4

HIGH PERFORMANCE COMPUTING & CYBER SECURTY

The importance of computation in science is well known. In contemporary research, the capability of a scientific organization is judged by the computational facility it has access to. CSIR-4PI (erstwhile C-MMACS) has been given the responsibility to provide world-class computational facility to the computational scientists and researchers of CSIR to address Grand Challenge problems in their frontier areas of science and engineering. The facility at CSIR-4PI (erstwhile C-MMACS) is one of the fastest in the country and is aimed at providing multiple architectures suitable for domain specific applications. The HPC group at CSIR-4PI (erstwhile C-MMACS) is also involved in research on cyber security, the importance of which can be only realized when computing infrastructure is provided as a centralized service over Internet. Under the 12th five year plan CSIR has sponsored a sub-project "CySeRO" to carryout research in the field of Cryptography and Cyber Security.

Inside

- *Cyber Security Inference through Unsolicited Traffic Analysis*
- *Automation of Security Assurance Process based on PCI-DSS for Cloud Computing*
- *Fully Homomorphic Encryption*
- *Cloud based High Performance Computing and its Security*
- *Method and Device for Categorizing a Stream Control Transmission Protocol (SCTP) Receiver Terminal as a Malicious SCTP Receiver (International Patent, filed on 27.03.2014)*
- *High Performance Computing*

4.1 Cyber Security Inference through Unsolicited Traffic Analysis

Unsolicited traffic continues to occupy a portion of the overall Internet traffic. It was found that during November 2010, about 5.5 Gigabits of unsolicited traffic was generated every second on the Internet. Also, a modem user would lose about 20 bits per second of his bandwidth to the unsolicited traffic. These are typically TCP/IP (Transmission Control Protocol/Internet Protocol) packets addressed to globally routable IP addresses, which are not assigned to any network devices.

Unsolicited network traffic can be used for remote inference of cyber security incidents. For example, one of the major sources of unsolicited traffic is Internet worm propagation. A worm-infected machine, in the process of spreading the worm to other hosts on the Internet, generates unsolicited packets. Likewise, Internet Botnets or Zombies, trigger unsolicited networks as part of network scanning to identify vulnerable hosts on the Internet. Misconfigured hardware, reflections from IP spoofed Denial-of-Service attacks and leaks from private networks are other potential sources of unsolicited traffic. By capturing, analyzing and actively responding to unsolicited traffic, significant insight can be gained to the nature and prevalence of various malicious incidents.

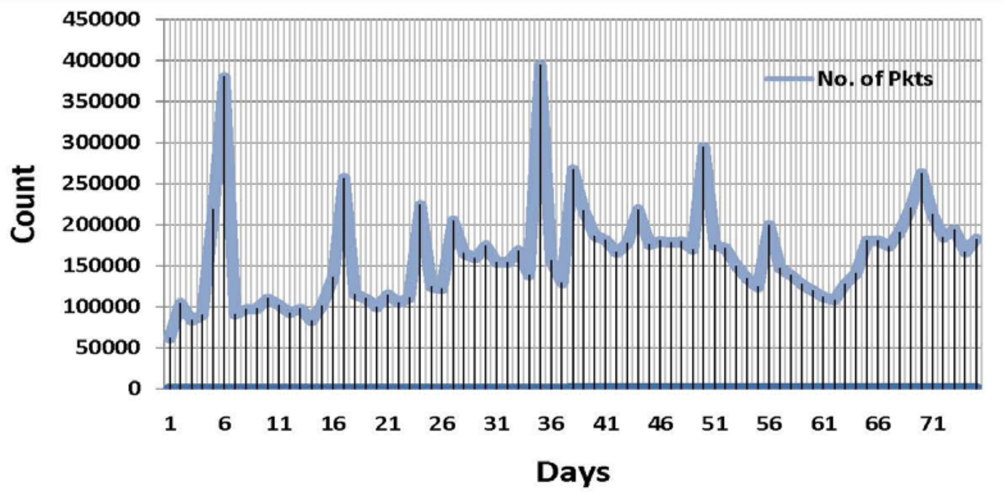

Figure 4.1 Number of unsolicited packets collected at CSIR-4PI (erstwhile C-MMACS) network.

CSIR-4PI (erstwhile C-MMACS) has been focusing on unsolicited network traffic analysis for past several years. In order to strengthen this activity, a comprehensive software framework is being designed and developed to capture and analyze such traffic. Our broad objective is to gain better understanding of the security dynamics on the Internet as whole and regional networks of particular interest to us. The tool is being implemented in python with several modules for various analyses.

The software is capable of identifying different type of network packets and accordingly parsing the packet headers at different layers such as Ethernet, ARP, IP, TCP, UDP, SCTP and other application protocols. Figure 4.1 shows the number of unsolicited packets collected at CSIR-4PI (erstwhile C-MMACS) network during a time span of about 75 days.

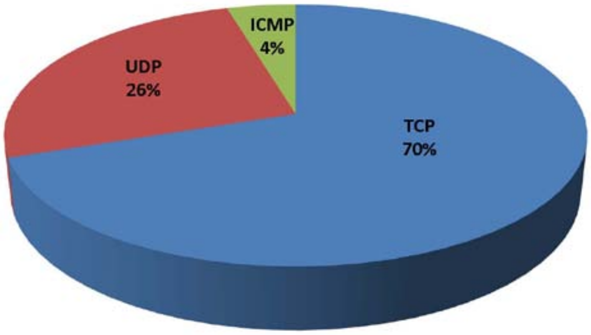

Figure 4.2 Protocol level break-up of unsolicited traffic.

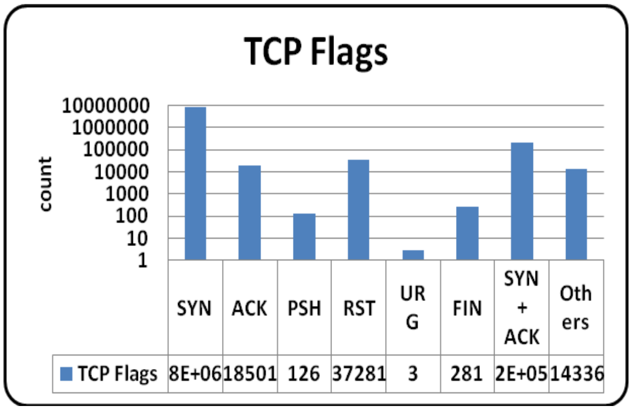

Figure 4.3 Flag-level break-up of TCP packets.

Figure 4.2 provides a protocol level break-up of unsolicited packets. As expected TCP, being the most popular transport layer protocol, accounts for the maximum portion of the unsolicited traffic.

The tool can further penetrate into the minute protocol details. For example, Figure 4.3 provides a flag level break-up of all unsolicited TCP packets received. Majority of them are connection initiation request (SYN packets). This could be either a active port scan or worm propagation attempt over TCP protocol.

Other features are being designed and integrated into the tool. Our immediate focus is on a trend analysis module to automatically generate trends from short and long-term data. Other feature enhancements include parallelizing the code with MPI for faster execution on supercomputer class of machines and development of GUI for data visualization.

Anil Kumar V, Sudeep Nesakumar S [*] and Patra G K
[*] SPARK (VIT University, Chennai Campus)

4.2 Automation of Security Assurance Process Based on PCI-DSS for Cloud Computing

Security and privacy have always been a challenge and are of primary concern with computing resources, and with the invention of cloud computing the need to secure the data stored in the cloud becomes much more challenging. As the cloud appears to be a black box, the user of the cloud is completely unaware of the security of the data residing in the cloud servers. To address such security concerns, standard organizations have devised a number of guidelines that are needed to be followed by the Cloud Service Providers to ensure secured and reliable services to the users. There are a number of standard security compliance guidelines available and a vendor is expected to comply with at least one of them to assure security if their services need to gain user trust. Some of the standard compliance guidelines are ISO 270001/2, PCI DSS, HIPAA, FedRAMP, SOC 1/ SOC 2/ SOC 3 etc. In this work, we have focused primarily on PCI DSS which is the mostly followed standard for security assurance. PCI DSS stands for Payment Card Interface - Data Security Standard and it was originally developed to enhance and ensure data security for the cardholder users. It is comprised of twelve requirements and any organization that needs to be PCI DSS certified must follow all the twelve guidelines. The certification is obtained by a rigorous auditing process that is carried out by PCI Qualified Security Assessors (QSA). Since the audit process is being carried out manually, there is a huge gap between security requirements and the actual measurements taken to provide the security. Moreover, the audit needs to be carried out after every three months to ensure the service is still according to the guidelines and incorporate any update to the original guidelines. To ease the process, it is best if the whole process can be automated which will reduce the requirement of human intervention making the audit process much faster. In this work we have proposed an automated system for the audit process of PCI DSS. Instead of performing it manually, we install a newly designed dedicated software agent authenticated by PCI Approved Scanning Vendor (ASV). This software agent monitors the CSP's services and notifies the security manager of the CSP if there is some mismatch in their service from the compliance guideline. The software can also be programmed to notify the manager about any changes or modifications made in the original guideline such that the CSP incorporates those changes in their operations.

*Tejas N Rao, Patra G K and Nilotpal Chakraborty**
* Devi Ahilya Vishwavidyalaya, Indore

4.3 Fully Homomorphic Encryption

Computing in an encrypted domain has been an intense area of research for cryptographers and computer scientists for quite a long time now and the scheme that supports such operations is fully homomorphic encryption (FHE). Fully homomorphic encryption provides a third party with the ability to perform simple computations on encrypted data. Typically, a third party can calculate one of the encrypted sums or the encrypted product of two encrypted messages. It can be pictorially represented with the help of figure 4.4.

The concept was originally devised by Rivest et al in 1978 and since then rigorous efforts have been rendered to devise a practical FHE. In 2009, Craig Gentry, an IBM researcher, showed the first plausible construction of fully homomorphic encryption based on the hardness problems of ideal lattices. But eventually, the scheme remained purely theoretical due to its inefficiency, and it cannot be implemented for real life applications.

$$\begin{tikzcd}
\text{Message } (m) \arrow[r, "f"] \arrow[d, shift left=2, "Enc"] & f(m) \arrow[d, shift right=2, "Enc"'] \\
Enc\,(m) \arrow[u, shift left=2, "Dec"] \arrow[r, "Eval"'] & Enc\,(f\,(m)) \arrow[u, shift right=2, "Dec"']
\end{tikzcd}$$

Figure 4.4 Fully Homomorphic Encryption

In this work, our primary focus has been to devise an efficient FHE scheme that would be practical for implementing in the real life scenario. The existing FHE schemes work only bit by bit and due to this they have been inefficient. Our motivation have been to develop the FHE scheme that can take decimal inputs directly, process it, and produce ciphertexts that can be homomorphically evaluated. We started with analyzing the FHE construction as proposed by Van Dijk et. al. and implemented it in C programming language. The scheme at first constructs a somewhat homomorphic encryption that can evaluate polynomials up to a certain limit, after which the decryption does not result in obtaining the actual plain text. To make it fully homomorphic, they followed Gentry's blueprint and suggested to apply the squashing technique. However eventually, we have observed a number of applications where the operations on the ciphertexts are limited and hence we can work with the somewhat homomorphic encryption only. Applications such as image processing, video processing, ASCII character encoding etc. have limited operations, with limited ciphertext space. Thus we can efficiently implement somewhat homomorphic encryption scheme for these applications. We have modified the DGHV scheme to process decimal inputs directly without converting them to their corresponding binary equivalents, thus increasing the efficiency up to a lot of extent as compared to the original scheme. The construction of the modified scheme is shown as follows—

KeyGen: The key is a large prime integer *'p'*.

Encrypt (p,m): Takes an input as a bit m={0,n-1}, and produces the cipher text as $c:=m+nr+pq$, where integers q and r are chosen at random and r is essentially very small than p such that $r<p/2$ in absolute value.

Decrypt (p,c): Takes cipher text c as input and produces the original pain text $m:= (c \bmod p) \bmod n$.

The above scheme works perfectly fine because by removing the p from the ciphertext, we remain with the noise part and the original message. By removing the mask of n, we obtain the original message. This scheme supports homomorphic additions and multiplications till the limit for the noise r is not crossed.

Here, the data can either be decrypted by administration or by Students and Faculty taken together. It cannot be the case that the same data is being accessed both by students/faculty and by the administration section. Functional encryption promises to provide such delegate cryptographic security to the systems using it.

Nilotpal Chakraborty, Patra G K and Anil Kumar V*
* Devi Ahilya Vishwavidyalaya, Indore

4.4 Cloud based High Performance Computing and its Security

In cloud computing, almost all the resources are delivered as a service. It would have been valuable if the same can be incorporated for High Performance Computing (HPC) services. As an HPC system incurs a lot of financial investment as well as manpower needed to run and maintain such systems, it would be very helpful for medium scale enterprises to leverage the benefits of HPC for various scientific and engineering problems in a very cost effective manner if the same can be provided as a service model of cloud computing. In general, supercomputers and HPC systems are used by the scientists for their researches and some limited number of organizations for gaining an upper hand against their competitors and maximizing their profits. By implementing such HPC services through cloud environment, if at all possible, such services can be provisioned for the general public and HPC can be made accessible to a large class of users. As general users, would use it like just another cloud service models, it would free them from the pain of huge investment and tedious maintenance job, as the same would be taken care of by the HPC Cloud Service Provider. The vendors, on the other hand would gain a lot, as now cloud services would not only be used for storage and simple compute operations, it would now be used to perform various important scientific and engineering modeling and simulations including data analytics and data processing in an HPC infrastructure.

It would both be beneficial and challenging for the service providers, as it would increase the number of customers in one hand as well as increase the concerns of security on the other hand. Though various security measurements have been proposed and introduced, only a limited number of techniques have actually succeeded in mitigating all the issues and security challenges. We in our work have proposed to use an advanced cryptographic scheme, called Functional Encryption, which not only allows users to perform computations on encrypted data, but also to provide user specific credentials for encryption/ decryption and operations on the data. The decryption operation in Functional encryption happens as depicted in Figure 4.5.

Functional encryption basically is a public key cryptosystem that provides fine grained access control to the encrypted data. The scheme is based on Fully Homomorphic Encryption which is ultimately the reason behind the possibility of computations on the encrypted information and over to that restricts the user to perform operations based on the public key issued to him.

Functional encryption is a new way towards public key cryptography, where data is not only encrypted and decrypted, we can specify what can be decrypted by whom and what operations can be performed based on the user credential. Such a scheme, if possible to implement, would certainly mitigate, if not all, most of the security problems of cloud computing, high performance computing and big data analytics. As now the owner of the data can decide who can do what and provide specific keys to the users. We have already implemented and tested the use of fully homomorphic encryption for specific applications and further work is being carried out to implement functional encryption. We believe that, such a scheme would be possible to implement in the HPC systems of CSIR Fourth Paradigm Institute and thus the same infrastructure can be provisioned to provide services to greater number of organizations and individuals through the use of cloud computing.

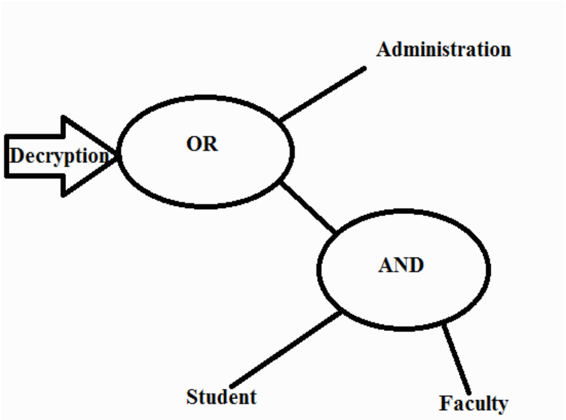

Figure 4.5 Functional Encryption

Patra G K, Nilotpal Chakraborty*, Anil Kumar V
* Devi Ahilya Vishwavidyalaya, Indore

4.5 Method and Device for Categorizing a Stream Control Transmission Protocol (SCTP) Receiver Terminal as a Malicious SCTP Receiver (International Patent, filed on 27.03.2014)

The invention provides a novel method and a system to detect and eliminate optimistic Selective Acknowledgement (SACK) spoofing in a Stream Control Transmission Protocol (SCTP) based communication consisting of a SCTP data sender and a SCTP data receiver with an established association. The said optimistic SACK spoofing can be remotely performed either by a malicious SCTP receiver for exploiting a SCTP sender as a flood source for Denial-of-Service (DoS) attacks or by a greedy SCTP receiver for downloading data from a SCTP data sender faster than normal SCTP receivers. The invention consists of a ***data enriched SACK generation*** *at* the SCTP data receiver side and a ***data enriched SACK validation*** at the SCTP data sender side. The SCTP data sender generates and sends SCTP data packets in the standard header format. The SCTP data receiver generates data enriched SACK packet, which contain a fixed size ***Cumulative Payload Essence*** of application data. Upon receiving a data enriched SACK, the SCTP data sender performs a data enriched SACK validation in which it locally computes the Cumulative Payload Essence using the data packets stored in its retransmission buffer, and compares the locally computed value with the Cumulative Payload Essence received through the data enriched SACK. SACKs whose Cumulative Payload Essences do not match with the locally computed value are marked maliciously spoofed optimistic SACKs even if the Cumulative TSN Ack of SACKs pretends to acknowledge data packets sent by the SCTP data sender. The SCTP data sender discards maliciously spoofed optimistic SACKs and an early termination of SCTP association is performed through SCTP ABORT packet to rescue the data sender from the exploitation.

*Anil Kumar V and Debabrata Das**
* International Institute of Information Technology (IIIT), Bangalore

4.6 High Performance Computing

Computational Scientists of CSIR have been provided with access to one of the largest supercomputing facilities of the country. 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 99th fastest in the world. The supercomputer was released to the CSIR community on 6th September 2013, after extensive testing and optimization to maximize the capability of the system.

Figure 4.6 CSIR centralized 360TF High Performance Computing Facility.

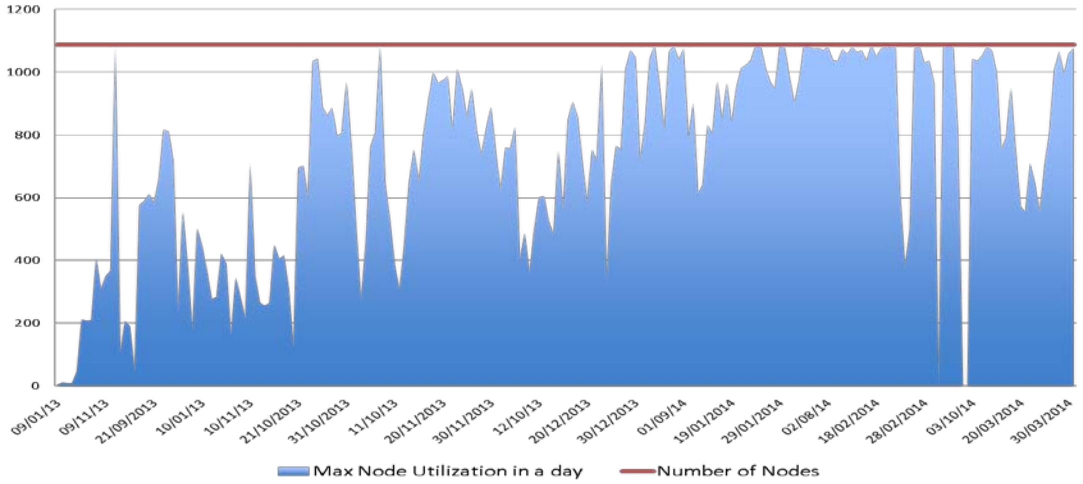

Figure 4.7 Intra-day maximum node used since the 1st September 2013 till 31st March 2014

The system is a cluster of 1088 computing nodes distributed over 17 numbers of 42U 600mm width racks. Each node is a HP Blade server, with two Intel Xeon E-5 2670 (8 cores, Sandy bridge) processors each. Hence, the system has 2176 physical processors and 17408 processing cores. The nodes are connected to each other using high speed FDR infiniband interconnect from Mellenox in a FAT tree topology, which is capable of providing a dedicated 56 Gbps interconnect bandwidth. The high throughput is achieved through a number of 16 port leaf switches in the computing racks connected with two 648-port core switch in a well designed redundancy with respect to availability of all the nodes. The memory per core is one of the important aspects of the system. The memory of about 68TB is distributed across the nodes, and the memory inside a node (48 GB) can be used in a shared memory architecture.

Online Storage, where jobs are run plays an important role in the performance of the system. The system has a high performance parallel file system. The size of the storage is about 2.1 Peta Byte (3 Peta Byte unformatted) capable of providing more than 20 Gbps read and write throughput. This storage is designed using the popular open source LUSTRE file system, with optimized performance tuning. It provides hardware RAID in a RAID6 configuration. The parallel I/O is achieved through 8 numbers of object servers and is controlled through two numbers of redundant meta-data servers.

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.

The users with large computational requirements have been shifted to this new supercomputer. The maximum nodes used in a day from the date of releasing the computational facility is presented in the Figure 4.7. It is worth noting that, the usage has touched the maximum in many days during the period, indicating the need of scientists for such computational capability. The distribution of computing usage by different CSIR Laboratories is shown in Figure 4.9. This indicates the usage of the system in different fields of computational sciences, such as Biological, Chemical, Engineering, Earth and Atmosphere, Physical and Information Sciences. The system has been used in both capacity and capability mode of computing.

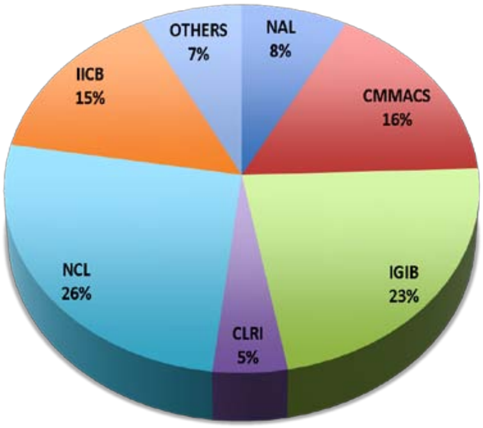

Figure 4.9 Distribution of usage of 360TF system in percentage by Major CSIR labs from 1st September 2013 till 31st March 2014

The computational scientists, who do not have large computational needs still use the Altix ICE cluster (Figure 4.8). The system with 2304 numbers of processing cores distributed over 192 nodes interconnected in the form of an enhanced hypercube using the QDR (32Gbps) infiniband interconnect. The system is equipped with Intel Westmere-EP Hex core processors running at 2.93/3.06 GHz frequencies. Each node has 12 processing cores with 24 GB of memory in a shared memory form, while the system as a whole has 4608 GB of memory across the 192 nodes in a distributed architecture. The peak performance the system is 27 TFLOPS. A lustre parallel file system of 30TB handles the storage requirements for the computing system.

All the computing systems use PBSPro, workload manager, compilers and other essential software to manage smooth transition from one to other. The workload manager ensures the efficient usage of the system and also provides an easy user interaction and submission process.

High Performance archival Storage

While, the working area to run jobs are provided through either a high performance parallel file system or a direct attached scratch file system for better read/write performance, the final storage is done on as archival system using a high performance SAN (Storage Area Network). 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 520 TB of offline storage. The home areas of all the users are centralized on a Network Attached Storage (NAS) of 200TB. Procurements of additional Tapes and archival facility are under process and likely to be installed by June end.

Data Center

The Tier-3 equivalent state-of-the-art data center along with the associated energy farm has been a constant support for the HPC system. The cooling infrastructure is the highlight of the complete setup. The cooling through water-cooling mechanism using Rear Door Heat Exchangers (RDHX) makes the datacenter, one of the high density datacenters in the country. This cooling infrastructure ensures highly efficient cooling with less power consumption..

The Power Usage Efficiency (PUE) of less that 1.5 is one of the best achieved PUEs. 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 the building management service. 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

Thanks to the National Knowledge Network (NKN), the accessibility to the centralized computing facility from other CSIR laboratories has been through a reliable and high speed communication medium. Currently the communication speed is 1 Gbps with a backup connectivity of 8 Mbps through ERNET. Scientists and researchers of CSIR-4PI (erstwhile C-MMACS) and CSIR-NAL (all the three campuses) use the facility from their desktops through a high speed local network interconnected using a 10 Gbps 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 provided for efficient communication and data dissemination.

Software Enhancements

Application software were maintained and upgraded to keep pace with hardware enhancements. The heavily used software are ABAQUS, CFD-ACE+, IDL, GAMIT/GLOBK, Tecplot, S-Plus, Hyperworks, Fluent, ANSYS, OpenFOAM etc. However, CSIR 4PI supports the open source movement and most of the open source software generally required for modelling and simulation are made available in the HPC systems. The systems are used extensively for running complex models in the field of ocean, atmosphere, earth and aerospace.

Technical support was provided to a large number of users from CSIR-4PI (erstwhile C-MMACS) & 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 (erstwhile C-MMACS) under the SPARK program. Technical advices and consultancies were provided to various institutions within and outside CSIR.

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