DeepLM Cluster Baseline Report

Session	20260819_0616
Date	2026-08-19 06:16:04
Duration	265m 43s
Coordinator	node04
Nodes tested	node05, node07, node08

ACTION REQUIRED

All 3 nodes have hardware issues that must be resolved before this cluster is ready for production workloads.

Node Health at a Glance

Node	GPU Health	Host Connection	Ethernet	High-Speed Netwo	Multi-GPU	Overall
node05	FAIL	FAIL	PASS	PASS	FAIL	Needs Attention
node07	FAIL	PASS	PASS	PASS	FAIL	Needs Attention
node08	FAIL	PASS	PASS	PASS	FAIL	Needs Attention

What Was Tested

Test Area	What It Checks
GPU Health Check	Checks that each GPU can run at full speed without overheating.
Memory & Host Connection	Measures how fast the GPU moves data to and from the server.
Ethernet Network	Confirms all nodes communicate over standard Ethernet at rated speed.
High-Speed Network Fabric	Validates the high-speed network used for GPU-to-GPU data transfer during training.
Multi-GPU Performance	Measures how efficiently GPUs across all nodes can coordinate — directly impacts distributed training speed.

Action Items

1. node05 — GPU 0 (NVIDIA RTX A2000 12GB) overheating: GPU temperature was already 83°C at idle — above the 83°C safety limit. The GPU throttled its own speed to protect itself. Inspect chassis fans, rack airflow, and thermal paste before running workloads.

2. node05 — GPU 0 (NVIDIA RTX A2000 12GB) compute performance is below specification. This is most likely caused by the overheating issue above — resolve cooling first and re-test.

3. node07 — GPU 0 (NVIDIA RTX A4000) overheating: GPU temperature was already 93°C at idle — above the 83°C safety limit. The GPU throttled its own speed to protect itself. Inspect chassis fans, rack airflow, and thermal paste before running workloads.

4. node07 — GPU 0 (NVIDIA RTX A4000) compute performance is below specification. This is most likely caused by the overheating issue above — resolve cooling first and re-test.

5. node08 — GPU 0 (NVIDIA RTX A4000) overheating: GPU temperature was already 93°C at idle — above the 83°C safety limit. The GPU throttled its own speed to protect itself. Inspect chassis fans, rack airflow, and thermal paste before running workloads.

6. node08 — GPU 0 (NVIDIA RTX A4000) compute performance is below specification. This is most likely caused by the overheating issue above — resolve cooling first and re-test.

7. node08 — GPU 1 (NVIDIA RTX A4000) overheating: GPU temperature was already 93°C at idle — above the 83°C safety limit. The GPU throttled its own speed to protect itself. Inspect chassis fans, rack airflow, and thermal paste before running workloads.

8. node05 — GPU-to-host connection running at 6.7 GB/s (21% of the 32 GB/s rated speed). The GPU is not negotiating the correct PCIe generation. Reseat the GPU in its slot and check PCIe settings in the server BIOS.

9. Multi-node GPU synchronization is below threshold for: node05 ↔ node07, node05 ↔ node08. This will reduce distributed training throughput. Resolve any overheating or connection issues on the affected nodes first.

Hardware Inventory

Node	GPUs	GPU Model	VRAM	PCIe Gen	IB Adapter
node05	1	NVIDIA RTX A2000 12GB	12 GB	DEFAULT	—
node07	1	NVIDIA RTX A4000	16 GB	GEN4	—
node08	2	NVIDIA RTX A4000	16 GB	GEN4	—

Technical Detail — for engineering reference

GPU Health Check — FAIL

Node: node05 | 2026-08-19 06:22:36

GPU Test Results

Test	GPU 0 NVIDIA RTX A2000 12GB
Compute Performance	FAIL
Compute Utilization	PASS
Memory Speed	PASS
Memory Capacity	PASS
GPU Cooling	FAIL
Throttle Detection	PASS
Overall	FAIL

GPU Cooling Detail

GPU	Model	Start °C	Stable °C	Limit	Clock Drop	Result	Diagnosis
GPU 0	NVIDIA RTX A2000 12GB	83°C	84.0°C	≤ 83°C	12.4%	FAIL	GPU temperature was already 83°C at idle — above the 83°C safety limit. The GPU throttled its own speed to protect itself. Inspect chassis fans, rack airflow, and thermal paste before running workloads.

GPU Health Check — FAIL

Node: node07 | 2026-08-19 06:21:53

GPU Test Results

Test	GPU 0 NVIDIA RTX A4000
Compute Performance	FAIL
Compute Utilization	PASS
Memory Speed	PASS
Memory Capacity	PASS
GPU Cooling	FAIL
Throttle Detection	PASS
Overall	FAIL

GPU Cooling Detail

GPU	Model	Start °C	Stable °C	Limit	Clock Drop	Result	Diagnosis
GPU 0	NVIDIA RTX A4000	93°C	90.5°C	≤ 83°C	3.1%	FAIL	GPU temperature was already 93°C at idle — above the 83°C safety limit. The GPU throttled its own speed to protect itself. Inspect chassis fans, rack airflow, and thermal paste before running workloads.

GPU Health Check — FAIL

Node: node08 | 2026-08-19 06:26:58

GPU Test Results

Test	GPU 0 NVIDIA RTX A4000	GPU 1 NVIDIA RTX A4000
Compute Performance	FAIL	PASS
Compute Utilization	PASS	PASS
Memory Speed	PASS	PASS
Memory Capacity	PASS	PASS
GPU Cooling	FAIL	FAIL
Throttle Detection	PASS	PASS
Overall	FAIL	FAIL

GPU Cooling Detail

GPU	Model	Start °C	Stable °C	Limit	Clock Drop	Result	Diagnosis
GPU 0	NVIDIA RTX A4000	93°C	93.0°C	≤ 83°C	3.2%	FAIL	GPU temperature was already 93°C at idle — above the 83°C safety limit. The GPU throttled its own speed to protect itself. Inspect chassis fans, rack airflow, and thermal paste before running workloads.
GPU 1	NVIDIA RTX A4000	93°C	92.0°C	≤ 83°C	1.6%	FAIL	GPU temperature was already 93°C at idle — above the 83°C safety limit. The GPU throttled its own speed to protect itself. Inspect chassis fans, rack airflow, and thermal paste before running workloads.

Memory & Host Connection — FAIL

Node: node05 | 2026-08-19 06:27:28

GPU-to-Host Connection Speed (PCIe)

Host to GPU	GPU to Host	PCIe Gen	Rated Speed	Threshold	Efficiency	Status	Result
6.7 GB/s	6.8 GB/s	DEFAULT	32 GB/s	>= 8 GB/s	21%	Degraded	FAIL

Warning: GPU is connected at 21% of rated speed. The most common causes are: BIOS PCIe speed setting not set to Auto/Gen4, or GPU seated in an x8 slot instead of x16. Reseat the GPU and check BIOS.

GPU-to-GPU Connection Speed

only 1 GPU on node

Memory Architecture

NUMA Nodes	Threshold	Status	Result
1	Same memory domain	GPU and network card on same memory domain	PASS

Memory & Host Connection — PASS

Node: node07 | 2026-08-19 06:27:21

GPU-to-Host Connection Speed (PCIe)

Host to GPU	GPU to Host	PCIe Gen	Rated Speed	Threshold	Efficiency	Status	Result
25.0 GB/s	23.1 GB/s	GEN4	32 GB/s	>= 20 GB/s	80%	Degraded	PASS

GPU-to-GPU Connection Speed

only 1 GPU on node

Memory Architecture

NUMA Nodes	Threshold	Status	Result
1	Same memory domain	GPU and network card on same memory domain	PASS

Memory & Host Connection — PASS

Node: node08 | 2026-08-19 06:27:41

GPU-to-Host Connection Speed (PCIe)

Host to GPU	GPU to Host	PCIe Gen	Rated Speed	Threshold	Efficiency	Status	Result
26.8 GB/s	27.1 GB/s	GEN4	32 GB/s	>= 20 GB/s	85%	Full Speed	PASS

Values shown are per-GPU. Aggregate: H→D 53.6 GB/s, D→H 54.1 GB/s across 2 GPUs.

GPU-to-GPU Connection Speed

Direction	Speed	Link Type
gpu0→gpu1	26.3 GB/s	PCIe
gpu1→gpu0	26.3 GB/s	PCIe

GPU-to-GPU BW 26.3 GB/s via PCIe — NVLink not supported on this GPU model, PCIe is the maximum interconnect.

Memory Architecture

NUMA Nodes	Threshold	Status	Result
1	Same memory domain	GPU and network card on same memory domain	PASS

Ethernet Network — PASS

Coordinator: node04 | Nodes: node05, node07, node08 | 2026-08-19 06:30:06

Network Throughput — TCP

Node Pair	Measured Speed	Minimum Required	Retransmits	Result
node05 → node07	935.9 Mbit/s	>= 800 Mbit/s	0	PASS
node05 → node08	935.9 Mbit/s	>= 800 Mbit/s	0	PASS
node07 → node08	936.0 Mbit/s	>= 800 Mbit/s	0	PASS
node07 → node05	935.9 Mbit/s	>= 800 Mbit/s	0	PASS
node08 → node05	942.6 Mbit/s	>= 800 Mbit/s	0	PASS
node08 → node07	942.1 Mbit/s	>= 800 Mbit/s	0	PASS

Packet Loss Check — UDP

Node Pair	Speed	Max Allowed Loss	Actual Loss	Jitter	Result
node05 → node07	949.2 Mbit/s	<= 0.1%	0.000%	0.015 ms	PASS
node05 → node08	949.2 Mbit/s	<= 0.1%	0.000%	0.012 ms	PASS
node07 → node08	949.1 Mbit/s	<= 0.1%	0.020%	0.012 ms	PASS
node07 → node05	949.2 Mbit/s	<= 0.1%	0.000%	0.008 ms	PASS
node08 → node05	949.9 Mbit/s	<= 0.1%	0.000%	0.021 ms	PASS
node08 → node07	949.9 Mbit/s	<= 0.1%	0.000%	0.009 ms	PASS

Network Packet Size (MTU)

Node	Interface	MTU	Type
node05	eno1	1500	Standard
node07	eno1	1500	Standard
node08	eno1	1500	Standard

Jumbo frames (MTU 9000) are recommended for GPU workloads. Contact your network administrator to enable them.

High-Speed Network Fabric — PASS

Coordinator: node04 | Nodes: node05, node07, node08 | 2026-08-19 06:30:29

Network Port Status

Node	State	Required State	SM LID	Port LID	Result
node05	PORT_ACTIVE	PORT_ACTIVE	4	1	PASS
node07	PORT_ACTIVE	PORT_ACTIVE	4	3	PASS
node08	PORT_ACTIVE	PORT_ACTIVE	4	2	PASS

Network Manager Check (Subnet Manager)

Node	SM LID	Port LID	Status	Result
node05	4	1	SM detected (lid=4)	PASS
node07	4	3	SM detected (lid=4)	PASS
node08	4	2	SM detected (lid=4)	PASS

High-Speed Transfer Rate (RDMA Bandwidth)

Node Pair	Measured	Threshold	Result
node05 ↔ node07	11,507.5 MiB/s	>= 10,000 MiB/s	PASS
node05 ↔ node08	11,508.2 MiB/s	>= 10,000 MiB/s	PASS
node07 ↔ node08	11,439.3 MiB/s	>= 10,000 MiB/s	PASS

Network Response Time (RDMA Latency)

Node Pair	Avg Response	Min Response	Max Allowed	Result
node05 ↔ node07	1.14 us	1.11 us	<= 5.0 us	PASS
node05 ↔ node08	1.03 us	1.00 us	<= 5.0 us	PASS
node07 ↔ node08	1.14 us	1.10 us	<= 5.0 us	PASS

Single-Node GPU Sync Speed (NCCL AllReduce)

Node	Measured Speed	Threshold	Errors	Result
node05	128.9 GB/s	>= 50 GB/s	0	PASS
node07	195.2 GB/s	>= 50 GB/s	0	PASS
node08	195.2 GB/s	>= 50 GB/s	0	PASS

Multi-GPU Performance — FAIL

Coordinator: node04 | Nodes: node05, node07, node08 | 2026-08-19 06:34:11

Node-to-Node GPU Sync Speed — Transfer Speed Matrix (peak busbw, GB/s)

	node05	node07	node08
node05	—	4.09	4.10
node07	4.09	—	7.92
node08	4.10	7.92	—

Threshold: >= 5.0 GB/s per pair. Green = above threshold, Red = below threshold. Transport: IB.

Pair	Peak Speed	1MB Speed	Network	Min Required	Result
node05 ↔ node07	4.09 GB/s	3.37 GB/s	IB	>= 5.0 GB/s	FAIL
node05 ↔ node08	4.10 GB/s	3.36 GB/s	IB	>= 5.0 GB/s	FAIL
node07 ↔ node08	7.92 GB/s	7.46 GB/s	IB	>= 5.0 GB/s	PASS

Scaling Projection — estimated throughput at larger node counts

Nodes	Est. Throughput	Ring Efficiency	Note
2	7.92 GB/s	100%	measured average
4	11.88 GB/s	150%	projected
8	13.86 GB/s	175%	projected
16	14.85 GB/s	188%	projected
32	15.35 GB/s	194%	projected

Projection uses ring-allreduce formula: throughput x 2(n-1)/n. Assumes uniform fabric with no congestion.

All-Node GPU Sync Speed (N-way AllReduce)

Total GPUs	Peak Speed	1MB Speed	Min Required	Network	Result
3	4.85 GB/s	3.91 GB/s	>= 5.0 GB/s	IB	FAIL
