Roar and Roar Collab Paid Allocation Service & Support Plans

Paid allocations offer researchers guaranteed access to compute and storage resources. With a minimum 36 month term length, purchase 40 cores of Basic Memory, 20 cores (or equivalent) of Standard Memory, High Memory, or GPUs, and receive 5TiB of group storage, free.

Roar Collab Resources

Roar Collab is not to be used with restricted (Level-3 or Level- 4) data. Researchers who must comply with restricted data should utilize the Roar cluster instead.

Type	Description	Cost/Month
Basic Memory	For minimal memory and high throughput jobs. • 5.33/4 GB/core • 24/64 cores and 128/256 GB RAM per node	$80 per 20 core block
Standard Memory	For moderate memory jobs and multicore processing • 8/10 GB/core, Infiniband • 20/48 cores and 256/384/512 GB RAM per node	$140 per 20 core block
High Memory	For large memory jobs and multicore processing • 25 GB/core, Infiniband • 40 cores and 1 TB RAM per node	$420 per 20 core block
Single GPU (A100)	For GPU enabled compute. • NVIDIA A100 GPU, Infiniband, 24 cores and 192 GB RAM per card • 2 GPU, 48 cores, and 384 GB RAM per node	$400 per GPU card
Single GPU (P100)	For GPU enabled compute. • 1 NVIDIA P100 GPU, Infiniband, 24 cores and 256/512 GB RAM per card/node	$300 per GPU card
Multi-Instance GPU (1/2^{h} Slice)	A "slice" of an NVIDIA A100 card; allows for GPU enabled compute without purchasing a whole card. • ≈3600 CUDA core and 20GB VRAM per slice • 12 CPU cores and 96 GB RAM per slice	$200 per MIG slice
Multi-Instance GPU (1/8th Slice)	A "slice" of an NVIDIA A100 card; allows for GPU enabled compute without purchasing a whole card. • ≈900 CUDA core and 5GB VRAM per slice • 3 CPU cores and 24 GB RAM per slice	$57 per MIG slice
Active Storage	High performance storage mounted to compute resources - designed for actively used files and datasets.	$33.35 per 5TiB
Archive Storage	Archival storage separate from compute resources - designed for long term storage and infrequently used files and data sets.	$6.25 per 5TiB

Purchasers of cores will access servers with at least these specifications.

Roar Resources

The Roar cluster is maintained for researchers who utilize restricted data.

Type	Description	Cost/Month
Standard Memory	For moderate memory jobs and multicore processing. • 10 GB/core, Infiniband • 20 cores and 256 GB RAM per node	$140 per 20 core block
Single GPU (P100)	For GPU enabled compute. • 4 NVIDIA P100 GPUs, Infiniband, 24 cores and 256 GB RAM per node • 1 GPU, 6 cores, and 64 GB RAM per card	$300 per card
Active Storage	High Performance storage mounted to compute resources - designed for actively used files and datasets.	$33.35 per 5TiB

Purchasers of cores will access servers with at least these specifications.

Special Services

Your service level agreement may include special services such as RISE consultations, science gateways, or hosted compute. Contact us at icds@psu.edu to discuss your needs and pricing.

Payment for Services

ICDS utilizes iLab for monthly invoicing of our Service Level Agreements.

Contact ICDS to Get Started

Call or email to discuss your research goals so we can create a package that meets your needs and your budget. icds@psu.edu 814-867-1467

Note: *Prices reflect discounts made possible through substantial funding from the Provost to encourage the use of ICDS services for University-sponsored research.*