

Good Practices in HPC Management

CARLA Conference

Preguntas...

1. ¿Como haces para llenar tu clúster?
2. Clasifica a los usuarios de tu centro
3. Si tuvieras \$10K, \$100K, \$1M, \$10M para invertir en tu centro, ¿qué harías?
4. Califica entre 1-5 cuál es tu injerencia en el proceso de diseño, licitación, compra y deployment de recursos. Justifica tu respuesta.
5. ¿Porque HPC es diferente de IT?

Good Practices in HPC Management

Experiences at Temple University

Dr. Fernando Posada

Assoc. Research Professor

Temple University

Temple University and HPC

- TU is located in North Philadelphia, PA.
- CST is the largest HPC user.
- Currently around 270 active users.
- Molecular Mechanics, Computational Chemistry, Bioinformatics, and Machine Learning workloads.



The HPC Team



Axel
Kohlmeyer



Richard Berger
(LANL)



Fernando
Posada

Systems Operated by the HPC Team

Resource	Name	Nodes	Cores	Memory	GPUs	Storage
Interactive	Compute	1	88	1.5 T	No	ZFS
	Machine Learning	1	16	512 G	4 NVIDIA Tesla V100	NFS
Batch-operated	Owl's Nest	254	6,464	47.1 T	32 NVIDIA Tesla P100	GPFS
	ICMS	39	2,016	16 T	30 NVIDIA A100	NFS
	CB2RR	66	1,392	9.1 T	12 NVIDIA A100	ZFS
	EFRC	72	1,440	9.1 T	No	ZFS
Total		433	11,416			



Shared Resources

Main HPC Cluster

- 6,464 CPU Cores
- 1.5 PT Parallel Storage (GPS)
- Nodes with RAM (96GB - 3.0 TB)
- 16 dual Tesla P100 GPU
- ~270 users
- ~200 million SU

Compute Interactive

- 88 CPU Cores
- 500 TB Storage (ZFS)
- 1.5 TB of RAM
- 8 batch-operated nodes
- ~243 users
- ~1 million SU

Machine Learning

- DGX-1
 - 40 CPU Cores
 - 521 GB RAM
 - 8x Tesla V100 GPU
- GPU Server
 - 16 CPU Cores
 - 512 GB RAM
 - 4x Tesla V100 GPU

Systems Facing Good Practices

Procurement

- Where is the funding coming from?
- The HPC Cluster design should match research's needs (not researcher's).
- Contact vendors directly, not resellers! (Have a plan before contacting).
- Review quotes line per line.
- **Do not forget datacenter limits and specifications.**

Systems Facing Good Practices

Deployment and Management

- More anticipation, less reaction.
- Use tools (like Ansible and the scheduler) to monitor the node's health.
- Anticipation reduces the number of reported incidents!
- Use each crash event to improve your health check system.
- Keep the scratch area clean (purge regularly).
- Never perform updates or massive hardware changes while on-production.
- Test, test and TEST!

Research Facing Good Practices

HPC Outreach

- Via training.
- Office hours.
- Website-hosted documentation.
- Proactively look for opportunities to facilitate.
- Encourage “oral tradition” or “residual knowledge” within groups.
- Again! Use the reported incidents to improve your material.

Software Facing Good Practices


- “We” support only one MPI Library.
- Encourage software compilation in \$HOME. (Documentation).
- Encourage the use of Singularity containers!



Networking!

- Join a research computing professionals' network!
- Campus Champions model.
- Americas HPC Collaboration.
 - <https://americas-hpc-collab.slack.com>

Thanks!



HPC
High-Performance Computing



CARLA 2023

LATIN AMERICA HIGH PERFORMANCE
COMPUTING CONFERENCE
Cartagena de Indias, Colombia September 18-22



College of Science
and Technology