सुस्वागतम्

நல்வரவு

ସୁସ୍ୱାଗତମ

సుస్వాగతం

ಸುಸ್ವಾಗತ

സുസ്വാഗതം

সুস্বাগতম

ਸੁਆਗਤਮ

સુસ્વાગતમ્

خوش آمدید

Agenda

- PARAM-Ganga Architecture
- Technical Specification of PARAM-Ganga
- PARAM Ganga Software Stack
- How to access Param-Ganga.
- Ticketing Tool and Helpdesk

PARAM GANGA

PARAM Ganga Architecture Diagram

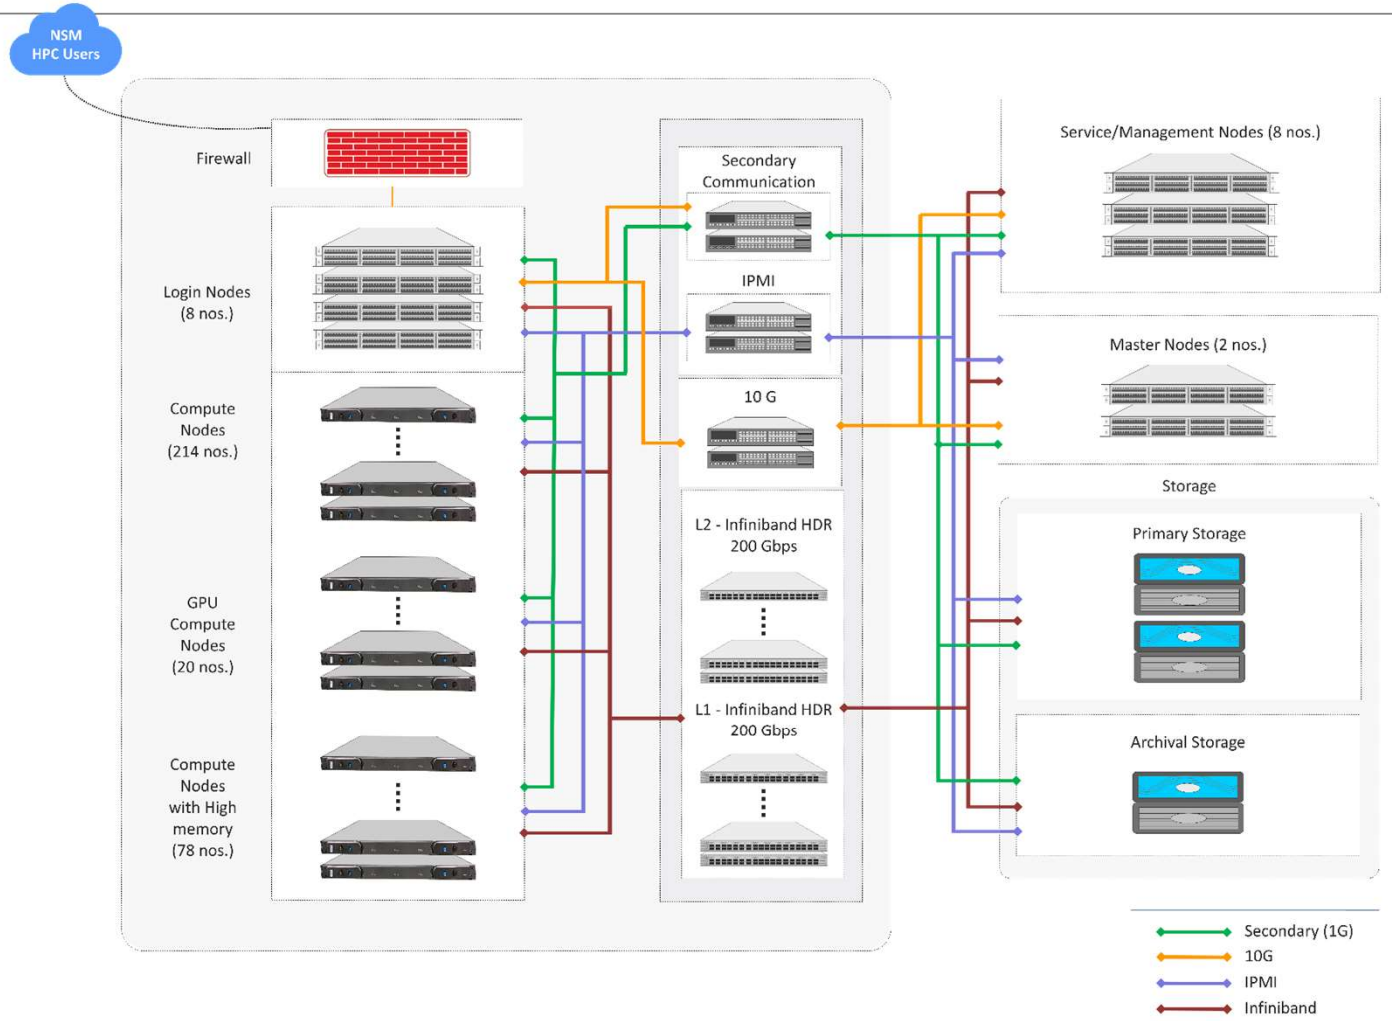

PARAM Ganga - System Configuration

System Specifications

Theoretical Peak Floating-point Performance Total (Rpeak)	1.67 PFLOPS
Base Specifications (Compute Nodes)	2 X Intel Xeon Cascadelake 8268, 24 Cores, 2.9 GHz, Processors per node, 192 GB Memory, 480 GB SSD
Master/Service/Login Nodes	20 nos.
CPU only Compute Nodes (Memory)	214 nos. (192GB)
GPU Compute Nodes (Memory)	20 (192 GB)
High Memory Compute Nodes	78 nos. (768GB)
Total Memory	104.832 TB
Interconnect	Primary: 100Gbps Mellanox Infiniband Interconnect network 100% non blocking, fat tree topology Secondary: 10G/1G Ethernet Network Management network: 1G Ethernet

CPU Only Compute Nodes

- 214 Nodes
- 10272 Cores
- Compute power of Rpeak 953.2 TFLOPS
- Each Node with
  - 2 X Intel Xeon Cascadelake 8268, 24 cores, 2.9 GHz, processors
  - 192 GB memory
  - 480 GB SSD

GPU Compute Nodes

- 20 Nodes
- 800 CPU Cores
- 204800 CUDA Cores
- Rpeak CPU 64 TFLOPS + GPU 312 TF
- Each Node with
  - 2 X Intel Xeon Cascadelake 6248, 20 cores, 2.5 GHz, processors
  - 192 GB Memory
  - 2 x NVIDIA V100 SXM2 GPU Cards
  - 480 GB SSD

High Memory Compute Nodes

- 78 Nodes
- 3744 Cores
- Compute power of Rpeak 347.33 TFLOPS
- Each Node with
  - 2 X Intel Xeon Cascadelake 8268, 24 cores, 2.9 GHz, processors
  - 768 GB Memory
  - 480 GB SSD

Technical Specifications

CPU only Compute Nodes

- 214 Nodes
- 10272 Cores
- Compute power of Rpeak 953.156 TFLOPS
- Each Node with
  - 2* Intel Xeon Platinum 8268, 24 cores, 2.9GHz, processors
  - 192 GB memory
  - 480 GB SSD

High Memory Compute Nodes

- 78 Nodes
- 3744 Cores
- Compute power of Rpeak 347.412 TFLOPS
- Each Node with
  - 2* Intel Xeon Platinum 8268, 24 cores, 2.9GHz, processors
  - 768 GB memory
  - 480 GB SSD

GPU Compute Nodes

- 20 Nodes
- 800 CPU cores
- 204800 GPU Cores
- Compute power of Rpeak 64 TFLOPS + 280 TF =344
- Each Node with
  - 2* Intel Xeon Gold G-6248, 20 cores, 2.5 GHz, processors
  - 192 GB Memory
  - 480 GB SSD
  - 2xNvidia V100 SXM2 GPU cards each with 5120 CUDA cores

Software Stack

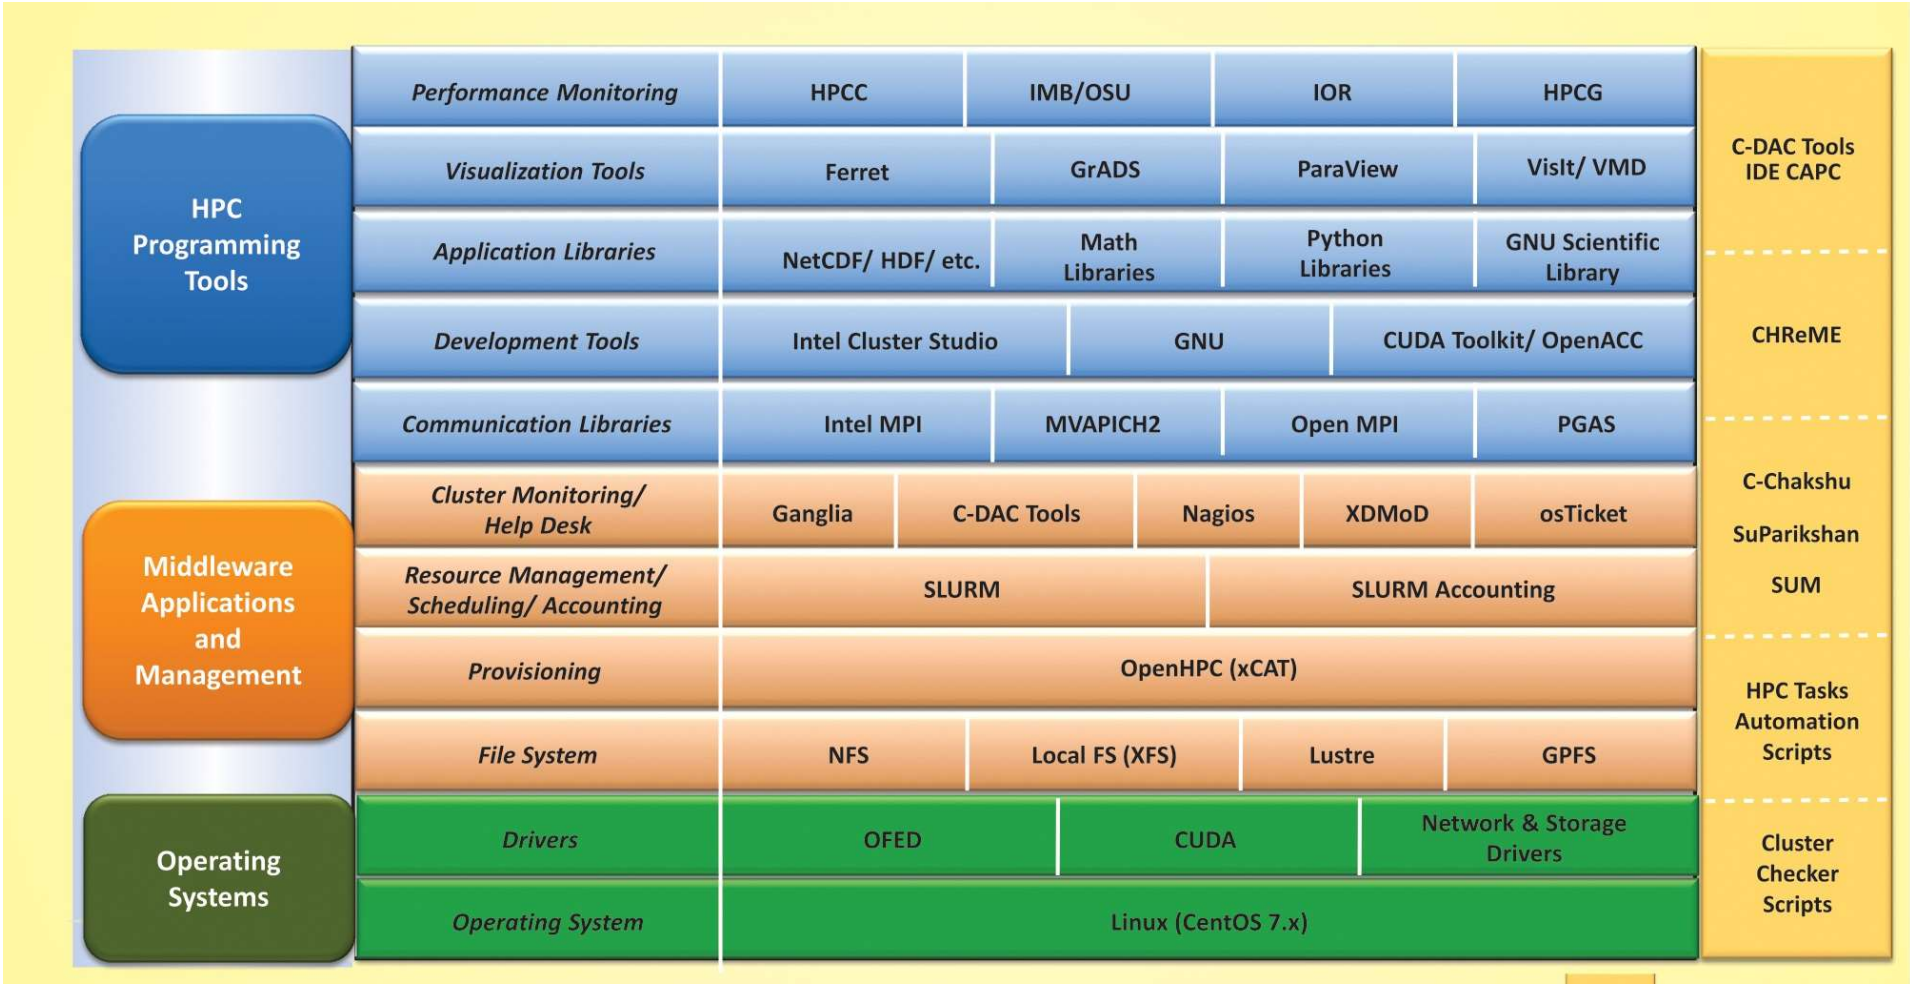

Software Components

Operating System –

- HPC clusters generally are build with Linux operating system as a base OS (Centos7.6)
- It includes all the device drivers for the H/W connected to each node.

Cluster Manager/Orchestrator

- Tools in this category builds a centralized architecture where a controller node builds and manages the cluster.
- xCAT – Is an open-source cluster Manager, Maintained by community, is the widely used tool for HPC as well as cloud clusters.
- It provides flexibility to handle objects within the cluster with its easy manageable methods
- It provide methods to deploy nodes with a very light weight stateless images.

Software Components

Resource Manager (SLURM)

- As there are a lot of resource within a cluster like : CPU-Cores, Memory banks, GPU accelerator cards managing which becomes a tedious task for a user and a system administrator.
- Resource manager with in “slurm” tool helps to manage and represent resources to the users in a simplest way.

Job Schedular (SLURM)

- A scheduler checks the available resources within a cluster and manages which jobs run where and when.
- Allocating resources to each users for optimal utilization of system resources.
- Provides multiple algorithm, which provides different ways to initiate jobs on the resources.
- BACKFILL scheduling is the widely used and the most efficient algorithm.
- Provider batch jobs as well as Interactive jobs submission methods.

NSM Clusters – Applications, Tools, Programming Models

HPC Applications	**Bio-informatics**	MUMmer, HMMER, MEME, PHYLIP, mpiBLAST, ClustalW
	Molecular Dynamics	NAMD (CPU & GPU), LAMMPS(CPU & GPU), GROMACS
	Material Modeling, Quantum Chemistry	Quantum-Espresso, Abinit, CP2K, NWChem,
	CFD	OpenFOAM, FDS, SU2
	Weather, Ocean, Climate	WRF, RegCM, MOM, ROMS
	Disaster Management	ANUGA Hydro

Visualization Programs	GrADS, ParaView, VisIt, VMD
Dependency Libraries	NetCDF, PNETCDF, Jasper, HDF5, Tcl, Boost, FFTW
Programming Models	MPI, OpenMP, OpenACC, CUDA, PGAS, Pthreads

Additional applications, libraries, tools on different NSM systems as per requirements from users of respective systems

AI/ ML/ DL Tools/ Technologies	**DL Frame work:** TensorFlow , keras, theano, pytorch, scikit-learn,scipy, cuDNN
	Data Science: Numpy , RAPIDS
	Distributed DL Framework: TensorFlow with Horovod
	Container Technology: enroot
	JupyterHub: DL application development platforms and web based IDE

Accessing the cluster

Login Environment

- The cluster can be accessed through 8 general login nodes.
- The login nodes is primary gateway to the rest of the cluster.
- All libraries, compilers, preinstalled applications, user installed application are available over login nodes.

Remote Login

- You may access login node through ssh.
- Using SSH in Windows (Putty,Moab-xterm,etc).
- Using SSH in Linux via terminal.
- For example, to connect to the PARAM Ganga Login Node, with the username.
- **Ex**: ssh –p 4422 testuseruser@paramganga.iitr.ac.in

Access Policy

- **Access to Login nodes are in Round-Robin Mode.**
- **Users are not allowed to access Master/Management Nodes**

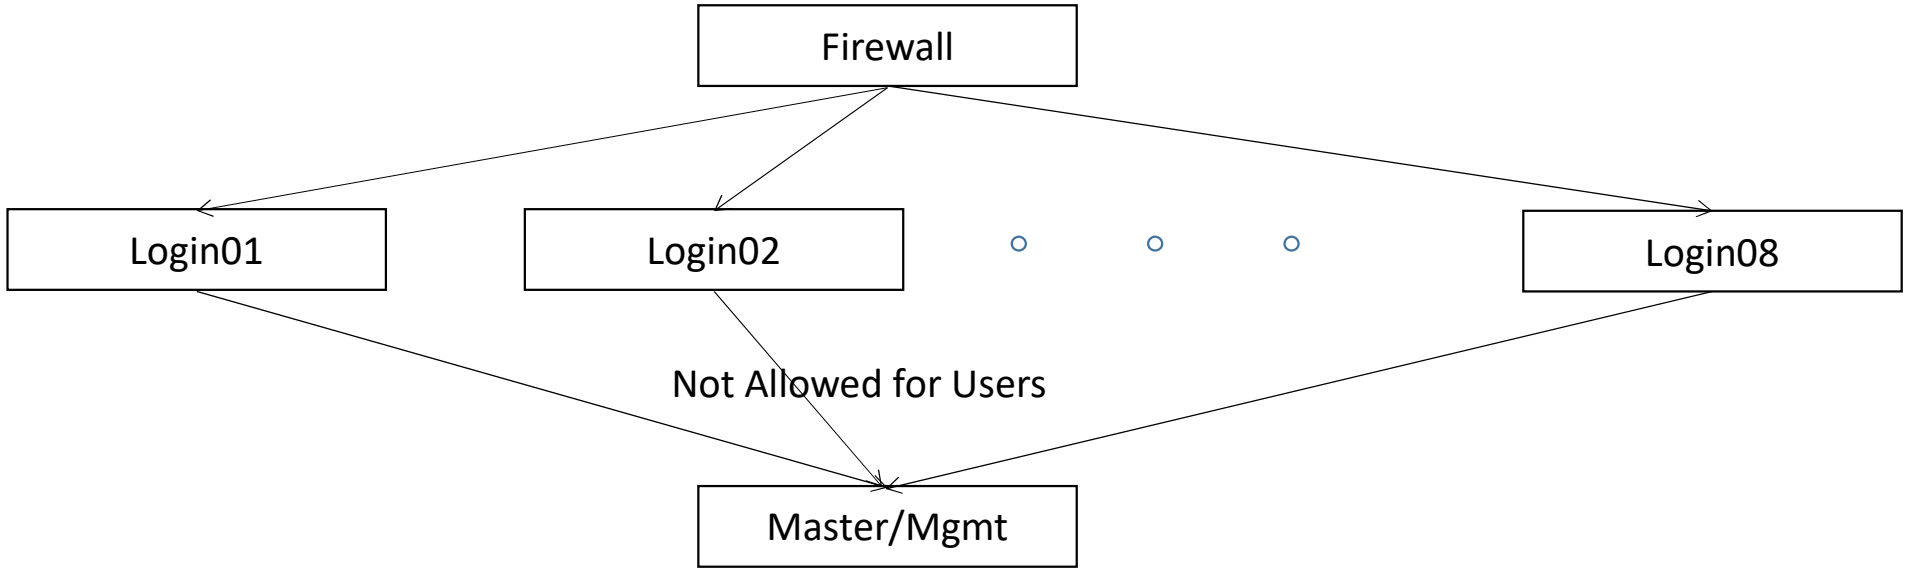

How to Access PARAM Ganga ?

- **If you are using windows you can access via(SSH Clients):**
  - MobaXterm
  - Putty, etc
- **Within IIT Roorkee Campus:**

  ssh username@paramganga.iitr.ac.in
- **Outside IIT Roorkee Campus:**

  ssh username@paramganga.iitr.ac.in -p 4422

Ticketing Tool

- **A Support Portal will be created for Assisting the Users.**

  https://paramganga.iitr.ac.in/support

National Supercomputing Mission

Centre for Development of Advanced Computing

General FAQ

How to get account on HPC cluster ?

- Get 'User Account Creation Form'
- Fill the relevant details.
- Get the signatures of your Head of the Department and the 'approving authority'.
- You will receive an Email in your official Email ID intimating the creation of your account

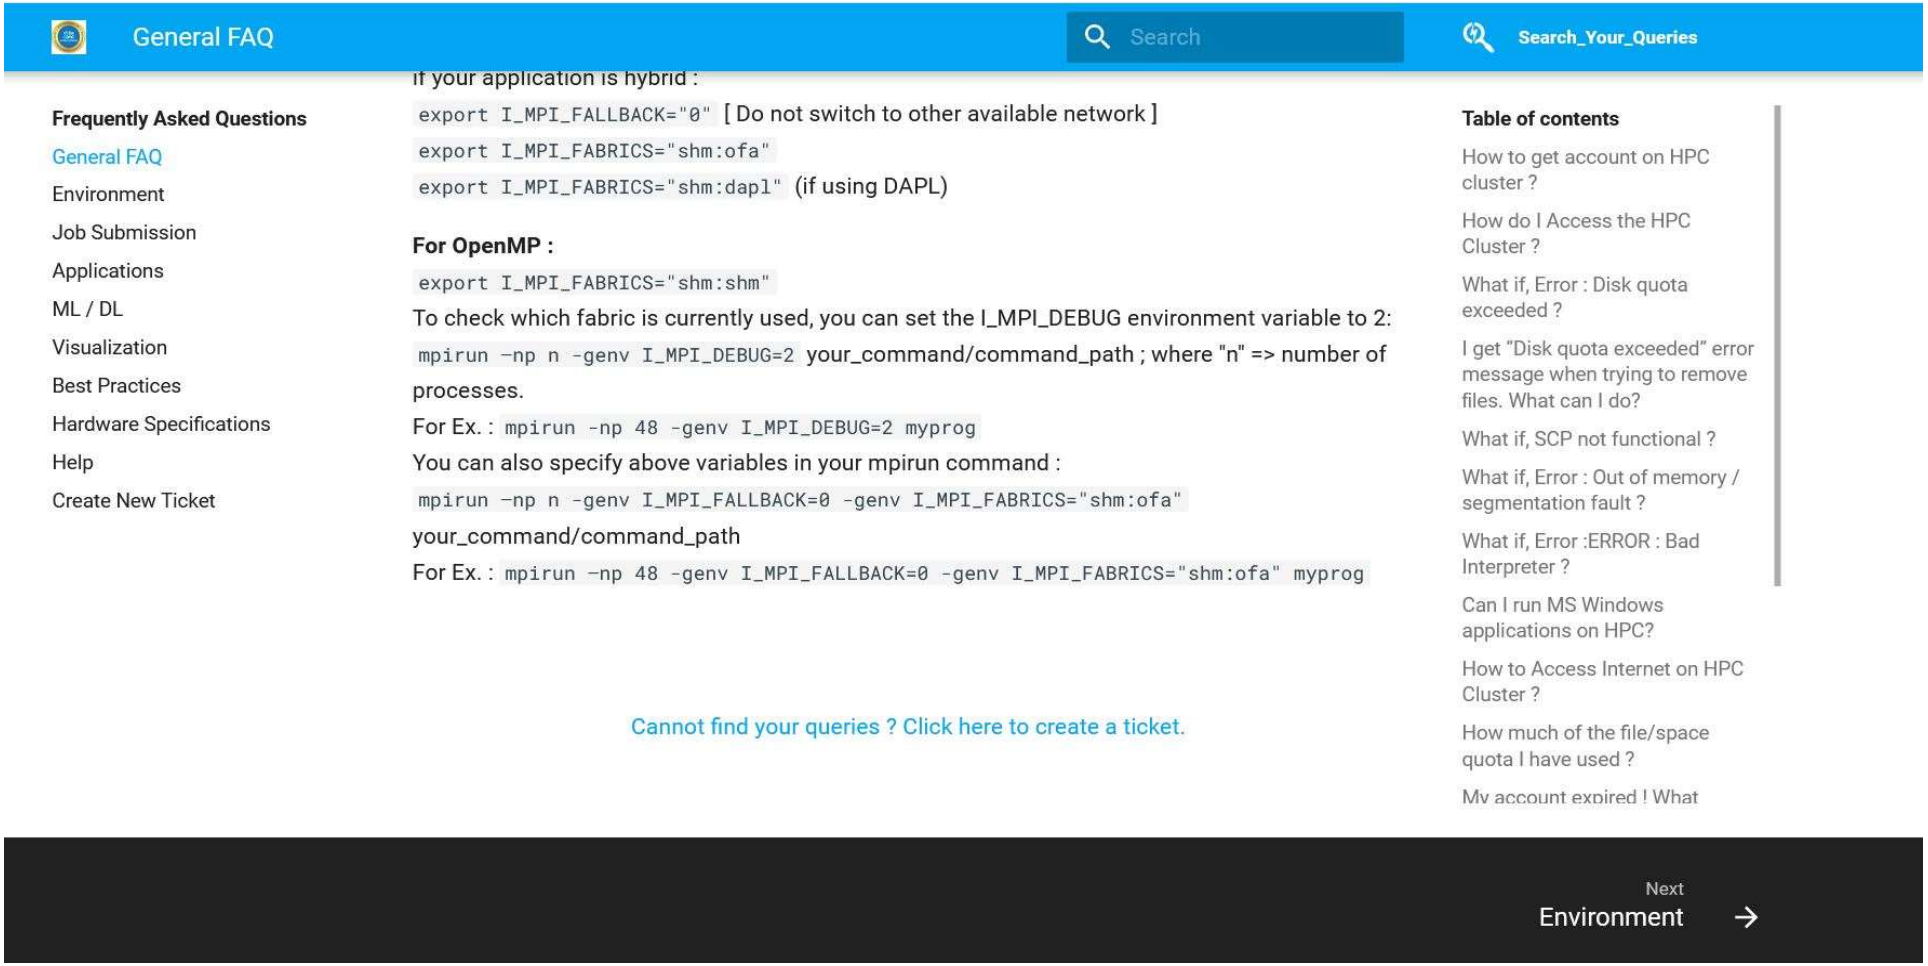
General FAQ
Search
Search_Your_Queries
Frequently Asked Questions
General FAQ
Environment
Job Submission
Applications
ML / DL
Visualization
Best Practices
Hardware Specifications
Help
Create New Ticket
if your application is hybrid :
export I_MPI_FALLBACK="0" [Do not switch to other available network]
export I_MPI_FABRICS="shm:ofa"
export I_MPI_FABRICS="shm:dapl" (if using DAPL)
For OpenMP :
export I_MPI_FABRICS="shm:shm"
To check which fabric is currently used, you can set the I_MPI_DEBUG environment variable to 2:
mpirun –np n -genv I_MPI_DEBUG=2 your_command/command_path ; where "n" => number of processes.
For Ex. : mpirun -np 48 -genv I_MPI_DEBUG=2 myprog
You can also specify above variables in your mpirun command :
mpirun –np n -genv I_MPI_FALLBACK=0 -genv I_MPI_FABRICS="shm:ofa"
your_command/command_path
For Ex. : mpirun –np 48 -genv I_MPI_FALLBACK=0 -genv I_MPI_FABRICS="shm:ofa" myprog
Cannot find your queries ? Click here to create a ticket.
Table of contents
How to get account on HPC cluster ?
How do I Access the HPC Cluster ?
What if, Error : Disk quota exceeded ?
I get "Disk quota exceeded" error message when trying to remove files. What can I do?
What if, SCP not functional ?
What if, Error : Out of memory / segmentation fault ?
What if, Error :ERROR : Bad Interpreter ?
Can I run MS Windows applications on HPC?
How to Access Internet on HPC Cluster ?
How much of the file/space quota I have used ?
My account expired ! What
Next
Environment

Thank You !

darbari@cdac.in;
www.cdac.in

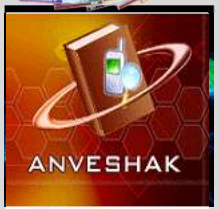

One Vision.
One Goal..
Advanced Computing
for Human Advancement...

धन्यवाद आभार આભાર ಧನ್ಯವಾದ ধন্যবাদ நன்றி നന്ദി ధన్యవాదములు