

Workshop on Performance and Architecture of Web Servers (PAWS-2000)

June 17-18, 2000, Santa Clara, CA

(Held in conjunction with SIGMETRICS-2000)

1 Motivation

After having grown briskly in the last several years, Internet services have not only entered the mainstream of society but also moved into areas where a plain best-effort service model is no longer adequate. This phenomenon is well illustrated by two major thrust areas: Electronic commerce where poor performance or unavailability could be very expensive, and streaming media services (including voice over IP) where quality of service is a fundamental requirement. These areas have brought to light a host of performance, availability and architectural issues that must be resolved effectively in order to prevent widespread customer dissatisfaction that could adversely affect the long-term growth of online services. In particular, the unresponsiveness and apparent failures of e-commerce servers, and consequent loss in revenue for e-commerce industry is well noted. Similarly, the difficulties in providing adequate quality of service have stunted the proliferation of streaming media services on the Internet. The advances needed for handling these issues fall into the following four areas: better architectural support (both hardware and software), better workload characterization and development of robust capacity engineering methods, application level overload control, and ensuring quality of service. The workshop was intended to bring together researchers and practitioners to exchange ideas as they attempt to cope with many of these challenges for Internet servers.

2 Keynote Speech

The keynote address entitled, "Network versus Server Issues in End-to-End Performance", was delivered by Christian Huitema from Microsoft Corporation. Huitema projected the growth rate of Internet usage and quantified the delay components in web services. Based on his de-

tailed measurements, he suggested that the servers seem to contribute to about 40% of the delay. He then discussed about the state of the network in terms of packet losses and round-trip time distribution. The performance issues of DNS server were discussed in the context of scalability. Huitema also raised the issue of the processor keeping up with increasing LAN speeds, which are increasing much faster than Moore's Law. In particular, the deployment of 10 Gb ethernet in near future would most certainly tax the fastest processors available, yet there are applications that can utilize such bandwidths. Handling such high speed networks would require innovations in a number of areas including processor and memory architecture, I/O processing, lightweight protocols, and offloading packet processing. At the system level, he proposed large clusters that use efficient resource management techniques and pushing the content nearer to the clients. At the end, Huitema discuss the "Gnutella" model and suggested that the approach could be a viable solution for several web applications. He suggested to distribute the load as much as possible to the client's side and placing in the servers only the functions that cannot be placed in the clients. Several interesting questions, including the feasibility and applicability of peer-to-peer computing, were raised by a large audience and discussions continued over the break that followed the keynote address.

3 Technical Paper Sessions

There were two technical sessions, each consisting of four papers. The presenter of the first paper of the first session proposed the concept of crawler-friendly web servers that can save significant amount of bandwidth by discriminating between crawlers and surfers. They proposed ideas in which the crawlers could be helped in identifying unchanged pages and in discovery of new links using metadata assists. These aspects will avoid unnecessary updates

and thereby reduce the bandwidth usage. The second paper was on consistency and locking for distributed updates to the web servers using a file system. The consistency protocols used in file systems demonstrate poor performance when used in distributed web server environment. This work introduced a publish consistently model that is well-suited for web workloads. The third paper provided the justification of the service differentiation at the server level. They proposed architecture and algorithms for supporting differentiated QoS in a web server. The last paper of the first session unveiled the overload issues involved in upload applications. They have developed a framework called, Bistro, that can be used for handling the upload problems during overload. In the proposed framework the actual data transfer is decoupled from the deadline by using shared resources.

The second session started with a paper on minimizing the impact of orphan requests in e-commerce services. They proposed a kernel service which keeps track of requests being performed and notify processes of any cancellations. They discussed an experimental implementation and showed a performance improvement of about 30% for a real workload. The second paper was based on improving web responsiveness by pushing politely. They presented a framework that exploits idle periods by pushing future requests speculatively, while giving priority to the on-demand requests. The third paper gave a detailed characterization of user sessions by analyzing the data collected from the 1998 World Cup Web site. The paper also discussed how these characterization could be used for improving the performance of web servers. The last paper of the session was on the end to end performance modeling of web servers. The paper presents a queuing model and also identified bottleneck in web server performance.

Both the sessions were very well attended and almost all the papers generated a lot of discussions.

4 Panel Session

A panel session on the architecture and performance of Internet servers was moderated by Arun Iyengar (IBM T.J. Watson Research Center). The panelists were: Martin Arlitt (Hewlett Packard Labs), Peter Druschel (Rice University), Jim Pitkow (GroupFire Inc.), Irving Robinson (Intel Corporation), and Brian Snitzer (IBM Global Services).

Iyengar introduced the panel members and discussed the factors affecting the web performance, which included the nature of traffic, caching approaches, server hardware, back-end databases, operating system, TCP/IP implemen-

tation, and other networking aspects. Druschel focussed on the impact of server delay on the client's response time. He discussed the factors affecting the performance of servers, and focussed on the solutions at the operating system level. For future directions, he listed service differentiation, scalable, and distributed web servers. Arlitt discussed about the issues involved in predicting the resource and architectural requirements of future Internet workloads by understanding current and emerging workloads. He outlined some of his recent work on workload characterization, and pointed out the hurdle in getting companies to share the interesting data they have. Snitzer is involved in managing the IBM servers used for special events. He showed several data related to these busy servers, and discussed issues involved in managing these servers that encounter very high volume of data. In addition to high performance, these servers also need to be highly available. Robinson addressed the issues of scalability, I/O, reliability, and availability of Internet servers. He discussed about the scalability in the networking of servers and the I/O technology needed to support such scalability. Pitkow presented some interesting results about user behavior in web access and presented a model to characterize this behavior, which has a significant impact on the server design process. The presentations by the panelist was followed by several intriguing questions from the audience, ranging from traffic characterization, parallel computing, quality of service, I/O, and denial of service attacks.

5 Issues and Challenges

The internet research community is likely to face a number of challenges at both the network and server side to satisfy the fast growth of internet services. With the network bandwidth increasing much faster than the server capacity, the bottlenecks will be observed on the server side. As reported by several PAWS-2000 panelists, the web-site hit rate at major events is already approaching 1 million per minute. As Internet usage skyrockets throughout the world, this number could easily increase by more than an order of magnitude for major world-wide events such as the Olympics. However, the hit count doesn't tell the entire story — as the richness, perceptual quality, dynamism, and security requirements of the content increase, the size of the web-pages and the processing time required to render them could itself increase by an order of magnitude. Studies have shown that the use of HTTPS to retrieve static web pages could increase processing requirements by an order of magnitude. Thus, even from a rather short-term perspective, we need to be able to design Internet servers that can handle sustained

10 million hits/minute with 10 times the bandwidth and processing requirements of today's servers. At the same time, the increasing criticality of Internet applications demands that the servers exhibit a very high level of robustness and availability while delivering transaction response times under a few hundred milliseconds with a high probability. Based on our experiences and discussions at the PAWS workshop, we enumerate the major technical problems that need to be solved to enable the scalability of Internet servers.

1. *The servers need to keep up with the increasing LAN speed by avoiding inefficiencies in the chip-set, protocol stack, and the operating system.*
2. *The servers need to be efficient and provide high performance for the increasing dynamic content, especially in e-commerce environments.*
3. *We need to explore scalable solutions to enable scaling up web server farms for handling huge content and hit rates.*
4. *With the advent of peer-to-peer computing, there will be several issues in supporting distributed internet services.*
5. *The management and engineering of large servers and server farms is becoming a critical issue with the increase in volume and reliability requirements.*
6. *The architectural, O/S, and I/O issues of web servers needs to be revisited in the context of Internet workload.*
7. *Support for quality of service is a critical issue because of the increasing usage of continuous media applications and e-commerce transactions.*
8. *Scaling and performance issues for back-end systems has not yet been studied well, although it is likely to be a bottleneck in e-commerce environment.*
9. *Overload control, which is strongly imposed in telecommunications industry, is almost non-existent in Internet services. The traffic and load characterization of Internet requests does indicate the need for good overload control mechanisms.*
10. *Secure transactions, which form the core of e-commerce traffic, have severe resource requirements; thus needing efficient mechanisms for higher throughput.*
11. *There is a strong need for a data repository of server performance and traffic logs. Several industry could provide valuable resources for this effort.*

A more detailed version of these issues and challenges can be downloaded from the PAWS website "<http://kkant.ccwebhost.com/PAWS/>".

6 Acknowledgements

We express our gratitude to the SIGMETRICS organizing committee for their help in taking care of all the mechanics of the workshop organization. Our thanks to ACM and Intel Corporation for sponsoring the workshop. It was a pleasure to have Dr. Christian Huitema as our keynote speaker. We appreciate the efforts of Dr. Arun Iyengar for assembling and organizing an excellent panel. We also thank all the authors and workshop participants. Above all, our sincere thanks goes to the following program committee members for their time and effort in putting up the technical sessions.

Program Committee:

Nina Bhatti (Hewlett Packard Labs, Palo Alto, CA)
Ian Leslie (University of Cambridge, Cambridge, UK)
Richard Muntz (University of California, Los Angeles)
Venkat Padmanabhan (Microsoft Research, Redmond, WA)
Avi Rubin (AT&T Labs – Research, Florham Park, NJ)
Michael Rumsewicz (Ericsson Australia, Melbourne, Australia)
Prashant Shenoy (University of Massachusettes, Amherst)
Edmundo de Souza e Silva (Universidade Federal, Rio de Janerio, Brazil)
Hui Zhang (Carnegie Mellon University, Pittsburgh, PA)

Program Chairs:

Krishna Kant, Intel Corporation, Beaverton, OR
krishna.kant@intel.com

Prasant Mohapatra, Michigan State University, Lansing, MI
prasant@cse.msu.edu