A Low Powered NAS-based Web Server for Small and Medium Enterprises

Surasak Srisawan

Rajabhat Rajanagarindra University, Thailand

0179

The Asian Conference on the Social Sciences 2013

Official Conference Proceedings 2013

Abstract

Small and Medium Enterprise or SME are companies that has small number of personnel. In Thailand, SMEs outnumber large enterprise significantly. Both large and small enterprises have equal footing in term of their identities on the Internet. Some successful Internet-based firms do not operate brick and mortar stores and doing business solely through their web sites. However, due to their lack of funding, SMEs usually have no resources for Information Technology infrastructure such as dedicate servers and an air conditioning system to support the underlying infrastructure. This is impossible for most SMEs especially those in agriculture and handicraft areas, which are the majority of SMEs in Thailand. This study focused on identifying a low powered, low cost and always on solution viable for SMEs to operate their e-commerce web sites without premium infrastructure cost. Network-attached storage or NAS is suitable in this situation due to its low power consumption property. Its ARM-based processor has low TDP and can operate 24/7 without air conditioning environment. A Linux-based NAS with a 300MHz single core ARM CPU was modified to serve as a web server with PHP support. The result shown that it was able to serve static content to multiple clients concurrently but the performance was dropped dramatically once the connections contain CPU intensive calculations. Using faster and/or multi-core ARM CPU may improve this situation and should be study further.

Keywords: SME, E-commerce, Web Server, Embedded System, Energy Saving

1. INTRODUCTION

The advancement of information technology especially the technology related to the Internet has propelled the popularity of e-commerce all over the world. With e-commerce, even small enterprises or individuals can sell and deliver products to their customers directly through web sites. In Thailand, Small and Medium Enterprises or SMEs outnumber large enterprise significantly. In order to compete with larger companies, SMEs need to set up their e-commerce web sites. There are many third party web hosting options available but hosting their own web sites give the benefit of flexibility in term of software and database that can be used in the system.

Although the cost of computer, even the server specific hardware, was decreased and can be acquired easily, the power consumption of these general purpose computers is relatively high. Furthermore, Thailand is located near the equator thus make the average temperature above thirty degree Celsius in many areas. Therefore, air conditioning systems are usually required in order to operate web servers 24/7. However, SMEs usually have no resources to support the underlying infrastructure especially those in agriculture and handicraft areas, which are the majority of SMEs in Thailand.

Unlike x86 architecture, which is used in general purpose computers, ARM architecture used in embedded computer systems focuses on optimizing power consumption. Network-Attached Storage or NAS is an example of the systems using ARM processors. These systems usually run modified versions of Linux, making it possible to install additional services or making changes to the operating system itself. Thus, running a web server on one of the ARM-based embedded systems is possible.

To evaluate the suitability of a NAS-based system when it is used as a web server platform for SMEs, this research focuses on the comparison of the performance and power consumption ratio between a general purpose computer, which is used as a reference system, and a NAS-based system. Series of tests were performed and the result scores were analyzed in order to evaluate if the NAS-based system is a viable alternative web server for SMEs' workload.

2. RELATED WORK

A. Web Server Hardware

Benini and Micheli (2000) mentioned that it is important to design computer hardware to have a balance in computer power and performance. For example, a data center of an Internet Service Provider which has 8,000 servers requires two million watts of power or 250 watts per server on average. However, with recent improvement in term of computer architecture, recent embedded system architecture requires only 7-10 watts while retaining the processing ability similar to x86 architecture (Domeika 2008) although it is considerably slower than the general purpose processors available at the same generation.

Despite the fact that x86 architecture is generally available in the form of personal computers, the largest numbers of processors are embedded in electronic devices such as household appliances, cars and mobile devices. ARM is an embedded architecture that gained popularity since 2001 due to is low power consumption nature allowing it to run on battery (Sloss et al. 2004). Furthermore, a research by Wu Min-hua (2008) shown that the web server based on ARM architecture performed well on low complex tasks and had advantages on cost and power consumption over general purpose computers as well.

B. Web Server Software

This research compared the performance and power consumption of two systems running on x86 and ARM architectures. Linux was chosen as the operating system of choice because it is a widely used operating system running on web servers and support both x86 and ARM systems. The web server software used was Apache HTTP server, which was widely used both in term of the number of sites on the Internet and in term of researches (Robertsson et al. 2004; Liu, Niclausse, and Jalpa-Villanueva 2001; Bosque et al. 2007). In term of cost, Apache HTTP server is also an open source software that can be freely modify and distribute making it a good choice for a low cost web server.

C. Techniques Used to Evaluate the Performance and Power Consumption of Computer Systems

Generally, the performance of a computer system can be evaluated by measure the period of time used to compute a specific workload or the amount of work that can be completed during a specific period of time. To evaluate the performance of a web server, Barford and Crovella (1998) developed a method called Scalable URL Reference Generator: SURGE which can be used to evaluate real-life performance. A more recent method called Web traffic GeneratOr and benchmark: WAGON by Liu, Niclausse, and Jalpa-Villanueva (2001) focused on workload generation and the analysis of the parameters passing between servers and browsers. These methods were used as a framework for the tools used in this research.

The other aspect in this research was the power consumption measurement. Although it can be measured using a watt meter, software alternatives are also available. In this research, a software suite called Lesswatts.org by Gray (2008) particularly PowerTOP was used to collect the power consumption data from both systems. Hence, allowing the data collection process to be completed entirely using software.

3. EXPERIMENT

A. Hardware Components

Two computers were used in the experiment in order to evaluate the effectiveness and viability of the NAS-based web server. A general purpose computer constructed with commodity parts was used as the reference. The NAS box used in the experiment was AgeStar NSB3AS1T, which can be acquired locally at low cost. The detail of the hardware components are listed in table 3.1.

Table 3.1 Hardware components

	Reference System	NAS-based System
Processor		
Model	Intel® Celeron® 2.0GHz	FA526id(wb) rev 1 (v4l)
Number of Core	1	1
Frequency	2,000 MHz	300 MHz
BogoMIPS	3,999.91	153.19
Architecture	x86	ARM
Memory		
RAM	1,024 MB	32 MB
Swap	2,048 MB	100 MB
Auxiliary Storage		
Model	WD10EARS	WD10EARS
Capacity	1 TB	1 TB
Cache	64 MB	64 MB
Interface	SATA 3 GB/s	SATA 3 GB/s
Network		
Type	Ethernet	Ethernet
Speed	100 Mbps	100 Mbps
Power Consumption		
Overall System	235 Watts	13 Watts

In term of raw performance, the NAS-based system was obviously inferior to the reference system with the processor operated at 300 MHz or 15% comparing to the Celeron processor. The BogoMIPS (Bogus Million Instructions per Seconds), which represents the relative operating cycle of Linux kernel, also shown that the NAS-based system was operating at 3.75% relative performance. However, due to the differences in architecture, the operating frequency and BogoMIPS are not reliable means to evaluate the real-world performance.

On the other hand, the NAS-based system had a clear advantage in term of power consumption. At overall power consumption rate of 13 watts, it consumed only 5.5% of the reference system.

B. Software Components

Unlike the hardware part, the software components installed onto both systems are relatively similar. Debian Linux was used as the operating system and the web server software was Apache HTTP server with PHP5 enabled.

The reference system is a standard general purpose x86 computer. Therefore, Debian Linux was installed conventionally. The Apache HTTP server and related programs were installed via apt, which is Debian package management system. On the NAS-based system, a method to install a full version of Debian Linux on Agestar devices by Yohanes (2011) was followed. The standard apt-sources were added and all the software required to operate a web server were installed in the same manner as the reference system.

C. Data Collection Methods

To evaluate the workload of web servers, a load generator program was developed using the models from SURGE (Barford and Crovella 1998) and WAGON (Liu, Niclausse, and Jalpa-Villanueva 2001) as a guideline. The first part of the experiment was to evaluate the performance of the systems categorized by file types. This part was separated into two phases, static evaluation and dynamic evaluation.

1. Static evaluation.
   By using HTML documents, the input/output subsystem of the web servers was evaluated. Both small and large files were used in this phase to perform the I/O intensive test on the systems.

2. Dynamic evaluation.
   The second phase was designed to evaluate the computing power of the web servers using PHP scripting language to generate dynamic documents. Two kinds of PHP documents, low complexity and high complexity scripts were used in this experiment as the CPU intensive test.

The second part of the experiment was to evaluate the real-world performance of the systems. Normally, a web page consists of multiple components such as HTML files, images and scripts and the web server transfers multiple files simultaneously to the user requesting a web page. Furthermore, the nature of the Internet allows multiple users to connect to the web servers at the same time. Therefore, this part focused on evaluating multiple thread performance of the systems.

4. FINDINGS AND ANALYSIS

A. Access time categorized by file types

Four file types were used in the evaluation, small file, large file, simple script and complex script. The tests were performed three times for each file type and the access time required to complete the transfer was recorded as shown in table 4.1.

Table 4.1 Access time categorized by file types (Unit: Millisecond)

	Reference System	NAS-based System
Small files ± SD		
Round 1	22.56 ± 4.76	177.28 ± 137.39
Round 2	19.73 ± 0.62	212.68 ± 160.41
Round 3	19.69 ± 0.95	220.94 ± 192.10
Average	20.66 ± 3.12	203.63 ± 165.37
Large files ± SD		
Round 1	81.45 ± 14.58	173.45 ± 46.91
Round 2	85.05 ± 6.44	172.36 ± 39.75
Round 3	84.48 ± 2.52	180.51 ± 60.23
Average	83.66 ± 9.42	175.44 ± 49.66
Simple Script ± SD		
Round 1	6.44 ± 6.14	20.43 ± 12.87
Round 2	4.98 ± 10.25	19.82 ± 7.32
Round 3	3.89 ± 0.76	18.98 ± 7.56
Average	5.10 ± 6.97	19.74 ± 9.58
Complex Script ± SD		
Round 1	8.16 ± 1.61	241.72 ± 25.06
Round 2	7.23 ± 0.80	241.35 ± 22.83
Round 3	7.12 ± 0.59	240.10 ± 22.03
Average	7.50 ± 1.19	241.06 ± 23.27

The access time for each transaction shown above represented the period between the client start sending a request until the client finished receiving the requested file. On average, the NAS-based system took 203.63 milliseconds to complete transferring small files while the reference system took 20.66 milliseconds. In term of response time, it took 985.62% or almost 10 times longer to transfer a small file.

B. Access time categorized by number of threads

To evaluate real-life performance where multiple files are transferred simultaneously, the performance of the system based on the number of web server threads running at the same time was measured. The result is shown in table 4.2.

Table 4.2 Access time categorized by number of threads (Unit: Millisecond)

	Reference System	NAS-based System
2 Threads ± SD		
Round 1	154.33 ± 14.05	550.00 ± 17.44
Round 2	158.33 ± 15.53	587.00 ± 11.14
Round 3	159.00 ± 15.72	550.67 ± 25.15
Average	157.22 ± 13.27	562.56 ± 24.52
5 Threads ± SD		
Round 1	335.33 ± 114.03	1,530.67 ± 7.09
Round 2	330.67 ± 0.58	1,402.33 ± 6.66
Round 3	353.33 ± 0.58	1,568.67 ± 6.11
Average	330.27 ± 50.98	1,471.53 ± 86.74
10 Threads ± SD		
Round 1	512.40 ± 46.46	2,707.70 ± 86.86
Round 2	531.20 ± 3.49	2,718.70 ± 250.51
Round 3	505.40 ± 12.88	2,651.40 ± 237.26
Average	516.33 ± 29.12	2,692.60 ± 200.46
20 Threads ± SD		
Round 1	799.25 ± 49.33	11,373.15 ± 436.23
Round 2	811.40 ± 22.52	8,426.05 ± 2,157.16
Round 3	692.10 ± 88.60	17,203.45 ± 3,299.76
Average	767.58 ± 79.98	12,334.22 ± 4,312.08

In term of real-life performance, the access time of NAS-based system grew significantly when the number of threads increased. At 20 threads, the system took more than 10 seconds on average to complete the execution and transferred all files. Considering at the performance alone, it is obvious that the number of threads should not be expanded further.

C. The performance per power consumption ratio

It was obvious that the performance of the reference system was superior to the NAS-based system. However, when taking the power consumption into account, it turned out that the performance per power consumption ratio of NAS-based system was generally better than the reference, as shown in table 4.3.

Table 4.3 Performance per power consumption ratio (higher is better)

	Reference System	NAS-based System
Small files	1	1.83
Large files	1	8.67
Simple scripts	1	4.69
Complex scripts	1	0.57
2 Threads simultaneously	1	5.08
5 Threads simultaneously	1	4.08
10 Threads simultaneously	1	3.49
20 Threads simultaneously	1	1.13

Taking power consumption into consideration changed the overall impression of the NAS-based system where it won all tests but one. The strength of the system lied in the large file transfer department where it has 8.67 times efficiency comparing to the reference system. However, with the efficiency rate of 0.57, the complex script test clearly showed the weakness of the system.

Also, the efficiency of the NAS-based system was significantly reduced with the increasing number of threads. From the test, twenty was the maximum number of threads where the efficiency rate of the NAS-based system was higher than the reference system.

5. CONCLUSION AND SUGGESTIONS

In most tests, the NAS-based system required less power to complete the same task performed on a general purpose computer. However, the response time of the system was also higher across all tests. For small file transfer and simple script tests, the response time was higher but stayed under one second on average and should not affect the usability of the web site. On the other hand, the complex script test and running a large number of threads simultaneously greatly affect the responsiveness of the system.

Despite its weakness in term of raw performance, a NAS-based system can operate without air conditioning infrastructure. Its low power consumption requirement leads to low electricity cost. For SMEs in remote area where electricity system is sometimes unreliable, these properties allow the system to operate longer on a backup battery. Furthermore, the number of transactions performed each day for small and medium size business is relatively low and the overall responsiveness of the system should not be greatly compromised. Thus, a NAS-based web server could be categorized as a viable solution for SMEs with low daily transaction.

Currently, it is clear that ARM processors are gaining popularity. Recent processors used in hi-end smartphones and connected devices have multiple cores and operate at higher frequency. It is obvious that the technology will continue to evolve to the point where the computing power of a low-powered system exceeds the requirements for the web server operated in medium and large enterprises. For future researches, alternative forms of embedded computer systems such as mobile phones or tablets may be studied further in order to identify alternate solutions that can benefit SMEs while maintaining a sustainable ecology as a whole.

6. ACKNOWEDGEMENT

This study was financed by the division of research and development, Rajabhat Rajanagarindra University.

REFERENCES

Barford, Paul, and Mark Crovella. 1998. “Generating Representative Web Workloads for Network and Server Performance Evaluation.” SIGMETRICS Perform. Eval. Rev. 26 (1) (June): 151–160. doi:10.1145/277858.277897.

Benini, Luca, and Giovanni de Micheli. 2000. “System-level Power Optimization: Techniques and Tools.” ACM Trans. Des. Autom. Electron. Syst. 5 (2) (April): 115–192. doi:10.1145/335043.335044.

Bosque, Ana, Pablo Ibañez, Víctor Viñals, Per Stenström, and Jose M. Llabería. 2007. “Characterization of Apache Web Server with Specweb2005.” In Proceedings of the 2007 Workshop on MEmory Performance: DEaling with Applications, Systems and Architecture, 65–72. MEDEA ’07. New York, NY, USA: ACM. doi:10.1145/1327171.1327179. http://doi.acm.org/10.1145/1327171.1327179.

Domeika, Max. 2008. "Chapter 10 - Getting Ready for Low Power Intel Architecture." In Software Development for Embedded Multi-core Systems, 347–383. Burlington: Newnes. http://www.sciencedirect.com/science/article/pii/B978075068539900010X.

Gray, James. 2008. "Go Green, Save Green with Linux." Linux J. 2008 (168) (April). http://dl.acm.org/citation.cfm?id=1371526.1371527.

Liu, Zhen, Nicolas Niclausse, and César Jalpa-Villanueva. 2001. "Traffic Model and Performance Evaluation of Web Servers." Performance Evaluation 46 (2–3) (October): 77–100. doi:10.1016/S0166-5316(01)00046-3.

Nugroho, Yohanes. 2011. "Installing Debian on Agestar Without Serial Port." http://tinyhack.com/2008/07/30/installing-debian-on-agestar-without-serial-port/.

Robertsson, A., B. Wittenmark, M. Kihl, and M. Andersson. 2004. "Design and Evaluation of Load Control in Web Server Systems." In American Control Conference, 2004. Proceedings of the 2004, 3:1980–1985 vol.3. IEEE.

Sloss, Andrew N., Dominic Symes, Chris Wright, and John Rayfield. 2004. "Chapter 1 - ARM Embedded Systems." In ARM System Developer's Guide, 3–16. Burlington: Morgan Kaufmann. http://www.sciencedirect.com/science/article/pii/B9781558608740500022.

Wu Min-hua. 2008. "Research for the Embedded WEB Server." In Microwave Conference, 2008 China-Japan Joint, 776–779. IEEE. doi:10.1109/CJMW.2008.4772543.

iafor

iafor