Thin Client Cluster Server Computing: 21st Century Enablers for Schools

Ingrid Siebörger, Alfredo Terzoli and Cheryl Hodgkinson-Williams
Department of Computer Science
Rhodes University, Grahamstown, South Africa
Tel: 0466038291 Fax: 0466361915
email: {i.sieborger, a.terzoli}@ru.ac.za, cheryl.hodgkinson-williams@uct.ac.za

***Abstract*— Worldwide computers are being used in schools for, amongst other things, developing the knowledge and skills required for citizens to be able to operate within the 21st Century Information Age. In addition, the use of cluster computers are becoming more commonly used as a result of the increasing ease with which supercomputer systems can be built using commodity-off-the-shelf hardware often using free software, such as Linux and Unix operating systems and third party applications. This work explores whether cluster thin client computer servers can provide robust and cost effective ICT infrastructure models for use in previously disadvantaged schools in South Africa. The paper presents the initial experimentation that has been conducted with a DNS round robin cluster. Further investigation is needed to determine whether cluster servers will adequately support pedagogy and provide robust and easy to maintain solutions that are sustainable in previously disadvantaged schools with clear advantages over single server models.** [1]

***Index Terms*— Cluster computing, DNS round robin cluster, SSI cluster, ICTs in Education, thin client computing**

I. Introduction

Worldwide computers are being used in schools for, amongst other things, developing the knowledge and skills required for citizens to be able to operate within the 21st Century Information Age. Many governments around the world have put in place specific policies to encourage the use of Information Communication Technology (ICT) in education [1], [2] . The use of ICTs in education "is now a political orthodoxy, seen by many politicians and educators as a ready means of widening participation to those social groups traditionally excluded from learning" [3]. Investment in ICTs in education may also be seen as an attempt to address the so-called "digital divide" [3].

According to the white paper on e-Education in South Africa [4], nationally, 39.2 % of schools have computers, while 26.5 % have computers for the purposes of teaching and learning [4]. In schools the traditional ICT infrastructure consisted of a computer lab housing thick client desktop computers, possibly with a central server for storage. In developing countries this model can be undesirable as the cost of desktop computers and a server can be prohibitively expensive. In response, schools and donor organisations such as the TuxLab project have deployed thin client computer labs to schools. In this model the greatest cost is the server, while the thin clients are older refurbished computers that can be cheaply acquired. However, for some previously disadvantaged schools even the cost of a single server can be prohibitively high. Furthermore, the central server, upon which all the clients rely, is a single point of failure for the ICT facilities in a school. For these two reasons this paper investigates the implementation of cluster server solutions for thin client computer labs in schools. The remainder of the paper is divided into three sections: the first discusses relevant background literature, followed by a description of initial experimentations. Finally, future work and conclusions are discussed.

[1]This work was undertaken in the Distributed Multimedia CoE at Rhodes University, with financial support from Telkom SA, Comverse, Stortech, Tellabs, Amatole Telecommunications Services, Mars Technologies, Bright Ideas 39 and THRIP.

II. Related Work

A common and easy method of clustering computers together is to use a Domain Name System (DNS) round robin. DNS round robin allows a computer cluster to have multiple IP addresses associated with one host name. When the DNS server receives a query for the host name it returns a number of IP addresses [5]. The requesting computer will then randomly pick one of the IP addresses and initiate communication. As a result of the relatively random nature with which the computers choose an IP address from the list they are presented with, the communication between the computers and the cluster nodes are balanced. Google, perhaps the most famous service provider on the Internet, employs the use of DNS round robin in order to load balance search requests [6], [7]. DNS round robin implementation of clusters can also be beneficial on a much smaller scale. An example of this is the cluster at Rhodes University which provides email and proxy authentication services on the campus.

Unfortunately DNS round robin solutions do not maximise the efficiency of aggregation because not all of the resources of each node are available to all the other nodes, meaning that it is not possible to migrate running processes to other nodes during execution. This is because the cluster is not able to present or behave as though it were a single large computer. In order to present a cluster of computers as a single computer, Single System Image (SSI) architectures must be employed [8], [9]. A SSI architecture hides the heterogeneous and distributed aspects of the available resources and presents them to users and applications as a single computing resource. This means that users of the system have a global view of all the resources that are available within the system from any one of the nodes from which they are operating. SSI solutions can be constructed at both the hardware and software levels. SSI systems can ensure

high availability – the system will continue to operate if one node fails – load balancing and communal multiprocessing [9]. "SSI design goals for cluster-based systems are mainly focused on complete transparency of resource management, scalable performance, and system availability in supporting user applications" [9, p124]. The concept of SSI clustering is not new. The Locus Distributed System produced in the 1980s provided an SSI environment [8]. There are a number of SSI clustering systems available, for example, Kerrighed [10], OpenSSI [8], Amoeba [11], and Mosix [12].

III. Initial Experimentation

Standard thin client computer labs that we deploy in schools use a single, central server and numerous thin clients. The server runs Edubuntu Linux [13], a flavour of Ubuntu that comes standard with the Linux Terminal Server Project and a host of educational games and programs. The server, however, is a single point of failure and the largest financial cost. In this section we present one of the four identified clustering methods that could be employed as possible solutions to the short comings of a the single server model.

The first cluster solution implemented was that of a DNS round robin cluster. Ideally this implementation should take the form of high availability; the thin clients should be able to use Trivial File Transfer Protocol (TFTP) to retrieve their kernel and ramdisk image from any of the servers in the cluster regardless of which Dynamic Host Configuration Protocol (DHCP) server responds. However, because of how LTSP operates (the link between the DHCP server and the TFTP server are tightly coupled) the initial communication required to boot the thin clients and the mounted Network Block Device (NBD) all take place between the thin client and the original DHCP server. It is only at the point at which the client connects to the X display manager that any one of the LTSP servers on the network can be utilised. The display server can be set (changed) by using the LDM_SERVER parameter in the lts.conf file. Using DNS, the LDM_SERVER parameter can be set to a host name that has multiple A-records, balancing the load of connections between the clients and the display servers. In order to provide a uniform environment and access to all the users files and folders, regardless of which display server the client connects to, winbind is used for central authentication. Using one of the cluster servers, the user's home directory on that server is mounted. Consequently, while this does balance the display load of the LTSP servers (and other services such as DHCP and DNS that can be portioned out to various servers to further, albeit crudely, balance their loads) the system is not as robust as there are still single points of failure for certain services within the system.

IV. Future Work and Conclusion

The aim of this research is to improve upon the thin client model of computer labs for previously disadvantaged schools in South Africa, and indeed in many other parts of the world. The central, single thin client server imposes cost complications for some schools as they can not afford a server of the required specifications to support their thin client computers. In addition, the server is a single point of failure and can result in complete ICT infrastructure failure should it become nonoperational. For these reasons, it is desirable to investigate the use of cluster computers as thin client servers in school computer labs using commodity computers.

There are a number of methods that have been identified, three methods within a DNS round robin based solution and an SSI solution, of which just one is reported here. Each implementation needs to be deployed to a school, as well as a traditional single server model and compared and contrasted.

References

[1] L. Dawes, "First connections: teacher and the National Grid for Learning," *Computers and Education*, vol. 33, pp. 235–252, 1999.

[2] M. Warschauer, M. Knobel, and L. Stone, "Technology and equity in schooling: Deconstructing the Digital Divide," *Educational Policy*, vol. 18, no. 4, pp. 562–588, 2004.

[3] N. Selwyn, S. Gorard, and S. Williams, "Digital Divide or Digital Opportunity? The Role of Technology in Overcoming Social Exclusion in U.S. Education," *Educational Policy*, vol. 15, no. 2, pp. 258–277, 2001.

[4] Department of Education, "White Paper on e-Education: Transforming learning and teaching through Information and Communication Technologies," Sept. 2004.

[5] O. P. Damani, P. E. Chung, Y. Huang, C. Kintala, and Y.-M. Wang, "ONE-IP: techniques for hosting a service on a cluster of machines," *Computer Networks and ISDN Systems*.

[6] D. F. Carr. (2006) How Google Works. Web article. Last Accessed May 2008. [Online]. Available: http://www.baselinemag.com/c/a/Projects-Networks-and-Storage/How-Google-Works-[1]/

[7] L. AndrBarroso, J. Dean, and U. Hzle, "Web search for a planet: The google cluster architecture," *IEEE MICRO*, vol. 23, no. 2, pp. 22–28, 2003.

[8] B. J. Walker, "Open Single System Image (openSSI) Linux Cluster Project," 2004, last Accessed May 2008. [Online]. Available: http://openssi.org/ssi-intro.pdf

[9] R. Buyya, T. Cortes, and H. Jin, "Single System Image (SSI)," *The International Journal of High Performance Computing Applications*, vol. 15, no. 2, pp. 124–135, 2001.

[10] R. Lottiaux, B. Boissinot, P. Gallard, G. Vall, and C. Morin, "OpenMosix, OpenSSI and Kerrighed: A Comparative Study," in *CCGrid 2005. IEEE International Symposium on Cluster Computing and the Grid, 2005.*, vol. 2, 2005, pp. 1016 – 1023, last Accessed May 2008. [Online]. Available: http://hal.inria.fr/action/open_file.php?url=http://hal.inria.fr/docs/00/07/06/04/PDF/RR-5399.pdf&docid=70604

[11] A. S. Tanenbaum and S. J. Mullender, "An Overview of the Amoeba Distributed Operating System," *ACM SIGOPS Operating Systems Review*, vol. 15, no. 3, pp. 51–64, 1981.

[12] A. Barak, O. La'adan, and A. Shiloh, "Scalable Cluster Computing with MOSIX for LINUX," in *Proceeding of Linux Expo '99*, 1999, pp. 95–100, last Accessed May 2008. [Online]. Available: http://www.mosix.org/pub/mosix4linux.pdf

[13] Canonical. (2006) Edubuntu. Last accessed Oct. 2006. [Online]. Available: http://www.edubuntu.org/

Ingrid Siebörger is currently reading for her PhD Degree in Computer Science at Rhodes University. She is working in the area of ICT in Education integration focusing on what ICT architectures would be the most suitable, cost effective way to meet pedagogical, technical and economic needs of South African schools. She is also working full time for the Centre of Excellence (CoE) at Rhodes University as a Research Assistant.