


ZCOMPUTE GPU NODES

Flexible, Scalable Cloud Power

With Zadara’s zCompute offering, you can develop, deploy, run and virtualize any application. It offers pay-as-you-go, consumption based (100% OpEx) pricing model* to help control and optimize your spending (zero hidden transport fees; free ingress and free egress helps here, too!). Technology wise, you get easy, self service provisioning using predefined & fully-customizable instance types, auto-scaling to adjust resources to meet evolving requirements, virtual machine (VM) snapshots for easy, convenient backups of VM based data, and EC2 compatible APIs.



* Other commitment based models available

Available instance types

Product type	Instance Type	vCPU	Memory [GIB]	NVMe [GIB]	GPUs
ZGL4 Family					
ZGL4	zgl4.7large	14	112	1920	1
ZGL4	zgl4.7xlarge	28	224	3840	2
ZGL4	zgl4.14xlarge	56	448	7680	4
ZGL40S Family					
ZGL40S	zgl40s.7large	14	112	1920	1
ZGL40S	zgl40s.7xlarge	28	224	3840	2
ZGL40S	zgl40s.14xlarge	56	448	7680	4

The GPU based services require Zadara GPU nodes** in the cloud

	NVIDIA L40S GPUs	NVIDIA L4 GPUs
Form Factor	2U Rackmount	
Server model	SupeMicro SYS-220GP-TNR	
CPU	2x Intel Xeon Gold 6326 (16C/32T, 2.9GHz, 185W)	
Memory	16x 32GB DDR4-3200 RDIMM - 512GB Total	
Storage	4x 1.9TB NVMe PCIe 4.0 SSDs for VMs + 2x 480GB SATA SSDs for boot	
Networking	2x 10GBase-T + 2x 100GbE QSFP28	
Power Consumption (Base Server)	400-600W	
GPU Configuration	4x NVIDIA L40S (48GB GDDR6, 350W each)	4x NVIDIA L4 (24GB GDDR6, 70W each)
GPU Architecture	Ada Lovelace	
Power Consumption (GPUs)	1400W Total (350W x4)	280W Total (70W x4)
Common Use Case	AI Training, Inference, HPC, CAD, VDI	AI Inference, Edge AI, Light-weight Compute, CAD, VDI

** A minimum of 2 GPU nodes of the same type is required for redundancy



Fully managed cloud services: compute, networking, storage and more; designed for service providers and the modern enterprise.

Contact us at:

www.zadara.com
info@zadara.com