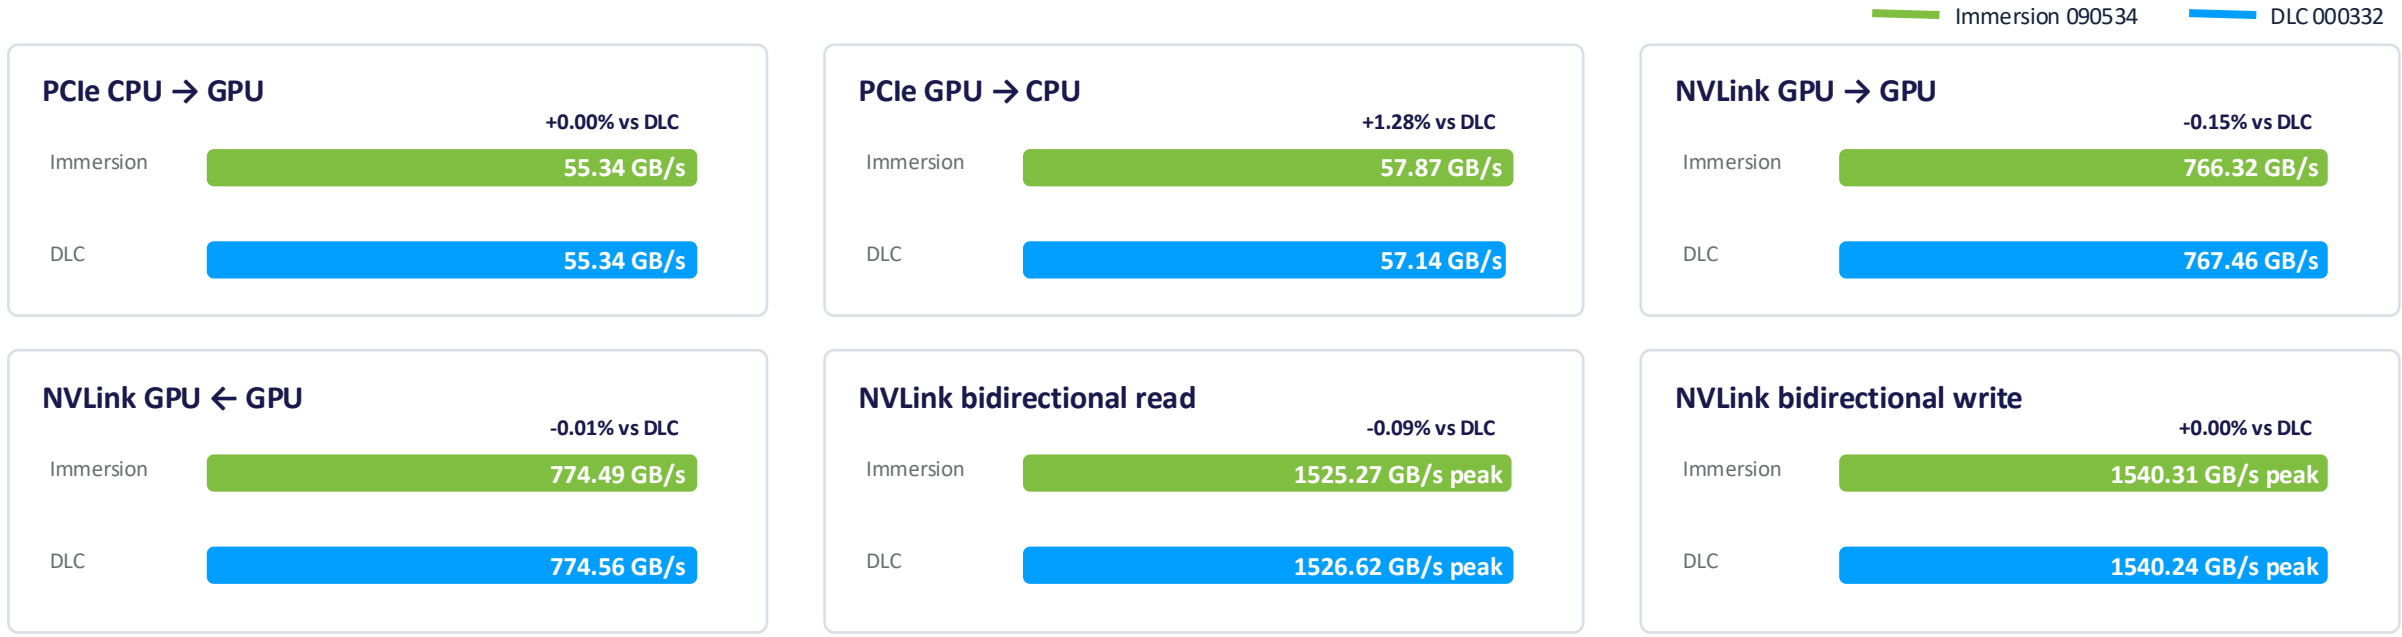


Immersion 68–77°C DLC 55–64°C

Takeaway: Both cooling approaches passed all 8 GPUs and sustained at least 99.5% of the 1,000 W target on average. Immersion matched DLC power delivery while remaining below the 83°C Tcore limit.

Immersion Matches DLC PCIe & NVLink Bandwidth

Like-for-like DCGMI NVBandwidth results • CPU↔GPU represents PCIe • GPU↔GPU represents NVLink



Result: All six bandwidth tests passed in both cooling configurations. Immersion remained within 0.15% of DLC on directional NVLink and within 1.27% on directional PCIe.

PCIe replay events	0
EDATL error events	0
Immersion samples	28,360

Takeaway: Immersion bandwidth matched the DLC baseline, while separate monitoring observed zero PCIe replay-counter events and zero EDATL error events.

Call to Action: Enable Innovation Through Collab!

- **This work was enabled by OCP:**
  - Power distribution & open rack designs
  - Immersion requirements
  - Material compatibility
  - Signal Integrity
  - Safe handling and cleaning
  - FMEA and operational guidance
- **Engage with OCP technical committees to help shape what comes next!**

UNICOM® Engineering, Inc
A Division of UNICOM Global

DENVR

After Hours discussion – Ask Me Anything!!

Register and submit your questions in advance to help shape the conversation.

Slides and Technical Deep Dive

For deeper insights on the topic, scan the QR code here!

Thank You!

Access these slides and other technical deep dive content.

Schedule a meeting with our team.

After-hours discussion topics – Ask Me Anything!

