




We Will Win by

Embracing Cloud Architectural Revolution

Smart PBS Scheduling



Nvidia Branded Infiniband EDR

- PCIe3.0 x16 required
- Speed up to 100 GBps

Cutting-edge Cloud GPU Service

- Nc24-series Azure Service
- equipped with 2 node 8 card K80
- 600GB/s memory bandwidth
- 250 Watts TDP
- RDMA Inside

Zen-2 assured AMD Processors

- equipped with 2 node 60 cores 7552 & 1 node 96 7702 vCPU
- up to latest PCIe4.0 x16 Bandwidth
- up to 384GiB Memory
- hexa channel to achieve 126 GB/s

More than Greatest SSDs & File Systems

- Lustre + NFS over raid on Premium SSDs & Ultra SSDs

Not so huge Cost

- Total Cost Master(login node) + execute + storage = $0.5 + 4 \times 2 \times 6 \times 2 + 8 + 0.5 = 29$ USD/hour



We Will Win Endowed with

Strong Software Supports

Operating System

- CycleCloud version 8.0.0 front end
- Nvidia Node for RedHat RHEL 7 Based System
- stable and reproducible
- always applies stable releases without expensive subscription
- friendly to HPC applications with kernel-level support

Package Management

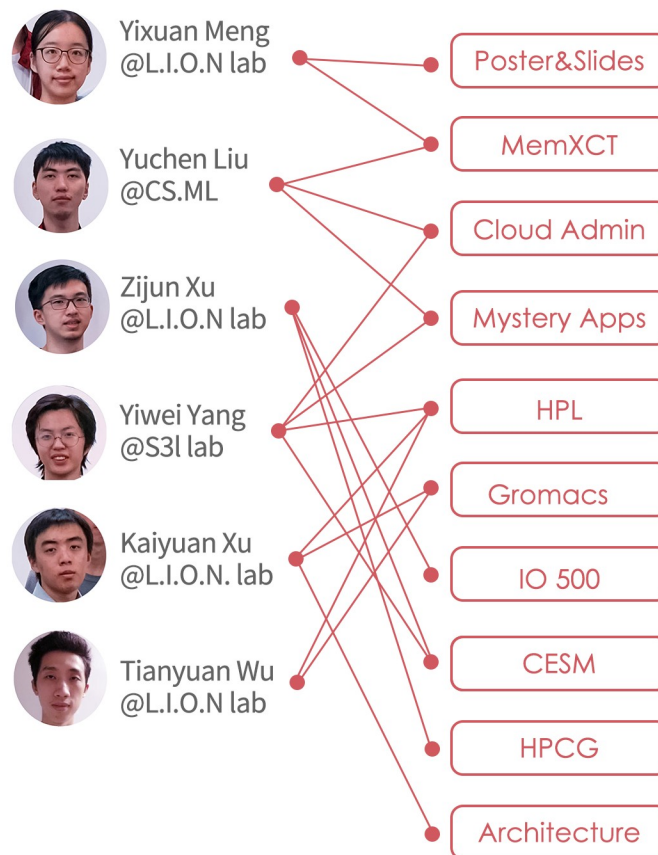
- Azure Spack flexible package management
- Fast to deploy compiler & libraries
- Intel® Parallel Studio Cluster 2019.5 over gcc 7.5.0 & 2018.4 over gcc 6.5.0
- Cuda Aware OpenMPI 3.1.6 with hca support
- Nvidia® Cuda Library & CuDNN 11.1

System Monitoring

- Restful API assured Telegraf & Money API
- Transparent Monitoring and Error recovery
- Grafana Panel & Slack Alert Tailored for Azure

We Will Win with

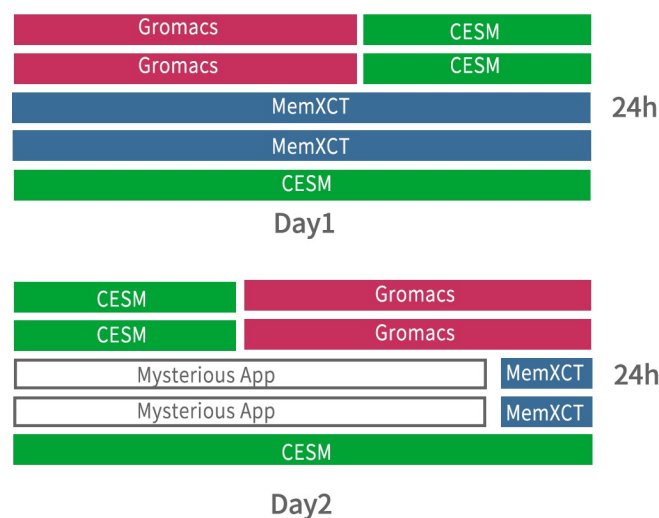
Collaboration among Most Powerful Teammates



GeekPie_HPC backup Members Ziyue Yang, Feiran Qin & Siyuan Zhang - Slack, Website, Cloud Admin

We Will Win for

Sophisticated Tactic (Task scheduling)



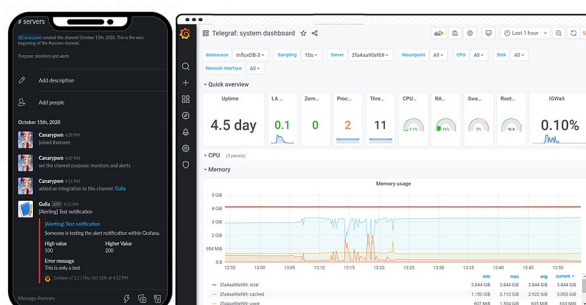
We Will Win for

Strict Panel Control

1. PBS Auto deallocating VMs.

All VM switch on mode & economic mode fast transition.
Inside VM Persistent-state & Computing-state

2. Slack alert and Grafana panel show all parameter.



3. Cloud virtual machine tweak

- Turbo Ponnst : disable by the sys config
- GPU Over clock : utilize computing power if needed.
- Down throttling : if over heat, down clock.

We Will Win over

Tweaked Applications

HPCG & HPL

- Script Driven parameter optimization
- Optimize communication and balance threads to GPUs ratio

IO 500

- Ultra SSD provided by Azure
- Parameter tuning for IOR and MD test

CESM® 2

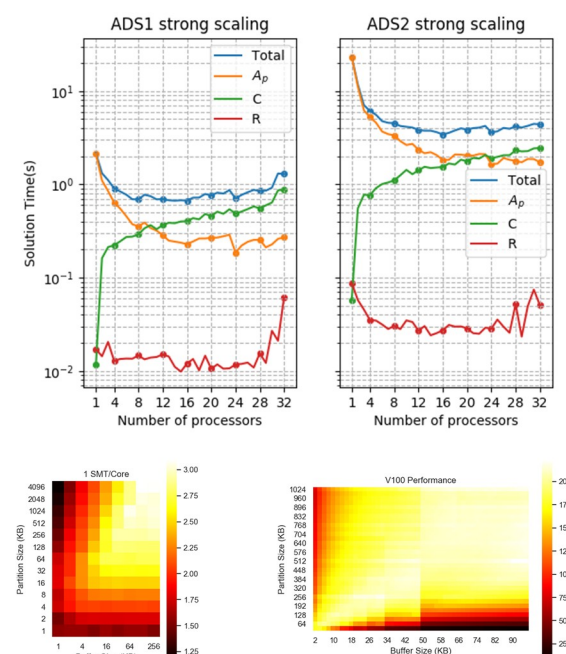
- Tuned PE layout for cloud devices.
- Compiler Option Switched **Maximum** guided by AMD® uProf Profiling Result.
- Fast deployment using pre-compiled **Static Library**.

Gromacs®

- Balanced CPU & GPU **Heterogeneous** Computing.
- **RDMA** speed up for data passing.
- Fine tuned GPU code and parameter over **SIMT** architecture.


MemXCT

- Modify the code best for **MPI** multi-processing.
- Both CPU & GPU code **fully** tested.
- Strong scaling and weak scaling test are reproduced on CPU(line plot)
- Buffer size and block size are tuned on CPU and GPU(heatmap)



Mysterious Apps

- Spack for fast deployment and different compilers
- Flexible Architecture for any Apps



Cost & Performance & Power Consumption & Flexibility Figure

We Will Win with

Support from ShanghaiTech



上海科技大学
ShanghaiTech University

信息科学与技术学院
School of Information Science and Technology