

ARC COMPUTE

HIGH PERFORMANCE GPU CLOUD

Justin Ritchie - CEO, Arc Compute
justin@arccompute.io

ASK ABOUT OUR 1 WEEK FREE TRIAL

[ARCCOMPUTE.IO](https://arccompute.io)

ABOUT ARC COMPUTE

THE COMPANY

Arc Compute is a next generation High Performance Compute Cloud Provider, powered by in-house proprietary GPU software that provides up to 80% better performance than instances from AWS, GCP and Azure, at 30%-60% less cost. With fixed monthly costs and unlimited usage, you always know what to expect.

OUR CUSTOMERS

Arc Compute's customers have one thing in common; they are all large consumers of GPUs who are tired of the current cloud business models and are looking for better, transparent pricing and better performance and security. Our clients are leaders in artificial intelligence and machine learning, pharmaceutical, media and entertainment, defence and more.

BENEFITS OF ARC COMPUTE

- 01 Up to 80% better performance over AWS, GCP and Azure thanks to proprietary GPU software
- 02 30%-60% lower cost than AWS, GCP and Azure
- 03 Fixed monthly pricing for dedicated GPU compute
- 04 Only cloud provider with the option of Composable Disaggregated Infrastructure
- 05 Highly customizable and easily scalable infrastructure
- 06 Global network of energy efficient green data centers
- 07 No Ingress or Egress fees
- 08 Fixed Monthly Pricing

COMPUTING WITH ARC



SUPERIOR PERFORMANCE


PROPRIETARY SOFTWARE

Unlike other cloud providers, Arc Compute utilizes its own proprietary, next-generation GPU software in its cloud services. Compared to software used by the vast majority of cloud providers (vGPU, VMware, etc), Arc Compute's GPU software boosts GPU performance by up to 80%. Additionally, since we own the software, we are able to pass on savings of up to 60% (over AWS) to our customers.

C2 - Pass Through (2 Arc GPUs)

10h 52m 14s

A single VM with 2 "ARC A100 (40 GB)" passed through to a VM which is run by LibVF.IO.



In the above scenario, multi-tenancy is used to provide over 70% performance boost by allowing under-utilized VM's to allocate their GPU cores, execution capabilities and VRAM at run-time to active VM's.


ZERO TRUST ISOLATION

Arc Compute provides discrete virtual machines with native performance. This is possible thanks to our proprietary software, which uses IOMMU for hardware enabled mediation, ensuring that your workloads and data are safe from viruses, ransomware, bad actors, data breaches and more.

C5 - Multiplexed (Reduced VRAM)

8h 21m 44s

A single VM with 4 "ARC A100 (20 GB)" passed through to a VM which is run by Hyperborea.



01

Hardware enforced isolation via Intel Virtualization Technology (Intel VT-d) and AMD Virtualization (AMD-V) Technology

02

Dedicated IP addresses to virtual machines

03

GPU Quality of Service (QoS) guaranteed via Arc Compute's LibVF.IO



SUPERIOR PRICING

Arc Compute offers the latest NVIDIA, AMD and Intel GPUs as a service and works directly with the GPU and server manufacturers to create next-generation HPC cloud services.

Below is a comparison of Amazon's main HPC instance with Arc's 8x A100 instance.

Arc Compute can offer mix and match instances (even different GPUs in the same instance) but the following are readily available:

NVIDIA A100 (1X,2X,4X,8X,16X, 32X)

NVIDIA A40 (1X,2X,4X,8X,16X, 32X)

NVIDIA A16 (1X,2X,4X,8X,16X, 32X)

NVIDIA A2 (1X,2X,4X,8X,16X,32X)



Better Performance, Better Price

	Monthly 1-year commit	Monthly 2-year commit	Monthly 3-year commit
Arc Compute 8 X A100 Instance Dedicated	\$1,875 per GPU	\$1,315.56 per GPU	\$964.70 per GPU
AWS 8 X A100 EC2 P4d Instance Dedicated	\$2,532.57 per GPU	N/A	\$1,607.93 per GPU

AWS prices are based on the available prices on their website as of April 2022.

BUSINESS IMPACT

GLOBAL RESOURCE OPTIMIZATION

- Optimize your infrastructure when workforce is across multiple timezones
- Reduce your infrastructure without sacrificing performance; thanks to Arc's proprietary software, you can have a smaller infrastructure footprint and still get the performance that you are used to, all while spending less

JUST IN TIME RESOURCE ALLOCATION

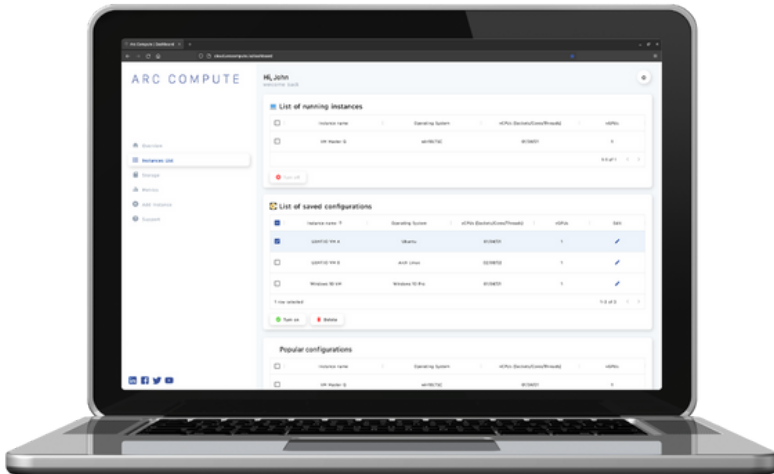
- Pool GPUs and memory using composable disaggregated infrastructure for use when demand is present and make your infrastructure multi-purpose
- Drive cost efficiency by scheduling resources to consider peak usage rather than individual GPU needs
- Run-time allocation of GPU cores, execution capabilities and VRAM providing increased performance beyond minimum guarantee of service

INFRASTRUCTURE INVESTMENT OPTIMIZATION & EFFICIENCY

- Address the total needs of your organizations GPU compute at peak usage; reimagine how organizational demand is forecasted and acquired as a whole, rather than siloed resource procurement



SIMPLE USER EXPERIENCE



```
dopefish — arc@dev1: ~ — ssh arc@9.0.0.56 — 79x26
arc@dev1:~$ nvidia-smi
Tue May 3 21:14:04 2022

+-----+
| NVIDIA-SMI 510.47.03      Driver Version: 510.47.03   CUDA Version: 11.6     |
+-----+-----+
| GPU   Name                               Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap       Memory-Usage | GPU-Util  Compute M. |
|=====+=====+
| 0.  NVIDIA A100-SXM...                  Off           | 00000000:4B:00:00 Off |   0%       Default  |
| N/A   24C    P0     51W / 400W           0MiB / 40960MiB |           Disabled   |
+-----+-----+
| 1.  NVIDIA A100-SXM...                  Off           | 00000000:65:00:00 Off |   0%       Default  |
| N/A   25C    P0     51W / 400W           0MiB / 40960MiB |           Disabled   |
+-----+-----+
| 2.  NVIDIA A100-SXM...                  Off           | 00000000:CA:00:00 Off |   0%       Default  |
| N/A   27C    P0     53W / 400W           0MiB / 40960MiB |           Disabled   |
+-----+-----+
| 3.  NVIDIA A100-SXM...                  Off           | 00000000:E3:00:00 Off |   0%       Default  |
| N/A   28C    P0     53W / 400W           0MiB / 40960MiB |           Disabled   |
+-----+-----+
```



LOW IMPACT TRANSITIONING

With the Arc Cloud compatible with today's leading orchestration tools, schedulers, containers and more, transitioning workloads to Arc Compute is a breeze.

CASE STUDY

DATA MACHINES CORP.

Data Machines Corporation is a leading defence contractor based in Ashburn Virginia, who leverages machine learning to provide solutions to various US government agencies.

THE GOAL

Data Machines came to Arc Compute looking to scale up one of its research projects, trying to enable shorter training times at scale for a rapidly deployable solution.

THE SOLUTION

Arc Compute was able to provide Data Machines with a rapidly scalable solution that decreased workload training times using half the hardware that they were previously using. This was made possible by Arc Compute's proprietary GPU software which enabled better performance and utilization by allocating GPU cores and VRAM between Data Machines researcher's virtual machines. This saved them a tremendous amount of time and about 75% over what they were paying with AWS.

"ARC COMPUTE'S SERVICES HAVE ALLOWED US TO ADVANCE OUR RESEARCH AT A FASTER RATE, THANKS TO THE BOOST IN PERFORMANCE AND COST SAVINGS WE'VE EXPERIENCED."

Dr Brian Dennis, PhD - CTO, Data Machines Corporation



CLOUD VS ON-PREM

On-Premises

9%
software licenses

Customization & implementation

Hardware

IT personnel

Maintenance

Training




Ongoing costs

- Apply filters, patches, upgrade
- Downtime
- Performance tuning
- Rewrite integrations
- Upgrade dependent applications
- Ongoing burden on IT
- Maintain/upgrade hardware
- Maintain/upgrade network
- Maintain/upgrade security
- Maintain/upgrade database

Cloud Computing

68%
subscription fee

Implementation, Customization & training




Ongoing costs

- Subscription fee

Both cloud computing and on-premises IT systems each have their own advantages, along with a range of disadvantages.

Both need regular investment, but one option is light years ahead when it comes to the ability for companies to scale and compete with larger organizations, which is the cloud platform.





ARC COMPUTE