

# Introduction to High-Performance Computing

# What is High Performance Computing?


- **There is no clear definition**
  - Computing on high performance computers
  - Solving problems / doing research using computer modeling, simulation and analysis
  - Engineering design using computer modeling, simulation and analysis
- **My understanding**
  - A huge number of computational and memory requirements
  - Cannot be afforded by a PC efficiently
  - Speeds and feeds are the keywords
- **Who uses High-Performance Computing**
  - Research institutes, universities and government labs
    - Weather and climate research, bioscience, energy, military etc.
  - Engineering design: more or less every product we use
    - Automotive, aerospace, oil and gas explorations, digital media, financial simulation
    - Mechanical simulation, package designs, silicon manufacturing etc.
- **Similar concepts**
  - Parallel computing: computing on parallel computers
  - Super computing: computing on world 500 fastest supercomputers

- **Case 1: Complete a time-consuming operation in less time**
  - I am an automotive engineer
  - I need to design a new car that consumes less gasoline
  - I'd rather have the design completed in 6 months than in 2 years
  - I want to test my design using computer simulations rather than building very expensive prototypes and crashing them
- **Case 2: Complete an operation under a tight deadline**
  - I work for a weather prediction agency
  - I am getting input from weather stations/sensors
  - I'd like to predict tomorrow's forecast today
- **Case 3: Perform a high number of operations per seconds**
  - I am an engineer at Amazon.com
  - My Web server gets 1,000 hits per seconds
  - I'd like my web server and databases to handle 1,000 transactions per seconds so that customers do not experience bad delays

- **High-performance computing is fast computing**
  - Computations in parallel over lots of compute elements (CPU, GPU)
  - Very fast network to connect between the compute elements
- **Hardware**
  - Computer Architecture
    - Vector Computers, MPP, SMP, Distributed Systems, Clusters
  - Network Connections
    - InfiniBand, Ethernet, Proprietary (Myrinet, Quadrics, Cray-SeaStar etc.)
- **Software**
  - Programming models
    - MPI (Message Passing Interface), SHMEM (Shared Memory), PGAS, etc.
  - Applications
    - Open source, commercial

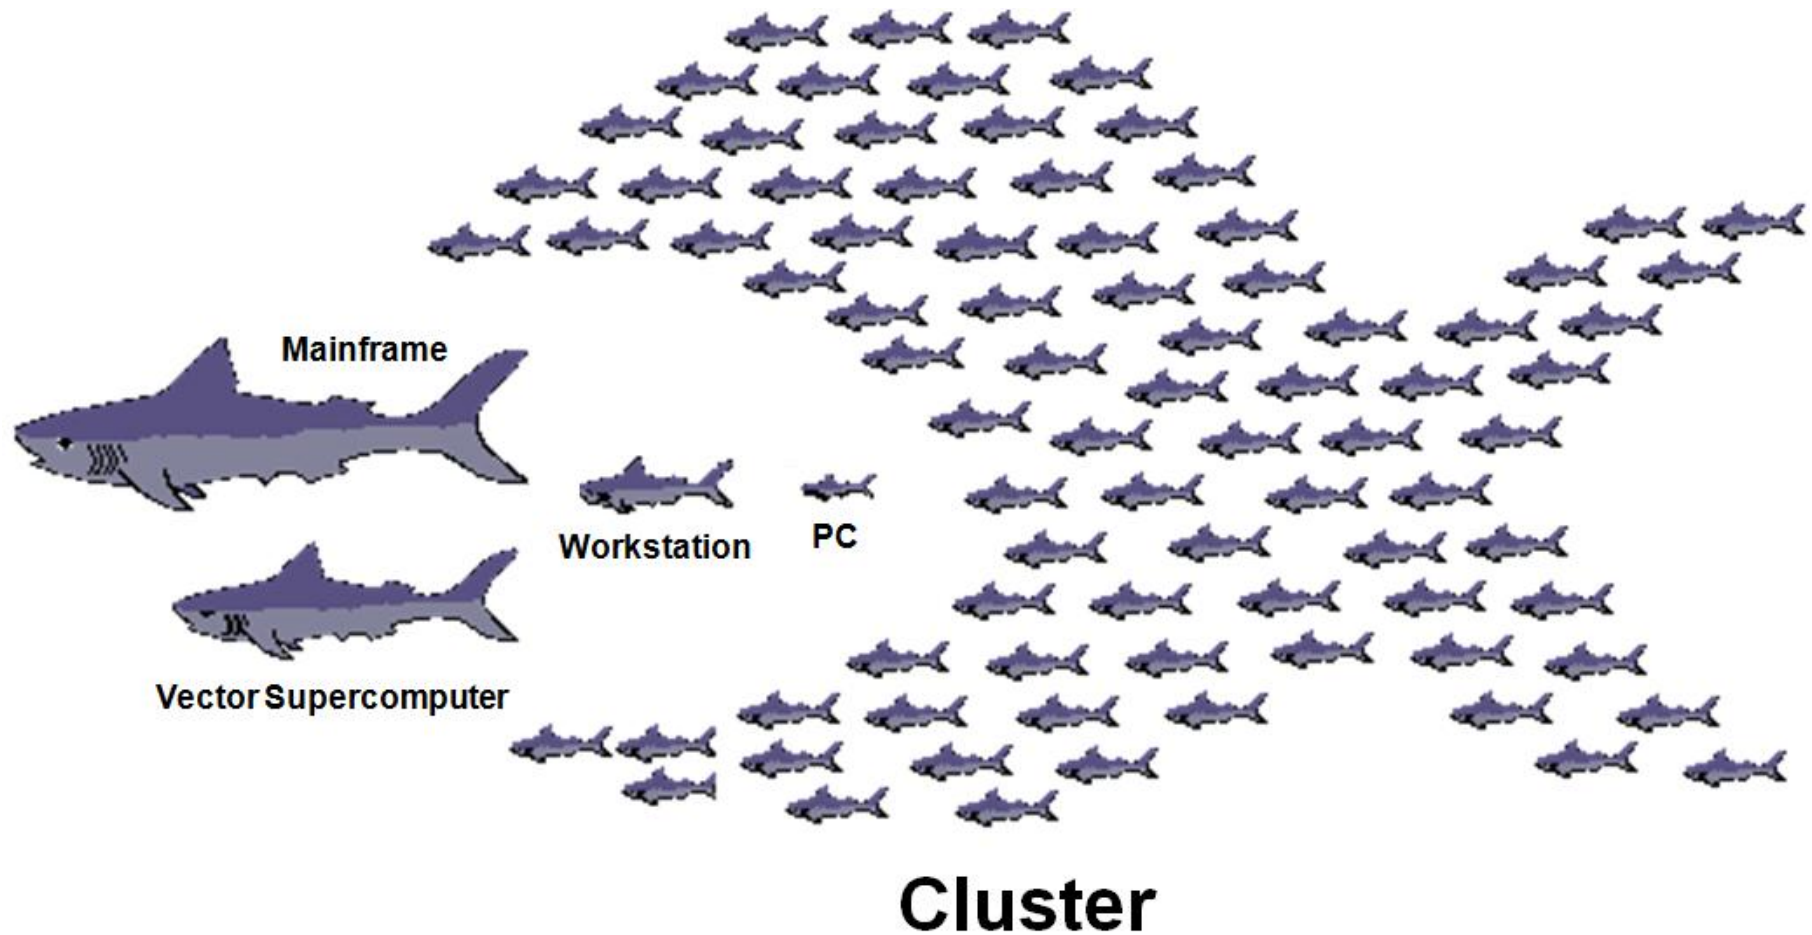
- **Vector Computers (VC) - proprietary system**
  - Provided the breakthrough needed for the emergence of computational science, but they were only a partial answer
- **Massively Parallel Processors (MPP) - proprietary systems**
  - High cost and a low performance/price ratio.
- **Symmetric Multiprocessors (SMP)**
  - Suffers from scalability
- **Distributed Systems**
  - Difficult to use and hard to extract parallel performance
- **Clusters – commodity and highly popular**
  - High Performance Computing - Commodity Supercomputing
  - High Availability Computing - Mission Critical Applications

Architecture Share Over Time  
1993-2010



Clusters have become the most used HPC system architecture

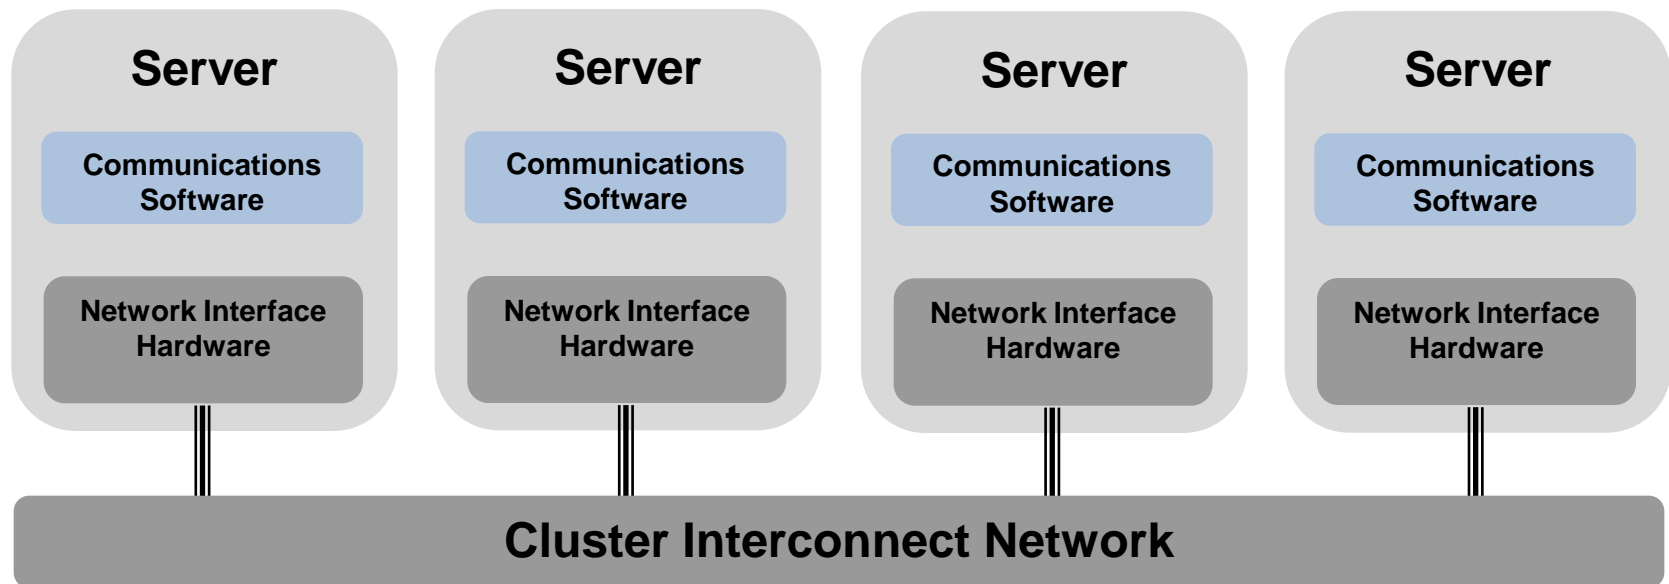
More than 80% of Top500 systems are clusters



**Parallel Computing on a Large Number of Servers is More Efficient than using Specialized Systems**

- **Since the 1990s, there has been an increasing trend to move away from expensive /specialized proprietary parallel supercomputers to clusters of computers**
  - From specialized supercomputers to cost effective, general purpose systems
- **So What's So Different about Clusters?**
  - Commodity, standard, affordable, cost effective, scalable and reliable architecture


## Cluster Architecture






- **Commoditization/standardization** are the clustering and interconnect driving forces
- **InfiniBand and Ethernet** are the most used interconnect solutions for HPC systems

### TOP500 Clustering Interconnect Trends






## Top100 Systems



## Product Development Process



- **From concept to engineering, from design to test and manufacturing, from weather prediction to medical discoveries, our day to day life depends more and more on HPC simulations**
  - Safer products, accurate predictions, research, etc.
- **High-performance compute clusters provide the most efficient, flexible, cost effective HPC environment for any HPC simulation**

# Thank You

[www.hpcadvisorycouncil.com](http://www.hpcadvisorycouncil.com)

[info@hpcadvisorycouncil.com](mailto:info@hpcadvisorycouncil.com)

