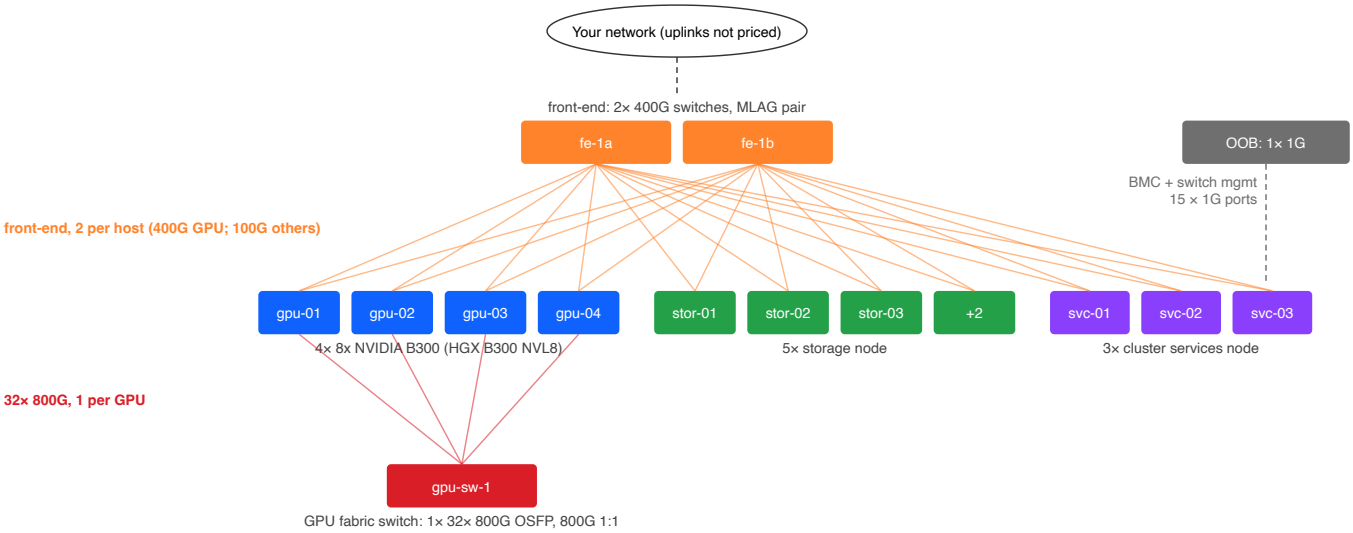
UpStation

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

Network topology — budgetary design, not a quote

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

Topology



Fabrics

Fabric	Speed	Switches	Ports used / total	Oversubscription
gpu fabric	800G	1× 32× 800G OSFP	32 / 32	1:1
front end	400G	2× 32× 400G QSFP-DD	18 / 64	3:1 or better
oob	1G	1× 48× 1G RJ45 + 4× 10G	15 / 48	n/a

How it was designed

1. 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).
2. 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).
3. OOB: 12 hosts + 3 switches = 15 × 1G ports → 1 × 48× 1G RJ45 + 4× 10G.

Network bill of materials

Item	Qty	Unit (USD)	Extended
32× 800G OSFP switch (GPU fabric)	1	$37,500.00	$37,500.00
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
32× 400G QSFP-DD switch (front-end pair, MLAG)	2	$20,000.00	$40,000.00
+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
400G AOC 10 m — GPU nodes to the pair, 2 per node	8	$1,055.00	$8,440.00
400G→4×100G breakout 2 m — storage nodes to the pair (4 nodes per cable)	4	$199.00	$796.00

400G→4×100G breakout 2 m — cluster services nodes to the pair (4 nodes per cable)	2	$199.00	$398.00
400G 2 m passive DAC — peer link	2	$174.00	$348.00
48× 1G RJ45 + 4× 10G switch (OOB)	1	$989.00	$989.00
Cat6 patch — 1G BMC/management (not priced)	15	$0.00	$0.00
Network total			**$165,102.76**

Port map and cable schedule (63 cables, 73 connections)

Rows that share a cable ID are lanes of one breakout cable. Device names match the rack elevation.

Cable	Fabric	From	To	Cable type
C0001	gpu fabric	gpu-01:GPU NIC 1	gpu-sw-1:1	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0002	gpu fabric	gpu-01:GPU NIC 2	gpu-sw-1:2	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0003	gpu fabric	gpu-01:GPU NIC 3	gpu-sw-1:3	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0004	gpu fabric	gpu-01:GPU NIC 4	gpu-sw-1:4	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0005	gpu fabric	gpu-01:GPU NIC 5	gpu-sw-1:5	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0006	gpu fabric	gpu-01:GPU NIC 6	gpu-sw-1:6	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0007	gpu fabric	gpu-01:GPU NIC 7	gpu-sw-1:7	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0008	gpu fabric	gpu-01:GPU NIC 8	gpu-sw-1:8	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0009	gpu fabric	gpu-02:GPU NIC 1	gpu-sw-1:9	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0010	gpu fabric	gpu-02:GPU NIC 2	gpu-sw-1:10	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0011	gpu fabric	gpu-02:GPU NIC 3	gpu-sw-1:11	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0012	gpu fabric	gpu-02:GPU NIC 4	gpu-sw-1:12	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0013	gpu fabric	gpu-02:GPU NIC 5	gpu-sw-1:13	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0014	gpu fabric	gpu-02:GPU NIC 6	gpu-sw-1:14	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0015	gpu fabric	gpu-02:GPU NIC 7	gpu-sw-1:15	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0016	gpu fabric	gpu-02:GPU NIC 8	gpu-sw-1:16	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0017	gpu fabric	gpu-03:GPU NIC 1	gpu-sw-1:17	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0018	gpu fabric	gpu-03:GPU NIC 2	gpu-sw-1:18	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0019	gpu fabric	gpu-03:GPU NIC 3	gpu-sw-1:19	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0020	gpu fabric	gpu-03:GPU NIC 4	gpu-sw-1:20	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0021	gpu fabric	gpu-03:GPU NIC 5	gpu-sw-1:21	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0022	gpu fabric	gpu-03:GPU NIC 6	gpu-sw-1:22	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0023	gpu fabric	gpu-03:GPU NIC 7	gpu-sw-1:23	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0024	gpu fabric	gpu-03:GPU NIC 8	gpu-sw-1:24	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0025	gpu fabric	gpu-04:GPU NIC 1	gpu-sw-1:25	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0026	gpu fabric	gpu-04:GPU NIC 2	gpu-sw-1:26	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0027	gpu fabric	gpu-04:GPU NIC 3	gpu-sw-1:27	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0028	gpu fabric	gpu-04:GPU NIC 4	gpu-sw-1:28	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0029	gpu fabric	gpu-04:GPU NIC 5	gpu-sw-1:29	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0030	gpu fabric	gpu-04:GPU NIC 6	gpu-sw-1:30	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0031	gpu fabric	gpu-04:GPU NIC 7	gpu-sw-1:31	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0032	gpu fabric	gpu-04:GPU NIC 8	gpu-sw-1:32	800G link: 2× SR8 + 2× MTP-12 OM4 10 m
C0033	front end	gpu-01:FE 1	fe-1a:1	400G AOC 10 m
C0034	front end	gpu-01:FE 2	fe-1b:1	400G AOC 10 m
C0035	front end	gpu-02:FE 1	fe-1a:2	400G AOC 10 m
C0036	front end	gpu-02:FE 2	fe-1b:2	400G AOC 10 m

C0037	front end	gpu-03:FE 1	fe-1a:3	400G AOC 10 m
C0038	front end	gpu-03:FE 2	fe-1b:3	400G AOC 10 m
C0039	front end	gpu-04:FE 1	fe-1a:4	400G AOC 10 m
C0040	front end	gpu-04:FE 2	fe-1b:4	400G AOC 10 m
C0041	front end	stor-01:FE 1	fe-1a:5/1	400G→4×100G breakout 2 m
C0042	front end	stor-01:FE 2	fe-1b:5/1	400G→4×100G breakout 2 m
C0041	front end	stor-02:FE 1	fe-1a:5/2	400G→4×100G breakout 2 m
C0042	front end	stor-02:FE 2	fe-1b:5/2	400G→4×100G breakout 2 m
C0041	front end	stor-03:FE 1	fe-1a:5/3	400G→4×100G breakout 2 m
C0042	front end	stor-03:FE 2	fe-1b:5/3	400G→4×100G breakout 2 m
C0041	front end	stor-04:FE 1	fe-1a:5/4	400G→4×100G breakout 2 m
C0042	front end	stor-04:FE 2	fe-1b:5/4	400G→4×100G breakout 2 m
C0043	front end	stor-05:FE 1	fe-1a:6/1	400G→4×100G breakout 2 m
C0044	front end	stor-05:FE 2	fe-1b:6/1	400G→4×100G breakout 2 m
C0045	front end	svc-01:FE 1	fe-1a:7/1	400G→4×100G breakout 2 m
C0046	front end	svc-01:FE 2	fe-1b:7/1	400G→4×100G breakout 2 m
C0045	front end	svc-02:FE 1	fe-1a:7/2	400G→4×100G breakout 2 m
C0046	front end	svc-02:FE 2	fe-1b:7/2	400G→4×100G breakout 2 m
C0045	front end	svc-03:FE 1	fe-1a:7/3	400G→4×100G breakout 2 m
C0046	front end	svc-03:FE 2	fe-1b:7/3	400G→4×100G breakout 2 m
C0047	front end	fe-1a:31	fe-1b:31	400G 2 m passive DAC
C0048	front end	fe-1a:32	fe-1b:32	400G 2 m passive DAC
C0049	oob	gpu-01:BMC	oob-1:1	Cat6 patch
C0050	oob	gpu-02:BMC	oob-1:2	Cat6 patch
C0051	oob	gpu-03:BMC	oob-1:3	Cat6 patch
C0052	oob	gpu-04:BMC	oob-1:4	Cat6 patch
C0053	oob	stor-01:BMC	oob-1:5	Cat6 patch
C0054	oob	stor-02:BMC	oob-1:6	Cat6 patch
C0055	oob	stor-03:BMC	oob-1:7	Cat6 patch
C0056	oob	stor-04:BMC	oob-1:8	Cat6 patch
C0057	oob	stor-05:BMC	oob-1:9	Cat6 patch
C0058	oob	svc-01:BMC	oob-1:10	Cat6 patch
C0059	oob	svc-02:BMC	oob-1:11	Cat6 patch
C0060	oob	svc-03:BMC	oob-1:12	Cat6 patch
C0061	oob	gpu-sw-1:mgmt	oob-1:13	Cat6 patch
C0062	oob	fe-1a:mgmt	oob-1:14	Cat6 patch
C0063	oob	fe-1b:mgmt	oob-1:15	Cat6 patch

Topology (Mermaid source)

```
flowchart TB
  GF["GPU fabric: 1× 32× 800G OSFP<br/>800G 1:1, 32 links"]
  GN["4× 8x NVIDIA B300 (HGX B300 NVL8)"] -- "32× 800G" --> GF
  FE["Front-end MLAG pair: 2× 32× 400G QSFP-DD<br/>400G, storage converged"]
  GN -- "2× 400G per node" --> FE
  SN["5× 2U all-NVMe storage node"] -- "2× 100G per node" --> FE
  MG["3× cluster services node"] -- "2 ports per node" --> FE
  FE -. "uplinks to your core (not priced)" .-> CORE(("Your network"))
  OOB["OOB: 1× 1G switch"] -. "BMC, 15 ports" .- GN
```

Assumptions

- None beyond the catalog conventions.

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