

SANDIA REPORT

SAND2011-7132
Unlimited Release
Printed October 2011

Experimental Evaluation of the Impact of Packet Capturing Tools for Web Services

Chao-Chih Chen, Yung Ryn Choe, Chen-Nee Chuah, and Prasant Mohapatra

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.osti.gov/bridge>

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online ordering: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



Experimental Evaluation of the Impact of Packet Capturing Tools for Web Services

Chao-Chih Chen◇
Yung Ryn Choe★
Chen-Nee Chuah○
Prasant Mohapatra◇

- ◇ Department of Computer Science. University of California, Davis, CA 95616
Email: {cchchen,pmohapatra}@ucdavis.edu
- ★ Sandia National Laboratories. Livermore, CA 94551
Email: yrchoe@sandia.gov
- Department of Electrical Engineering. University of California, Davis, CA 95616
Email: chuah@ucdavis.edu

Abstract

Network measurement is a discipline that provides the techniques to collect data that are fundamental to many branches of computer science. While many capturing tools and comparisons have been made available in the literature and elsewhere, the impact of these packet capturing tools on existing processes has not been thoroughly studied. While not a concern for collection methods in which dedicated servers are used, many usage scenarios of packet capturing now require the packet capturing tool to run concurrently with operational processes.

In this work we perform experimental evaluations of the performance impact that packet capturing processes have on web-based services; in particular, we observe the impact on web servers. We find that packet capturing processes indeed impact the performance of web servers, but on a multi-core system the impact varies depending on whether the packet capturing and web hosting processes are co-located or not. In addition, the architecture and behavior of the web server and process scheduling is coupled with the behavior of the packet capturing process, which in turn also affects the web server's performance.

Contents

1	Introduction	7
2	Packet Capturing Process	9
3	Evaluation	11
3.1	Experimental Set-up.....	11
3.2	Web Servers	13
4	Related Works	17
5	Conclusion	19
	References	20

List of Figures

3.1	Apache Results	12
3.2	Nginx Results	14
3.3	Web server performance in fetches/second averaged over 24 runs, with 95% confidence interval shown	15

Chapter 1

Introduction

Network measurement is a discipline that provides the foundation for many studies in networked systems. From capacity planning to anomaly detection to network security, being able to measure and collect data from the network is crucial for the success in these tasks. One such popular method to collect data from the network is packet capturing. Because the collected data (the packet) contains application-invariant and application-specific information, it is a good candidate for one-time data collection that can support various types of analysis. In addition, packet capturing tools are widely available (e.g., Wireshark, TCPDump for Linux, NetMon for Windows), and there are mature libraries for custom codes to tap into the packet monitoring process.

In the past network measurements are often collected by means of port mirroring at the router and dedicated machines to collect the packets. Another way to monitor the network is to do so at the edge of the network (i.e., capture at the server machines). This reduces the need to have a dedicated machines and also amortizes the cost of packet capturing over all the machines.

Although not a concern if dedicated machines are used to capture packets, the performance impact of capturing process becomes important if it is to co-locate in the same physical machine as running processes. In fact, in some cases monitoring at the server machines is preferred, if not essential. For example, there are works that attempt to discover application-level dependency [17, 6] while others try to localize the source of faults from information derived from captured packets [4]. Without capturing the network information at the server machines, details such as application-level dependency is either impossible or much more difficult to capture elsewhere.

Recognizing the deficiency of research work in this area, we carry out experiments to examine the impact that packet capturing process has on web-based services. In particular, we test packet capturing process' performance impact on web servers at their saturation point. This gives us some insight into the maximum performance achievable for web services when packet capturing process is also running, and whether it adversely impacts existing services.

The contributions of this work are:

- Experimental evaluation on the performance impact of capturing process to co-located web-based services.
- Deployment of two web servers of different architecture to validate that results are consistent across different web server architectures.

- Measurement of both system-level statistics and user-perceived statistics to observe correlation between the two.

In Section 2 we briefly go over the packet capturing process; Section 3 presents the evaluation methodology and the results obtained; Section 4 discuss the related works; Section 5 discusses future works and concludes the work.

Chapter 2

Packet Capturing Process

This section provides a high-level overview of the packet capturing process in Linux, as to make the discussion in this work complete. The goal of this section is not to provide a complete detail of the internals of packet capturing, but to illuminate on the steps involved in delivering the packet from the network card to the kernel and to the user application. This will help towards understanding the behaviors observed in Section 3; details of the packet capturing process described here can be found in [10, 11].

When packets are transmitted over the wire, the network interface card (NIC) normally picks up the packet if the packet is destined for it; under promiscuous mode it will pick up all packets sensed. Once the packet is recognized for reception, the NIC's interrupt routine is invoked, in which the routine allocates some space in memory and copies the packet into the allocated memory. The packet is not immediately processed after moving to the memory, as the interrupt routine is intended to perform as little operation as possible. When the packet is picked up later by the software interrupt handler, it passes the packet upwards to the appropriate protocol handler based on the packet's protocol type.

For packets destined for packet capturing tools, a special handler is used so that all the packets can be handled and subsequently forwarded to the capturing process. Corresponding to this special handler is a special protocol family called `PF_PACKET`, and the packet is copied¹ and delivered to a socket created specifying the `PF_PACKET` family. Copying is needed because the packet might be actually destined for an application at the local machine, so the packet must be copied for separate consumption by the capturing process and the application.

While this overview is brief, it illuminates the many transactions involved in capturing the packet, and these transactions will result in the use of CPU resource. The experiments to be discussed in Section 3 are aimed at observing how its uptake of CPU resource affects the performance of web servers.

¹The way in which the packet is copied is that an internal structure, `sk_buf`, that holds information regarding the data packet is copied. But only the fields are copied – the packet data itself is referred by pointers in the duplicate and the original structure.

Chapter 3

Evaluation

In this section we present the experimental results obtained when running the packet capturing tool under various application scenarios. We will first discuss the set-up of our experiments and the metrics we set out to collect; then we will discuss the results of the experiments and various inferences drawn from them.

3.1 Experimental Set-up

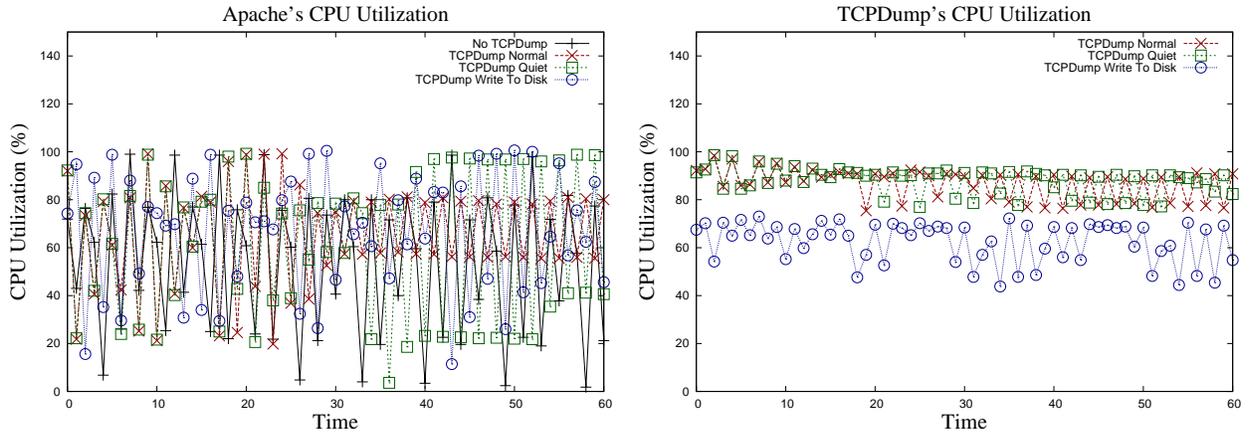
To observe the effects of packet capturing on applications, we deployed two web servers and collected various performance metrics. The web servers used are Apache version 2.2.17 with worker MPM[3] and Nginx version 0.8.54[15]; the kernel used is 2.6.35.11. We choose to use Apache because it is rated the most-used web server according to the latest survey by NetCraft (February 2011 at the time of writing) [14]; while Nginx is also used because its architecture is fundamentally different from Apache and is also rated high on the NetCraft survey.

The Apache architecture offers various Multi-Processing Modules (MPM) as a way to scale the web server with increasing user demand. These MPMs are either multi-processed, multi-threaded, or both. In the model we use, the worker module¹ is selected. Apache uses a parent process to accept incoming connections and distributes them to multiple processes/threads, where each process/threads handles one request at a time [9]². One implication of such an architecture means that in order to scale, the number of threads and process needed to spawn needs to increase. On the other hand, Nginx [15] is a highly-scalable web server developed to address the ability to serve high number of simultaneous user connections, also known as the C10K problem [12]. Nginx operates under the asynchronous call model, so a single process can scale quite well against increasing concurrent request volume.

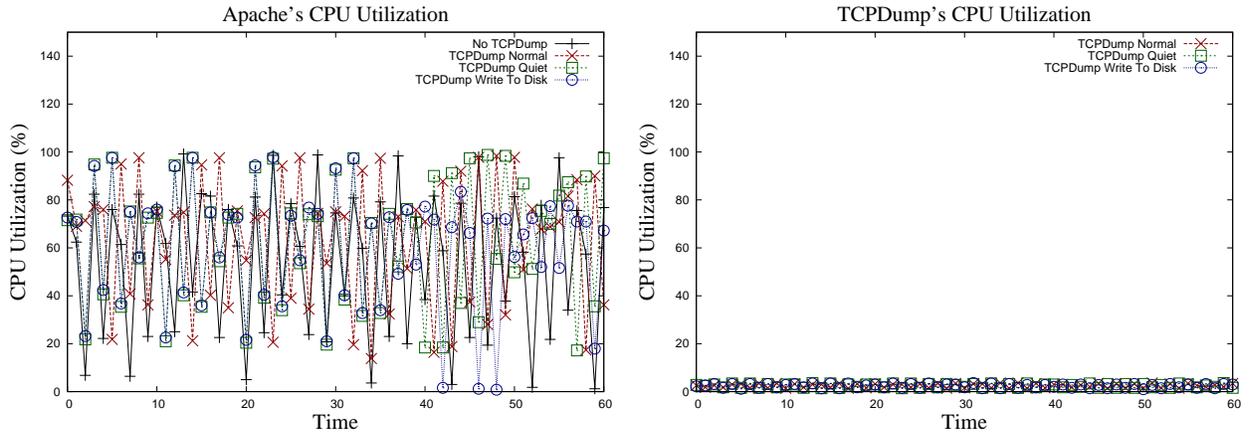
To measure the performance impact of packet capturing, we monitor the CPU and bandwidth utilization of all applications, and application-specific metrics. The packet capturing tool we used in this experiment is a vanilla TCPDump, as we would like to observe the impact of a free and widely available tool without any performance modification to it (e.g., mmap extension, PF_RING extension [7]). To stress the web servers, we use HTTP_LOAD [13] to constantly fetch a small

¹We choose this model based on the scalability and smaller memory footprint than pure process-based module

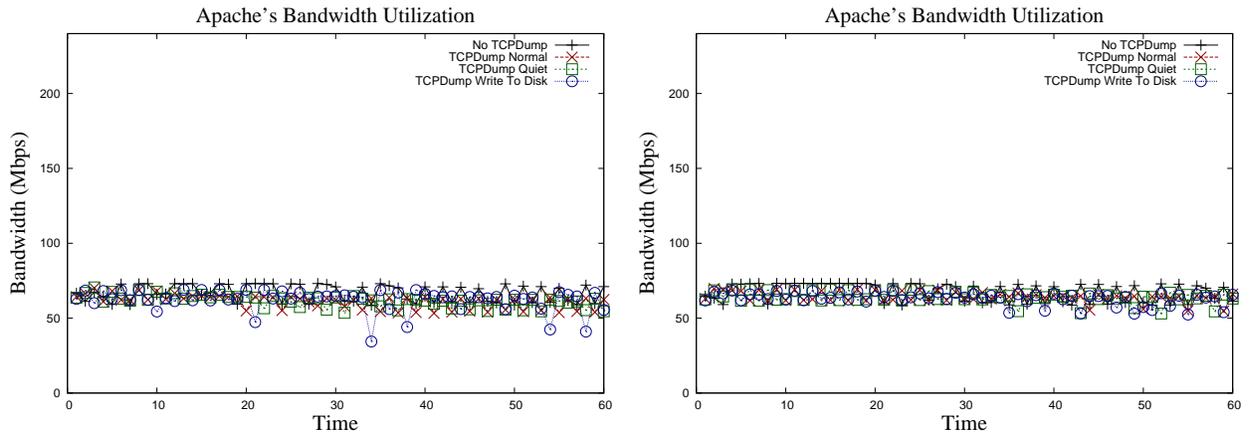
²We also consulted the documentation on the popular prefork module and it works in similar manner.



(a) CPU Utilization For Apache and TCPDump When Pinned To Different CPU Core



(b) CPU Utilization For Apache and TCPDump When Pinned To Same CPU Core



(c) Bandwidth Utilization When Apache and TCPDump Are Pinned To Different (left) Or The Same CPU Core (right)

Figure 3.1. Apache Results

static HTML page from the servers. In addition, for all the experiments the web server is hosted on a Duo Core system, but we confine the web server to a single core. This allows us to experiment

with co-locating and separating the packet capturing and web server process. By doing so, we can observe the effect that sharing CPU resource has on the performance degradation of the web servers.

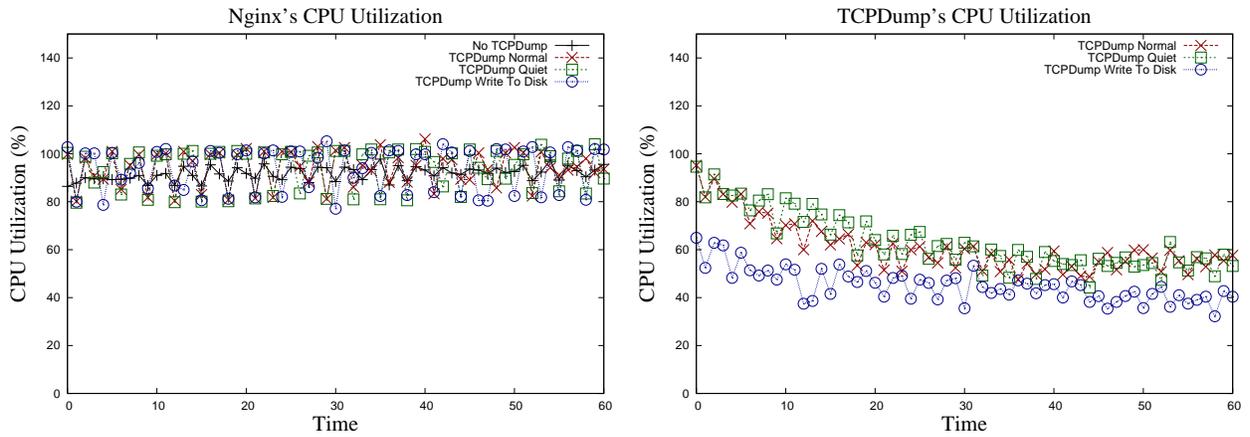
For each web server, we test the co-location factor by moving TCPDump and the web server processes onto the same or different CPU core, and we also vary TCPDump's behavior in four ways: no TCPDump running, normal TCPDump, quieter TCPDump, and TCPDump writing captured packets to disk. We have a total of eight settings, two of which are repeats (no TCPDump and co-locate is the same as no TCPDump and do not co-locate). For each setting we use HTTP_LOAD on two different clients, where each client set the HTTP_LOAD concurrency to 15, to fetch continuously from the server for five minutes; metrics are collected every five seconds to minimize its impact on the experiment.

3.2 Web Servers

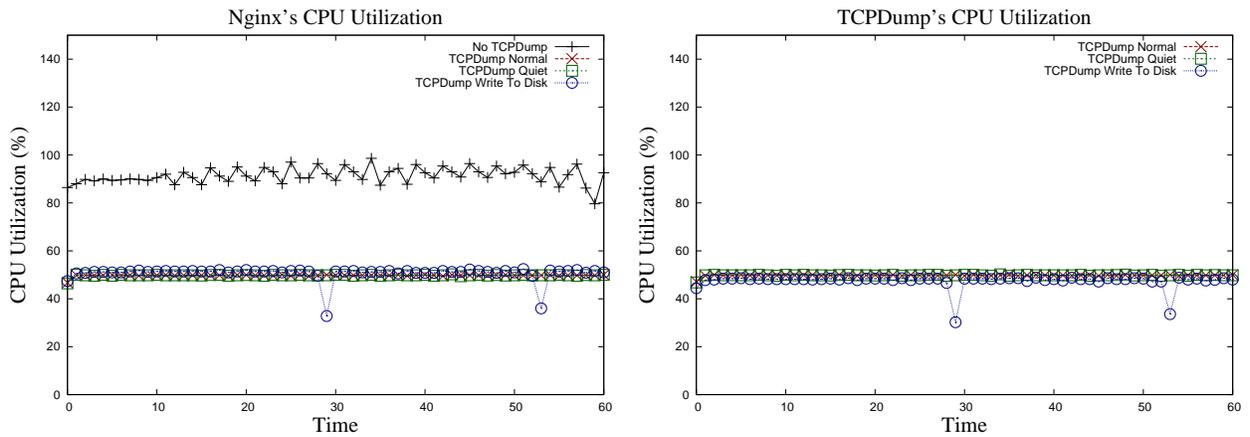
Figure 3.1 and Figure 3.2 show the CPU and bandwidth utilization for one run of the experiment, when running the HTTP_LOAD to retrieve files from the web servers.

Upon first glance, we note that the behavior of Apache and Nginx are visibly different, with Apache more prone to CPU fluctuation, while Nginx is more stable in CPU usage. This could be attributed to the fact that Apache's method of scaling with demand results in much more context switches between the various worker threads, resulting in the performance fluctuation. On the other hand, Nginx is using much less number of processes, and is able to serve requests without getting constantly interrupted.

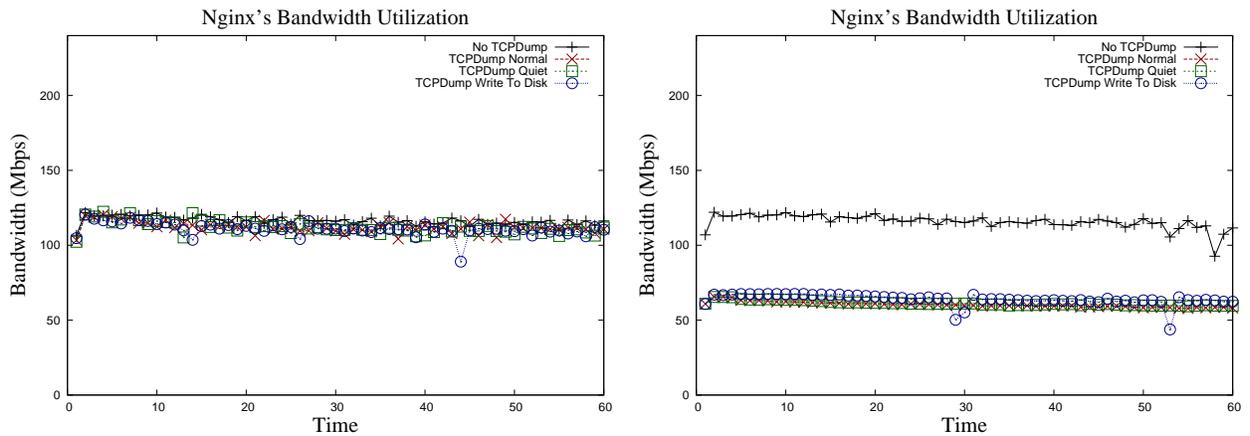
With the web server and TCPDump pinned to different CPU core, both processes are shown to take up significant amount of CPU resource. In both types of web servers, the web server process takes up nearly one hundred percent of the process while the TCPDump process also consumes significant CPU resource. When the server and TCPDump process are bound to the same CPU core, the Apache server seems to be holding more share of the CPU resource while Nginx predictably shares the resource equally with TCPDump. In fact, TCPDump has almost no access to the CPU, using only a few percent of the CPU resource at any given time. This phenomenon can be explained by understanding the method in which Apache scales with increasing user demand. To maintain scalability, the worker MPM in Apache has a parent process that monitors and distributes incoming load. The parent process spawns a number of child process that actually serve the request, with the maximum number of child processes constrained by the ServerLimit directive (defaults to 16, which is our setting). The kernel scheduling algorithm would then try to equally distribute the available CPU resource to all active process, majority of which belongs to Apache. On the other hand, Nginx and TCPDump predictably shares the CPU resource equally, due to the fact that only one Nginx process is actively serving incoming requests. This result has significant implication to the efficiency of the web server and packet capturing process. To ensure capturing process has access to enough CPU to process the captured packet, it should be scheduled on a separate core. If the capturing process is to be co-located with the web server, care must be taken to understand



(a) CPU Utilization For Nginx and TCPDump When Pinned To Different CPU Core



(b) CPU Utilization For Nginx and TCPDump When Pinned To Same CPU Core



(c) Bandwidth Utilization When Nginx and TCPDump Are Pinned To Different (left) Or The Same CPU Core (right)

Figure 3.2. Nginx Results

the structure and behavior of the web server, to ensure the capturing process also has access to the CPU resource. It is interesting to also note that, in some cases the recorded CPU utilization is

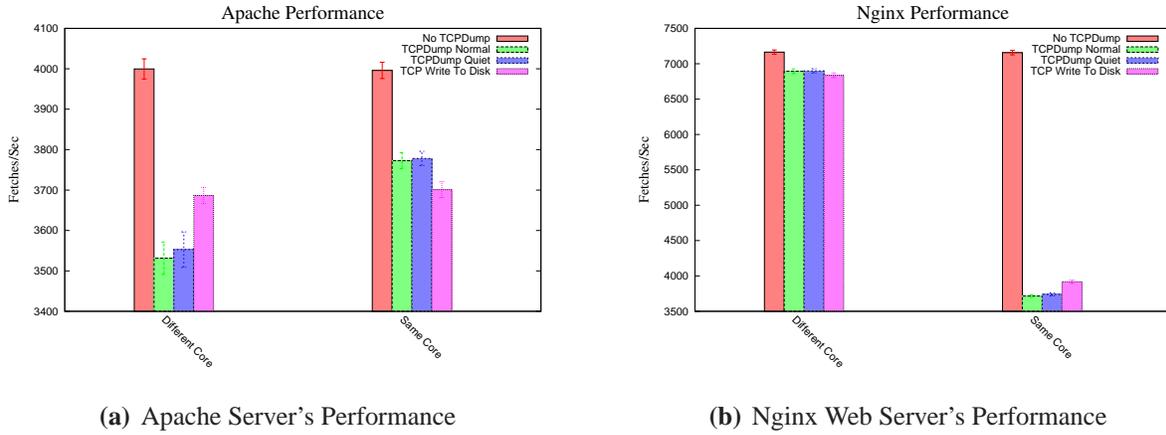


Figure 3.3. Web server performance in fetches/second averaged over 24 runs, with 95% confidence interval shown

over one hundred percent. This is impossible because we confine the processes to a single core, so the maximum possible is one hundred percent. After inspecting the source code of the tool we used to collect the CPU utilization, we noticed the tool reads two pseudo-files to calculate the CPU utilization for a process, and the two reads are not atomic. We believe that this behavior, in addition to the CPU collection process running on another thread, means that there is a potential race condition, and the calculation could be slightly off. However, we note that despite the observed error, the general trend is still significant enough for us to make the above observations.

Next we look at the bandwidth utilization result. For Nginx the achieved bandwidth utilization is lower in the case when TCPDump is running and co-locating with the web server. On the other hand, the bandwidth utilization for apache is about the same as before. Comparing Figure 3.1(c) and Figure 3.2(c) with Figure 3.1(b), Figure 3.2(b), the difference in bandwidth usage is correlated with the difference in CPU usage: a higher CPU usage corresponding to higher bandwidth utilization, and vice versa; however the variability in the bandwidth/CPU utilization are different, as the bandwidth utilization remains more stable than CPU utilization.

While these metrics shed some light on the resource consumption and possible performance of the system, they do not explicitly tell us the performance that users can expect. To gain such an insight, we scrape the reports generated by HTTP_LOAD at each client machine at the end of the experiments. Figure 3.3 shows the aggregate number of fetches per second that are observed from all the clients. We note that the baseline experiment (i.e., TCPDump is not running) shows the performance of Apache and Nginx is quite good. With the presence of TCPDump, the performance of the web server varies depending on the co-location. When TCPDump is not co-located with the server process, TCPDump seems to decrease the average performance of Apache more significantly (up to 10%). However, when TCPDump co-locates with the server process on the same core, the result is more dramatic. Correlating the results in Figure 3.3 with Figure 3.1(b) and Figure 3.2(b), we can see that the CPU share TCPDump has obtained is directly proportional to the decrease in the average fetches per second achievable. The significant performance degradation for the case where TCPDump and Nginx are on same core – but little performance degradation when

on different cores – suggests that CPU utilization could be a major source of the web servers' performance degradation.

In summary, the experiments carried out in this section implies that CPU utilization can tell us that performance degradation has occurred, but performance degradation can still occur even if CPU utilization looks normal. In the case for Apache, even though it dominates the CPU when co-located with TCPDump, the fetches per second achievable is considerably lower; while Nginx consumes much less CPU but has similar performance degradation as Apache. However it is undeniable that the presence of TCPDump has negative performance impacts to the web servers, so care should be taken when running packet capturing process such as TCPDump, as to ensure the performance impact to the web servers is minimal. In addition, when TCPDump has equal opportunity to contend for the CPU resource it does consume a non-trivial amount, and this has performance impact for admins looking to consolidate different types of task onto a single machine. For tasks that are CPU-bound, consolidating it with machines running packet capturing processes could elongate the task completion time as well as diminishing the number of packets captured.

Chapter 4

Related Works

For evaluation of packet capturing tools, the closest works to this are works that either explicitly evaluate packet capturing performance or attempts to improve the packet capturing performance. This is because in both types of work, an evaluation of the various aspects of packet capturing tools such as CPU utilization and packets captured are usually presented. Below we briefly describe both types of work.

Deri [7] has suggested that the packet capturing process is inefficient due to overhead involved in copying the packet. The work proposes a new socket type, PF_RING, in which the packet can be copied directly from the device driver buffer to user-accessible memory, drastically reducing memory allocation and copying operations. A later work improved upon PF_RING by proposing a new architecture in which multi-core processor can be utilized to increase the monitoring capability of the system [8]. In both of these works, Deri et al. discovered that the capturing process do not handle high traffic volume well due to the memory operations from device driver to kernel and from kernel to user level, as well as sub-optimal utilization of resource available at device and kernel level.

In [18, 5], the authors investigate the performance of packet capturing tools in various software and hardware platform. The metrics investigated in these works are the packets captured [18, 5], with [18] having some emphasis on the CPU utilization and [5] focusing on the percentage of packets captured. Both works are important because they evaluate the performance of packet capturing tools using common platforms, so the valid conclusions can be drawn regarding the hardware or software stacks involved in packet capturing.

Our work differs from these works in that we do not emphasize on the performance of the packet capturing tool, but whether the packet capturing tool affects existing applications, and if so to what degree. Even though the hardware used to host the capturing process and web server is multi-core, we only utilize one core for either the web server or the packet capturing process. This is so we could monitor the effect of co-locating the two processes, and have shown that co-location causes dramatic performance degradation to the web servers.

Chapter 5

Conclusion

In this work we examine the performance of web servers in the presence of packet capturing process. We find that CPU sharing is directly proportional to the performance degradation experienced by the web server, and separating the two processes onto different cores still has some performance impact on the web server.

This work is a good start, but more environments can be considered:

- **Serving larger web pages:** We need to repeat the experiments for the case when web servers are serving larger web pages. This would make the web server more I/O bound, and having the capturing process write to disk should create another venue for resource contention.
- **Serving dynamic pages:** In this work we have looked at the case when web servers host static pages, dynamic page would put more CPU demand on the server, and the performance impact of such needs to be investigated.
- **Caching pages:** When serving static pages, the web page can be cached in memory, thus avoiding the disk completely. More experiments should be performed to investigate the effect of such a strategy.
- **Monitoring technology:** In this work we do not take advantage of the prototypes made available from prior research works (e.g., PF_RING), other techniques (e.g., sampling), or other types of packet capturing tools (e.g., dumpcap), future work will investigate these varieties.

We believe this work is the first step towards thoroughly understanding the behavior of co-locating capturing process and web servers. From these experiments, we can understand how to best capture packets when the capturing process has to be co-located with the on-line service, and whether new techniques can be applied to perform network measurement.

References

- [1] Apache. Apache MPM Worker.
- [2] Apache. Apache Performance Tuning.
- [3] Apache. Apache Software Foundation.
- [4] Paramvir Bahl, Ranveer Chandra, Albert Greenberg, Srikanth Kandula, David A. Maltz, and Ming Zhang. Towards highly reliable enterprise network services via inference of multi-level dependencies. In *Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications*, SIGCOMM '07, pages 13–24, New York, NY, USA, 2007. ACM.
- [5] Lothar Braun, Alexander Didebulidze, Nils Kammenhuber, and Georg Carle. Comparing and improving current packet capturing solutions based on commodity hardware. In *Proceedings of the 10th annual conference on Internet measurement*, IMC '10, pages 206–217, New York, NY, USA, 2010. ACM.
- [6] Xu Chen, Ming Zhang, Z. Morley Mao, and Paramvir Bahl. Automating network application dependency discovery: experiences, limitations, and new solutions. In *Proceedings of the 8th USENIX conference on Operating systems design and implementation*, OSDI'08, pages 117–130, Berkeley, CA, USA, 2008. USENIX Association.
- [7] Luca Deri, Netikos S. P. A, Via Del Brennero Km, and Loc La Figuretta. Improving passive packet capture: Beyond device polling. In *In Proceedings of SANE 2004*, 2004.
- [8] Francesco Fusco and Luca Deri. High speed network traffic analysis with commodity multi-core systems. In *Proceedings of the 10th annual conference on Internet measurement*, IMC '10, pages 218–224, New York, NY, USA, 2010. ACM.
- [9] Ibrahim F. Haddad. Apache 2.0: The Internals of the New, Improved.
- [10] Gianluca Insolvibile. Inside the Linux Packet Filter.
- [11] Gianluca Insolvibile. Inside the Linux Packet Filter, Part II.
- [12] Dan Kegel. The C10K problem. 2006.
- [13] ACME Laboratories. `http_load` - Multiprocessing HTTP Test Client.
- [14] NetCraft. February 2011 Web Server Survey. 2011.
- [15] Nginx. Nginx.

- [16] Nginx. Nginx.
- [17] Lucian Popa, Byung-Gon Chun, Ion Stoica, Jaideep Chandrashekar, and Nina Taft. Macro-scope: end-point approach to networked application dependency discovery. In *Proceedings of the 5th international conference on Emerging networking experiments and technologies*, CoNEXT '09, pages 229–240, New York, NY, USA, 2009. ACM.
- [18] Fabian Schneider and Jörg Wallerich. Performance evaluation of packet capturing systems for high-speed networks. In *Proceedings of the 2005 ACM conference on Emerging network experiment and technology*, CoNEXT '05, pages 284–285, New York, NY, USA, 2005. ACM.

DISTRIBUTION:

2 MS 9018 Central Technical Files, 8944
1 MS 0899 Technical Library, 9536 (electronic copy)

