

Glossary

Units of Measurement:

1TB (Terabyte) = 1,000GB (Gigabyte)

1GB (Gigabyte) = 1,000MB (Megabyte)

Batch processing

A way of running jobs on a computer. Instead of executing them straight away, the jobs will join a queue. A scheduler will determine when they can run, based on resource availability. The user has no way to directly interact with a batch job.

Core

an alternative term for a central processing unit (CPU), especially used when several are present in a single processor chip.

Embarrassingly parallel

This term describes a problem that is able to be broken down into many small parts, each of which runs on a single node (or even a single core), entirely independently of the other parts. interactive processing: The opposite of batch processing. Instead of going through a scheduler, the job will be run by the operating system, and the user can often interact with it directly. Most user-space applications on a personal computer are run this way.

Job

A piece of work that a user requests of a computer. A job consists of one or more tasks, which may be run sequentially, in parallel, or a mix of both.

MPI

An abbreviation of Message Passing Interface, a way of running the same job across multiple nodes (or maybe the same node). An MPI job will spawn multiple tasks that exchange control signals and data by "messages", hence the name.

Node

A single computing device within a cluster. A node is the equivalent of a personal computer or laptop. Every node in a NeSI cluster has many cores.

Scheduler

A piece of software that runs on a shared cluster. Its role is to determine in what order queued jobs should run.

Task

A single running instance of a computer program. Each task consists of at least one thread.

Thread

A sequence of instructions that can be managed independently by the operating system. Threads are grouped into tasks. Ideally, each thread should have its own core to run on.