

PARAM padma

C-DAC's Terascale Supercomputing Facility (CTSF) at C-DAC Knowledge Park, Bangalore



Centre for Development of Advanced Computing

India

C-DAC's HPCC (High Performance Computing and Communication) initiatives are aimed at designing, developing and deploying advanced computing systems, tools and technologies that impact strategically important application areas.

Fostering an environment of innovation and dealing with cutting edge technologies, C-DAC's PARAM series of supercomputers have been deployed to address diverse applications in science and engineering, and business computing at various institutions in India and abroad.

C-DAC's commitment to the HPCC initiative has once again manifest as a deliverable through the design, development and deployment of PARAM Padma, a terascale supercomputing system.

PARAM Padma is C-DAC's next generation high performance scalable computing cluster, currently with a peak computing power of One Teraflop. The hardware environment is powered by the Compute Nodes based on the state-of-the-art Power4 RISC processors, using Copper and SOI technology, in Symmetric Multiprocessor (SMP) configurations. These nodes are connected through a primary high performance System Area Network, PARAMNet-II, designed and developed by C-DAC and a Gigabit Ethernet as a backup network.

The PARAM Padma is powered by C-DAC's flexible and scalable HPCC software environment. The Storage System of PARAM Padma has been designed to provide a primary storage of 5 Terabytes scalable to 22 Terabytes. The network centric storage architecture, based on state-of-the-art Storage Area Network (SAN) technologies, ensures high performance, scalable and reliable storage. It uses Fibre Channel Arbitrated Loop (FC-AL) based technology for interconnecting storage subsystems like Parallel File Servers, NAS Servers, Metadata Servers, Raid Storage Arrays and Automated Tape Libraries, achieving an I/O performance of upto 2 Gigabytes/Second.



The Secondary backup storage subsystem is scalable from 10 Terabytes to 100 Terabytes with an automated tape library and support for DLT, SDLT and LTO Ultrium tape drivers. It implements a Hierarchical Storage Management (HSM) technology to optimize the demand on primary storage and effectively utilize the secondary storage.

The PARAM Padma system is also accessible by users from remote locations.

PARAMNet-II

PARAMNet-II, a low latency, high bandwidth interconnect, provides data rates of 2.5 Gigabits/sec in full duplex over fiber. The message latency is as low as 10 μ sec.

PARAMNet-II uses a 16 port switch and a Network Interface Card (NIC) alongwith an Application Programming Interface i.e. C-DAC's Virtual Interface Provider Library (C-VIPL). The non-blocking architecture of the switch allows multilevel switching for realizing a large cluster. The switch offers very low latency of the order of 0.5 μ sec. Use of an Interval routing scheme and group adaptive routing based on Least Recently Used (LRU) algorithm, ensures uniform bandwidth distribution. The NIC is



PARAMNet-II Switch

based on C-DAC's Communication Co-Processor-III (CCP-III) chip based on 0.15 micron 1 million gate technology. Implementation of packetisation & reassembly, flow control, protection mechanism, address translation and error recovery in CCP-III, results in low latencies and very low overheads for the CPU.

PARAM Padma has 12 PARAMNet-II switches connected in two level configurations to form a 64-node CLOS network.



PARAMNet II boards

HPCC Software



C-DAC's HPCC software on PARAM Padma effectively addresses the performance and usability challenges of clusters through a high performance flexible software environment. It provides low overhead communication, optimized MPI and a parallel file system with MPI-IO interface to enable applications to scale on large clusters.

C-PFS, a client-server and user-level parallel file system, addresses the high I/O throughput requirements of scientific and engineering applications. Exporting MPI-IO interface for parallel programming and UNIX interface for system management, the C-PFS fully exploits the concurrent data paths between the

compute nodes and the terabytes of storage in PARAM Padma.

The Fortran programming environment includes a Fortran 90 compiler, a parallelizing compiler for Fortran 77/90 for SMP, an integrated development environment and Fortran 77 to Fortran 90 converter.

The DIViA debugging environment comprises of a coherent set of tools that help programmers in both correctness and performance debugging of C or Fortran language parallel programs for MPI environment. PARMON, the system management tool, enables the administrator to monitor activities and resource utilization of various cluster components. The RMS resource management tool enables users to optimally schedule jobs across the cluster nodes.

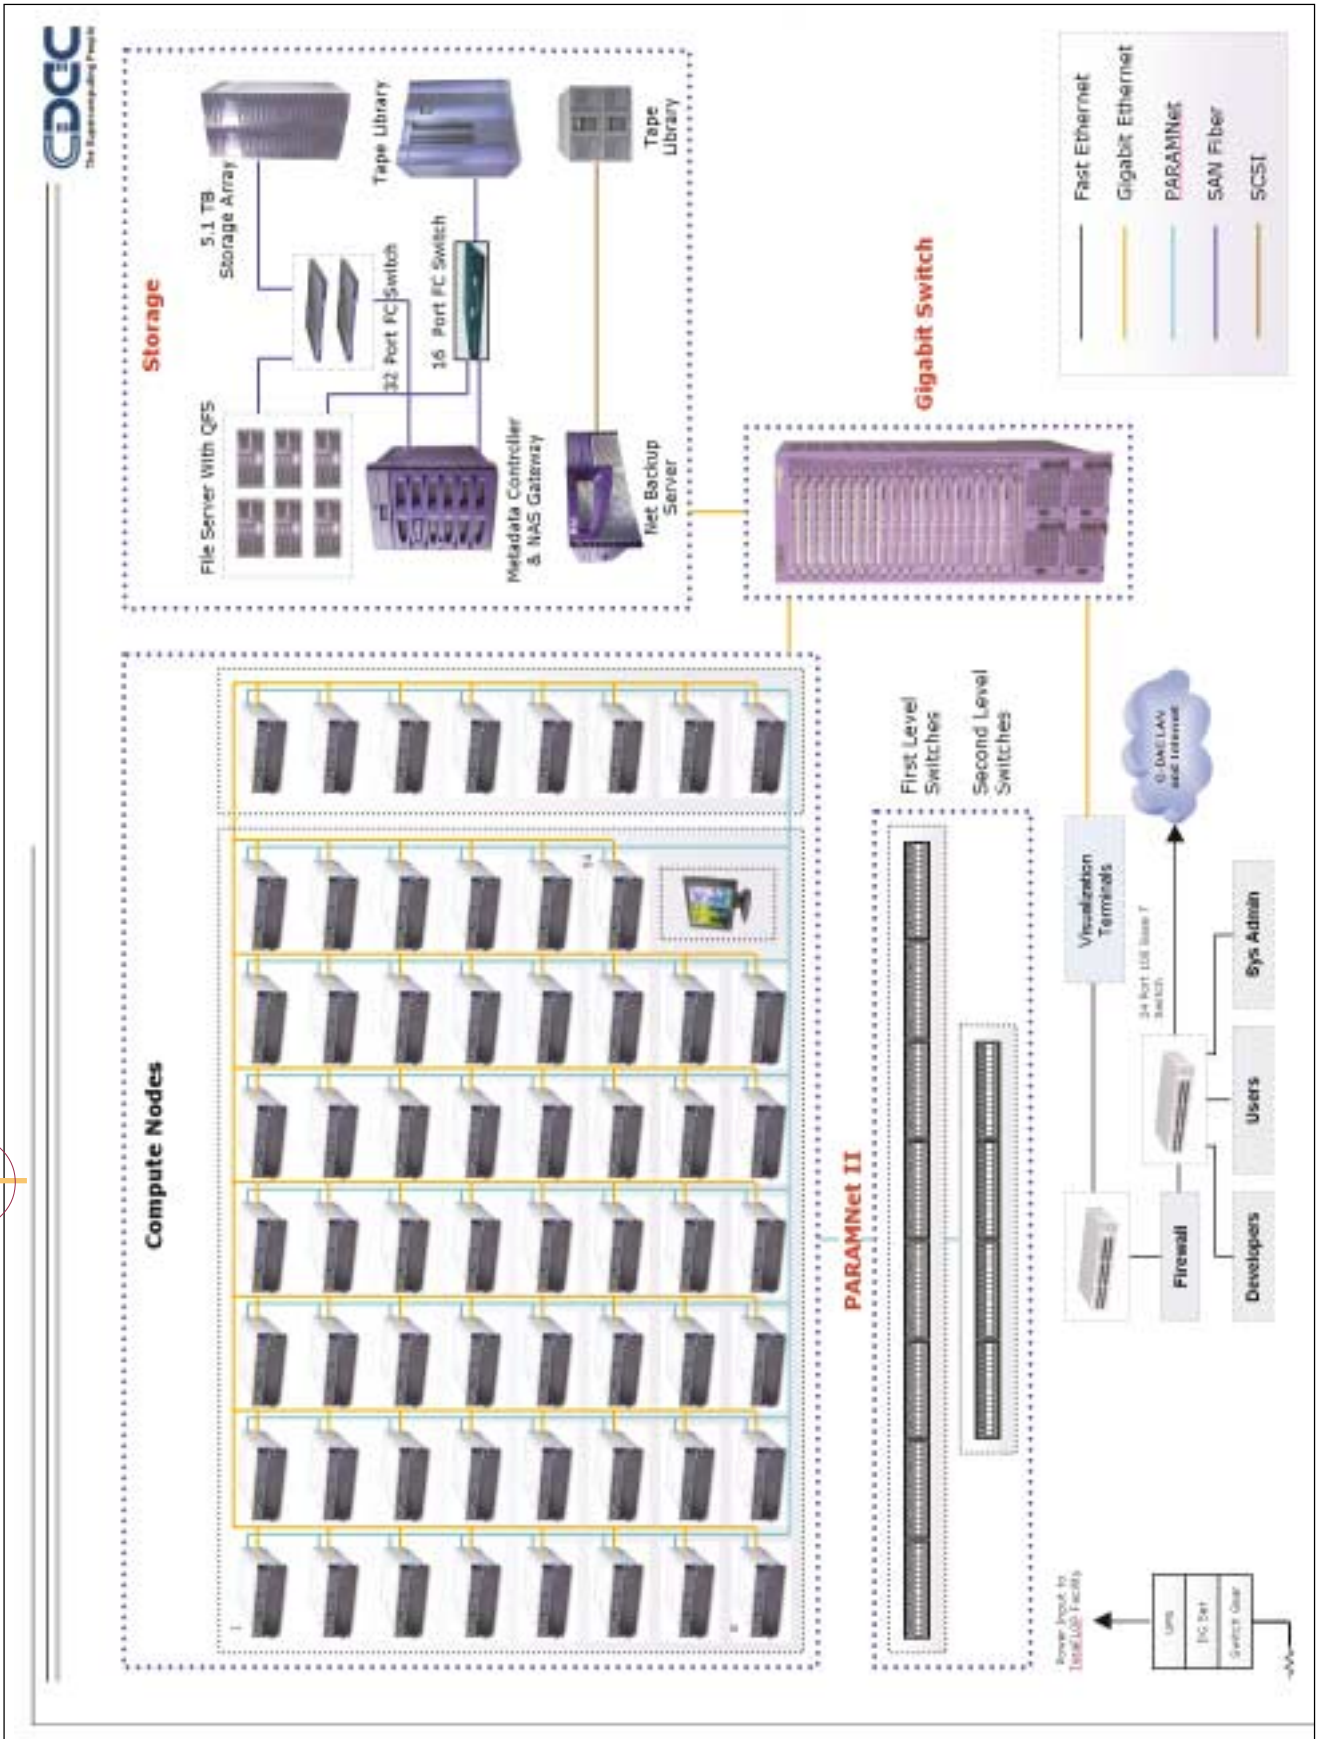


Industrial Design and Packaging

The industrial design & packaging of PARAM Padma offers flexibility to scale from a system with a few nodes to systems having a large number of nodes. The PARAM Padma enclosure has been designed taking into consideration environmental requirements of high performance computing sub-assemblies. It is a 19 inch 38 U standard rack with shielded cable trays and accessories for housing Compute Nodes, File Servers, Network Switches and Cables.



configuration





specifications

PARAM PADMA: One TF Cluster

1. Compute Nodes:	· Configuration	:	54 Nos. of 4 Way SMP & 1 No. of 32 Way SMP	
	· No. of Processors	:	248 (Power 4@1GHz)	
	· Aggregate Memory	:	0.5 TeraBytes	
	· Internal Storage	:	4.5 TeraBytes	
	· Operating System	:	AIX / LINUX	
	· Peak Computing Power	:	1005 GF (~1 TF)	
2. File Servers	· Configuration	:	6 Nos. of 4 Way SMP	
	· No. of Processors	:	24 (UltraSparc-III@900MHz)	
	· Aggregate Memory	:	96 GigaBytes	
	· Internal Storage	:	0.4 TeraBytes	
	· File System	:	QFS	
	· Operating System	:	Solaris	
3. Networks	· Primary	:	PARAMNet-II @ 2.5 Gbps Full Duplex	
	· Backup	:	Gigabit Ethernet @ 1 Gbps Full Duplex	
4. External Storage	· Storage Array	:	5 TeraBytes	
	· Tape Library	:	12 TeraBytes	
5. System Software	:	C-DAC's HPC Suite of System Software		
		I)	C-MPI	
		II)	C-PFS	
		III)	DIViA	
		IV)	F90IDE	
		V)	KSHIPRA	
		VI)	MetricAdvisor	
		VII)	PARMON	
		VIII)	PCF90	
		IX)	RMS	
6. Mechanical	:	No. of Racks	:	15 (8 Compute Nodes, 2 File Servers, 3 Networks, 2 Storage)
7. Total Weight	:		:	9 Tons
8. Total Power	:		:	160 KVA
UPS	:		:	2 * 200 KVA
DG Set	:		:	500 KVA
Battery BackUp	:		:	30 Minutes
9. Cabling	(i)	Networks	:	12Kms (CAT6) & 2.5Kms(Fibre Optic)
	(ii)	Power	:	5.5Kms
10. Air Conditioning	:		:	39 Tons
11. Temperature /Humidity	:		:	18±2 Degree Celsius/50±3% RH
12. Total Machine Floor Area	:		:	168 Sq.metre (1800 sq.ft.)

applications



APPLICATIONS on PARAM Padma

Real life complex problems and scientific and engineering research are the driving force behind the development of PARAM Padma. C-DAC's expertise also lies in developing, porting and demonstrating applications in critical scientific and engineering fields such as computational atmospheric sciences, bioinformatics, computational structural mechanics, evolutionary computing, computational chemistry, computational fluid dynamics and seismic data processing.

Bioinformatics



- Expertise in analyzing large sequence databases using popular tools like BLAST, FASTA, Smith-Waterman for studying the functions of genes and proteins to identify drug targets.
- Realistic simulations of large biomolecules using molecular modeling codes like AMBER, CHARMM & GROMACS.
- Research initiatives aimed at utilizing the right combination of Genomics, Bioinformatics and High Performance Computing power for better health care.

Computational Structural Mechanics



- Specialized software for structural analysis, tuned for high-end parallel architectures, modeling and visualization tools.
- FEMCOMP, a stress analysis package for fibre reinforced composite structures based on higher-order shear deformation theory.
- INTEGRA, a platform independent, modeling and visualization package for structural and fluid mechanics.
- NONLIN (nonlinear stability analysis of shells), a collaborative project for studying thin walled structures.
- FRACT3D (fracture analysis), a parallel software based on domain decomposition.

Computational Atmospheric Science




- Medium range (3-10 days) weather forecasts and long range (~100 years) climate forecasts.
- Pre-processors to ingest the Indian meteorological data into the MM5 modeling system for regional weather forecast.
- Climate System Model consisting of computationally intensive components: Atmosphere model, Ocean Model, Land Model, Sea-Ice Model and Coupler.
- World ocean circulation model based on the system of primitive equations for momentum, stream function, heat and salinity.

meteorological

Seismic Data Processing




- WAVES, a parallel seismic modeling and migration package for seismic data processing in oil and natural gas exploration.
- High precision 2-D and 3-D migration and modeling algorithms.
- 2-D and 3-D migration codes using wavelets.



Evolutionary Computing




- Parallel genetic algorithms based methodologies for protein structure prediction and multiple sequence alignment.
- Financial modeling for analyzing stock market trends.
- Internal Combustion Engine design optimization by modeling and simulation of fuel flow characteristics.
- A parallel genetic algorithms library for general optimization problems.



Computational Chemistry




- INDMOL, a package developed for electronic structure and molecular properties based on theoretical quantum chemistry.
- MP2 level calculations using GAMESS with much less turn around time.
- Large-scale calculations of electronic structure and properties of molecules.
- *ab initio* molecular dynamics simulation software developed and optimized.



Computational Fluid Dynamics



- Simulation of external and internal flows.
- Finite volume code for re-entry vehicles at hypersonic flows.
- General purpose navier stokes code for large simulation.





C-DAC Knowledge Park, Bangalore



Centre for Development of Advanced Computing

Pune University Campus, Ganeshkhind, Pune - 411 007, India

Tel: +91-20-569 4000/01/02/03

Fax: +91-20-569 4059

website: <http://www.cdacindia.com>

Bangalore

C-DAC Knowledge Park, No. 1, Old Madras Road,
Byappanahalli, Bangalore - 560 038, India

Tel: +91-80-534 1874, 534 1909

Fax: +91-80-524 7724

e-mail: bdm@cdacindia.com

New Delhi

A 335, Shivalik Enclave,
Near Malviya Nagar,
New Delhi - 110 017

Tel/Fax: +91-11-667 4689/91/97

e-mail: bd@cdacindia.com

Hyderabad

2nd Floor, Delta Chambers,
Ameerpet,
Hyderabad - 500 016

Tel: +91-40-340 1331/32

Fax: +91-40-340 1531

Centres:

Chennai +91-44-3719326

Kolkata +91-33-3212357

Thiruvananthapuram - +91-471-554681

<http://www.cdacindia.com>