Cluster Computing

Prof. Anil Mishra
Department of Chemistry
University of Lucknow
Lucknow

www.anilmishra.name 1

Computing

- *Computing* is the study of systematic processes that describe and transform information:
  - Their theory
  - Analysis
  - Design
  - Efficiency
  - Implementation
  - Application

www.anilmishra.name 2

Computing

- **Computing** is any goal-oriented activity that
  - Requires Computers
  - Benefits from Computers
  - Creating Computers

www.anilmishra.name 3

Earlier Computing Devices

- Slide Rule

- Punch Cards

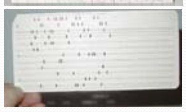

- Digital Calculators

www.anilmishra.name 4

Computing

- Computing includes
  - Designing, developing and building hardware and software systems
  - Processing, structuring, and managing various kinds of information;
  - Doing scientific research on and with computers;
  - Making computer systems behave intelligently

www.anilmishra.name 5

Computing

- What is Process?
  - How we do things
  - How we specify what we do
  - How we specify what kinds of things we're processing

www.anilmishra.name 6

Computing

- What is an Algorithm?
  - Precise description of a process
  - Specifies exactly what is to be done in what order
  - Uses terms that can be completely defined and understood
  - Similar to a recipe
    - Word *algorithm* is derived from name of Persian textbook author al-Khowârizmî (approx. A.D. 825)
    - Word originally referred to process of using arithmetic computations using Arabic numerals

www.anilmishra.name 7

The Goal of Computing

- To develop a *computing machine*
- Specify a precise *algorithm,* represent the data used, and devise a way to translate all of this into a language (encoding) that our computing machine can "understand"

  Then,
- Our *computing machine* can accurately, consistently, and more quickly carry out our computation for us

www.anilmishra.name 8

Purpose of Computing

- What do we want to achieve?
  - To Run Applications Faster
  - To complete the job Sooner
- There are 3 ways to improve performance
  - Work Harder
  - Work Smarter
  - Get Help

www.anilmishra.name 9

Faster Computing

- To achieve faster results use a device....
  - Calculator
  - Computer
  - Any other
- In the Computer Analogy
  - Using faster hardware
  - Optimized algorithms and techniques used to solve computational tasks
  - Multiple computers to solve a particular task

www.anilmishra.name 10

Faster Computing

- Need more computing power
  - Improve the operating speed of processors & other components
    - Constrained by the speed of light, thermodynamic laws, & the high financial costs for processor fabrication
  - Connect multiple processors together & coordinate their computational efforts
    - Parallel computers
    - Allow the sharing of a computational task among multiple processors

www.anilmishra.name 11

Super Computers

- A **supercomputer** is a computer with great speed and memory.
- This kind of computer can do jobs faster than any other computer of its generation.
- They are usually thousands of times faster than ordinary personal computers.
- In 2015, such machines can perform quadrillions of floating point operations per second.

www.anilmishra.name 12

What is Supercomputing About?

- **Size**:
  - Many problems that are interesting to scientists and engineers can't fit on a PC – usually because they need more than a few GB of RAM, or more than a few 100 GB of disk.

What is Supercomputing About?

- **Speed**:
  - Many problems that are interesting to scientists and engineers would take a very very long time to run on a PC: months or even years.
  - A problem that would take a month on a PC might take only a few hours on a supercomputer.

Motivation for Clusters

- Many science and engineering problems today require large amounts of computational resources and cannot be executed in a single machine.
- Large commercial supercomputers are very expensive…
- A lot of computational power is underutilized around the world in machines sitting idle.

Motivating factors

- **High cost of 'traditional' High Performance Computing:**
  - Clustering using Commercial off the Shelf is way cheaper than buying specialized machines for computing.
- **Increased need for High Performance Computing:**
  - As processing power becomes available, applications which require enormous amount of processing, like weather modeling are becoming more common place requiring the high performance computing provided by Clusters

Motivating factors

- Hard to find money to buy expensive systems
- Parallel and Distributed computing are solutions for low cost supercomputing.
- **Network of Workstations** (NOW) - We can use local networked resources to achieve better performance for large scale applications.

Motivating factors

- Just as we learned to fly, not by constructing a machine that flaps its wings like birds, but by applying aerodynamics principles demonstrated by the nature…
- Computational calculations are also modeled after those of biological species.

Computational Power Improvement

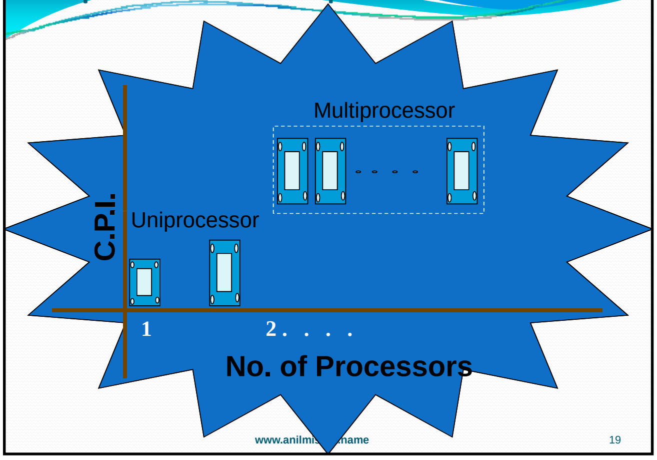

High Performance Computing

- HPC can mean high flop count
  - Per processor
  - Totaled over many processors working on the same problem
  - Totaled over many processors working on related problems
- Can mean faster turnaround time
  - More powerful system
  - Scheduled to first available system(s)
  - Using multiple systems simultaneously

High Performance Computing

- HPC: *any* computational technique that solves a large problem faster than possible using *single, commodity* systems
  - Custom-designed, high-performance processors (e.g. Cray, NEC)
  - Parallel computing
  - Distributed computing
  - Grid computing

Some Definitions

- **Parallel Computing**: single systems with many processors working on the same problem
- **Distributed Computing**: many systems loosely coupled by a scheduler to work on related problems
- **Grid Computing**: many systems tightly coupled by software and networks to work together on single problems or on related problems

What is a Parallel Computer?

- It is the use of multiple computers or processors working together on a common task
- Using a computer that contains multiple processors:
  - Each processor works on its section of the problem
  - Processors are allowed to exchange information with other processors

Parallel computing

- Splitting problem in smaller tasks that are executed concurrently
- Why?
  - Absolute physical limits of hardware components
  - Economical reasons –more complex = more expensive
  - Performance limits –double frequency <> double performance
  - Large applications –demand too much memory & time

Parallel Computing

- Advantages:
  - Increasing speed & optimizing resources utilization
- Disadvantages:
  - Complex programming models –difficult development

Parallel vs. Serial Computers

- Some benefits of parallel computing include:
  - More data points
  - Bigger domains
  - Better spatial resolution
  - More particles
- More time steps
  - Longer runs
  - Better temporal resolution
- Faster execution
  - Faster time to solution
  - More solutions in same time
  - Larger simulations in real time

Types of Parallel Computers

- The simplest and most useful way to classify modern parallel computers is by their memory model:
  - Shared memory
  - Distributed memory

Parallel Architecture 1

Shared Vs Distributed Memory

- **Vector machines**
  - CPU processes multiple data sets
  - shared memory
  - advantages: performance, programming difficulties
  - issues: scalability, price
    - examples: Cray SV, NEC SX, Athlon3/d, Pentium- IV/SSE/SSE2

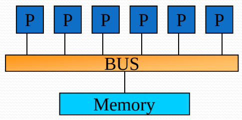

- **Massively parallel processors (MPP)**
  - large number of CPUs
  - distributed memory
  - advantages: scalability, price
  - issues: performance, programming difficulties
    - examples: ConnectionSystemsCM1 i CM2, GAAP (GeometricArrayParallel Processor)

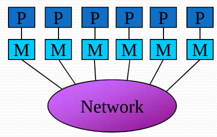

Parallel Architecture 2

Shared Memory: UMA vs. NUMA

- **Uniform Memory Access (UMA)**
  - Two or more processors
  - Shared memory
  - Advantages: price, performance, programming difficulties
  - Issues: scalability
  - examples: UltraSparcII, Alpha ES, Generic Itanium, Opteron, Xeon

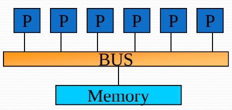

- **Non Uniform Memory Access (NUMA)**
  - Solving SMP's scalability issue
  - Hybrid memory model
  - Advantages: scalability
  - Issues: price, performance, programming difficulties
  - examples: SGI Origin/Altix, Alpha GS, HP Superdome

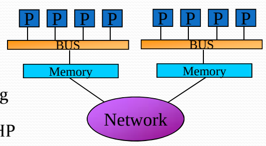

Distributed Computing

- Basic idea:
  - Run scheduler across systems to runs processes on least-used systems first
    - Maximize utilization
    - Minimize turnaround time
  - Have to load executables and input files to selected resource
    - Shared file system
    - File transfers upon resource selection

Distributed vs. Parallel Computing

- Different
  - Distributed computing executes independent (but possibly related) applications on different systems;
  - jobs do not communicate with each other
  - Parallel computing executes a single application across processors, distributing the work and/or data but allowing communication between processes

What is Computer Cluster ?

- **Poor Man's Supercomputer….**
  - Collection of interconnected stand-alone computers working together as a single, integrated computing resource
- Cluster consists of:
  - Nodes
  - Network
  - OS
  - Cluster middleware
- Standard components
  - Avoiding expensive proprietary components

What is Computer Cluster ?

- A **computer cluster** is a group of linked computers, working together closely so that in many respects they form a single computer.
- A cluster is a type of parallel or distributed processing system, which consists of a collection of interconnected stand-alone computers cooperatively working together as a single, integrated computing resource.

Cluster Components

- **A Node**
  - A single or multiprocessor system with memory, I/O facilities, & OS
- **A Cluster**
  - Generally 2 or more computers (nodes) connected together
  - In a single cabinet, or physically separated & connected via a LAN
  - Appear as a single system to users and applications
  - Provide a cost-effective way to gain features and benefits

Cluster Computing

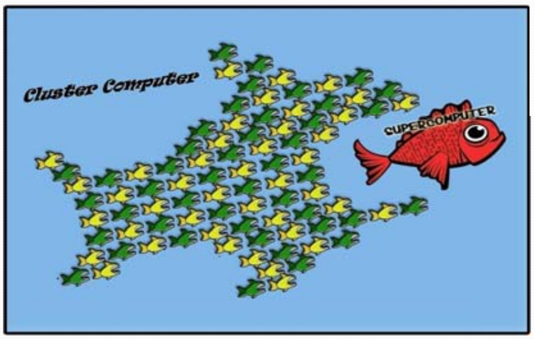

A Typical Cluster

Server Cluster

Cluster Definitions

- Group of computers and servers (connected together) that act like a single system.
- Each system called a Node.
- Node contain one or more Processor , Ram ,Hard disk and LAN card.
- Nodes work in Parallel.
- We can increase performance by adding more Nodes.

A Conventional Computer

- Speed is limited by the rate at which computer can transfer information internally.
  - Ex:PC, Macintosh, Workstations

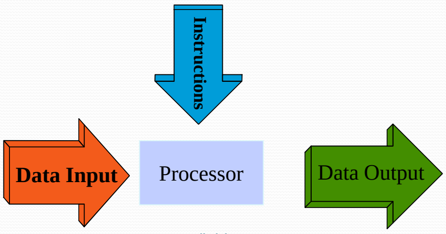

System Performance Improvements

- Multiple CPUs
- Faster clock speed, buses and circuits
- Wider instruction and data paths
- Faster disk access
- More and faster memory

Multiprocessing

- Reasons
  - Increase the processing power of a system
  - Parallel processing
- Types of multiprocessor systems
  - Tightly coupled systems
  - Loosely coupled systems

Cluster classification

- High performance clusters (HPC)
  - Parallel, tightly coupled applications
- High throughput clusters (HTC)
  - Large number of independent tasks
- High availability clusters (HA)
  - Mission critical applications
- Load balancing clusters
  - Web servers, mail servers, ...
- Hybrid clusters
  - Example: HPC+HA

Clusters Classification (I)

- Application Target
  - High Performance (HP) Clusters
    - Grand Challenging Applications
  - High Availability (HA) Clusters
    - Mission Critical applications

Clusters Classification (2)

- Node Ownership
  - Dedicated Clusters
  - Non-dedicated clusters
    - Adaptive parallel computing
    - Communal multiprocessing

Clusters Classification (3)

- Node Hardware
  - Clusters of PCs (CoPs)
    - Piles of PCs (PoPs)
  - Clusters of Workstations (COWs)
  - Clusters of SMPs (CLUMPs)

Clusters Classification (4)

- Node Operating System
  - Linux Clusters (e.g., Beowulf)
  - Solaris Clusters (e.g., Berkeley NOW)
  - NT Clusters (e.g., HPVM)
  - AIX Clusters (e.g., IBM SP2)
  - SCO/Compaq Clusters (Unixware)
  - Digital VMS Clusters
  - HP-UX clusters
  - Microsoft Wolfpack clusters

Clusters Classification (5)

- Node Configuration
  - Homogeneous Clusters
    - All nodes will have similar architectures and run the same OSs
  - Heterogeneous Clusters
    - All nodes will have different architectures and run different OSs

Cluster Classification

High availability clusters (HA) (Linux)
- Mission critical applications
- High-availability clusters (also known as Failover Clusters) are implemented for the purpose of improving the availability of services which the cluster provides.
- provide redundancy
- eliminate single points of failure.

Network Load balancing clusters
- operate by distributing a workload evenly over multiple back end nodes.
- Typically the cluster will be configured with multiple redundant load-balancing front ends.
- all available servers process requests.
- Web servers, mail servers,..

Science Clusters
- Beowulf

Load Balancing Clusters

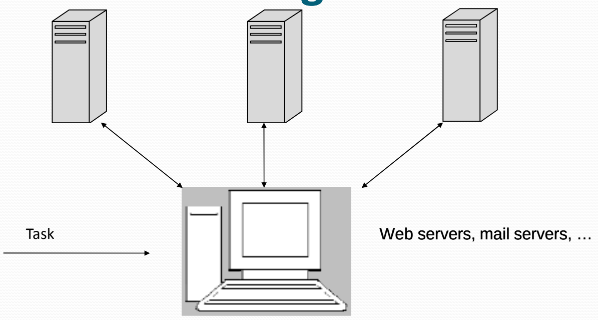

Computational Clusters

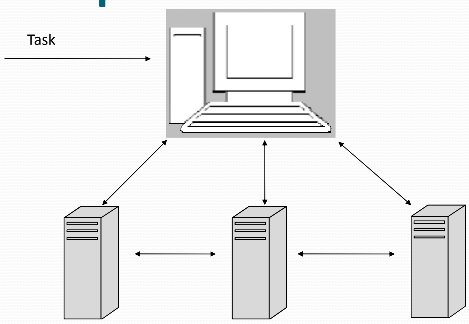

High Availability Clusters

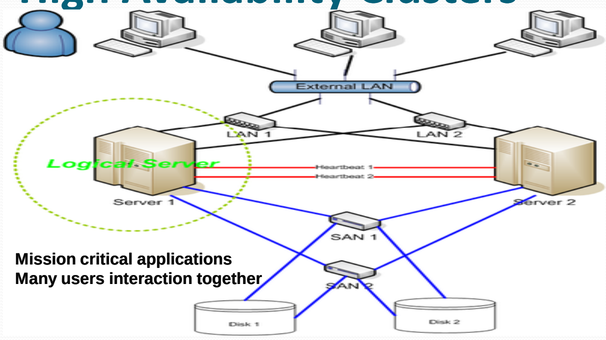

Beowulf Cluster

- A Beowulf Cluster is a computer design that uses parallel processing across multiple computers to create cheap and powerful supercomputers.
- A Beowulf Cluster in practice is usually a collection of generic computers, either stock systems or wholesale parts purchased independently and assembled, connected through an internal network.

Beowulf Cluster

- A cluster has two types of computers, a master computer, and node computers.
- When a large problem or set of data is given to a Beowulf cluster, the master computer first runs a program that breaks the problem into small discrete pieces
- it then sends a piece to each node to compute.
- As nodes finish their tasks, the master computer continually sends more pieces to them until the entire problem has been computed.

Beowulf Cluster

- **Master**: *or service node* or *front node* (used to interact with users and manage the cluster)
- **Nodes** : a group of computers (computing node s)(keyboard, mouse, floppy, video…)
- Communications between nodes on an interconnect network platform (Ethernet, Myrinet….)
- In order for the master and node computers to communicate, some sort message passing control structure is required. MPI,(Message Passing Interface) is the most commonly used such control.

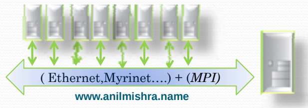

Beowulf Cluster

- Basically, the Beowulf architecture is a multi-computer architecture that is used for parallel computation applications.
- Therefore, Beowulf clusters are primarily meant only for processor-intensive and number crunching applications and definitely not for storage applications.

www.anilmishra.name 55

Typical Beowulf Cluster

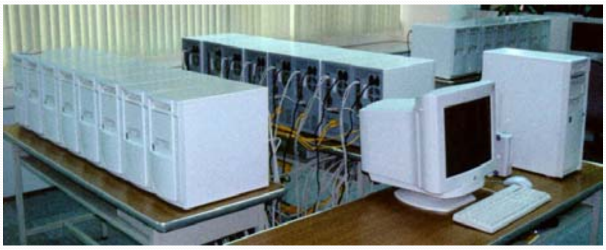

www.anilmishra.name 56

Running Program (Sequential)

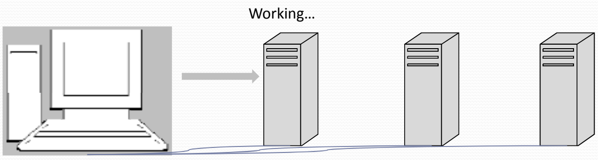

www.anilmishra.name 57

Running Program (Sequential)

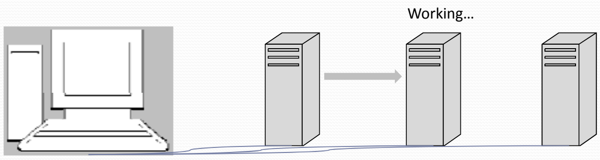

www.anilmishra.name 58

Running Program (Sequential)

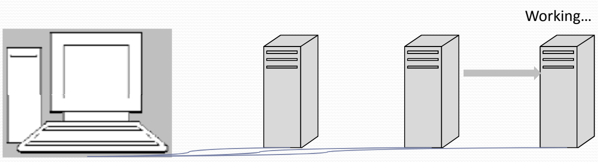

www.anilmishra.name 59

Running Program (Sequential)

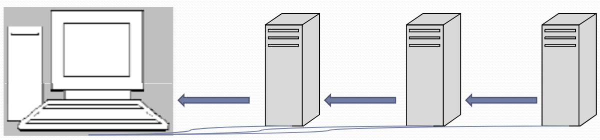

www.anilmishra.name 60

Running Program (Parallel)

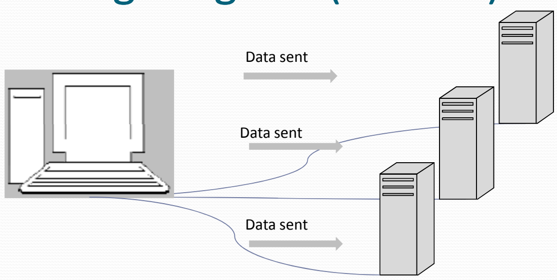

Running Program (Parallel)

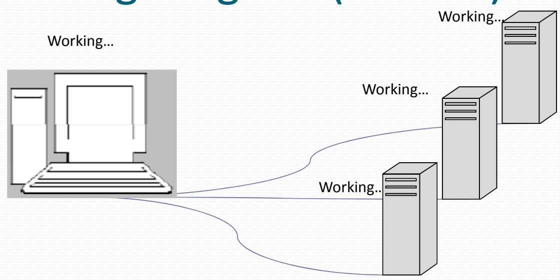

Running Program (Parallel)

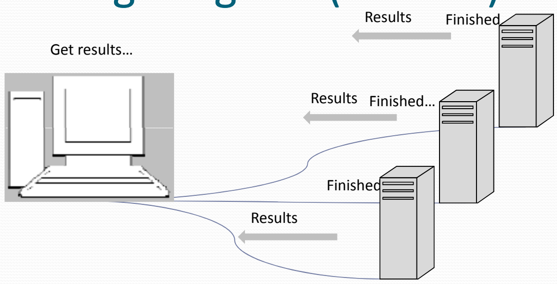

Cluster Architecture

- ***Close Clusters***
  - They hide most of the cluster behind the gateway node. Consequently they need less IP addresses and provide better security. They are good for computing tasks.

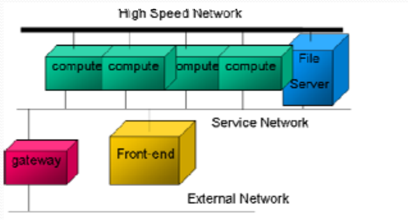

Cluster Architecture

- ***Open Clusters:***
  - All nodes can be seen from outside,and hence they need more IPs, and cause more security concern .But they are more flexible and are used for internet/web/information server task.

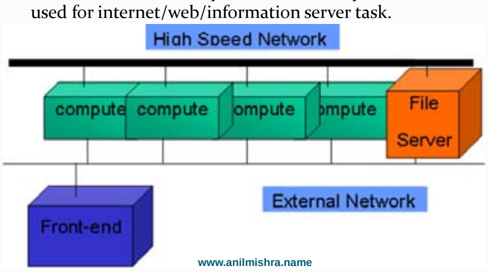

Computer Cluster Facility at University of Lucknow

Central Facility for Computational Research (CFCR)

CFCR Cluster

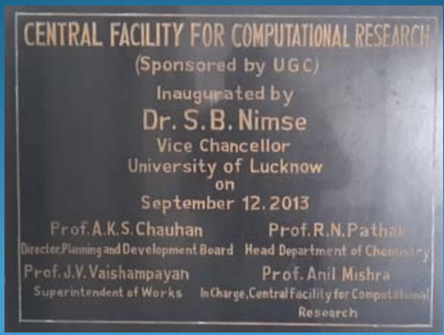

www.anilmishra.name 67

CFCR Cluster

www.anilmishra.name 68

CFCR Cluster

www.anilmishra.name 69

CFCR Cluster

- Name
  - Krishna
- Location
  - Department of Chemistry
- Access Through
  - GUI Web portal
  - Putty
  - VNC viewer
- Availability
  - On University LAN

www.anilmishra.name 70

Accessing the CFCR Cluster

Through Putty

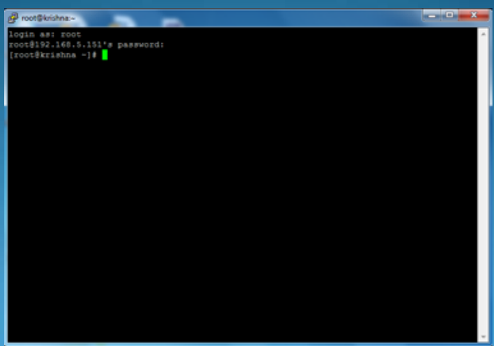

www.anilmishra.name 71

Accessing the CFCR Cluster

Through VNC Viewer

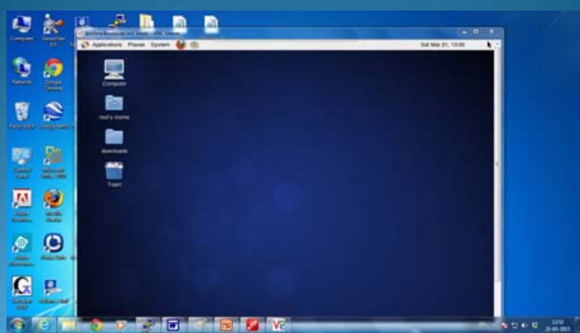

www.anilmishra.name 72

Accessing the CFCR Cluster

Through GUI Portal

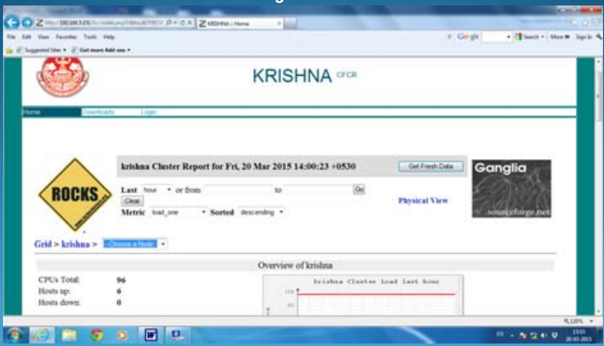

www.anilmishra.name 73

CFCR Cluster

- Total Nodes 6
- Total Processors 96
  - Master Node 1
    - 16 processors
  - Compute Node 5
    - 80 Processors

www.anilmishra.name 74

Configuration of CFCR Cluster

- **Master Node**
  - 2 * 2.4 GHz Intel Xeon Sandy bridge Processor E2665 (8 Cores Per Processor)
  - 4GB/Core Memory ECC Memory
  - 10TB usable HDD after RAID 5.
  - Internal DVD RW drive
  - Infini band QDR ports

www.anilmishra.name 75

Configuration of CFCR Cluster

- Compute Node
  - 2 * 2.4 GHz Intel Xeon Sandy bridge processor E2665 (8 Cores Per Processor)
  - 2GB/Core ECC Memory
  - 2*500GB Disk
  - Infini band QDR ports

www.anilmishra.name 76

Softwares on CFCR Cluster

- Gaussian 09 and Gauss view along with LINDA
  - Site license
- Mathematica
  - 2 Licenses
- MATLAB
  - 5 Licenses
    - Addons (2 Licenses each)
      - Simulink
      - Symbolic Math Toolbox
      - Partial Differential Equation Toolbox

www.anilmishra.name 77

Using the CFCR Cluster

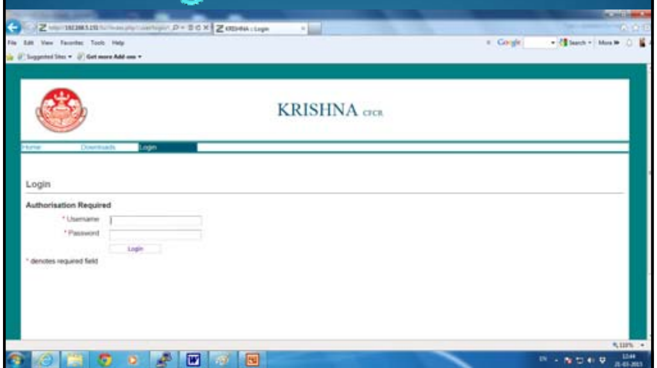

www.anilmishra.name 78

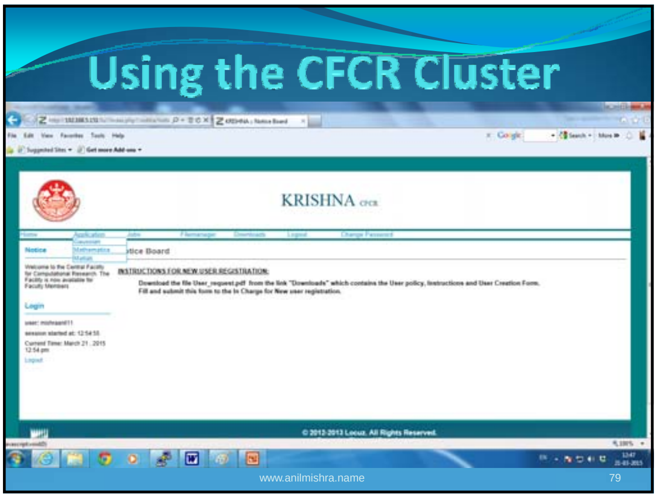
Using the CFCR Cluster
KRISHNA CFCR
www.anilmishra.name
79

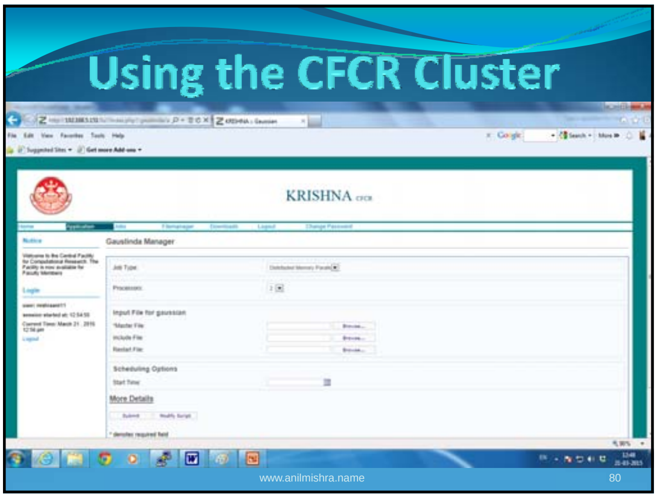
Using the CFCR Cluster
KRISHNA CFCR
www.anilmishra.name
80

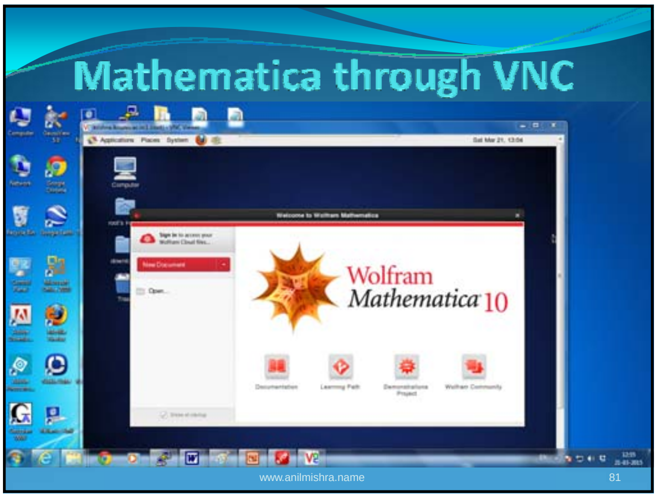
Mathematica through VNC
Wolfram Mathematica 10
www.anilmishra.name
81

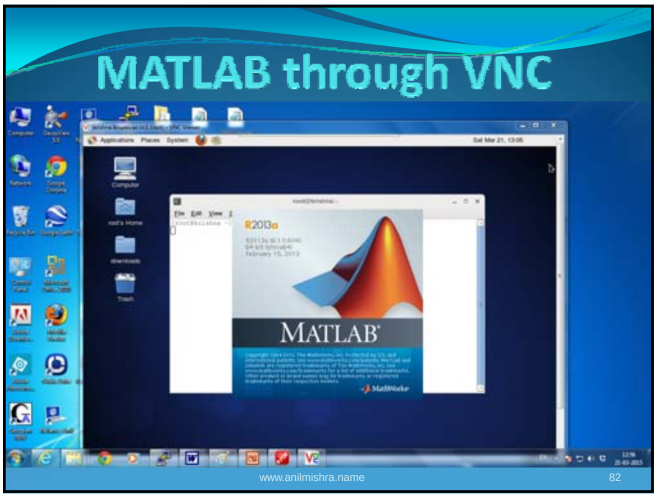
MATLAB through VNC
MATLAB
www.anilmishra.name
82

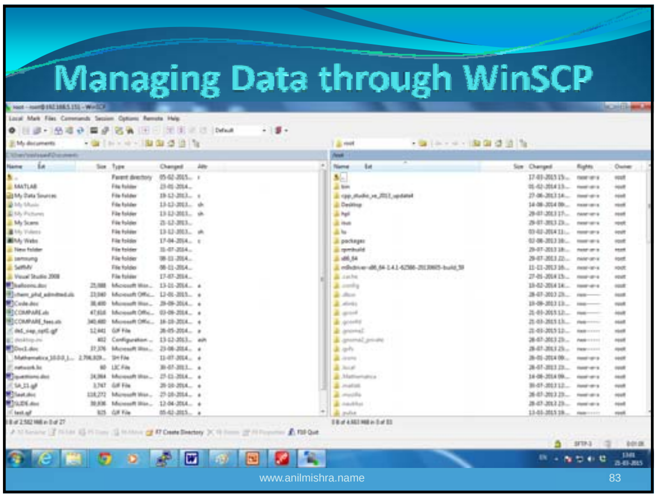
Managing Data through WinSCP
www.anilmishra.name
83

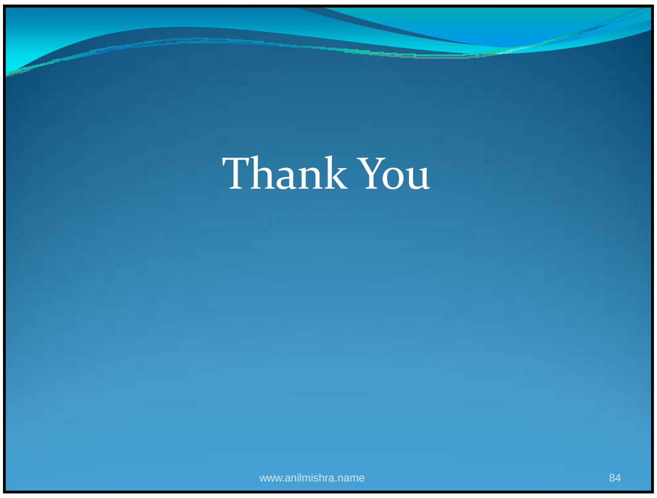
Thank You
www.anilmishra.name
84