



SNS COLLEGE OF TECHNOLOGY

(An Autonomous Institution)

Approved by AICTE, New Delhi, Affiliated to Anna University, Chennai

Accredited by NAAC-UGC with 'A++' Grade (Cycle III) &

Accredited by NBA (B.E - CSE, EEE, ECE, Mech&B.Tech.IT)

COIMBATORE-641 035, TAMIL NADU



19CSE310 GRID AND CLOUD COMPUTING

UNIT I INTRODUCTION

Evolution of Distributed computing: Scalable computing over the Internet – Technologies for network based systems – clusters of cooperative computers - Grid computing Infrastructures – cloud computing - service oriented architecture – Introduction to Grid Architecture and standards – Elements of Grid – Overview of Grid Architecture.

EVOLUTION OF DISTRIBUTED COMPUTING

(a) Define Distributed Computing (2 marks)

1. Describe about Evolution of Distributed computing. (8 marks)

- **Distributed computing** is a field of computer science that studies distributed systems. A distributed system is models in which components located on networked computers communicate and coordinate their actions by passing messages
- The components interact with each other in order to achieve a common goal.
- Three significant characteristics of distributed systems are: concurrency of components, lack of a global clock, and independent failure of components.
- Examples of distributed systems vary from SOA-based system to massively multiplayer online games to peer-to-peer applications.
- A computer program that runs in a distributed system is called a distributed program, and distributed programming is the process of writing such programs.
- There are many alternatives for the message passing mechanism, including pure HTTP, RPC-like connectors and message queues.

HISTORY

- The use of concurrent processes that communicate by message-passing has its roots in operating system architectures studied in the 1960s.
- The first widespread distributed systems were local area networks such as Ethernet, which was invented in the 1970s.
- ARPANET, the predecessor of the Internet, was introduced in the late 1960s, and ARPANET email was invented in the early 1970s. E-mail became the most successful application of ARPANET, and it is probably the earliest example of a large-scale distributed application.
- In addition to ARPANET, and its successor, the Internet, other early worldwide computer networks included Usenet and FidoNet from the 1980s, both of which were used to support distributed discussion systems.
- The study of distributed computing became its own branch of computer science in the late 1970s and early 1980s.
- The first conference in the field, Symposium on Principles of Distributed Computing (PODC),

dates back to 1982, and its European counterpart International Symposium on Distributed Computing (DISC) was first held in 1985.

SCALABLE COMPUTING OVER THE INTERNET

Part-A	<ul style="list-style-type: none"> ➤ List out the types of Computing paradigms. ➤ Highlight the importance of the term “Cloud Computing” (April/May 2017) ➤ Tabulate the difference between High Performance Computing and High Throughput Computing (April/May 2017)
Part-B	1) Explain in detail about Scalable computing over the Internet

Over the past 60 years, computing technology has undergone a series of platform and environment changes. Evolutionary changes in machine include architecture, operating system platform, network connectivity, and application workload. Instead of using a centralized computer to solve computational problems, a parallel and distributed computing system uses multiple computers to solve large-scale problems over the Internet. Thus, distributed computing becomes data-intensive and network-centric.

The Age of Internet computing

- The platform evolution
- High performance computing
- High throughput computing
- Three new computing paradigm
- Computing paradigm distinction
- Distributed system families

Scalable computing trends and new paradigms

- Degrees of parallelism
- Innovative applications
- The trend towards utility computing
- The hype cycle of new technologies

The internet of things and cyberphysical system

- The internet of things
- Cyberphysical system

The Age of Internet Computing

- Billion of people use the Internet every day.
- Supercomputers sites and large data centers must provide high-performance computing services to huge numbers of Internet users concurrently.
- The Linpack Benchmark for high-performance computing (HPC) applications is no longer optimal for measuring system performance.
- The emergence of computing clouds instead demand high-throughput computing (HTC) systems built with parallel and distributed computing technologies.
- The purpose is to advance network-based computing and web services with the emerging new technologies.

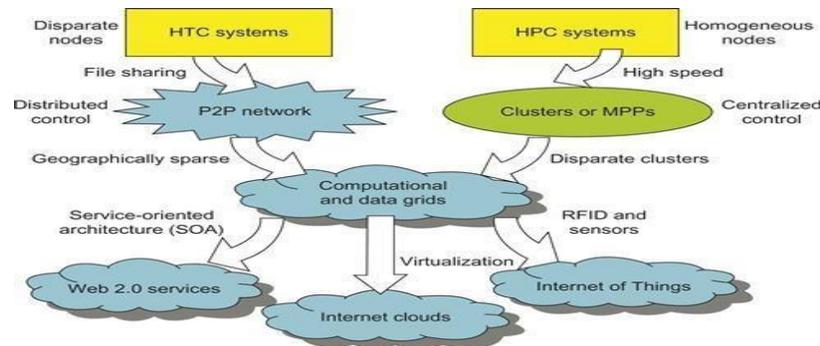
The Platform Evolution

- Computer technology has gone through five generations of development, with each generation lasting from 10 to 20 years.
- 1950 to 1970 – mainframes – eg: IBM 360 and CDC 6400.
- 1960 to 1980 – lower-cost mini-computers – eg: DEC PDP 11 and VAX Series.
- 1970 to 1990 – personal computers built with VLSI microprocessors.

- 1980 to 2000 - portable computers

High-Performance Computing

- For many years, HPC systems emphasize raw speed performance.
- The speed of HPC systems has increased from Gflops in the early 1990s to now Pflops in 2010.
- This improvement was driven mainly by the demands from scientific, engineering, and manufacturing communities.



High-Throughput Computing

- The development of market-oriented high-end computing systems is undergoing a strategic change from an HPC paradigm to an HTC paradigm.
- The main application for high-flux computing is in Internet searches and web services by millions or more users simultaneously.
- The performance goal thus shifts to measure high throughput or the number of tasks completed per unit of time.
- HTC technology needs to not only improve in terms of batch processing speed, but also address the acute problems of cost, energy savings, security, and reliability at many data and enterprise computing centers.

Three New Computing Paradigms:

- The maturity of radio-frequency identification (RFID), Global Positioning System (GPS), and sensor technologies has triggered the development of the Internet of Things (IoT).

Computing Paradigm Distinctions:

The following list defines these terms more clearly; their architectural and operational differences are discussed further in subsequent chapters.

- **Centralized computing:** This is a computing paradigm by which all computer resources are centralized in one physical system. All resources (processors, memory, and storage) are fully shared and tightly coupled within one integrated OS. Many data centers and supercomputers are centralized systems, but they are used in parallel, distributed, and cloud computing applications.
- **Parallel computing:** In parallel computing, all processors are either tightly coupled with centralized shared memory or loosely coupled with distributed memory. Some authors refer to this discipline as parallel processing. Inter processor communication is accomplished through shared memory or via message passing.
- **Distributed computing:** A distributed system consists of multiple autonomous computers, each having its own private memory, communicating through a computer network. Information exchange in a distributed system is accomplished through message passing. A computer program that runs in a distributed system is known as a distributed program. The process of writing distributed programs is referred to as distributed programming.

- **Cloud computing** An Internet cloud of resources can be either a centralized or a distributed computing system. The cloud applies parallel or distributed computing, or both. Clouds can be built with physical or virtualized resources over large data centers that are centralized or distributed. Some authors consider cloud computing to be a form of utility computing or service computing.

Distributed system families ---The system efficiency is decided by speed, programming, and energy factors (i.e., throughput per watt of energy consumed). Meeting these goals requires yielding the following design objectives:

- Efficiency
- Dependability.
- Adaptation in the programming model
- Flexibility in application deployment

Scalable computing trends and new paradigms

a) Degrees of Parallelism

- **Bit-level parallelism (BLP)** converts bit-serial processing to word-level processing gradually. Over the years, users graduated from 4-bit microprocessors to 8-, 16-, 32-, and 64-bit CPUs.
- **Instruction-level parallelism (ILP)**, in which the processor executes multiple instructions simultaneously rather than only one instruction at a time. For the past 30 years, ILP practiced through pipelining, super-scalar computing, VLIW (very long instruction word) architectures, and multithreading. ILP requires branch prediction, dynamic scheduling, speculation, and compiler support to work efficiently.
- **Data-level parallelism (DLP)** was made popular through SIMD (single instruction, multiple data) and vector machines using vector or array types of instructions.

b) Innovative Applications

- Both HPC and HTC systems desire transparency in many application aspects.
- For example, data access, resource allocation, process location, concurrency in execution, job replication, and failure recovery should be made transparent to both users and system management.
- Users must deal with multiple database servers in distributed transactions. Maintaining the consistency of replicated transaction records is crucial in real-time banking services.
- Other complications include lack of software support, network saturation, and security threats in these applications.

c) The Trend toward Utility Computing

Characteristics:

- Ubiquitous
- Reliability
- Scalability
- Autonomic
- Self-organized to support dynamic discovery.
- Composable with QoS and SLAs (service-level agreements).
- Utility vision - Utility computing focuses on a business model in which customers receive computing resources from a paid service provider. All grid/cloud platforms are regarded as utility service providers.

d)

e) **The hype cycle of new technologies** For example, at that time consumer-generated media was at the disillusionment stage, and it was predicted to take less than two years to reach its plateau of adoption. Internet micropayment systems were forecast to take two to five years to move from the enlightenment stage to maturity. It was believed that 3D printing would take five to 10 years to move from the rising expectation stage to mainstream adoption, and mesh network sensors were expected to take more than 10 years to move from the inflated expectation stage to a plateau of mainstream adoption.

THE INTERNET OF THINGS AND CYBERPHYSICAL SYSTEM

THE INTERNET OF THINGS (IOT)

- The traditional Internet connects machines to machines or web pages to web pages. The idea is to tag every object using RFID or a related sensor or electronic technology such as GPS.
- With the introduction of the IPv6 protocol, 2¹²⁸ IP addresses are available to distinguish all the objects on Earth, including all computers and pervasive devices.
- The IoT researchers have estimated that every human being will be surrounded by 1,000 to 5,000 objects.
- The IoT demands universal addressability of all of the objects or things.
- The IoT obviously extends the Internet and is more heavily developed in Asia and European countries.

CYBERPHYSICAL SYSTEM

- A cyber-physical system (CPS) is the result of interaction between computational processes and the physical world. A CPS integrates “cyber” (heterogeneous, asynchronous) with “physical” (concurrent and information-dense) objects.
- A CPS merges the “3C” technologies of computation, communication, and control into an intelligent closed feedback system between the physical world and the information world, a concept which is actively explored in the United States.



TECHNOLOGIES FOR NETWORK BASED SYSTEMS

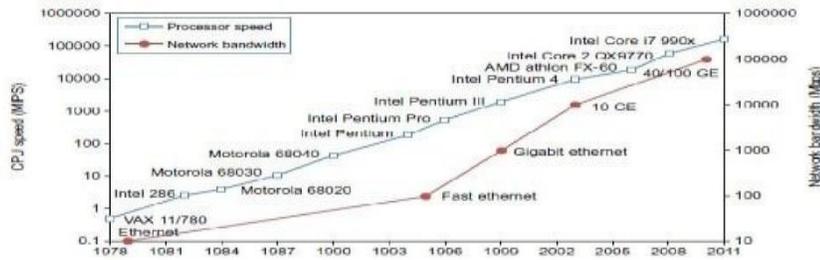
PART-A	1) Give the Basic Operation of a VM (AU/April/May 2017)
PART-B	<ol style="list-style-type: none"> 1) Explain in detail about Multicore CPUs and Multithreading Technologies 2) Explain in detail about GPU Computing to Exascale and Beyond 3) Brief the interaction between the GPU and CPU in performing Parallel execution of operation (AU/April/May 2017) 4) Illustrate the Architecture of Virtual machine and brief about the operations (AU/Nov/Dec 2016)

- o Multicore CPUs and multithreading technologies

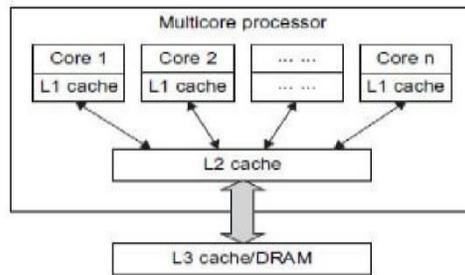
- Advances in CPU processors
 - Multicore CPU and many-core GPU architectures
 - Multithreading technology
- o GPU computing to exascale and beyond
 - How GPUs work
 - GPU programming model
 - Power efficiency of GPU
- o Memory, storage and Wide area networking
 - Memory technology
 - Disk and storage technology
 - System area interconnect
 - Wide area networking
- o Virtual machines and virtualization middleware
 - Virtual machines
 - VM primitive operations
 - Virtual infrastructure
- o Data center virtualization for cloud computing
 - Data center growth and cost breakdown
 - Low cost design philosophy
 - Convergence of technologies

1) Multicore CPUs and Multithreading Technologies

- A multicore architecture - with dual, quad, six, or more processing cores.
- These processors exploit parallelism at ILP and TLP levels.
- Processor speed growth is plotted in the upper curve in a cross-generations of microprocessors or CMPs.
- 1 MIPS - for the VAX 780 in 1978
- 1,800 MIPS - for the Intel Pentium 4 in 2002,
- 22,000 MIPS - the Sun Niagara 2 in 2008.
- As the figure shows, Moore's law has proved to be pretty accurate in this case.
- The clock rate for these processors: 10 MHz for the Intel 286, 4 GHz for the Pentium 4
- However, the clock rate reached its limit on CMOS-based chips due to power limitations.
- The clock rate will not continue to improve unless chip technology matures.
- This limitation is attributed primarily to excessive heat generation with high frequency or high voltages.
- Both multi-core CPU and many-core GPU processors can handle multiple instruction threads at different magnitudes today.
- Each core is essentially a processor with its own private cache (L1 cache).
- Multiple cores are housed in the same chip with an L2 cache that is shared by all cores. In the future, multiple CMPs could be built on the same CPU chip with even the L3 cache on the chip.
- Multicore and multi-threaded CPUs are equipped with many high-end processors, including the Intel i7, Xeon, AMD Opteron, Sun Niagara, IBM Power 6, and X cell processors. Each core could be also multithreaded.



Improvement in processor and network technologies over 33 years.



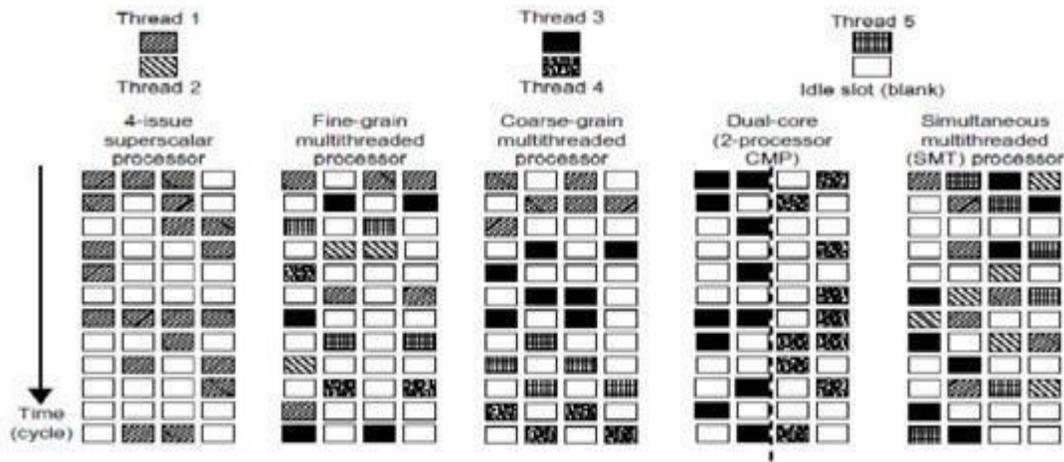
Schematic of a modern multicore CPU chip using a hierarchy of caches, where L1 cache is private to each core, on-chip L2 cache is shared and L3 cache or DRAM is off the chip.

Multicore CPU and Many-Core GPU Architectures

- Multicore CPUs may increase from the tens of cores to hundreds or more in the future.
- But the CPU has reached its limit in terms of exploiting massive DLP due to the memory wall problem.
- Many-core GPUs have with hundreds or more than cores.
- Both IA-32 and IA-64 instruction set architectures are rebuilt into commercial CPUs. Now, x-86 processors have been extended to serve HPC and HTC systems in some high-end server processors.

Multithreading Technology

The dispatch of five independent threads of instruction to four pipelined datapaths (functional units) in each of the following five processor categories from left to right: a



Five micro-architectures in modern CPU processors, that exploit ILP and TLP supported by multicore and multithreading technologies.

- **The superscalar processor** is single-threaded with four functional units. Each of the three multithreaded processors is four-way multithreaded over four functional data paths.
- **In the dual-core processor**, assume two processing cores, each a single-threaded two-way superscalar processor. Instructions from different threads are distinguished by specific shading patterns for instructions from five independent threads.
- **Fine-grain multithreading** switches the execution of instructions from different threads per cycle.
- **Course-grain multi-threading** executes many instructions from the same thread for quite a few cycles before switching to another thread.
- The multicore CMP executes instructions from different threads completely. These execution patterns closely mimic an ordinary program.
- The blank squares correspond to no available instructions for an instruction data path at a particular processor cycle.

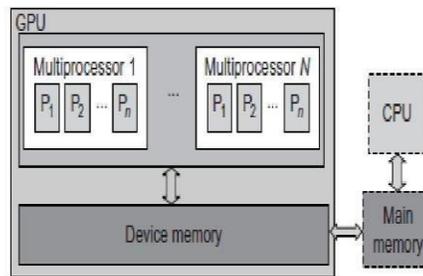
2) GPU computing to exascale and beyond

- AGPU is a graphics coprocessor or accelerator mounted on a computer's graphics card or video card.
- AGPU offloads the CPU from tedious graphic tasks in video editing applications.
- The world's first GPU, the GeForce 256, was marketed by NVIDIA in 1999.
- These GPU chips can process a minimum of 10 million polygons per second.
- Traditional CPUs are structured with only a few cores.
- For example, the Xeon X5670 CPU has six cores., a modern GPU chip can be built with hundreds of processing cores.
- GPUs have a throughput architecture that exploits massive parallelism by executing many concurrent threads slowly, instead of executing a single long thread in a conventional microprocessor very quickly.
- General-purpose computing on GPUs, known as GPGPUs, have appeared in the HPC field. NVIDIA's CUDA model was for HPC using GPGPU.

How GPUs Work

- Early GPUs functioned as coprocessors attached to the CPU.
- Today, the NVIDIA GPU has been upgraded to 128 cores on a single chip.
- Furthermore, each core on a GPU can handle eight threads of instructions.

- This translates to having up to 1,024 threads executed concurrently on a single GPU.
- Modern GPUs are not restricted to accelerated graphics or video coding.
- They are used in HPC systems to power supercomputers with massive parallelism at multicore and multithreading levels.
- GPUs are redesigned to handle large numbers of floating-point operations in parallel.
- In a way, the GPU offloads the CPU from all data-intensive calculations.
- GPUs are widely used in mobile phones, game consoles, embedded systems, PCs, and servers.
- The NVIDIA CUDA Tesla or Fermi is used in GPU clusters or in HPC systems for parallel processing of massive floating-pointing data.
- The interaction between a CPU and GPU for performing parallel execution of floating-point operations concurrently.
- The CPU is the conventional multicore processor with limited parallelism to exploit.
- The GPU has a many-core architecture that has hundreds of simple processing cores organized as multiprocessors.



The use of a GPU along with a CPU for massively parallel execution in hundreds or thousands of processing cores.

Power Efficiency of the GPU

- Bill Dally of Stanford University considers power and massive parallelism as the major benefits of GPUs over CPUs for the future.
- By extrapolating current technology and computer architecture, it was estimated that 60 Gflops/watt per core is needed to run an exaflops system
- Power constrains what we can put in a CPU or GPU chip.

3) MEMORY, STORAGE AND WIDE AREA NETWORKING

The upper curve in Figure 1.10 plots the growth of DRAM chip capacity from 16 KB in 1976 to 64 GB in 2011. This shows that memory chips have experienced a 4x increase in capacity every three years. Memory access time did not improve much in the past. In fact, the memory wall problem is getting worse as the processor gets faster. For hard drives, capacity increased from 260 MB in 1981 to 250 GB in 2004.

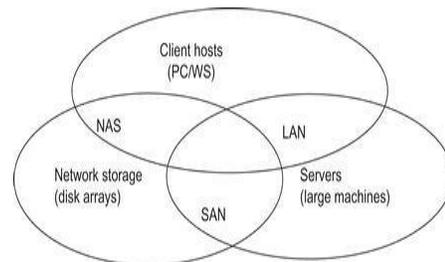
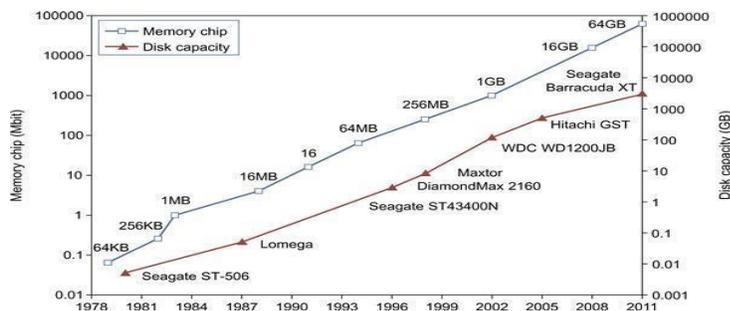


Fig1(a) Improvement in memory and disk technologies over 33 years

Fig1(B) Three interconnection networks for connecting servers, client hosts, and storage devices.

System-Area Interconnects:

The nodes in small clusters are mostly interconnected by an Ethernet switch or a *local area network* (LAN). As Figure 1(B) shows, a LAN typically is used to connect client hosts to big servers. A *storage area network* (SAN) connects servers to network storage such as disk arrays. *Network attached storage* (NAS) connects client hosts directly to the disk arrays. All three types of networks often appear in a large cluster built with commercial network components. If no large distributed storage is shared, a small cluster could be built with a multiport Gigabit Ethernet switch plus copper cables to link the end machines.

4) VIRTUAL MACHINES AND VIRTUALIZATION MIDDLEWARE

Virtual Machine Definition:

A Virtual Machine (VM) is an operating system OS or application environment that is installed on software which imitates dedicated hardware. The end user has the same experience on a virtual machine as they would have on dedicated hardware

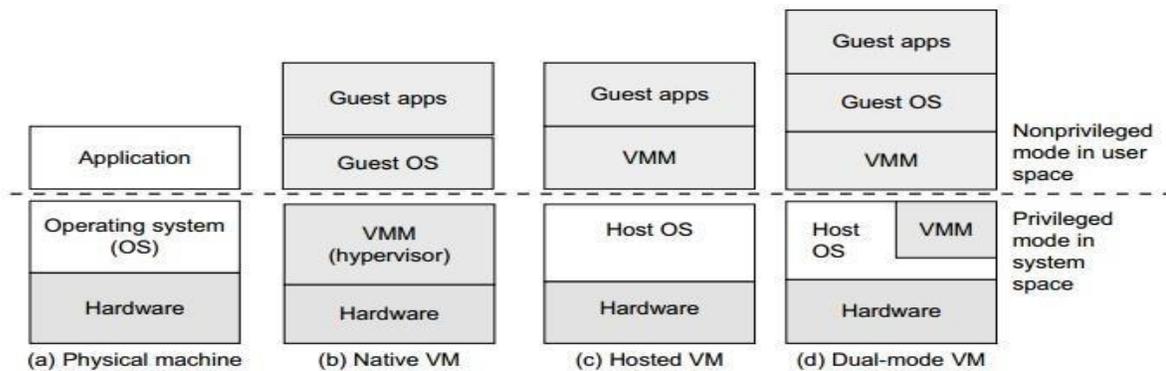


FIGURE 1.12 Three VM architectures in (b), (c), and (d), compared with the traditional physical machine shown in (a).

- The host machine is equipped with the physical hardware. The VM can be provisioned for any hardware system.
- The VM is built with virtual resources managed by a guest OS to run a specific application. Between the VMs and the host platform, one needs to deploy a middleware layer called a virtual machine monitor (VMM).
- A native VM installed with the use of a VMM called a hypervisor in privileged mode. For example, the hardware has x-86 architecture running the Windows system.

VM PRIMITIVE OPERATIONS:

With full virtualization, the VMM exports a VM abstraction identical to the physical machine so that a standard OS such as Windows 2000 or Linux can run as it would on the physical hardware.

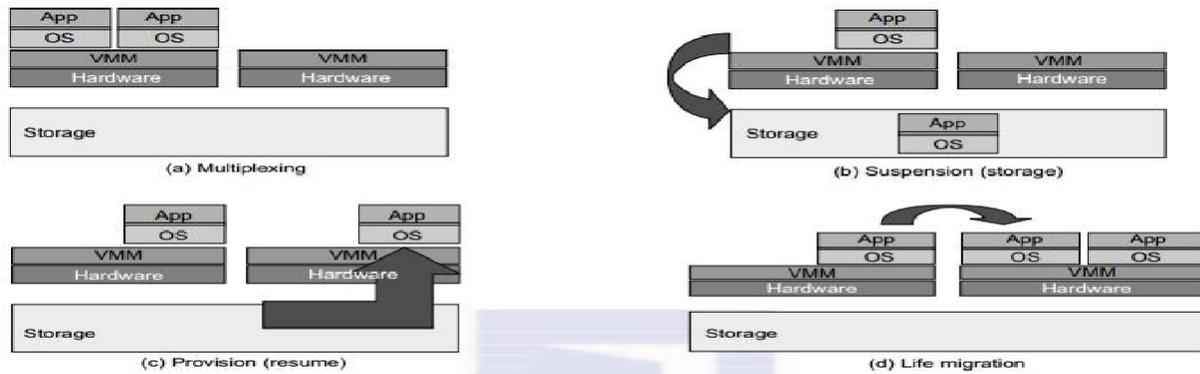


FIGURE 1.13

VM multiplexing, suspension, provision, and migration in a distributed computing environment.

- First, the VMs can be multiplexed between hardware machines.
- Second, a VM can be suspended and stored in stable suspension.
- Third, a suspended VM can be resumed or provisioned to a new hardware platform. Finally, a VM can be migrated from one hardware platform to another.

5) DATA CENTER VIRTUALIZATION FOR CLOUD COMPUTING

- Cloud architecture is built with commodity hardware and network devices.
- Low-cost terabyte disks and Gigabit Ethernet are used to build data centers.
- Data center design emphasizes the performance/price ratio over speed performance alone. In other words, storage and energy efficiency are more important than sheer speed performance.
- Worldwide, about 43 million servers are in use as of 2010. The cost of utilities exceeds the cost of hardware after three years.

CONVERGENCE OF TECHNOLOGIES

Essentially, cloud computing is enabled by the convergence of technologies in four areas:

- (1) hardware virtualization and multi-core chips,
- (2) utility and grid computing,
- (3) SOA, Web 2.0, and WS mashups,
- (4) autonomic computing and data center automation.

CLUSTERS OF COOPERATIVE COMPUTERS

PART-B

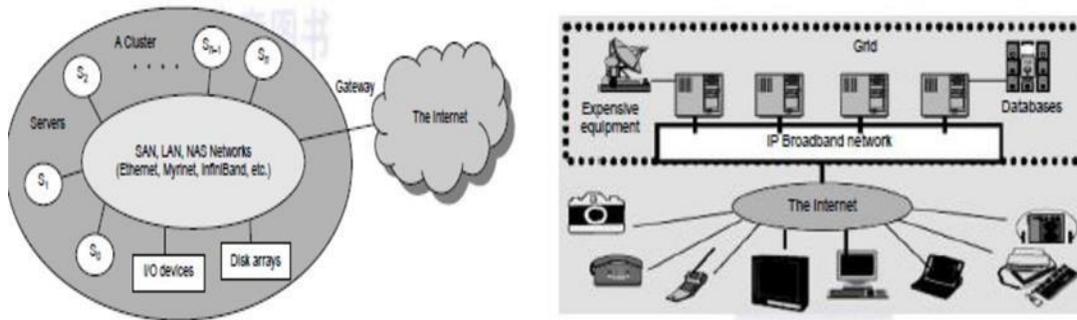
1) Write Short notes on Cluster of Cooperative computers (Nov/Dec 2016)

Definition:

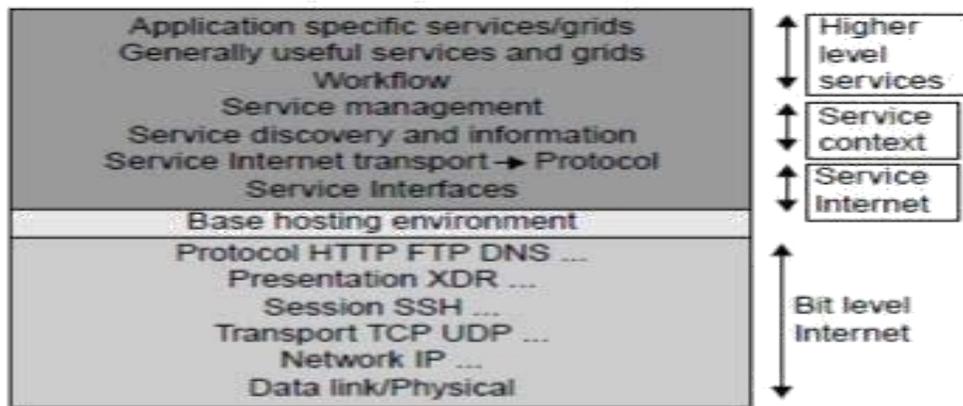
A computer cluster consists of a set of loosely or tightly connected computers that work together in many respects, they can be viewed as a single system. Clusters are generally deployed to improve speed or reliability.

Cluster Architecture:

- The architecture of a typical server cluster built around a low-latency, high bandwidth interconnection network. This network can be as simple as a SAN (e.g., Myrinet) or a LAN (e.g., Ethernet).
- To build a larger cluster with more nodes, the interconnection network can be built with multiple levels of Gigabit Ethernet, Myrinet, or InfiniBand switches.



Through hierarchical construction using a SAN, LAN, or WAN, one can build scalable clusters with an increasing number of nodes. The cluster is connected to the Internet via a virtual private network (VPN) gateway. The gateway IP address locates the cluster. The system image of a computer is decided by the way the OS manages the shared cluster resources.



Single-System Image

An ideal cluster should merge multiple system images into a single-system image (SSI). Cluster designers desire a cluster operating system or some middleware to support SSI at various levels, including the sharing of CPUs, memory, and I/O across all cluster nodes. An SSI is an illusion created by software or hardware that presents a collection of resources as one integrated, powerful resource. SSI makes the cluster appear like a single machine to the user. A cluster with multiple system images is nothing but a collection of independent computers.

GRID COMPUTING INFRASTRUCTURES

Part B Describe the Infrastructure requirements for grid Computing (AU/Nov/Dec2017)

- Users have experienced a natural growth path from Internet to web and grid computing services. Internet services such as the Telnet command enables a local computer to connect to a remote computer. Web service such as HTTP enables remote access of remote web pages.
- Grid computing is envisioned to allow close interaction among applications running on distant computers simultaneously.

Computational Grids

- Like an electric utility power grid, a computing grid offers an infrastructure that couples computers, software/middleware, special instruments, and people and sensors together.
- The grid is often constructed across LAN, WAN, or Internet backbone networks at a regional, national, or global scale Enterprises or organizations present grids as integrated computing resources.
- They can also be viewed as virtual platforms to support virtual organizations. The computers used in a grid are primarily workstations, servers, clusters, and supercomputers. Personal computers, laptops, and PDAs can be used as access devices to a grid system

Grid Families

- Grid technology demands new distributed computing models, software/middleware support, network protocols, and hardware infrastructures.
- National grid projects are followed by industrial grid platform development by IBM, Microsoft, Sun, HP, Dell, Cisco, EMC, Platform Computing, and others.
- New grid service providers (GSPs) and new grid applications have emerged rapidly, similar to the growth of Internet and web services in the past two decades.

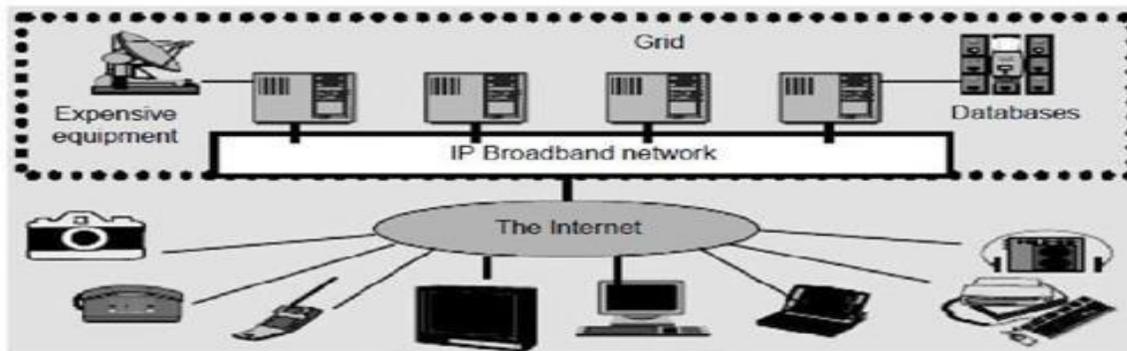


Table 1.4 Two Grid Computing Infrastructures and Representative Systems

Design Issues	Computational and Data Grids	P2P Grids
Grid Applications Reported	Distributed supercomputing, National Grid initiatives, etc.	Open grid with P2P flexibility, all resources from client machines
Representative Systems	TeraGrid built in US, ChinaGrid in China, and the e-Science grid built in UK	JXTA, FightAid@home, SETI@home
Development Lessons Learned	Restricted user groups, middleware bugs, protocols to acquire resources	Unreliable user-contributed resources, limited to a few apps

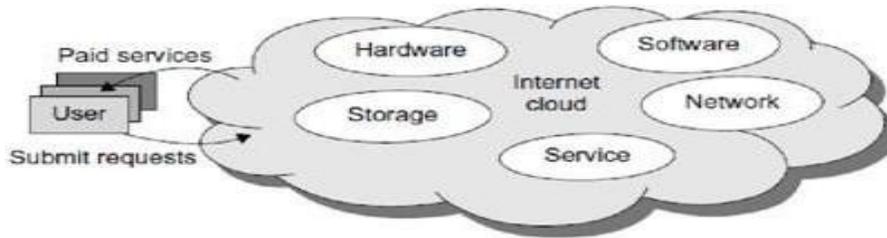
CLOUD COMPUTING OVER THE INTERNET

Part-A Highlight the importance of the term “Cloud Computing” (AU/Nov/Dec2016)

- **Definition** - “A cloud is a pool of virtualized computer resources. A cloud can host a variety of different workloads, including batch-style backend jobs and interactive and user-facing applications”
- The cloud allows workload to be deployed and scaled out quickly through rapid provisioning of virtual or physical machines.
- The cloud supports redundant, self-recovering, highly scalable programming models that allow workloads to recover from many unavoidable hardware/software failures.
- Finally, the cloud system should be able to monitor resource use in real time to enable rebalancing of allocations when needed.

Internet Clouds

- Cloud computing applies a virtualized platform with elastic resources on demand by provisioning hardware, software, and data sets dynamically.
- The idea is to move desktop computing to a service-oriented platform using server clusters and huge databases at data centers.
- Cloud computing leverages its low cost and simplicity to benefit both users and providers.



The Cloud Landscape

The traditional systems have encountered several performance bottlenecks: constant system maintenance, poor utilization, and increasing costs associated with hardware/software upgrades. Cloud computing as a non-demand computing paradigm resolves or relieves us from these **problems**.

CLOUD SERVICE MODELS:

- **Infrastructure as a Service (IaaS)**--- This model puts together infrastructures demanded by users — namely servers, storage, networks, and the data center fabric.
- **Platform as a Service (PaaS)**- This model enables the user to deploy user-built applications onto a virtualized cloud platform. PaaS includes middleware, databases, development tools, and some runtimes such as Web 2.0 and Java.
- **Software as a Service (SaaS)**- The SaaS model applies to business processes, industry applications, consumer relationship management (CRM), enterprise resource planning (ERP), human resources (HR), and collaborative applications.

The following list highlights eight reasons to adapt the cloud for upgraded Internet applications and web services:

- Desired location in areas with protected space and higher energy efficiency
- Sharing of peak-load capacity among a large pool of users, improving overall utilization
- Separation of infrastructure maintenance duties from domain-specific application development
- Significant reduction in cloud computing cost, compared with traditional computing paradigms
- Cloud computing programming and application development
- Service and data discovery and content/service distribution
- Privacy, security, copyright, and reliability issues

SERVICE ORIENTED ARCHITECTURE

Part-A	Differentiate between Grid and Cloud Computing (AU/Nov/Dec 2017)

Part-B Write Short notes on Service Oriented Architecture (AU/Nov/Dec 2016)

- In grids/webservices, Java, and CORBA, an entity is, respectively, a service, a Java object, and a CORBA distributed object in a variety of languages.
- These architectures build on the traditional seven Open Systems Interconnection (OSI) layers that provide the base networking abstractions.
- Eg.: .NET or Apache Axis for webservices, the Java Virtual Machine for Java, broker network for CORBA.

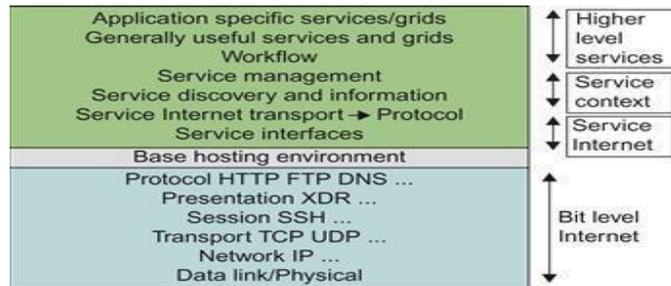


Fig: Layered architecture for webservices and the grids.

Layered Architecture for Web Services and Grids

- The entity interfaces correspond to the Web Services Description Language (WSDL), Java method, and CORBA interface definition language (IDL) specifications in these example distributed systems.
- These interfaces are linked with customized, high-level communications systems: SOAP, RMI, and IIOP in the three examples.
- These communications systems support features including particular message patterns

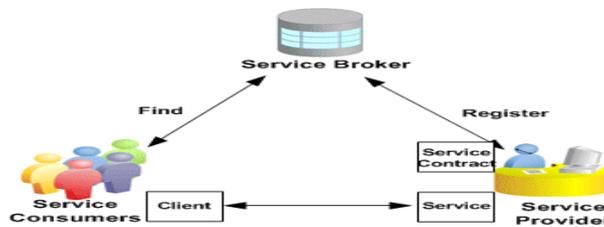


Fig: SOA Architecture

Web Services and Tools

- Loose coupling and support of heterogeneous implementations makes services more attractive than distributed objects.
- In webservices, one aims to fully specify all aspects of the service and its environment.
- This specification is carried with communicated messages using Simple Object Access Protocol (SOAP).
- The hosting environment then becomes a universal distributed operating system with fully distributed capability carried by SOAP messages.
- This approach has mixed success as it has been hard to agree on key parts of the protocol and even harder to efficiently implement the protocol by software such as Apache Axis.
- In the REST approach, one adopts simplicity as the universal principle and delegates most of the difficult problems to application (implementation-specific) software.
- For Java, this could be as simple as writing a Java program with method calls replaced by Remote Method Invocation (RMI), while CORBA supports a similar model with a syntax reflecting the C++ style of its entity (object) interfaces.

EVOLUTION OF SOA

- As shown in Figure 1.21, service-oriented architecture (SOA) has evolved over the years.
- SOA applies to building grids, clouds, grids of clouds, clouds of grids, clouds of clouds (also known as interclouds), and systems of systems in general.
- A large number of sensors provide data-collection services, denoted in the figure as SS (sensor service). A sensor can be a ZigBee device, a Bluetooth device, a WiFi access point, a personal computer, a GPA, or a wireless phone, among other things.
- Raw data is collected by sensor services.
- All the SS devices interact with large or small computers, many forms of grids, databases, the compute cloud, the storage cloud, the filter cloud, the discovery cloud, and so on. Filter services (fs in the figure) are used to eliminate unwanted raw data, in order to respond to specific requests from the web, the grid, or web services.

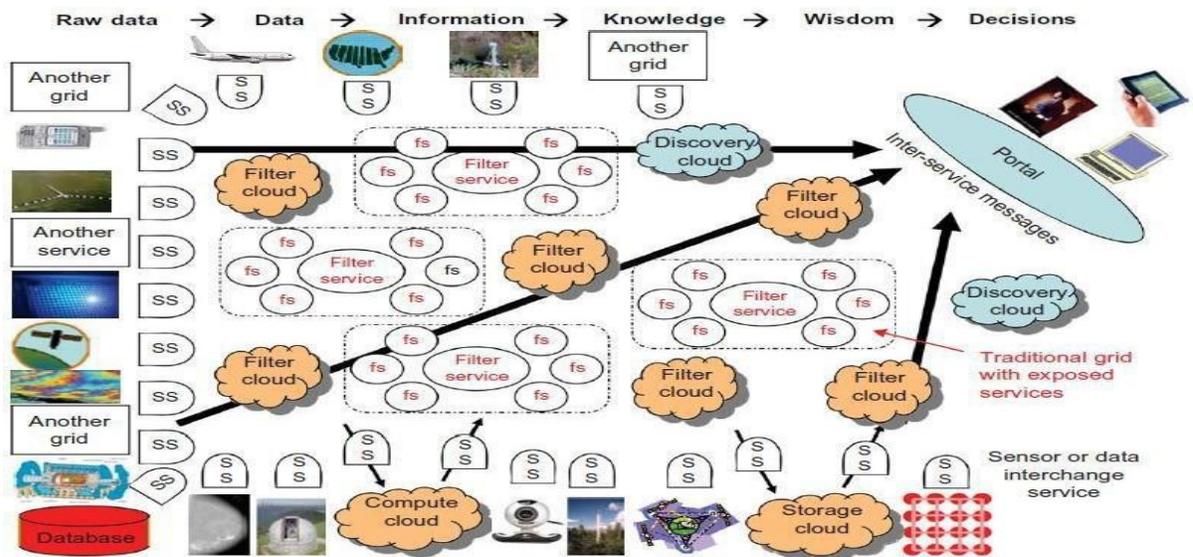


Fig: Evolution of SOA

Grid Vs Cloud

- In general, a grid system applies static resources, while a cloud emphasizes elastic resources.
- For some researchers, the differences between grids and clouds are limited only in dynamic resource allocation based on virtualization and autonomic computing. One can build a grid out of multiple clouds.
- This type of grid can do a better job than a pure cloud, because it can explicitly support negotiated resource allocation.
- Thus one may end up building with a *system of systems*: such as a *cloud of clouds*, a *grid of clouds*, or a *cloud of grids*, or *inter-clouds* as a basic SOA architecture.

INTRODUCTION TO GRID ARCHITECTURE AND STANDARDS – ELEMENTS OF GRID – OVERVIEW OF GRID ARCHITECTURE

PART-B1) Describe Layered Grid Architecture. How does it map onto internet Protocol architecture (AU/Nov/Dec 2017)

- Grid computing is applying the resources of many computers in a network to a single problem at the same time.

- Grid computing technologies & infrastructures support the sharing & coordinated use of diverse resources in dynamic, distributed virtual organization.
- Grid architecture identifies the fundamental system components, specifies purpose and function of these components and indicates how these components interact with each other.
- The grid can be thought of as a distributed system with non-interactive workload that involves a large number of files.
- The difference between grid computing and conventional high performance computing systems such as cluster computing is that grid is to be more loosely coupled, heterogeneous & geographically dispersed.
- Grids are often constructed with general purpose grid middleware software libraries.

GRID COMPUTING ARCHITECTURE:

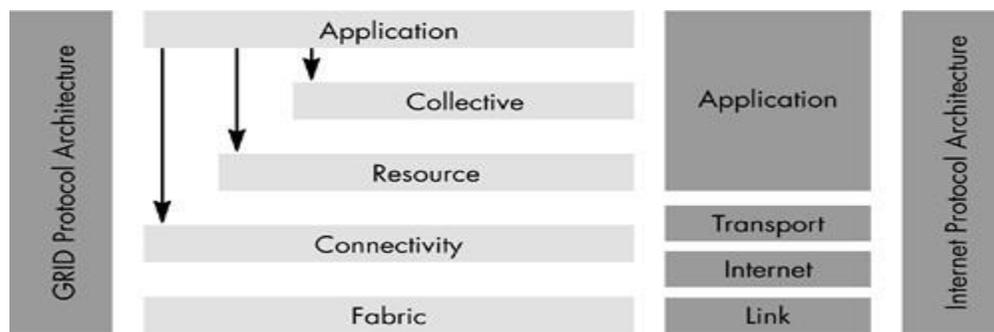


Figure: Grid Computing architecture

ADVANTAGES OF GRID COMPUTING:-

- No need to buy large symmetric multiprocessor server.
- Much more efficient use of idler resources.
- Grid environment are much more modular and don't have single point of failure.
- Policies can be managed by the grid software.
- Upgrading can be alone.
- Jobs can be executed in parallel speed.

DISADVANTAGES OF GRID COMPUTING:-

- For small memory applications memory passing interfaces you may need to have a fast interconnect between the computer resources.
- Some application may require full advantages of new model.
- Political challenges associated with sharing resources.

GRID COMPUTING STANDARDS:

- a) **OGSA (Open Grid Service Architecture):** The aim of OGSA is to standardize grid computing and to define a basic framework of a grid application structure. Some of the key concepts are first presented by Ian Foster who still leads the OGSA working group. This Architecture combines different aspects from grid computing with advantages from Web Services.
- b) **OGSA Services:** The OGSA specifies services which occur within a wide variety of grid systems. They can be divided into 4 broad groups: i) Core Services ii) Data Services iii) Program Execution Services and iv) Resource Management Services.
- c) **OGSI (Open Grid Service Infrastructure):** OGSA defines a Grid Application and what a Grid Service should be able to do. OGSI specifies a Grid Services in detail.

d) WSRF (Web Service Resource Framework): WSRF is a derivative of OGSI. A first implementation can be found in GT4 (Global Toolkit 4). The framework combines 6 different WS specifications “that define what is termed the WS-Resource approach to modeling and managing state in a Web services context. WSRF Specifications: i) WS-ResourceLifetime: mechanisms for WSResource destruction ii) WS-ResourceProperties: manipulation and definition of WS properties iii) WS-Notification: event management iv) WS-RenewableReference: defines updating proceeding v) WS-ServiceGroup: interface for by-reference collections of WSs vi) WS-BaseFaults: standardization of possible failures.

ELEMENTS OF GRID:

Grid computing combines elements such as distributed computing, high-performance computing and disposable computing depending on the application of the technology and the scale of operation. Grids can create a virtual supercomputer out of the existing servers, workstations and personal computers.

Present-day grids encompass the following types

- **Computational grids**, in which machines will set aside resources to “number crunch” data or provide coverage for other intensive workloads
- **Scavenging grids**, commonly used to find and harvest machine cycles from idle servers and desktop computers for use in resource-intensive tasks (scavenging is usually implemented in a way that is unobtrusive to the owner/user of the processor)
- **Data grids**, which provide a unified interface for all data repositories in an organization, and through which data can be queried, managed and secured.
- **Market-oriented grids**, which deal with price setting and negotiation, grid economy management and utility driven scheduling and resource allocation.

OVERVIEW OF GRID ARCHITECTURE:

The architecture of a grid system is often described in terms of “layers”, each providing a specific function as shown in the following figure. Higher layers are user centric, whereas the lower layers are hardware-centric. In Figure 1 a generic grid architecture showing the functionality of each layer.

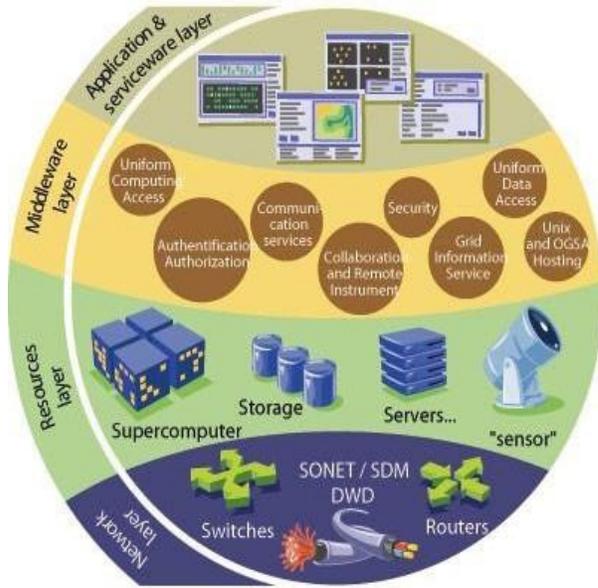


Fig: Grid Layer Architecture

NETWORK LAYER: It is the bottom layer which assures the connectivity for the resources in the grid.

RESOURCE LAYER: It is made up of actual resources that are part of the grid, such as computers, storage systems, electronic data catalogues, and even sensors such as telescopes or other instruments, which can be connected directly to the network.

MIDDLEWARE LAYER: It provides the tools that enable various elements (servers, storage, networks, etc.) to participate in a unified grid environment.

APPLICATION LAYER: In which includes different user applications (science, engineering, business, financial), portal and development toolkits-supporting applications.

Types of Layer & its Function

MAIN CHARACTERISTICS OF GRIDS:

The main characteristics of a grid computing environment can be listed as follows:

- **Large scale:** A grid must be able to deal with a number of resources ranging from just a few to millions.
- **Geographical distribution:** Grid resources may be spread geographically.
- **Heterogeneity:** A grid hosts both software and hardware resources that can be ranging from data, files, software components or programs to sensors, scientific instruments, display devices, personal digital organizers, computers, super-computers and networks.
- **Resource sharing and coordination:** Resources in a grid belong to different organizations that allow other organizations (i.e. users) to access them. The resources must be coordinated in order to provide aggregated computing capabilities.
- **Multiple administrations:** Each organization may establish different security and administrative policies under which resources can be accessed and used.
- **Accessibility attributes:** Transparency, dependability, consistency, and pervasiveness are attributes typical to grid resource access. A grid should be seen as a single virtual computing environment and must assure the delivery of services under established Quality of Service requirements.