

July 2001
155E-0701A-WWEN
Prepared by:
Industry Standard Server Division

Compaq Computer Corporation

Contents

Introduction	3
Benchmarking Cache Performance	4
Product Summary	5
Performance Summary	6
Discussions	6
Forward Proxy Test.....	6
Reverse Proxy Test.....	6
Streaming Test.....	7
Conclusion	7
Related Documents	7

Compaq TaskSmart C-Series Servers Performance Guide

Abstract: Compaq TaskSmart C-Series servers are based on industry-leading Compaq *ProLiant*[™] server platforms and the leading content-acceleration software suite, Inktomi Traffic Server. The TaskSmart C-Series servers are optimized to accelerate content serving and accesses. This white paper is to provide the reader with a brief summary of the performance results of the TaskSmart C-Series servers. The performance parameters used are also discussed to provide the reader with guidelines in deploying TaskSmart C-Series servers in forward proxy, reverse proxy, and streaming applications.

155E-0701A-WWEN © 2001 Compaq Computer Corporation

Compaq, Compaq Insight Manager, the Compaq logo, ProLiant, and TaskSmart Registered in U.S. Patent and Trademark Office.

Microsoft and Windows NT are trademarks of Microsoft Corporation in the United States and other countries. All other product names mentioned herein may be trademarks of their respective companies.

Intel is a trademark of Intel Corporation in the United States and other countries.

Inktomi Traffic Server is a trademark of Inktomi Corporation in the United States and other countries.

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is provided “as is” without warranty of any kind and is subject to change without notice. The warranties for Compaq products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

Introduction

A proxy cache can be deployed either as a client (forward) accelerator or a server (reverse) accelerator.

- Because a forward proxy is located closer to the clients than to the Web server, where the server caches commonly accessed page objects, the proxy server can serve these objects much faster to the clients than the origin Web server. The upstream network bandwidth utilization is significantly reduced when most clients are served by a forward proxy instead of a Web server.
- A reverse proxy is located closer to the origin Web server than to the clients. It intercepts and responds to clients' requests, off-loading the Web server TCP connections as well as serving already cached objects.

Regardless of the deployment choice, a proxy cache is located somewhere between clients and server to effectively improve the client access time (also commonly known as latency).

The effectiveness of a cache to improve the client access time depends on several factors:

- Hardware selection and software parameter configurations, which are primarily controlled by the IT department.
- Client browser applications using more HTTP/1.1 persistent connections, which increase server performance by reducing TCP connection processing.
- Fewer no-cache directives (in HTTP headers), which contribute more objects served by the proxy cache.
- Origin Web servers with fewer no-cache HTTP header directives, which allow a proxy cache to cache more objects. Servers with a greater percentage of objects in cache can increase both the effectiveness and the performance of a proxy by increasing the cache hit ratio.
- Problems associated with network architecture and traffic conditions, including cache server placement in relation to the client.
- Access patterns of multiple clients, which can have significant impact on the cache server effective latency.

Because the actual latency depends on these factors, throughput is commonly used as a performance indicator of a cache. Throughput is the amount of traffic that can pass through a proxy cache server, and is commonly represented in rps (requests per second)¹ and Mbps (million bits per second). Higher throughput can provide lower latency, which in turn translates to faster access times.

Related to TCP performance is the robustness of a TCP implementation inside a proxy cache server. Some vendors optimize the TCP performance to the point that it begins to downgrade their system reliability. Those vendors will have cache that can show attractive performance numbers, but eventually, the system user starts experiencing instability of the system in a heterogeneous network environment. Compaq uses proven networking stack that complies with all necessary standards in its proxy cache servers to enhance reliability.

¹ Some also use operations per sec (ops), transactions per sec (tps) or objects per sec (obj/sec).

Benchmarking Cache-Performance

The “Performance Summary” section in this white paper presents benchmark performance numbers for several models of the Compaq *TaskSmart*TM C-Series appliance servers. Each benchmark test was run in a condition that best matched a deployment scenario, such as forward/reverse proxy and number of network connections.

Benchmark programs are used to measure the performance of a cache. Commonly used benchmark programs are WebBench² and Web Polygraph³, as well as proprietary programs. Compaq uses the WebBench test suite to establish performance benchmarks for reverse proxy, and Web Polygraph establishes the forward proxy benchmarks.

Product vendors use benchmarked performance numbers in sales tools such as advertising and marketing. They also use benchmark performance numbers to provide useful information for customer use when making purchasing and deployment decisions. Most vendors publish one or two performance numbers to give the customer an idea of a product’s position in the market regarding its performance. When making purchase decisions, customers should not rely only on limited benchmarks published by different vendors but also consider benchmarking done in a neutral and consistent environment. An example of a neutral environment is the Cache-Off events⁴ held occasionally by Measurement Factory (www.measurement-factory.com). When compared to simulated conditions provided by a benchmark program, most proxy cache server traffic patterns present completely different workload conditions.

After design and implementation, performance evaluation can serve to alert support personnel to a potential problem that could be caught early. For example, if the MRTG graphs of the cache miss latency show a degrading trend, it can be concluded that bandwidth limits are being reached in the upstream Internet transit connections.

² From Ziff Davis. See www.WebBench.com for the details.

³ From Measurement Factory. See www.Web-Polygraph.org for the details.

⁴ Compaq hosted the second and the third Cache-Off events in Houston, Texas, in January 17, 2000 and in September 11, 2001. See these Cache-Off events reports at the Measurement Factory website at www.measurement-factory.com

Product Summary

The Compaq TaskSmart C4000 Model 30 server is best suited for a forward proxy deployment to improve Web browsers' access time, and to reduce upstream network bandwidth requirements. By nature, a forward proxy accesses origin Web servers in lieu of the individual client browsers.

The Compaq TaskSmart C4000 Model 40 server is best suited for a reverse proxy deployment to accelerate the origin Web servers. Because a reverse proxy, by nature, is memory-intensive and is not disk-intensive, 1 gigabyte of memory and two ATA drives are sufficient for most deployments.

Note the following reverse proxy sizing example:

With two 20-GB disk drives installed in the TaskSmart C4000 Model 40 server, about 30 GB will be available for the disk cache after factoring the disk space for OS, application, log files, and other user application tools. Assuming a 100-GB website, 30 percent of the most commonly accessed cacheable objects of the website can be stored in the disk cache.

The Compaq TaskSmart C4000 Model 50 server is best suited for a forward proxy deployment to accelerate Web browsers' access time and streaming media applications.

By deploying one or more of these TaskSmart C-Series appliance servers, website users will experience faster and better quality access to the Internet.

Table-1 A Brief System Configuration of TaskSmart C4000 Models 30, 40 and 50 Servers

Configuration	Model 30	Model 40	Model 50
Usage	Client-side acceleration (forward proxy)	Host-side acceleration (reverse proxy)	Client-side acceleration (streaming)
Software	Inktomi Traffic Server	Inktomi Traffic Server	Inktomi Traffic Server + Media IXT
Processor	(1) 933-MHz Intel PIII	(1) 933-MHz Intel PIII	(1) 933-MHz Intel PIII
RAM	384 MB	1 GB	1 GB
Storage	(2) 1-in, 20-GB ATA/100, 7200 rpm	(2) 1-in, 20-GB ATA/100, 7200 rpm	(2) 1-in 18-GB WU3 SCSI, 10K rpm, non-hot-plug
Network I/F (See Note)	(2) 10/100 Fast Ethernet	(2) 10/100 Fast Ethernet	(2) 10/100 Fast Ethernet
Chassis	1U rack-ready	1U rack-ready	1U rack-ready
Power Supply	(1) 180 W	(1) 180 W	(1) 180 W

Note: A gigabit network interface board or the Compaq Remote Insight Lights-Out Edition board (157866-001 NA) may be added as an option in the PCI slot.

Performance Summary

Table-2 Performance Parameters Achieved by the TaskSmart C-Series Appliance Servers

Use	Model 30	Model 40	Model 50
Forward Proxy	150 rps (15 Mbps) Response Time: 2 sec.	150 rps (15 Mbps) Response Time: 2 sec.	200 rps (20 Mbps) Response Time: 1.7 sec.
Reverse Proxy		1,200 rps (60 Mbps)	1,200 rps (60 Mbps)
Streaming			84 Mbps (840 streams @ 100 Kbps) 97 Mbps (130 streams @ 750 Kbps) Packet Loss: 0 @ <40 Mbps <1% @ 40-70 Mbps <2% @ >70 Mbps

Notes

1. For each model, a 100-Mbps Ethernet network port was connected to a Layer-4 switch to run all the benchmarks, unless otherwise stated.
2. All transactions in the WebBench workload used non-persistent connections. The standard for non-persistent connections states that for each object, a client opened a connection, received the object, and closed the connection. The requests per second (rps) throughput will be higher if persistent connections are used.

Discussions

Note in Table-2 that the forward and reverse proxy throughput numbers are not related because they were generated using different benchmark programs and payloads. Although TaskSmart C4000 Model 30 server is intended for forward proxy, and TaskSmart C4000 Model 40 servers for reverse proxy, both models may be deployed as forward or reverse proxy.

Forward Proxy Test

A forward proxy, by nature, serves a relatively small group of clients, and these clients can access fairly diverse websites. Therefore, a forward proxy normally caches a wide range of URLs and objects. It is not practical to provide large amount of RAM to have high hit ratio in the memory cache. Instead, providing a number of high performance disks is a more balanced approach in sizing a forward proxy. For this reason, in TaskSmart C4000 Models 30 and 40 servers, two optimized ATA drives are used to provide good economic models without sacrificing performance.

Reverse Proxy Test

A reverse proxy, by nature, will cache only from one website. Therefore, knowing the access pattern of a given website will help tremendously in selecting an appropriate reverse proxy server. For example, knowing the total storage requirement of a website, and taking a factor (for example, 10 percent) of the most commonly accessed objects will result in approximate disk capacity in a reverse proxy cache. Taking another factor of the disk capacity will result in

approximate memory capacity. TaskSmart C4000 Models 40 and 50 servers both have 1 GB of memory, and they have similar reverse proxy benchmark results because the benchmark load test is more memory-centric.

Streaming Test

For streaming throughput on the TaskSmart C-Series Model 50 server, Table-2 shows a set of numbers as guidelines for the number streams for different bit rates. Here, a stream corresponds to a client user. The actual number of streams will depend on the application, and users' tolerances on video and audio qualities. For example, telecasting of a CEO presentation will not require as stringent video quality (in terms of size, resolution, and frame rate) as a high-quality movie. How a video is compressed at the source also has significant impact on the streaming throughput. For variable bit-rate-encoded videos, the overall network bandwidth utilization and the cache server CPU utilization will vary according to the instantaneous bit-rate of the video traffic. When the peaks of the video traffic fluctuations are within the network and CPU utilization upper limits, there will be no packet loss. However, when the peaks of the video traffic fluctuations start getting over the network and CPU utilization, packet losses start occurring. The severity on the video and audio quality with respect to packet losses depend on the application, users' tolerance, and the packet loss interval and burstiness. For example, users may tolerate some infrequent video quality degradation as long as the audio quality is acceptable (for example, when viewing a CEO speech.) Streaming players also contribute to video and audio qualities perceived by users. Upon packet loss events, some players reduce the video frame rate, and the audio sampling rate necessary in dealing with variable bit-rate streams. Others freeze the video frame until there are virtually no packet errors.

Conclusion

TaskSmart C4000 Models 30, 40 and 50 servers provide good price performance choice for customers to deploy as a forward proxy, a reverse proxy, or a streaming proxy.

Related Documents

Compaq TaskSmart C-Series Content Acceleration Servers Deployment Guide, white paper, Document number 155D-0701A-WWEN.

Compaq TaskSmart C-Series Streaming Servers Deployment Guide, white paper, Document number 155C-0701A-WWEN.

Compaq TaskSmart C-Series Servers Feature Procedures Guide, white paper, Document number 158D-0701A-WWEN.

Filename: 155E-0701A-WWEN.doc
Directory: V:\cpq-ecg\Gazza\155E-0701A-WWEN\DOC
Template: C:\Program Files\Microsoft
Office\Templates\Compaq\ActiveAnswersTemplate.dot
Title: Compaq TaskSmart C-Series Servers Performance Guide
Subject:
Author: Mike Rometry
Keywords:
Comments: Released: 28-Feb-2000
Creation Date: 7/9/01 2:17 PM
Change Number: 43
Last Saved On: 7/12/01 2:27 PM
Last Saved By: monicae
Total Editing Time: 307 Minutes
Last Printed On: 7/13/01 11:06 AM
As of Last Complete Printing
Number of Pages: 7
Number of Words: 2,257 (approx.)
Number of Characters: 12,415 (approx.)