A Scalable HTTP Server:
The NCSA Prototype

Eric Dean Katz
Michelle Butler
Robert McGrath

National Center for Supercomputing Applications
University of Illinois, Urbana-Champaign, Il 61820
ekatz@ncsa.uiuc.edu
mbutler@ncsa.uiuc.edu
mcgrath@ncsa.uiuc.edu

Abstract

While the World Wide Web (www) may appear to be intrinsically scalable through the distribution of files across a series of decentralized servers, there are instances where this form of load distribution is both costly and resource intensive. In such cases it may be necessary to administer a centrally located and managed http server. Given the exponential growth of the internet in general, and www in particular, it is increasingly more difficult for persons and organizations to properly anticipate their future http server needs, both in human resources and hardware requirements. It is the purpose of this paper to outline the methodology used at the National Center for Supercomputing Applications in building a scalable World Wide Web server. The implementation described in the following pages allows for dynamic scalability by rotating through a pool of http servers that are alternately mapped to the hostname alias of the www server. The key components of this configuration include:

- Cluster of identically configured http servers
- Use of Round-Robin DNS for distributing http requests across the cluster
- Use of distributed File System mechanism for maintaining a synchronized set of documents across the cluster.
- Method for administering the cluster.

The result of this design is that we are able to add any number of servers to the available pool, dynamically increasing the load capacity of the virtual server. Implementation of this concept has eliminated perceived and real vulnerabilities in our single-server model that had negatively impacted our user community. This particular design has also eliminated the single point of failure inherent in our single-server configuration, increasing the likelihood for continued and sustained availability. While the load is currently distributed in an unpredictable and, at times, deleterious manner, early implementation and maintenance of this configuration have proven promising and effective.

Please note that it is neither the intention of the author nor of NCSA to either endorse or malign any particular architecture or hardware type. For this reason I will make no specific references to architecture types or equipment performance. Note also that any references to IP addresses or host names are not representative of the actual information and the reader should not attempt to make connections using this information.

1. Introduction

At the National Center for Supercomputing Applications we have seen a steady growth in requests to our http server from an initial weekly level of 91,000 to our most recent level of just over 1.5 million. That represents a consistent growth pattern of almost 11% per week for the 40 week reporting period. (See Figure 1.)

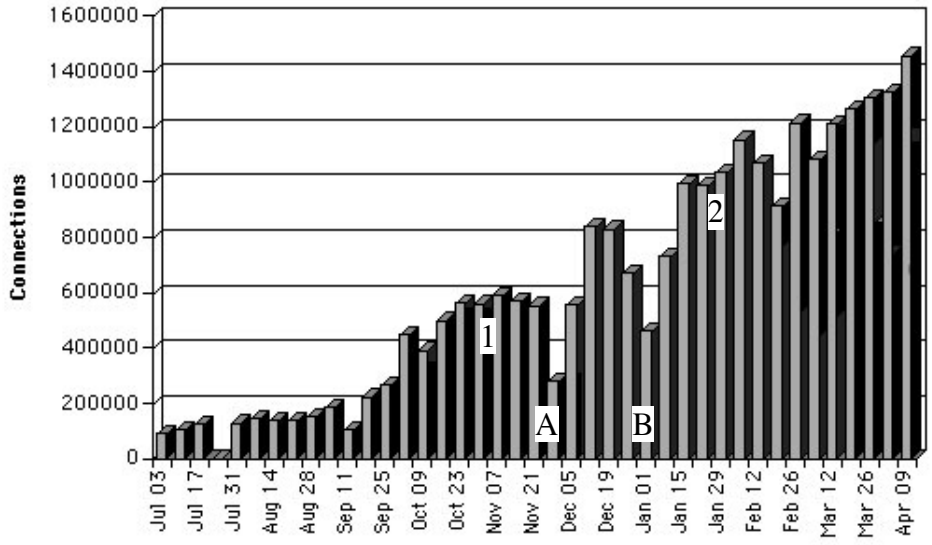

Figure 1. Growth in HTTP Requests to www.ncsa.uiuc.edu. Inception to Present.
A and B indicate reductions in connectivity commensurate with the observance of academic holidays.
On average 35% to 40% of all connections to our server are from the education domain (.edu)

There are two periods of time that are of particular interest in the context of this discussion. The first of these, between Oct 23 and Dec 05 of 1993 ([1]) is indicative of a time in which the number of http requests coming into our server exceeded its capacity to respond. Our solution was to migrate the http server software to a faster, more powerful architecture which solved our problem in the short term. The new server, however, was unable to respond to the needs of our user community as weekly accesses approached and eventually surpassed the 1 million mark ([2].) We determined that the only viable long term solution would embody some form of load distribution across a multi-server configuration.

This left us with two options; dividing our document tree amongst several systems each serving a portion of the document set and each answering to a unique host name, or sharing the entire document set with a number of servers each answering to a common host name. The former required a great deal of file manipulation and would have broken many existing links. This involved a prohibitive level of complexity and impact to the user community. An acceptable implementation of the latter would accommodate the growing demands placed upon the server and would impart virtually no hardship on the user community, thereby presenting the most favorable solution.

2. Overview

The manner in which the different network and computation components come together to define our scalable web are detailed in Figure 2. Understand that this simplified diagram is intended to illustrate the logical or virtual path of a request from the client to our server and is not indicative of the actual flow of information.

Figure 2. Logical Flow Chart of NCSA's ScalableWorld Wide Web Server