

Compute Power for Life Sciences Research

Balasubramanian.R @ Compaq.com

Compaq Computer Corporation, Bangalore, India



Central Paradigm of Bioinformatics

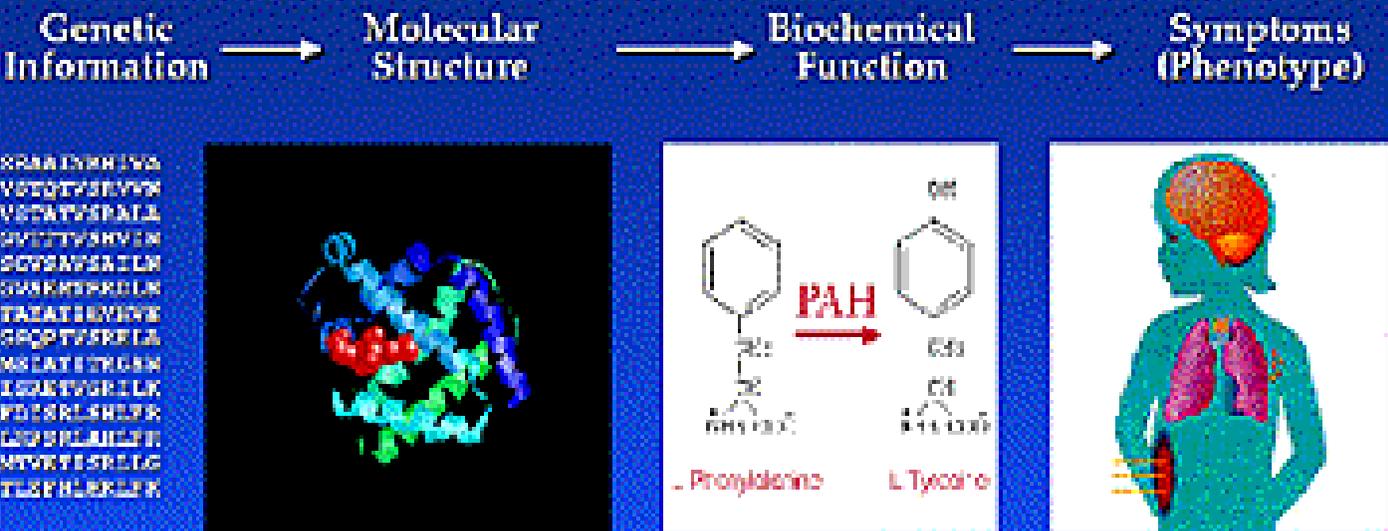
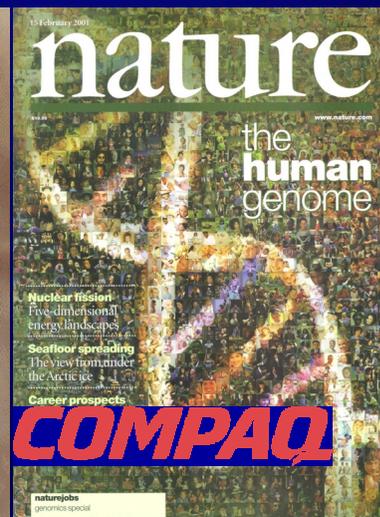


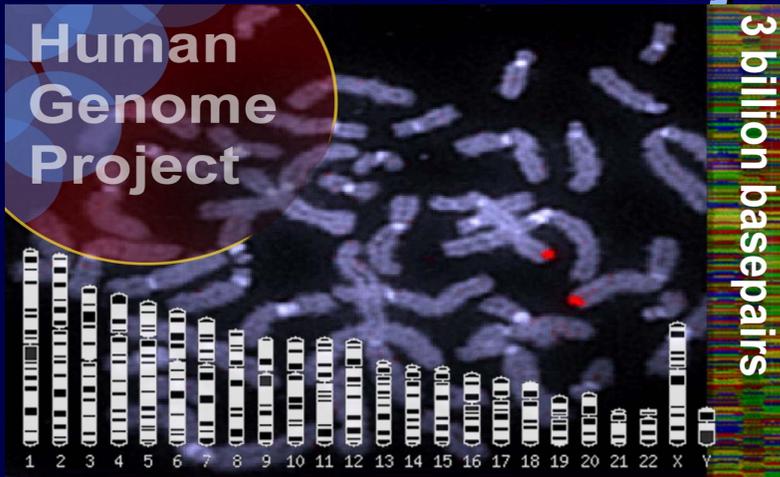
Image taken from Stanford Bioinformatics Group]

Genome has been Sequenced

Using Compaq Alpha Servers



Human Genome Project



- DOE Joint Genome Institute
- Baylor College of Medicine
- Sanger Centre
- Washington University Genome Sequencing Center
- Whitehead Institute/MIT Center for Genome Research

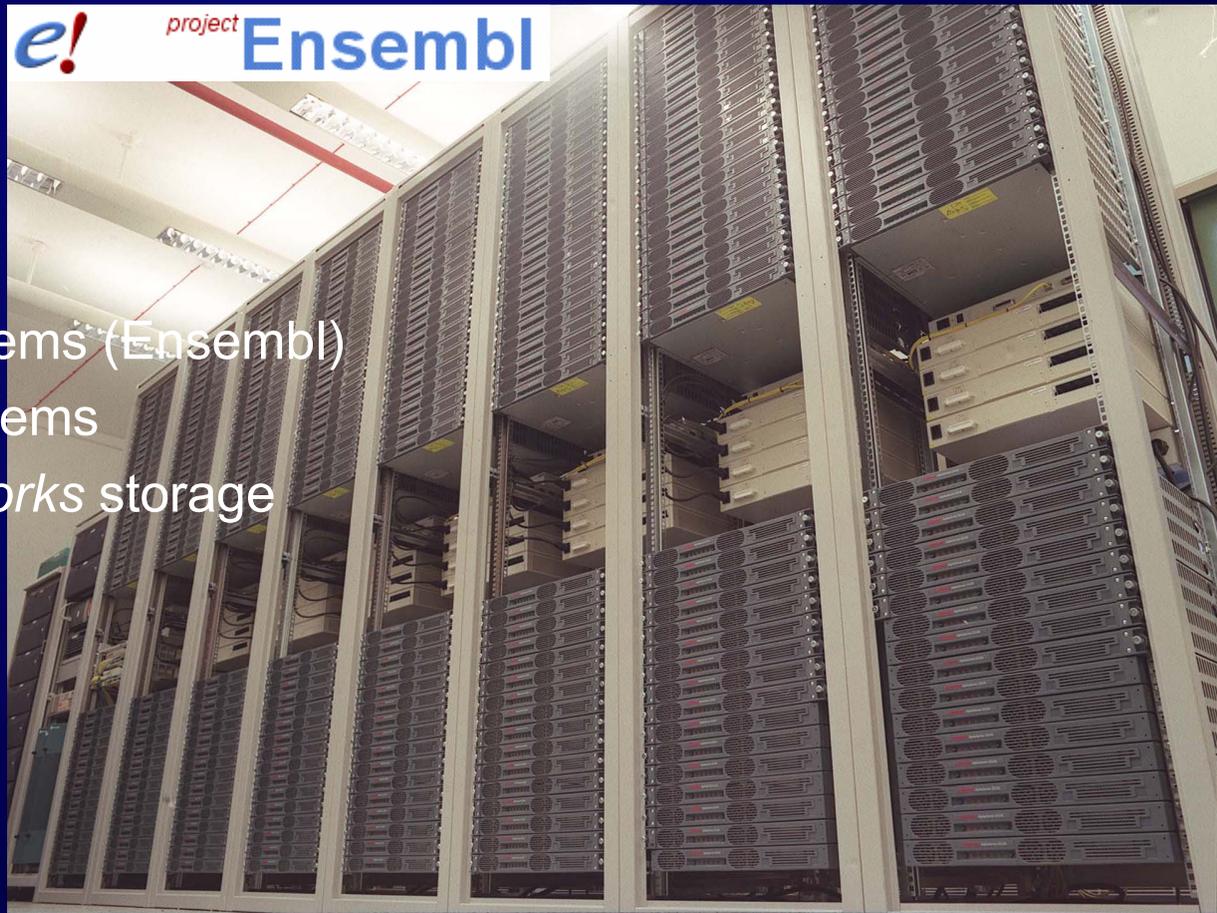
Supercomputer Facility @ Whitehead/MIT CGR

- 2 sites for Sequencing/Functional Genomics
- Standardized on Compaq Technology
- *AlphaServer* Systems
 - 20 Compaq *AlphaServer* 4-processor ES40 systems
 - 12 DS20s and DS10 systems
- Storage and Backup systems
 - Over 21 TB *StorageWorks*™ storage

Sanger's Supercomputer Facility

AlphaServer Systems

- Over 100 AlphaServer ES40 and DS20 systems
- 160 Alpha workstations in a clustered compute farm
- 120 node BLAST farm
 - 40 Alpha DS10L
 - 32 Alpha DS10
 - 48 Cpq Pentium
- 320 Alpha DS10L systems (Ensembl)
- Storage & Backup systems
- Over 20 TB StorageWorks storage



Celera Genomic's Computing Environment

The Virtual Compute Farm

440 x Alpha CPUs, includes model EV6 @500 MHz and EV67 @667 MHz.
Available memory ranges from 2 GB to 8 GB.
The VCF is used to manage trace file processing, and annotation.

Genome assembly was performed on a GS160

GS160 running 16 EV67s (731 MHz) and 64 GB of memory, and
10 x ES40s running 4 EV6s (500 MHz) and 32 GB of memory.
Total of 100 terabytes of SAN storage



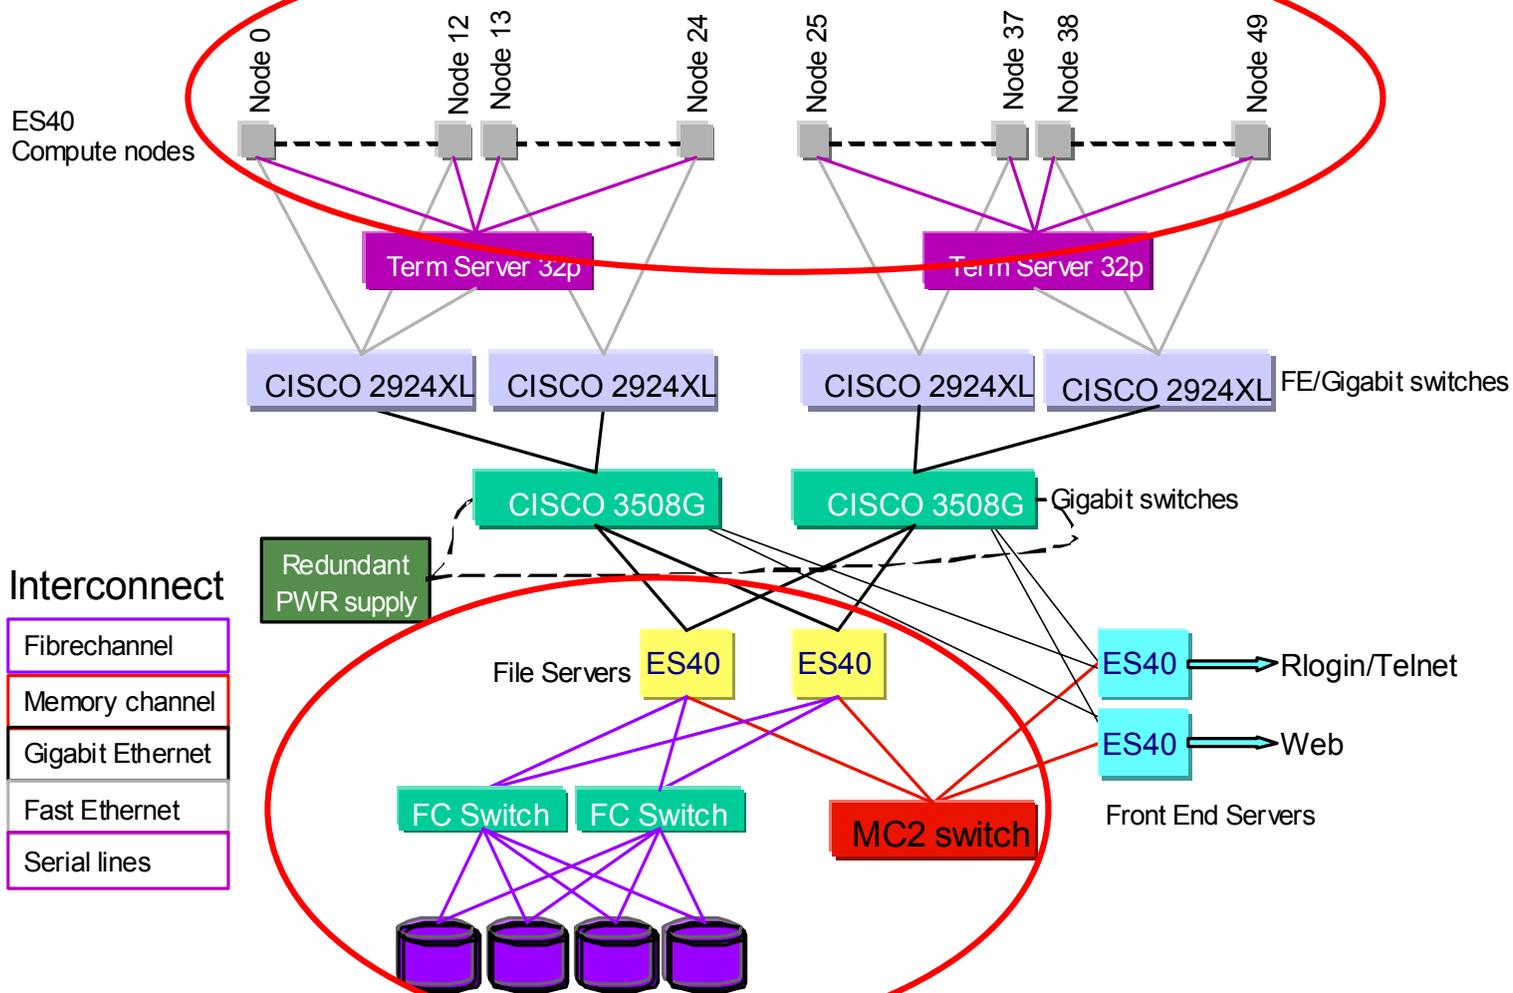
**Jun'00
Genome Assembly
= 20,000 CPU hours**

File and Database servers

Configured as 4-node Alpha *TruClusters* for High Availability.
Data Fault Resilience by RAID-0+1

Typical System Configuration for Bioinformatics

Bioinformatics Solution Architecture



Growth in selected configurations

	Starting Configuration	Current Configuration	Future Needs
Celera Genomics	<u>June 1999</u> <ul style="list-style-type: none"> ▪ 144 AlphaServer CPUs ▪ 14TB storage ▪ Sybase 	<ul style="list-style-type: none"> ▪ 800 Alpha CPUs ▪ 80TB storage ▪ Sybase & Oracle 	<p style="text-align: center;">Joint Celera-Sandia-Compaq 100 TeraOps Supercomputer</p>
Sanger Centre	<u>June 1999</u> <ul style="list-style-type: none"> ▪ 220 CPUs (AlphaServer, Alpha workstations and Intel PCs) ▪ 4TB storage ▪ Oracle 	<ul style="list-style-type: none"> ▪ 700 Alpha CPUs (a few Intel PCs) ▪ 20TB storage ▪ Oracle 	<ul style="list-style-type: none"> ▪ Need additional 100 CPUs just to annotate mouse genome ▪ In 2 years, need 5x CPUs and over 100TB of storage
MIT/Whitehead	<u>October 1997</u> <ul style="list-style-type: none"> ▪ 12 Alpha workstations ▪ 0.4TB storage ▪ Sybase 	<ul style="list-style-type: none"> ▪ 90 Alpha CPUs ▪ 9TB storage ▪ Oracle 	<ul style="list-style-type: none"> ▪ Triple computing resources in one year



GeneProt - Large Proteomic Discovery Center in Geneva, Switzerland

- Center runs 20 hours a day and use 1,420 Compaq Alpha-Tru64 UNIX and StorageWorks systems to capture, store, and analyze the terabytes of data generated by 51 Bruker Daltonics mass spectrometers.
- Compaq invest \$10 million in GeneProt.
- "What makes GeneProt unique among the companies seeking to commercialize proteins is the industrial-sized scale in which we work; " said Loiret-Bernal.
- GeneProt has already entered into a five-year proteomics analysis partnership with Novartis.
- Largest Protein Factory in the World !
- Largest Private-owned Computer Farm in the World !

The Compaq BioCluster

- Cluster of *AlphaServer ES40* systems
- Available to academic users for ...
- Annotation of the human genome

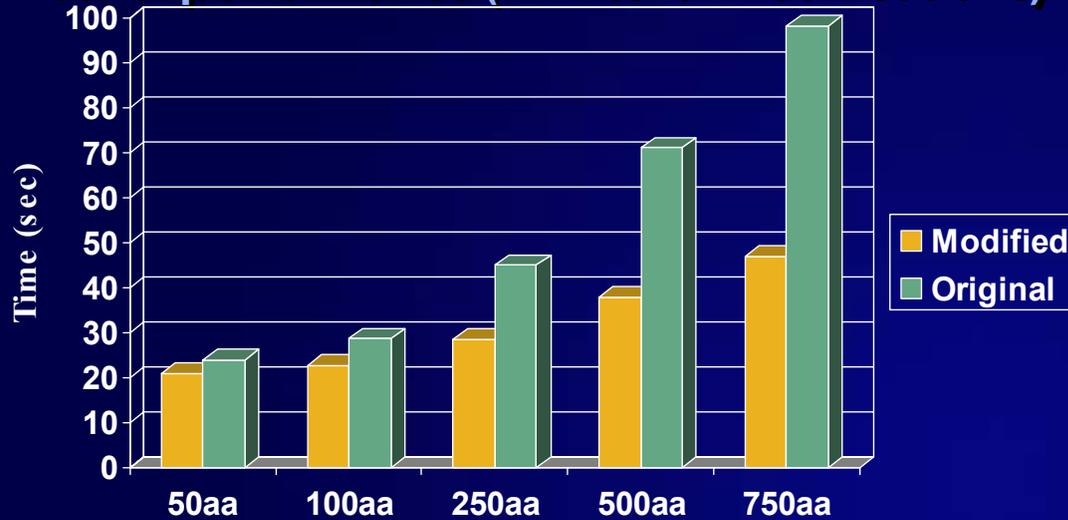


- *AlphaServer DS10* login node
- 25 compute nodes (*ES40*, 4GB, 54GB disk)
- 1 big compute node (*ES40*, 16GB, 54GB disk)
- *AlphaServer ES40* file server with 1TB of storage
- Dual 100Mb Ethernet
- LSF from Platform Computing used to distribute jobs

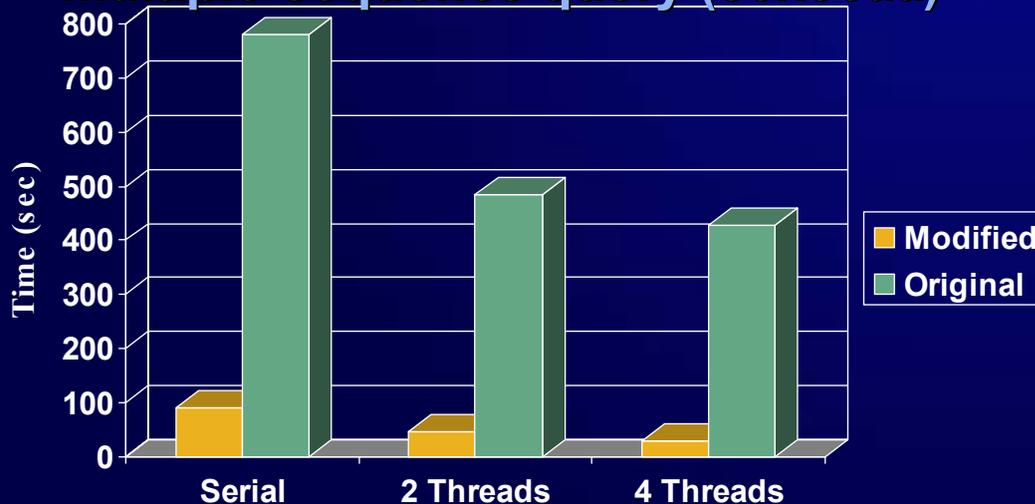
Application optimisation work

- We have a keen interest in application performance measurement and optimisation
- Primarily worked with public-domain developers
- Work is done in collaboration with original developers
 - Developers retain all IP
 - Work has benefited all system vendors
- Work done to date:
 - FASTA parallelisation
 - NCBI-BLAST optimisation
 - HMMER optimisation
 - GeneHunter parallelisation
 -

Modified HMMER vs original: serial performance (P7Viterbi modifications)



multiple-sequence query (30x56aa)



More Results from the BioCluster

AlphaServer cluster needed **25% less time to complete a run 2.5 times larger than all previous runs made on any system available from any vendor.** Overall system performance was 2.5 times greater than all the competitive systems”

Jean-Jacques Codani,
CEO of Gene-IT

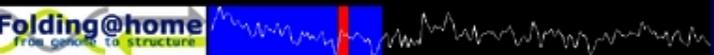
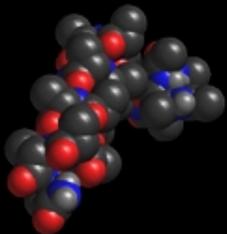


Entropia signs Novartis for Distributed Computing platform

- Novartis will evaluate its PC-based distributed computing technology—which harnesses spare cycles available from an existing network of PCs—to meet the computational demands of its drug-discovery efforts.
- Entropia estimated that the average PC sits idle more than 90 percent of the time and that its technology could provide a company an additional teraflop of compute power using as few as 2,500 processors with an average clock speed of about 500 MHz.

Folding @Home Stanford Pande-Group's Project

- Need more computing power to simulate Protein Folding.
- While the alpha helix folds in 100 nanoseconds, proteins just a little larger fold 100x slower (10 microseconds).
- While 10-100 processors were enough to simulate the helix, we will need many more to simulate these larger, more interesting proteins



Stanford's Pande group has developed a new way to simulate protein folding by dividing the work between multiple processors in a new way - with a near linear speed up in number of processors.

“Thus, with 1000's processors, we can break the microsecond barrier & unlock the mystery of how proteins fold”.

Challenge of Porting Codes

Moving Application Codes to New Hardware

Time it takes to move an application to a new generation of computers is time lost to scientific discovery.

- Scientific modeling codes endure much longer than the life of a generation of computers (~ 3 years)
 - The “Gaussian” code for which John Pople was awarded the 1998 Nobel Prize has been in existence (and evolving) since 1970.

Given the nature of the architecture, it is clear that a number of challenges must be faced in crafting an application to execute the science program on the target machine platform.

- F.Allen et.al. IBM Blue Gene Team

AlphaServer SC Customers



ASCI Q – Los Alamos



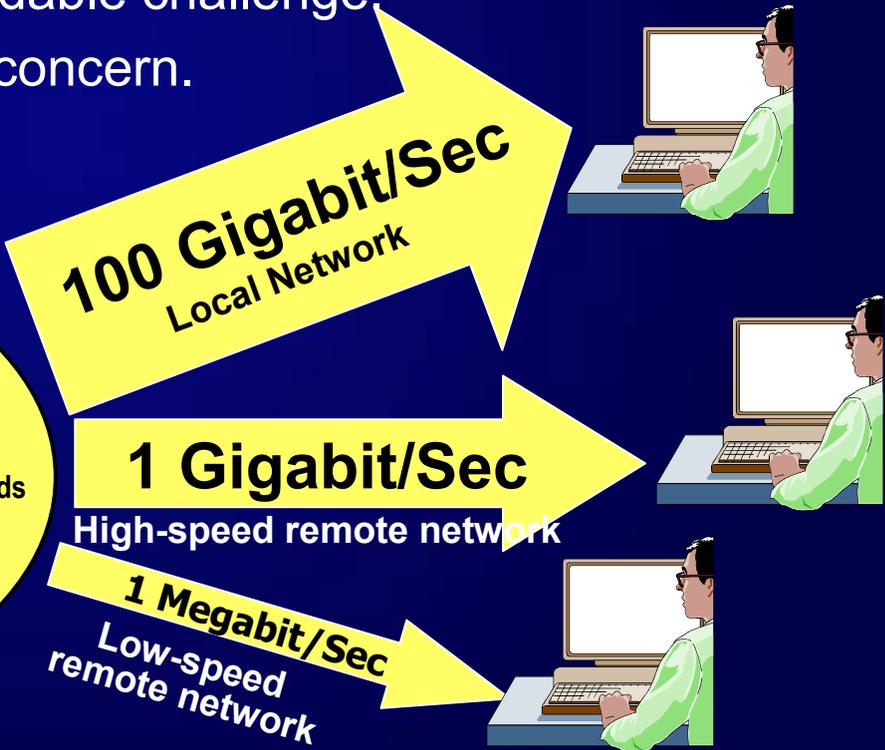
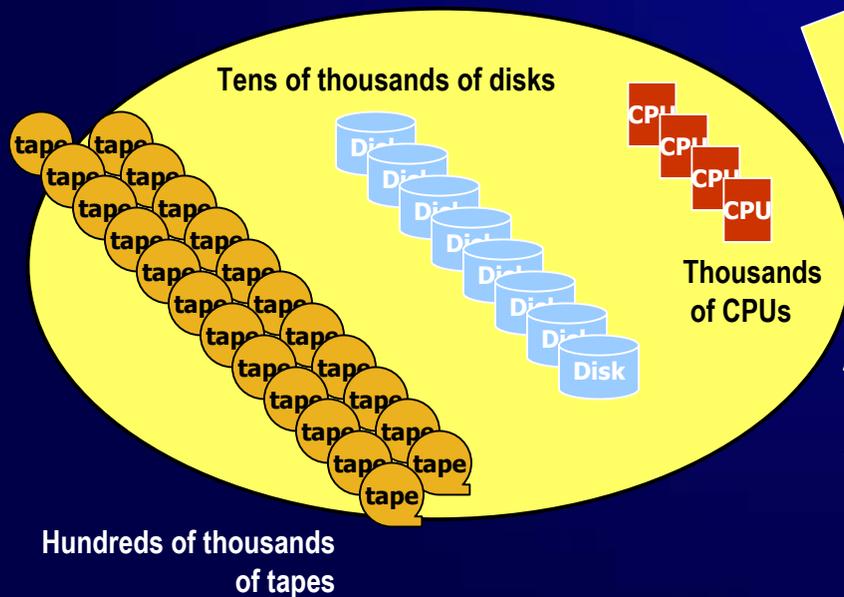
Pittsburgh TCS



Challenge of Using Petabytes of Data

Data, Data Everywhere

- High performance computers and new experimental facilities will produce millions of gigabytes of data.
- Enabling scientists to interact with this volume of data and find and analyze critical features is a formidable challenge.
- Remote access time is a serious concern.



Compaq demos Global Storage Network on Internet with FCIP

Data, Data Everywhere

- **World's first global storage network using Internet & FibreChannel.**
- **Global storage network links FC Storage Area Networks (SANs) in**
 - **Colorado Springs, Colorado;**
 - **Sydney, Australia; and**
 - **Nijmegen in the Netherlands**
- **SAN-connected devices situated across continents and replicate data using SANworks Data Replication Manager**
- **Demonstration of storage networking between continents using FCIP technology achieves the Global Replication Network defined in Compaq ENSA-2 vision**

“Grid” enables communities (virtual organizations) to share resources as they pursue common goals

Compute
Grid

Data Grid

AccessGrid

GRID

Distributed
Supercomputing

Distributed
Databases

On-demand
Collaborative
Multimedia

