UpStation

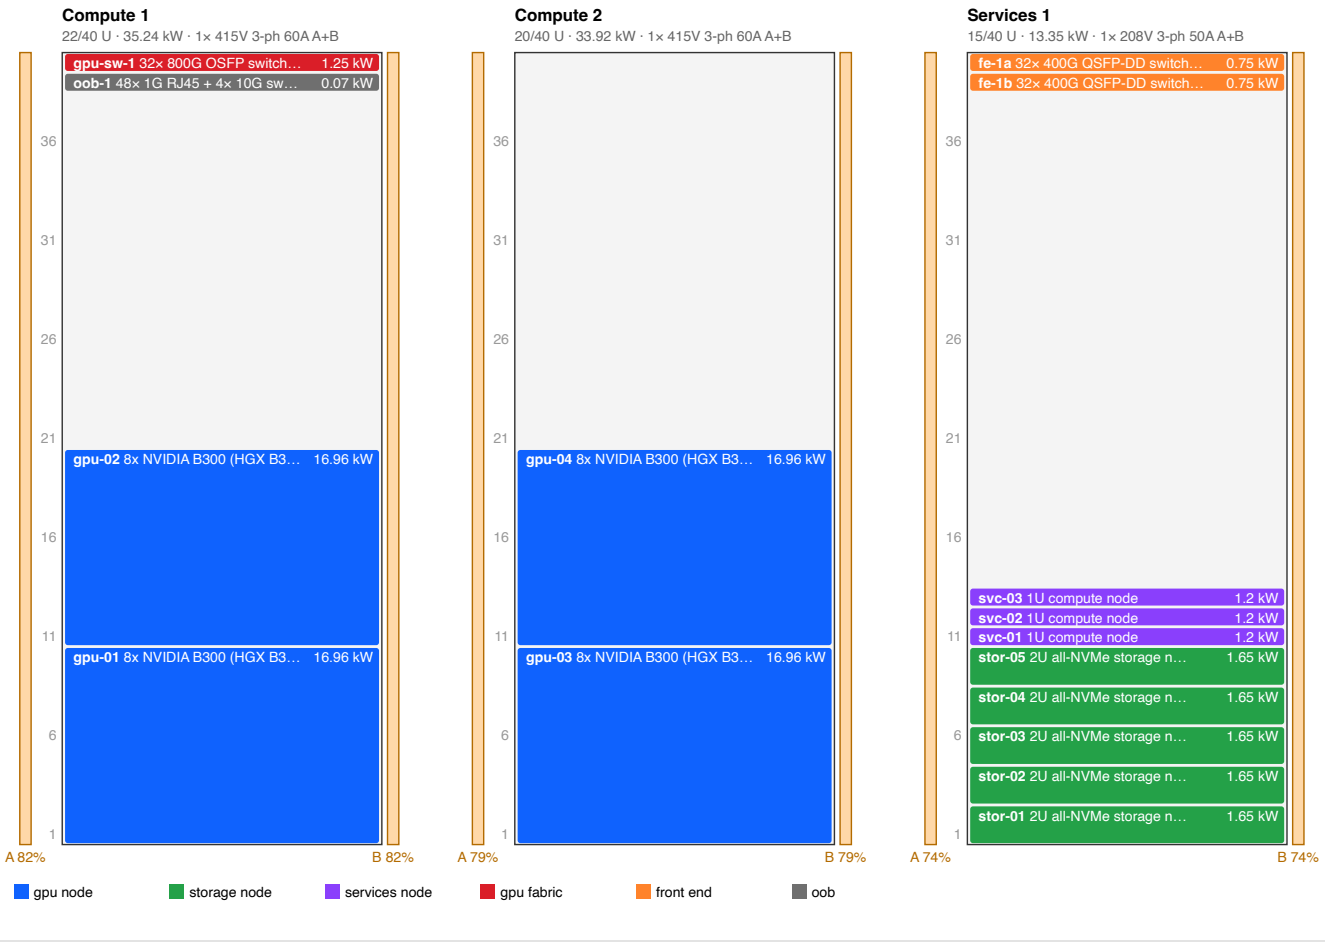
Design **UPS-DSN-20261103-G9YQ**
Prepared 2026-11-03 · prices as of 2026-11-03
Prepared for **Example AI Lab** · Head of ML Infrastructure

Rack elevation and power — budgetary design, not a quote

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

82.51 kW provisioned (nameplate ÷ 0.8) across the racks below, each held to 40 kW. Each rack takes an A and a B feed; either one alone carries the whole rack.

Rack elevation



Power per rack

Rack	Space	Provisioned	Circuits (A+B)	Load per feed	Cords per feed
Compute 1	22/40 U	35.24 kW	1× 415V 3-ph 60A	82%	8
Compute 2	20/40 U	33.92 kW	1× 415V 3-ph 60A	79%	6
Services 1	15/40 U	13.35 kW	1× 208V 3-ph 50A	74%	10

Colocation by location

Location	$/kW/month	Per month	Term total	Sales tax	Property tax
Montana (this design)	$260	$21,453.00	$772,308.00 (36 mo)	none	none on the first $1M of equipment, about 0.8%/yr above
Texas	$225	$18,565.00	$668,340.00 (36 mo)	8.25% (6.25% state + local); the data-center exemption needs $200M and 20 jobs	about 2.2%/yr of market value
Oregon	$175	$14,439.00	$519,804.00 (36 mo)	none	about 1.5-2%/yr of market value (ORS 308.290); abatements closed to data-center property since 2026
Pennsylvania	$250	$20,628.00	$742,608.00 (36 mo)	6% (8% Philadelphia, 7% Allegheny), or none if the facility is state-certified and the tenant holds its own exemption certificate (100 kW+, 2-year contract) before buying	none
Canada (Ontario, Quebec)	$225	$18,565.00	$668,340.00 (36 mo)	HST 13% (Ontario) or GST+QST (Quebec) on hardware and colo, recovered as an input tax credit by a registered business	none

How it was laid out

1. GPU nodes per 40 kW rack = min(floor(40 U ÷ 10 U), floor(40 ÷ 16.96 kW)) = 2 (power binds).
2. 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.
3. 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.
4. 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.

Devices

Rack	U	Name	Device	RU	kW	PSUs
Compute 1	40	gpu-sw-1	32× 800G OSFP switch — GPU fabric	1	1.25	2
Compute 1	39	oob-1	48× 1G RJ45 + 4× 10G switch — OOB	1	0.07	2
Compute 1	11	gpu-02	8x NVIDIA B300 (HGX B300 NVL8)	10	16.96	6
Compute 1	1	gpu-01	8x NVIDIA B300 (HGX B300 NVL8)	10	16.96	6
Compute 2	11	gpu-04	8x NVIDIA B300 (HGX B300 NVL8)	10	16.96	6
Compute 2	1	gpu-03	8x NVIDIA B300 (HGX B300 NVL8)	10	16.96	6
Services 1	40	fe-1a	32× 400G QSFP-DD switch — front-end pair, MLAG	1	0.75	2
Services 1	39	fe-1b	32× 400G QSFP-DD switch — front-end pair, MLAG	1	0.75	2
Services 1	13	svc-03	1U compute node	1	1.2	2
Services 1	12	svc-02	1U compute node	1	1.2	2
Services 1	11	svc-01	1U compute node	1	1.2	2
Services 1	9	stor-05	2U all-NVMe storage node	2	1.65	2
Services 1	7	stor-04	2U all-NVMe storage node	2	1.65	2
Services 1	5	stor-03	2U all-NVMe storage node	2	1.65	2
Services 1	3	stor-02	2U all-NVMe storage node	2	1.65	2
Services 1	1	stor-01	2U all-NVMe storage node	2	1.65	2

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

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