Preliminary Report on the nCUBE Parallel Web Server

E. E. Charrette

nCUBE Corporation
MIT Building E34-554
42 Carleton Street
Cambridge, MA 02142

November 23, 1994

The nCUBE Web Server

Overview

You are watching a prototype version of the nCUBE parallel World Wide Web server.

The server is based on the public domain version of the CERN Web server (Click here for more information). It contains all the functionality of the base server, plus support for a parallel server.

The nCUBE Parallel Web server was developed to address the needs of large web server site. In this context, large might mean either that the document volume of the server is large, or that the number of nearly-simultaneous connections is large enough to seriously degrade the ability of the server to serve the incoming requests.

The nCUBE Parallel Web server is really "n" independent Web servers running on "n" independent nCUBE processors. Each server has access to the entire content tree and a single ethernet controller with a unique IP address. For this reason, we will sometimes refer to the parallel server as the meta-server and each individual server running on a single nCUBE node as a server.

Advantages of the nCUBE Parallel Web Server

CERN Web Server plus Parallel Extensions

> The server is based on the public domain version of the CERN Web server (Click here for more information). It contains all the functionality of the base server, plus support for a parallel server.
>
> Parallelism is obtained in two ways.
>
> > Parallelism via striped disk storage
> >
> > > The server content can either be stored on the host file system, or the nCUBE parallel disk array. For performance reasons, the later file system is the preferred choice.
> > >
> > > The heart of the nCUBE parallel file system is a scalable array of SCSI disks. and a body of software which presents the distributed disks as a single logical file system. In general, all data written to the parallel file system is automatically striped across all the disk. This minimizes disk contention and allows the aggregate disk performance to far exceed the speed of accessing data from a single SCSI disk.
> >
> > Parallelism via Distributing the Hypertext Links
> >
> > > An important design criteria for the nCUBE Parallel Web server was to insure that the web server content must be identical in structure and protocol to the serial implementation. This allows for the possibility that the content tree could be moved between the nCUBE parallel file system and a serial UNIX environment and remain exactly and totally unchanged.
> > >
> > > This implies that any attempt to functionally decompose the process of the server must be done in the server program and not the content. We achieved this by adding a small body of code to the core distribution of the CERN web server. That code parses the outgoing HTML documents on the fly and

modifies the content so that the hypertext links to the same meta-server are randomly altered to have different server addresses. For large servers, this effectively splits further requests to the meta-server evenly between the n servers

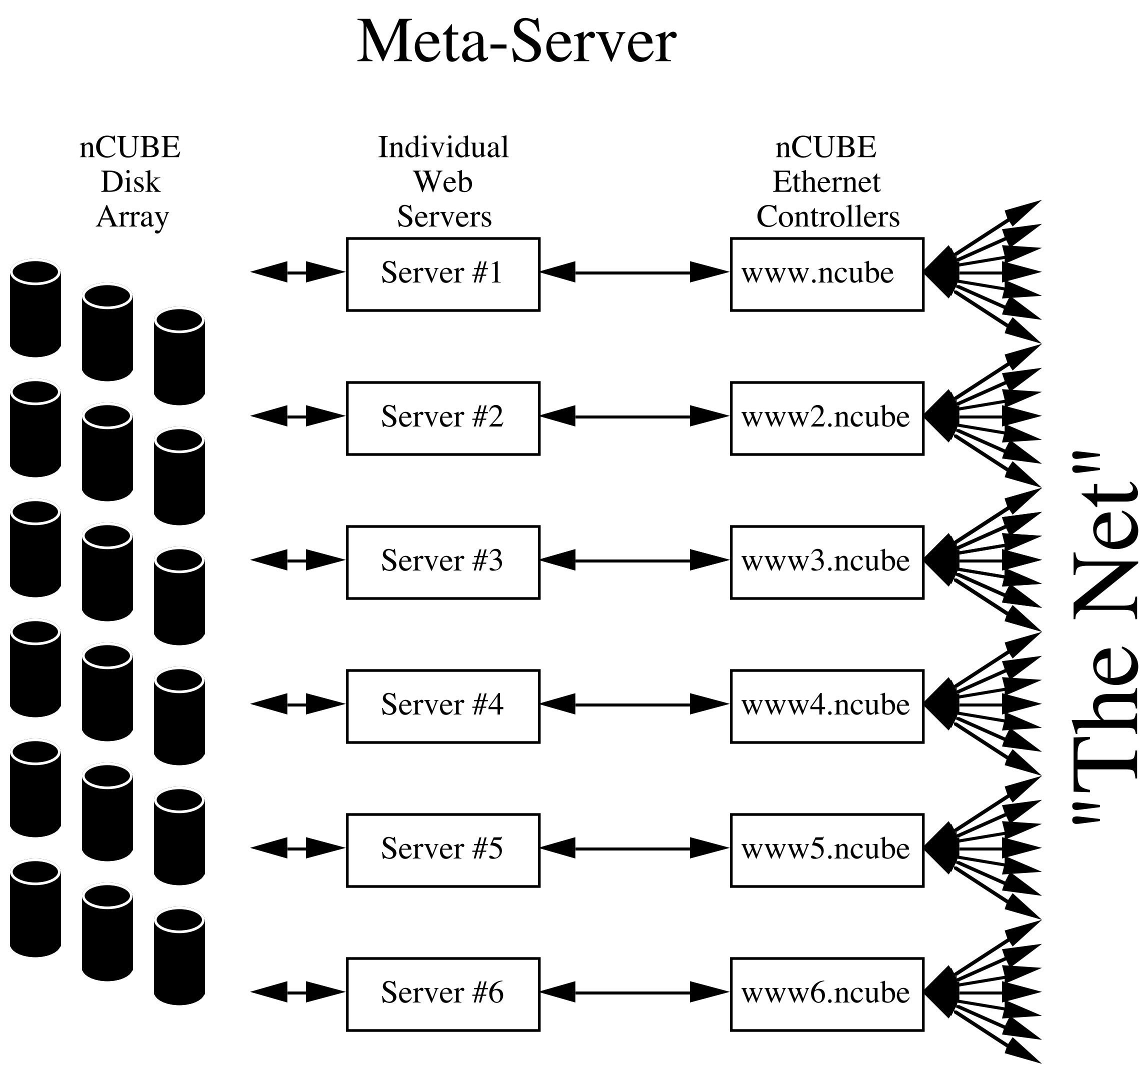

Scalability

By designing the meta-server so that each individual communicates with only one ethernet controller, we have minimized the number of possible bottlenecks in the system. Preliminary analysis has shown that the only possible point of contention might be access to the disk array. This can be minimized by correctly configuring the disk array so that the number of disk controllers is at least as large as the number of servers.

When configured in this way, we expect system performance to scale roughly linearly with the number of servers.

Shared File system

The nCUBE parallel web server can be configured so that the content resides on either the host file system or the nCUBE parallel file system. There are obvious performance gains from using the parallel file system and it is anticipated that all sites would use that approach.

Disks in the parallel file system are connected to controllers on an nCHANNEL board. The nCHANNEL board is an I/O board with 16 independent (8 on a nCUBE Model 5 system) I/O channels. Each channel is managed by an nCUBE-2 processor and can contain an I/O module, such as a SCSI controller or ethernet module. The nCHANNEL board itself is installed in an I/O card cage, which can hold up to 8 I/O boards. Each processor on the nCHANNEL board can connect directly to the processor array, other I/O processors, and its own I/O device.

nCUBE provides file system software which makes the distributed disk array appear as a single logical UNIX file system. This software automatically stripes files across multiple disk spindles to enhance performance and minimize disk contention. Striping is enacted automatically by the operating system, so the file system can be nfs mounted to the host computer, if desired.

Since the distributed disk array appears to be a single large UNIX file system, the server content in the nCUBE parallel web server can be stored in a manner consistent with traditional serial content trees. In fact, one of the design parameters for the parallel web server was to insure that it was possible to import and export content directly from/to a serial web server.