Implementations of the Integrated Linux, Windows and Macintosh Cluster for achieving the Personal High Performance Computing

Victor Chang
School of Electronics and Computer Science
University of Southampton, United Kingdom
icvc@ecs.soton.ac.uk

ABSTRACT

SUMMARY

The motivation and objective for this paper is to demonstrate (1) the integrated Linux (30 distributions, including Red Hat and SUSE), Windows (Cluster 2003 and Server 2003) and Macintosh (Tiger 10.4) cluster for running Web Services and jobs; and (2) the demonstration of the Personal High Performance Computing (PHPC), which only require a smaller number of computers, resources and space in the secure Gigabytes of 1,000 Mps and Wi-Fi/3G of 108 Mbps networking environments. This PHPC is based on a cluster of the 32-bit and 64-bit AMD and Intel dual-core machines, which can achieve the following: (a) jobs are shared amongst different platforms, which can run as a single job altogether or run different jobs concurrently; (b) running Java and .NET Web Services; (c) receiving up to a 60% better performance referring to multi-tasking performance; (e) ensuring a secure networking environment, which includes VPN, in-build firewall and integrated wired and wireless networks; (f) storing a massive amount of data (12 TB, or 12,000 GB) for database and server applications; and (g) successfully integrating with other technologies such as the integrated Java and .NET Web Services, X.509 security, Proxy Architecture and HTTPS .

Categories and Subject Descriptors

C.2.4 [**Distributed Systems**]: Client/server, Distributed applications, Distributed databases, Network operating systems.

General Terms

Personal High Performance Computing.

Keywords

Personal High Performance Computing.

1. INTRODUCTION

The 64-bit computing offers distinctive values for high performance computing, and this statement is strongly supported by the following experimental results. Firstly, a 64-bit computer performs much better than a 32-bit computer in dynamic multimedia and calculations, in a situation where high CPU and memory utilizations are required. A 64-bit computer has up to a 50% better performance than a 32-bit computer in stress testing [1]. Secondly, the "dual-core" 64-bit computer, a computer that has the capabilities for 2 CPUs, which performs between 6% - 42% better than a single-core 64-bit computer in high performance computing [1]. Thirdly, the dual-core 64-bit computers can successfully integrate with the latest emerging technologies, which include 3-D visualization, .NET and Java Web Services, broadband/wireless networks and home entertainment systems. Hence, a group of 64-bit dual-core machines are able to be further customized into the integrated Windows and Linux clusters, and this leads to an area known as Personal High Performance Computing (PHPC) [2]. Finally, there are four distinctive advantages offered by PHPC, which includes: (a) reducing CPU time; (b) improving memory and hard-disk workloads; (c) improving multi-tasking performance and finally (d) executing fast, robust and accurate calculations, visualization on the server-side.

2. EXPERIMENTAL SET-UPS

Implementations of the Integrated Cluster depend on the network infrastructure, which is built on top of the Gigabyte network of 1,000 Mbps and Wi-Fi/MIMO/3G network of 108 Mbps, which have the VPN, in-build firewall, WAP-PSK and .NET authentication and TKIP Data Encryption. This requires a robust and multi-functional wired/wireless router, a VPN, a MIMO router, a 108Mbps wireless router, 1,000 Mbps switches and network cards for every device in the Cluster. Each of the 32-bit machines is installed with eight 32-bit operating systems (OS), and each of the 64-bit machines is installed with sixteen OS, which includes 32-bit and 64-bit Windows and Linux. There are 30 Linux distributions altogether, including the big five: Red Hat WS/ES/AS 3/4, SuSE 9.3/10, Mandriva 2006, Debian 3.1 and Ubuntu 5.10/6.0.4. Each Linux distribution has its particular advantages for performing certain tasks, and therefore each Linux has its specific roles in the Integrated Cluster. Setting up Macintosh cluster requires a different approach. This requires the All-in-one "Apple iMac Duo Core Intel" Mac OS 10.4, which are connected to the Integrated Cluster via Gigabyte and Wi-Fi/MIMO/3G networks. Figure 1 below is the simplified architecture.

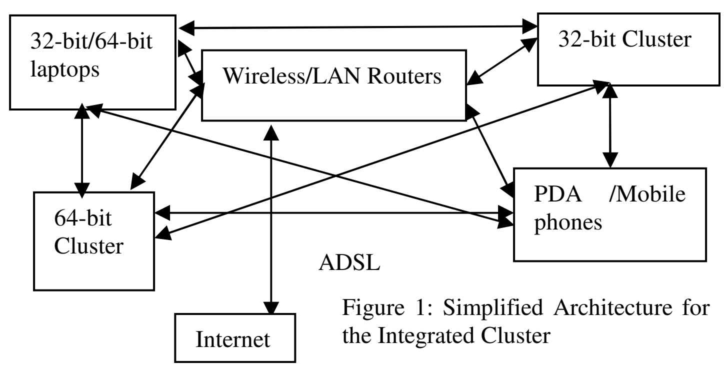

Figure 1: Simplified Architecture for the Integrated Cluster

3. EXPERIMENTS AND RESULTS

3.1 Investigation on Java Web Service

OMII_2.3.0 [3] is used for this part of research. The OMII client and servers were installed across a majority of Linux and Macintosh platforms, where interactions between client and server took place via X509 certificates, HTTP/HTTPS and Proxy Architecture. Both Sun JDK 1.4.2 – 1.5.0_06 and IBM JDK 1.4.2 were used for the tests. The test procedure included creating an account on the server, obtaining the account approval and ensuring the job submission was completed. The services included job submission (GridSAM), registry (Grimoires), PBAC services Java API (Jython), notification (FINS) and plots (PlotWS). Setting up OMII Services on Macintosh required setting up a remote database server on a Linux. The next task was to run all the jobs for the Client, Cauchy, GridSAM, Grimoires, Jython, FINS and PlotWS successfully on Macintosh with Apple JDK 1.4.2_09. Another test was to run GUI-Cauchy jobs 100 times with the time taken being recorded down. The results are discussed in Section 3.3.

3.2 Investigations on .NET Web Service

The author's 64-bit .NET e-portal and e-Services were used for this test, which were configured to be 32-bit compatible on the client side and 32/64 bit compatible on the server side. The procedure included creating an account on the client and recording the time of the server response for a .NET job request. Another test was to write a .NET script that informed the client to run a job on the server and then recorded down the CPU time required. The results are discussed in the next section.

3.3 CPU Time Investigation

"CPU time" is a common term used in high performance computing, which refers to either the amount of time taken for a job to be completed, or the duration of time that a task requires [3]. Results taken in Section 3.1 and 3.2 were recorded as below:

Java CPU time test	64-bit	32-bit
OMII Client to Server	20% quicker	20% slower
GUI-Cauchy (single)	10% quicker	10% slower
GUI-Cauchy (100 jobs)	30% quicker	30% slower

.NET CPU time test	64-bit	32-bit
Client-Server CPU Time	45% quicker	45% slower
CPU time to complete a task	50% quicker	50% slower

This does not imply .NET is a better Web Service than Java, as the 64-bit NET application is specifically written for 64-bit machines, operating systems and .NET API, which explains why it achieves a higher level of CPU-time reduction.

3.4 Multi-tasking performance

This set of experiments are designed to (a) understand how well the 64-bit and 32-bit clusters can cope whilst simultaneously running several applications, which require high CPU and memory consumption; and (b) compare the relative performance between these two models. In order to perform this scenario, two sets of tests were carried out. The CPU and memory utilizations were recorded every second, and the final results were taken thrice and were based on the combined outcome from the "FreeRAM XP Pro", "System Log" (Linux and Cluster 2003), and a .NET script. The 64-bit clusters completed all the tasks within 18 minutes. On the other hand, the 32-bit model could not complete the tasks within 25 minutes and at points, a system crash was encountered. To be equally fair, the comparison was taken up to 18 minutes and the better performance of the 64-bit model was up to 60%, as Figure 2 shown below.

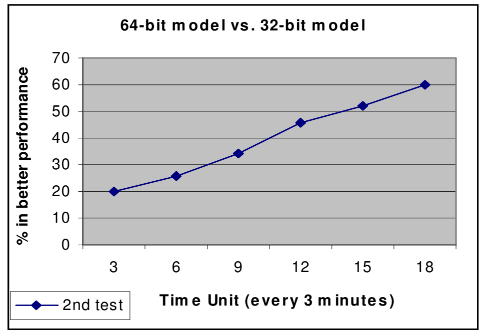

Figure 2: 64-bit vs. 32-bit model: % in better performance in multi-tasking

4. CONCLUSION AND FUTURE WORK

This paper demonstrates fundamental concepts and implementations of Personal High Performance Computing (PHPC), for which the specifications, functionalities and integrations with other technologies can be designed, developed and implemented based on individual needs and demands. PHCP includes the Integrated Cluster, Gigabyte and Wi-Fi/3G networks, integrated Java and .NET Web Services, 12,000 GB database server and Security, the later of which includes X.509 certificates, VPN, in-build firewall, WAP-PSK and .NET authentication and TKIP Data Encryption. The next stage of the PHCP is to integrate with the existing clusters in several parts of the world, which either the author owns or he has the access to. Therefore, this will form a global distributed computing system, where either a single job will be running concurrently, or different jobs will be running asynchronously, and the results of which will be analyzed and presented.

5. REFERENCES

[1] Chang, V., "Two research contributions in 64-bit computing": Testing and Applications. In Proceedings of IADAT-tcn2005 International Conference on Telecommunications and Computer Networks, Portsmouth, UK. 7– 9 September, 2005.

[2] Chang V., "Experiments and Investigations for the Personal High Performance Computing (PHPC) built on top of the 64-bit processing and clustering systems", In Proceedings of 13th Annual IEEE International Symposium and Workshop on Engineering of Computer, Based Systems (ECBS 2006), 27-30 March 2006, Potsdam, Germany.

[3] The OMII Software and User Guide, which is accessible via the OMII website: http://www.omii.ac.uk