UpStation

Design **UPS-DSN-20261103-G9YQ**
Prepared 2026-11-03 · prices as of 2026-11-03
Prepared for **Example AI Lab** · Head of ML Infrastructure

Budgetary design and estimate — not a quote

4× HGX B300 training cluster (32 GPUs) with NVMe storage and a services cluster, Montana colo

4× 8x NVIDIA B300 (HGX B300 NVL8) (32 GPUs, HGX class), 5× 2U all-NVMe storage node storage (208.3 TB plannable, replica_3), 3 cluster services nodes. Networks: gpu fabric 800G on 1 switch; front end 400G on 2 switches; oob 1G on 1 switch. 3 racks at ≤ 40 kW, 82.51 kW provisioned, in Montana colo. Every figure is computed from the ai.upstation.io catalog with the arithmetic shown below.

$3,307,252.84	**$21,453.00/mo**	**$3,564,688.84**	**$4,079,560.84**	**$4,594,432.84**
hardware (CapEx)	colocation, 82.51 kW	year one	3-year total	5-year total

Delivery: 20 weeks, set by 8x NVIDIA B300 (HGX B300 NVL8), Intel, 3TB (Aivres) — Aivres KR9288X3. **Colo:** Montana at $260/kW/month × 82.51 kW, 36-month term.

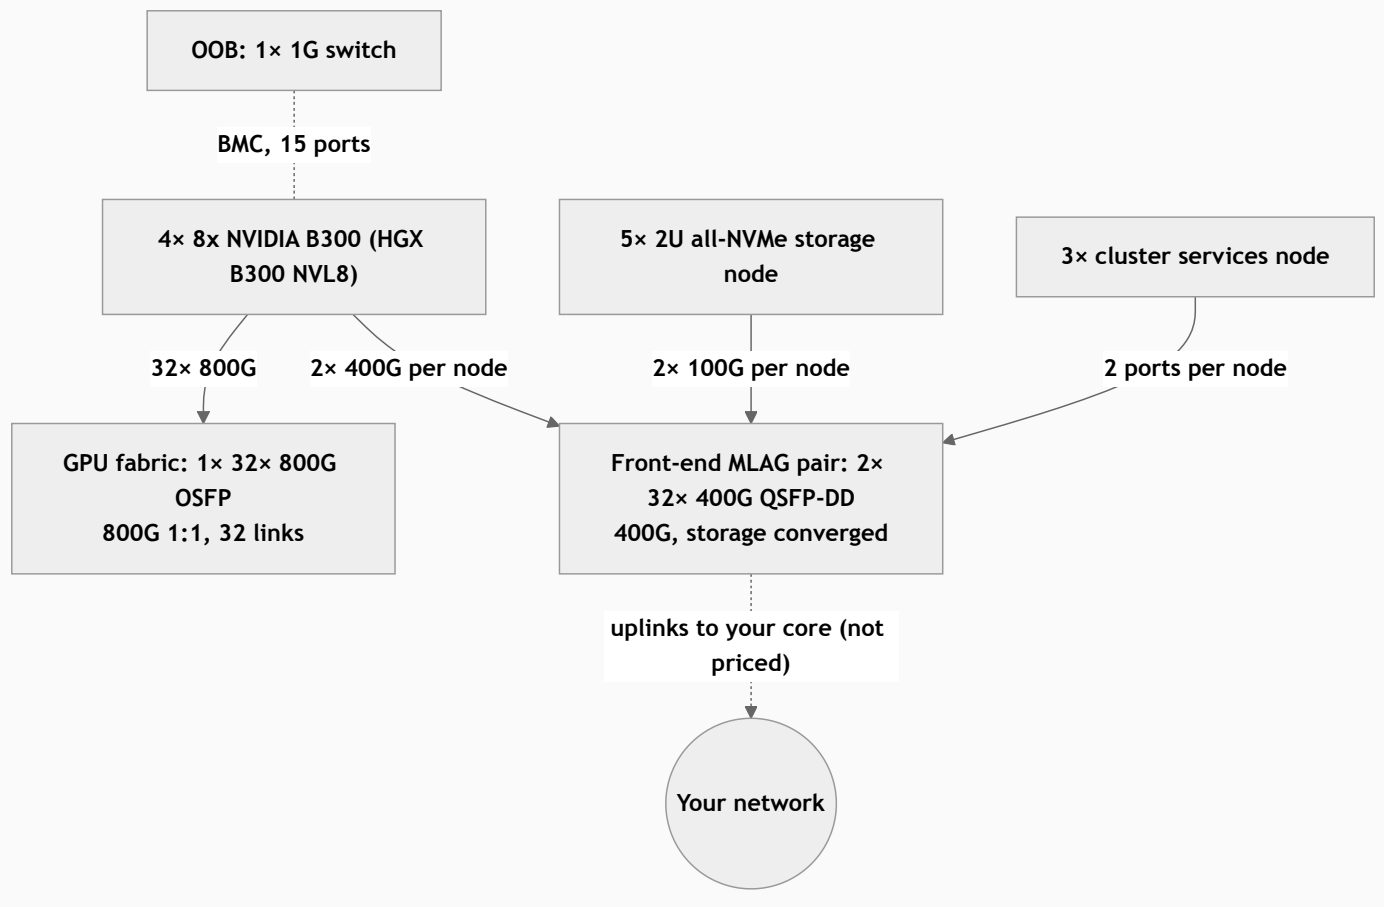
Requirements

Workload	training
GPUs	32
Storage	200 TB, NVME
Rack limit	40 kW
Colo	Montana
Timeline	production in Q2 2027

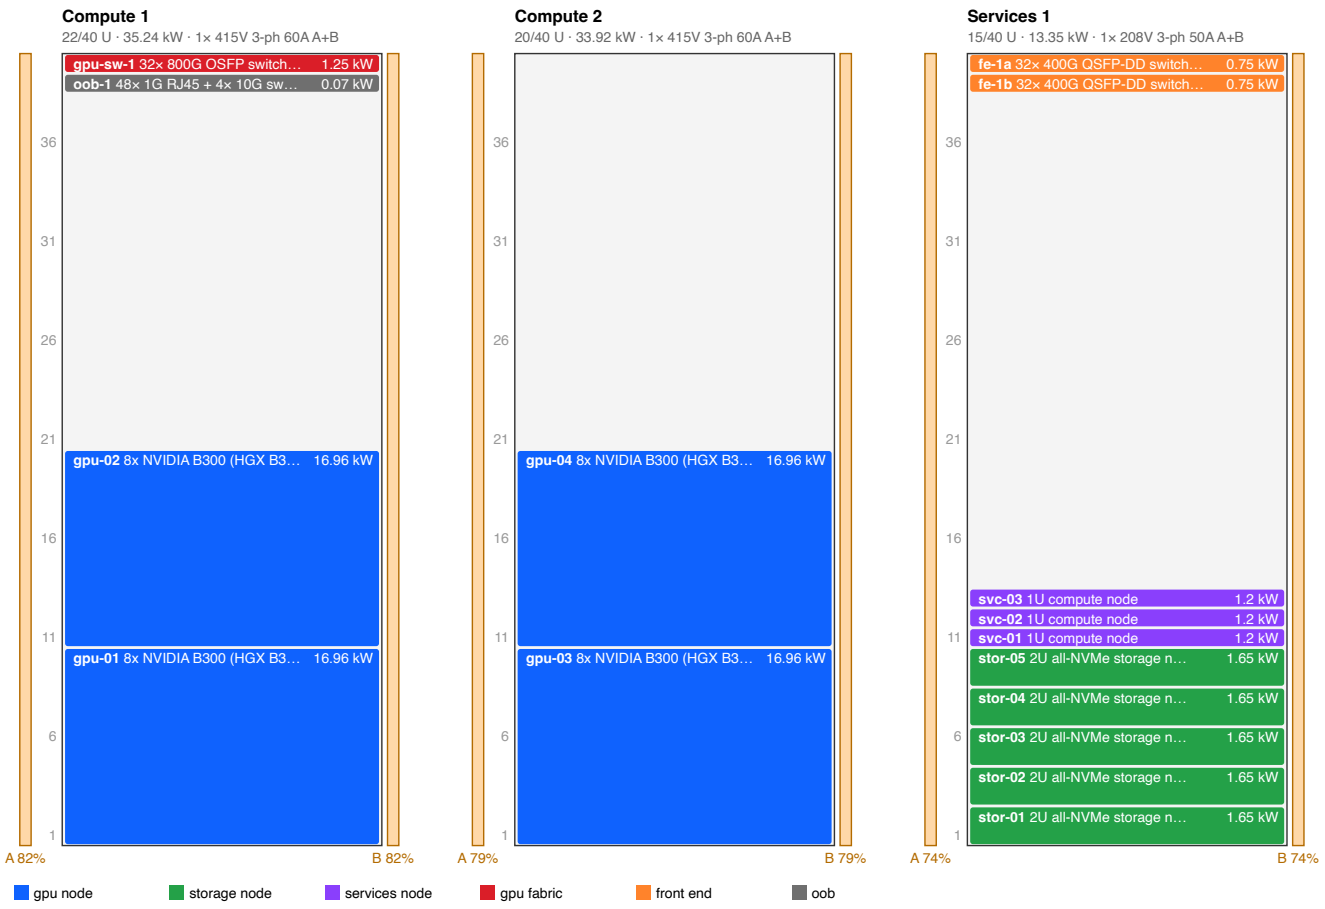
How it was designed

1. GPUs → nodes: ceil(32 ÷ 8) = 4 × 8x NVIDIA B300 (HGX B300 NVL8), Intel, 3TB (Aivres) (32 GPUs; nodes ship fully populated).
2. Class: HGX — the GPUs must act as one (training, multi-node). Does the model at its precision fit one GPU's memory, and can each GPU work alone? Yes: PCIe. No: HGX. Distributed training is never served by PCIe servers.
3. Storage (replica_3, Ceph across nodes, floor 3+1): plannable = 184 × N × (N–1)/N × 0.333 × 0.85; N = 5 gives 208.3 TB ≥ 204.5 TB and 5 × 20 = 100 GB/s ≥ 16.0 GB/s read.
4. Storage demand: read = 32 × 0.5 = 16.0 GB/s; checkpoint burst = 18.7 GB/s; write = 18.7 GB/s; capacity = 204.5 TB.
5. Cluster services (chosen with the buyer, optional): 3 × 1U CPU nodes for provisioning, scheduler and monitoring.
6. GPU fabric: 32 × 800G ports (1 per GPU) ≤ 32 → one 32× 800G OSFP switch, 1:1 non-blocking (single point of failure; a leaf-spine removes it).
7. Front-end: 2 ports per host, one to each switch of an MLAG pair of 32× 400G QSFP-DD switches; 7 ports per switch of 30 usable → 1 pair(s). Storage stays converged on this plane (well under 40 GB/s per node).
8. OOB: 12 hosts + 3 switches = 15 × 1G ports → 1 × 48× 1G RJ45 + 4× 10G.
9. GPU nodes per 40 kW rack = min(floor(40 U ÷ 10 U), floor(40 ÷ 16.96 kW)) = 2 (power binds).
10. Compute 1: 22 U, 35.24 kW provisioned → 1 × 415V 3-ph 60A circuit pair (A+B), 82% of each feed; 8 power cords per feed.
11. Compute 2: 20 U, 33.92 kW provisioned → 1 × 415V 3-ph 60A circuit pair (A+B), 79% of each feed; 6 power cords per feed.
12. Services 1: 15 U, 13.35 kW provisioned → 1 × 208V 3-ph 50A circuit pair (A+B), 74% of each feed; 10 power cords per feed.

Network topology



Rack elevation and power



Rack	Space	Provisioned	Circuits	Devices
Compute 1	22/40 U	35.24 kW	1× 415V 3-ph 60A (A+B)	2× 8x NVIDIA B300 (HGX B300 NVL8), 1× 32× 800G OSFP switch — GPU fabric, 1× 48× 1G RJ45 + 4× 10G switch — OOB
Compute 2	20/40 U	33.92 kW	1× 415V 3-ph 60A (A+B)	2× 8x NVIDIA B300 (HGX B300 NVL8)
Services 1	15/40 U	13.35 kW	1× 208V 3-ph 50A (A+B)	2× 32× 400G QSFP-DD switch — front-end pair, MLAG, 5× 2U all-NVMe storage node, 3× 1U compute node

Bill of materials

Item	Qty	Unit (USD)	Extended	Lead time	Source
GPU servers					
8x NVIDIA B300 (HGX B300 NVL8), Intel, 3TB (Aivres) — Aivres KR9288X3	4	537,196.40	2,148,785.60	Configure-to-order on Aivres B300 allocation (October 2026 s	/p/UPS-GPU-B300-8-T2-AIV
Storage					
2U all-NVMe storage node, dual socket, 24x 7.68TB NVMe (184TB raw) — SYS-222C-TN	5	177,958.91	889,794.55	Barebone in stock (2026-11-03); configure-to-order. Est. 3-5	/p/UPS-SN-INT-NVME
Cluster services					
1U compute node, dual socket, 64 cores, 512GB, diskless NVMe (virtualisation host) — SYS-122C-TN	3	34,523.31	103,569.93	Barebone in stock (2026-11-03); configure-to-order. Est. 3-5	/p/UPS-CN-INT-1U
GPU fabric					
32× 800G OSFP switch (GPU fabric) — FS N8650-32OD, PicOS	1	37,500.00	37,500.00	26 in stock (global warehouse, 2026-11-03). ~1 week	/p/UPS-NET-800G-32
GPU fabric cabling					

800G link: 2× SR8 + 2× MTP-12 OM4 10 m — node to leaf, 1 per GPU port (estimated optical)	32	2,058.00	65,856.00		/switches
Front-end network					
32× 400G QSFP-DD switch (front-end pair, MLAG) — FS N8610-32D, PicOS	2	20,000.00	40,000.00	33 in stock (global warehouse, 2026-11-03). ~1 week	/p/UPS-NET-400G-32
+2× 400G north-south (budgeted at the Supermicro equivalent option; Aivres fits its own 400G NICs and confirms them on its quote) (listed option, one per GPU node)	4	2,693.94	10,775.76		/p/UPS-GPU-B300-8-T2-AIV
Front-end network cabling					
400G AOC 10 m — GPU nodes to the pair, 2 per node	8	1,055.00	8,440.00		/switches
400G→4×100G breakout 2 m — storage nodes to the pair (4 nodes per cable)	4	199.00	796.00		/switches
400G→4×100G breakout 2 m — cluster services nodes to the pair (4 nodes per cable)	2	199.00	398.00		/switches
400G 2 m passive DAC — peer link	2	174.00	348.00		/switches
Out-of-band					
48× 1G RJ45 + 4× 10G switch (OOB) — FS S3410-48TS, PicOS	1	989.00	989.00	593 in stock (global warehouse, 2026-11-03). ~1 week	/p/UPS-NET-1G-48
Hardware total			**$3,307,252.84**		

Assumptions

- GPU server: 8x NVIDIA B300 (HGX B300 NVL8), Intel, 3TB (Aivres) (UPS-GPU-B300-8-T2-AIV), the default HGX line (proposed)
- Second source: 8x NVIDIA B300 (HGX B300 NVL8), Intel (UPS-GPU-B300-8), same price and design, 30+ weeks on NVIDIA allocation. Next generation: NVIDIA Rubin (HGX Rubin NVL8, liquid-cooled racks) follows B300, not priced

Not included

- Racks, PDUs, whips and cabling labour
- Uplinks to the buyer's core network
- Software (cluster manager, storage software), unless listed
- Taxes, shipping and duties
- Installation and on-site services

Next step. Book the intro (https://calendly.com/alecglover/intro-call) or email alec@upstation.io with the design ID; UpStation validates the design and introduces the supplier that fits.

Prepared by UpStation from public supplier prices published at ai.upstation.io (prices as of 2026-11-03). Budgetary only: the supplier confirms price, availability and lead time on order. UpStation does not resell; buyers transact directly with the supplier. alec@upstation.io