4

HIGH PERFORMANCE COMPUTING & CYBER SECURTY

Computation is the third pillar of scientific discovery. It is an in-expensive way to achieve high science, which complements theory and observations. It has become such a necessity that, the capability and credibility of a scientific organization is currently judged by the computational facility it has access to. CSIR 4PI provides state-of-the-art High Performance Computing facility to the computational scientists and researchers of CSIR to address Grand Challenge problems in their frontier areas of science and engineering through a centralized High Performance Computing facility. The facility at CSIR 4PI is one of the fastest in the country and is aimed at providing multiple architectures suitable for domain specific applications. The facility is accessed by all the CSIR laboratories, through the high speed National Knowledge Network. The group is involved in research on cyber security. Under the 12th five-year plan of CSIR work has been initiated to develop a "Cyber Security Research & Observation" abbreviated as "CySeRO" to carry out research in the field of Cryptography and Cyber Security.

Inside

4.1 Active TCP responder for cyber security inference

It has been well established that malicious activities on the Internet trigger certain special type of network traffic called 'unsolicited packets'. There exist a wide variety of malicious activities like worm propagation, port-scanning, Internet Protocol (IP) address spoofed Denial-of-Service attacks, etc. which generate unsolicited packets. CSIR Fourth Paradigm Institute, under its Cyber Security Research and Observation (CySeRO) programme, has been collecting and analyzing such unsolicited packets for better security inference.

It is a general perception that Internet hosts from where these unsolicited packets originate are hosting malware. Hence, the source addresses of these unsolicited packets are typically considered as the identity of infected hosts on the Internet. However, there is a potential risk of misidentification due to the prevailing IP address spoofing on the Internet. IP address spoofing is the process by which the actual source IP address of a packet is replaced with the IP identity of somebody else for obvious reasons. Hence the source IP address claimed in the unsolicited packets needs to be validated before using them for further inference.

We are developing a framework for source address validation of unsolicited packets. The software is developed in python using the scapy library. It consists of two tightly coupled modules: a passive packet listener to receive unsolicited packets and an active responder which prepares and sends appropriate TCP/IP packets as a response to the unsolicited packets. Acceptable and timely response, subsequent to this, from the remote host is analyzed to determine whether the initial unsolicited packet was IP address spoofed or not.

Considering the fact that majority of the unsolicited packets are based on the Transmission Control Protocol (TCP), the most popular and widely used transport layer protocol on today's Internet, our first version of the active responder validates only TCP packets. It listens to TCP connection requests (SYN packets) and responds to such requests with protocol compliant connection response (SYN/ACK). Further response to SYN/ACK, if any, is validated according to the TCP specification to ascertain the authenticity of the unsolicited SYN packet. This leads to the responder successfully establishing a TCP connection with the sender of the unsolicited packets through the well-known three-way handshake mechanism of TCP. The active responder is being integrated to the CySeRO testbed. Its preliminary deployment for a period of one-week in our NKN (National Knowledge Network) perimeter leads to interesting observations.

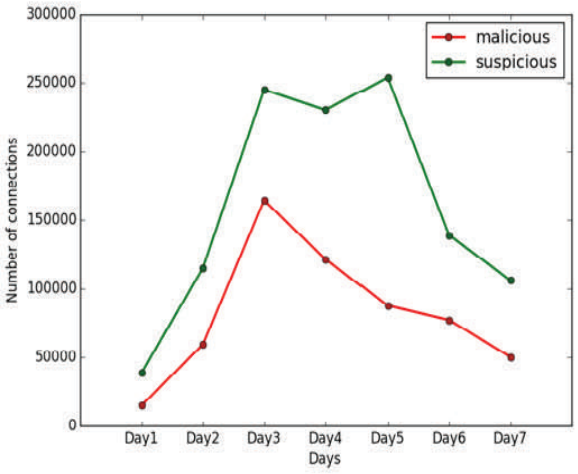

Figure 4.1 Classification of unsolicited packets through active TCP responder.

In Figure 4.1, suspicious indicates the number of SYN requests received and malicious indicates the source IP address validated connection by our active responder. It is important to note that only about half of the suspicious connection requests could be confirmed as malicious. The remaining connections either have spoofed source IP addresses or did not purposefully respond to the active responder's traffic.

We are currently strengthening our active responder with advanced classification schemes based on the further response received by our active responder. Our ongoing work uses TCP RST packets and values in the protocol fields such as TTL (Time to Live), Fragmentation Identify, Time Stamp Echo, etc. for further inference and classification. Active TCP responder will also be expanded to an active application responder (e.g. http responder) to respond to unsolicited application traffic for interacting with the *malicious hosts* until the host terminates the connection for timely extraction of *malicious payloads*. The responder can also be extended for other transport layer protocols like Stream Control Transmission Protocol (SCTP) and Multi Path Transmission Control Protocol (MPTCP).

V Anil Kumar, Sujata, Chinmaya Mohini and Kirthi Sagar

4.2 Leveraging bigdata technologies for cyber security inference

Cyber Security is a source of bigdata. For effective cyber security inference at regional, national and global scale, one needs to deal with massive amount of diverse data. Such data typically include application level logs from Internet servers, user access patterns, security logs from middle boxes such as firewalls, Intrusion Detection Systems (IDS) and Intrusion Prevention Systems (IPS), raw network packets consisting of multi-layer protocol information and application payload, etc. Among them, raw network packets, if collected, stored and analyzed, can lead to vital security inference.

Raw network packets on high-speed network can grow to several tera-bytes in a relatively short span of time. They are typically semi-structured with standard protocol headers and varying forms of payloads in each packet. Collecting, storing and effectively querying such volume of semi-structured data needs state-of-the-art technologies. CSIR-4PI, for its CySeRO (Cyber Security Research and Observation) programme, is exploring next generation database technologies to deal with the massive amount of raw network packets classified as unsolicited packets. In particular, we have designed a NOSQL (Not Only Structured Query Language) based database using the emerging bigdata tool called Cassandra.

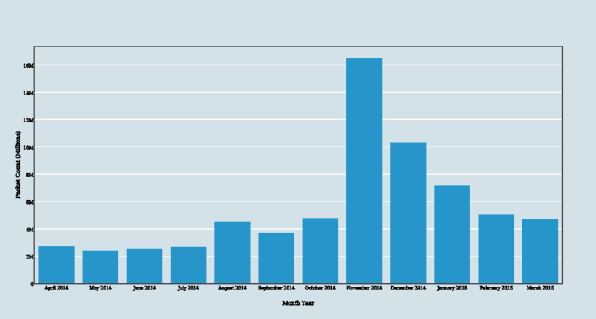

Figure 4.2 Volume of Unsolicited Packets

Our database currently consists of about 100 million raw unsolicited network packets. Each packet consists of multi-layer protocol headers (TCP/IP/ Ethernet) and application payload. Using a customized web interface, queries can be launched to extract any protocol information from the raw packets

in a timely manner. For time optimization of long-term queries, we have incorporated multi-thread based queries.

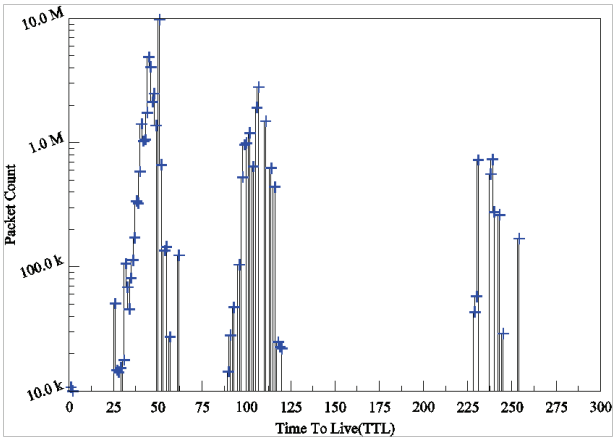

Figure 4.3 TTL Distribution of Unsolicited Packets

Figure 4.3 gives packet-count vs. time information extracted from the database for the past 12 month period. Figure 4.3 shows the extracted TTL (Time To Live) information from the Internet Protocol (IP) header of the raw packets. The three clusters below 64, 128 and 256 indicate the three widely used default TTL on the Internet.

V Anil Kumar, Navi Thejesh, Sudeep Nesakumar

4.3 Hierarchical trust model based on security to rate cloud service providers

In large scale distributed systems like cloud computing, customers need to interact with unknown cloud service providers (CSP) to carry out tasks or transactions. The ability to reason about and assess the possible risks in carrying out such transactions is necessary for providing a safe and trustworthy environment. Cooperative characteristics of distributed computing systems enforce a proper and secure trust management to be in place to minimize the risks posed by different malicious agents. Trust is the estimation of competency of a resource provider in completing a task based on dependability, security, ability and availability in the context of distributed environment. It enables users to select the best resources in the heterogeneous cloud infrastructure.

Typically a trust evaluation model comprises of two stages and three Key Performance Indicator (KPI). The first stage is the implementation with the help of Mamdani Fuzzy Inference System (FIS), which evaluates Performance, Financial and Agility parameters. The second stage implementation takes the output of the first stage FIS and helps to obtain the trust rating for each plan of the CSP.

In these models no emphasis is given to data security, including the risk of loss, unauthorized collection. So the model is extended to include two more parameters - Security and Usability. The Security parameter is described in terms of the Physical Security, Internal Security and Network Security levels available with the cloud provider, while the Usability parameter of the model is calculated based the contributions from the Understandability, Easability and Flexibility. Figure 4.4 shows the block diagram of the proposed model based on security.

Table 4.1 shows the trust values of the CSPs corresponding to the model in Figures 4.4. Here the Recommended Trust values are higher than the Direct Trust values except for Rackspace Performance Two due to higher total processing cost (in $) for the user requests. This is also reflected as a considerable reduction in Recommended Trust value for the Finance based model.

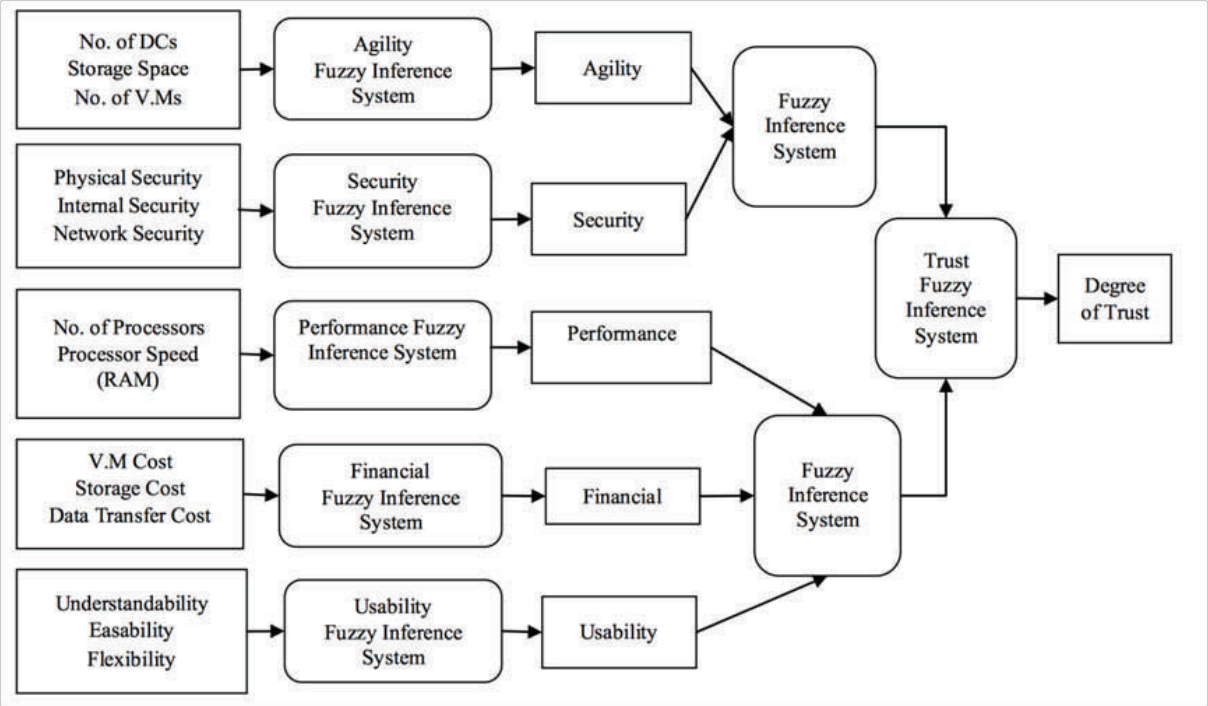

Figure 4.4 Hierarchical Model Based on Security

Table 4.1 Trust values of different CSPs with security as a parameter

CSP and Server type	Security based	
	Direct Trust	Recommended Trust
Gogrid Standard Dedicated Server	0.589	0.606
Gogrid Advanced Dedicated Server	0.585	0.607
Gogrid Ultra Dedicated Server	0.585	0.661
Gogrid Elite Dedicated Server	0.679	0.679
Rackspace Enhanced One	0.578	0.605
Rackspace Enhanced Two	0.585	0.585
Rackspace Performance One	0.67	0.755
Rackspace Performance Two	**0.755**	**0.67**
Amazon EC2 Small	0.471	0.578
Amazon EC2 Medium	0.586	0.619
Amazon EC2 Large	0.65	0.755
Cloud flare Pro	0.498	0.578
Cloud flare Business	0.64	0.649
Cloud flare Enterprise	0.67	0.67

Another important observation is that the priority based model is better in distinguishing between various plans. In the non-hierarchical model, where all parameters have equal weights, trust values of all the plans fall in a shorter range makeing it difficult to rank the CSPs. But in a priority based model with Finance / Security, the range varies from 0.471 to 0.755. Thus we can rank the various service provider plans.

Thus it is seen that with Security as the main requirement Rackspace Performance One and Amazon EC2 Large would be preferable. Such a conclusion cannot be arrived from a non-hirarichel model.

M Supriya*, Sangeeta K*, G K Patra
*Amrita School of Engineering, Bangalore

4.4 Secure key stream generation using particle swarm optimization

Neural Synchronization has been used to construct a cryptographic key agreement protocol using Particle Swarm Optimization (PSO) to accelerate the mutual learning between the sender and receiver in a Tree Parity Machine (TPM) neural network. PSO is a population based optimization algorithm that is motivated from the simulation of social behavior. Each individual in PSO flies in the search space with a velocity that is dynamically adjusted according to its own flying experience and its companions' flying experience. Compared with other evolutionary algorithms, such as GA, PSO algorithm possesses some attractive properties such as memory and constructive cooperation between individuals, so that it has more chance to "fly" into the better solution areas more quickly and discover reasonable quality solution much faster.

Neural cryptography defined by Kinzel and Kanter uses multilayer feed forward neural network called a tree parity machine. Here the two partners (A and B), receive the identical input vector and are trained, with the output of their partner. Each TPM consists of K hidden units and each hidden units receive different N inputs. Synchronization in this process is a stochastic, and is well balanced by attractive and repulsive forces. Though the method is very encouraging, it is weak in security, against Majority Flipping Attack (MFA).

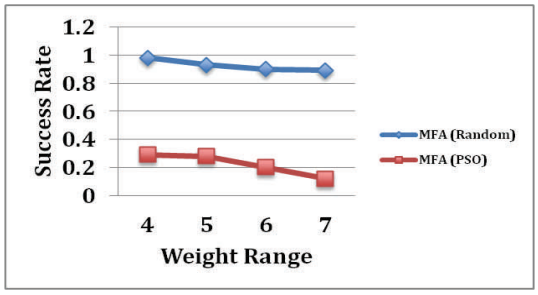

Figure 4.5 Probability of Success between a Random and PSO based initial weight vector against Majority Flipping Attack

PSO is used to find a best-fit weight vector, such that the synchronization happens faster. By doing so it does not give enough time for the attacker to synchronize. Figure 4.5 shows the probability of success by Majority Flipping attack with Random and PSO based weight vector. It can be seen that the possibility of success is very low in case of PSO based initial weights. Unlike the case of random initial weights, where the probability of success remains high even with increase in weight range, in PSO based initial weights, the probability becomes negligible with higher weight range. This promises to be a potential public-key key exchange mechanism.

S Santhanalakshmi*, Sangeeta K*, G K Patra
* Amrita School of Engineering, Bangalore

4.5 High Performance Computing

The centralized High Performance Computing facility located at CSIR 4PI is the main lifeline of the computational scientists across CSIR. The supercomputer with a peak computing power of 360TF and a sustained computing capability of 334 TF on a High Performance LINPACK (HPL) is currently the 3rd fastest system in the country and 155th fastest in the world. The supercomputer after its released to the CSIR community on 6th September 2013, has clocked more than 86% of average utilization during 2014-15.

Figure 4.6 CSIR centralized 360TF High Performance Computing Facility.

Each node in the system is a HP Blade server, with two Intel Xeon E-5 2670 (8 cores, Sandy bridge) processors. There are 1088 such computing nodes distributed over 17 numbers of 42U 600mm width racks, resulting in 2176 physical processors and 17408 processing cores. For communication among the nodes, the nodes are connected using high speed FDR infiniband interconnect (providing a dedicated 56 Gbps interconnect bandwidth) in a FAT tree topology and a high availability mode. To achieve the high throughput and high availability, the nodes are connected through two numbers of centralized Mellanox 648-port core switch and a number of 16 port leaf switches to the computing and storage racks in redundancy mode. Memory per core is one of the important aspects of the system. The nodes are designed with 4GB memory per core, which results in about 68TB of distributed memory in the total system. However, 48GB memory can be used for shared memory inside a single node for parallel applications.

A LUSTRE parallel, online storage, plays an important role in the overall performance of the system. The high performance computing system has a high performance parallel file system of about 2.1 Peta Byte (3 Peta Byte pre-formatted) capable of providing minimum 20 Gbps simultaneous read and write capability. The storage based on the popular open source LUSTRE file system, is optimized for performance and data availability, by using hardware RAID in a RAID6 configuration, parallel I/O through 8 numbers of object servers and two numbers of redundant metadata servers.

Figure 4.7 shows the intra-day maximum usage (number of nodes) for the year 2014-15. The size of the problems run by the CSIR computational scientists range from 1 node (16 cores) to about 250 nodes (4000 cores). It is interesting to note that the maximum nodes used in a day have reached 100% most of the days, indicating the heavy utilization of the systems by scientists and researchers. Figure 4.9 shows the percentage distribution of usage by different CSIR Laboratories in various fields of computational sciences, such as Biological, Chemical, Engineering, Earth and At-

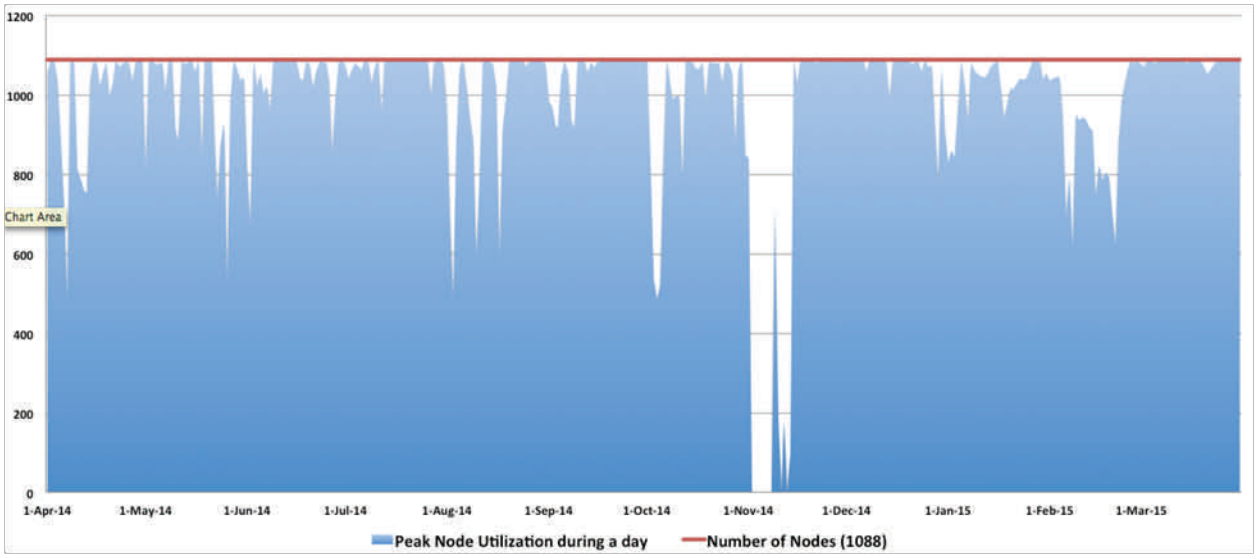

Figure 4.7 Intra-day maximum node used since the 1st September 2013 till 31st March 2014

mosphere, Physical and Information Sciences. The jobs are run in both capacity and capability mode for solving scientific problems, depending the nature of the models.

The demand of computing by CSIR community can be judge by the continuous heavy usage of Altix ICE cluster (Figure 4.8) even after being in service for more than 3 years. The system with 2304 numbers of processing cores distributed over 192 nodes inter-connected in the form of an enhanced hypercube using the QDR (32Gbps) infiniband interconnect was heavily utilized in the year 2014-15. The system has Intel Westmere-EP Hex core processors running at 2.93/3.06 GHz frequencies, with each node having 12 processing cores with 24 GB of memory in a shared memory configuration, while the system as a whole has 4608 GB of memory across the 192 nodes in a distributed architecture. The peak performance of the system is 27 TFLOPS. This system also uses a LUSTRE parallel file system of 30TB for high performance storage access during computation.

Figure 4.8 Altix ICE system with 2304 processing cores distributed over 192 nodes and 30 TB of parallel file system along with all associated hardware and software.

Efficient utilization of a system depends on maximizing the usage and minimizing the wait period. All the High Performance Computing systems at CSIR 4PI use PBSPro workload manager, Intel compilers and other essential software to efficiently manage and run jobs on the system. The workload manager not only ensures efficient usage of the system but also provides an easy user interaction and submission interface.

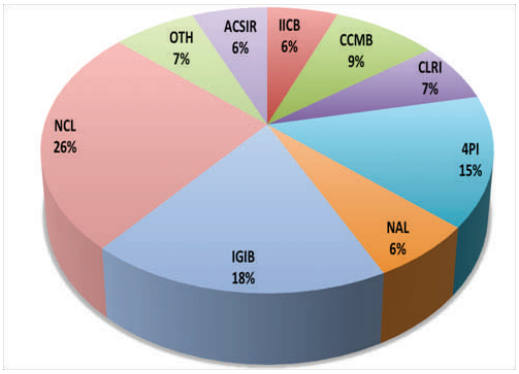

Figure 4.9 Distribution of usage of 360TF system in percentage by major CSIR labs from 1st April 2014 till 31st March

High Performance archival Storage

The parallel file systems are typically expensive, hence, normally used as scratch, to achieve performance, during job computation. To store and archive results, which need to be preserved for a longer period, an archival system based on a high performance SAN (Storage Area Network) is made available to the users. The archival system is upgraded regularly to support the growing need of data storage. The SAN archival system has four numbers of LTO Gen 5 drives. Currently the virtualized 3-tired storage solution has 6 TB online (FC), 20TB of near-line (SATA) and about 1.5 PB of offline storage. The home areas of all the users are centralized on a Network Attached Storage (NAS) of 200TB.

Data Center

The HPC systems are supported by a Tier-3 equivalent state-of-the-art data center along with an associated energy farm. The water based cooling mechanism using Rear Door Heat Exchangers (RDHX) makes the datacenter, one of the high density and high power efficient datacenters in the country.

The Power Usage Efficiency (PUE) of less that 1.5 is one of the best-achieved PUEs in a country like India. The energy farm consists of two numbers of compact substations of 1.25MVA each and backup power by using three numbers of diesel generators, an underground diesel yard of capacity more than 15000 liters, three numbers of UPS with battery backup etc.

The datacenter is monitored through a well-designed Building Management Service (BMS). The system, the electrical infrastructure, fire detection and suppression system, very early smoke detection system, water leakage system, CCTV, rodent repellant system is monitored continuously through an integrated building management services.

Network Facilities

National Knowledge Network (NKN), has greatly enabled high speed and reliable access to the centralized computing facility from other CSIR laboratories. Currently the NKN connectivity to CSIR-4PI is at 1 Gbps. The institute also has a backup connectivity of 8 Mbps through ERNET. Scientists and researchers of CSIR- 4PI and CSIR NAL (all the three campuses) use the facility from their desktops through a 10 Gbps high-speed backbone. All network services namely DNS (Domain Name Server), NIS (Network Information Services), WWW (World Wide Web), institutional repository, webmail, mail services, Intranet and Internet gateways (both for ERNET and NKN connections) have been shifted to newer systems, for efficient communication and data dissemination. An new

Unified Threat Management (UTM) has been procured and installed to address multiple security threats. This ensures safe communication with Internet through both the NKN and ERNET links.

Software Enhancements

To keep pace with the fast enhancement of hardware, application software were maintained and upgraded. Some of the heavily used software are ABAQUS, CFD-ACE+, IDL, GAMIT/GLOBK, Tecplot, S-Plus, Hyperworks, Fluent, ANSYS, OpenFOAM etc. CSIR 4PI also encourages use of open source software and most software required for modelling and simulations are made available on the HPC systems for users. The systems are used extensively for running complex models in the field of ocean, atmosphere, earth and engineering, biology, chemistry.

Other Technical Services

Technical support was provided to a large number of users from CSIR-4PI & CSIR NAL. The team also provided web-hosting facilities for organizing different workshops and conferences during this period. In addition, several students from academic institutions across the country have availed the computing services as part of their academic work at CSIR-4PI under the SPARK program. Technical advices and consultancies were provided to various institutions within and outside CSIR.

R P Thangavelu, G K Patra, V Anilkumar, Ashapurna Marndi, N Prabhu, Nagaraju V Y Mudkavi*, Premalatha*
*National Aerospace Laboratories