Internet2 Distributed Storage Infrastructure (I2-DSI) Project: Improving Global Access to Digital Collections

Bert J. Dempsey[1], Micah Beck[2], and Terry Moore[2]

[1]School of Information and Library Science
University of North Carolina at Chapel Hill
dempsey@ils.unc.edu

[2]Innovative Computing Laboratory
University of Tennessee, Knoxville
{mbeck, tmoore}@cs.utk.edu

The Internet2 Distributed Storage Infrastructure (I2-DSI) Project (http://dsi.internet2.edu/) [1] is developing a novel distributed network storage solution as part of Internet2 (http://www.internet2.edu) technical activities. Through the replication of application source objects to dedicated replication servers in the network, I2-DSI seeks to reduce the network distance between clients and servers and thereby address the performance challenges facing Internet-based applications. The approach will leverage current capabilities in mass storage and high-bandwidth wide-area networks to enable significant new applications today.

Replication-based hosting services will enable applications with large data sets, streaming media, and performance-sensitive interfaces. These applications include many digital library efforts with data-intensive or media-rich environments that are stymied by the lack of high-performance access for clients located across the wide-area Internet. In contrast to commercial companies developing similar solutions [4,9], the I2-DSI project is focused on serving innovative applications within the research and education community through the development of open solutions.

Replication

The I2-DSI project is creating middleware that will support robust, scaleable replication of server-side content with transparent resolution at the client. Content providers will create Internet content channels, [1] "a collection of content that can be transparently delivered to end user communities at a chosen cost/performance point through a flexible, policy-based application of resources." Source objects within a channel are replicated across a set of dedicated replication hosts in the network (I2-DSI servers). By controlling the set of servers to which a channel is mapped and other replication parameters, the system resources within I2-DSI dedicated to replicating a channel can be tuned. This situation is in contrast to the dynamic replication of data that takes place in the current Internet through WWW caching.

Like WWW caching, I2-DSI replication seeks to improve clients' performance and reduce the demands on network bandwidth through localized access to copies of the data. Server-based replication, however, enables controlled replication costs and the tuning of the replication service to individual channel needs. Also, server-based replication extends the advantages of distributed storage in the network to a broad range of application services (e.g., streaming media servers), not just HTTP as with conventional caching.

A key long-term challenge for effective replication is to develop models of portable content [8]. Even static file replication across heterogeneous WWW servers today is not straightforward due to filesystem differences, configurable WWW server parameters, and other platform-specific issues. Dynamic content introduces the further complexities of platform-specific binaries and execution environment assumptions. I2-DSI researchers are working long-term towards APIs for content developers that address what content is portable across DSI servers and how channel providers can create, access, and update content. This research will proceed incrementally with definitions of constrained WWW and media channels as the first step.

Transparent Resolution

Clients accessing the channel will require a resolution service to aid them in finding a channel on a "nearby I2-DSI server." By mapping content channels onto Internet domain names, the Domain Name System (DNS) can be used to enable transparent network-level resolution. In the DNS database, multiple IP addresses (those of the I2-DSI servers carrying the channel) are associated with the domain name of the channel. Special DNS resolver

DL 99, Berkeley, CA USA
 1-58113-145-3/99/08 . . . $5.00

software [6] or an intermediary [2, 7] gains control of the dynamic translation of the domain name to an IP address.

The DSI-aware resolver uses network measurements or routing table information to determine the "best" DSI server for the client, and the IP address of this server is returned to the application client. If the special resolver is placed on the server-side of the DNS resolution process, application clients, whether standard commercial WWW browsers or special-purpose application clients, need not be modified in order to access I2-DSI channels.

Project Status

Under the umbrella of the Internet2 project, I2-DSI is now pursuing a full implementation of the initial architecture, including academic developers with their applications, and industrial sponsorship valued at over $1.5M in equipment and $175K in development funding. Corporate participation comes from companies in all areas of the computer and networking industry, including Cisco, IBM, Novell, Ellemtel, StorageTek, Sun Microsystems, and Starburst Communications.

The initial WAN testbed will be operational by early Q3 1999, and Figure 1 shows the location of current high-capacity (~1-terabyte) storage servers donated by IBM (large towers) and potential sites for additional servers (stars). As of May 1999, the potential sites shown in Colorado and Texas have commitments to receive high-capacity Linux-based servers from StorageTek. Additional sites are anticipated in the near future, and, also as shown, content exchange with network groups outside the U.S. are under active negotiation.

Applications

At a recent workshop at the University of North Carolina, I2-DSI developers brought together a number of leading application groups to explore collaboration in developing early I2-DSI channels [3]. These efforts will lead to the development of approximately 6-8 Internet content channels working on the I2-DSI testbed by the end of 1999. Active project participants include (see workshop URL for on-line white papers of the projects):

- digital library efforts such as the Indiana Variations project with its large collection of digital audio,
- document repositories including the HTTP and FTP-based Linux repository at UNC MetaLab and a database-backed repository of Internet standards documents developed by Normos.org, and
- image databases such as the medical imaging group with Vanderbilt University.

These and other innovative content channels will be focused on giving feedback to the I2-DSI developers from real content providers while building applications that demonstrate the power of the I2-DSI replication services architecture.

References

[1]. Beck, M. and Moore, T., "The Internet2 Distributed Storage Infrastructure Project: An Architecture for Internet Content Channels", in Computer Networking and ISDN Systems, 1998, 30(22-23): pp. 2141-2148.

[2]. G. Carpenter, G. Goldszmidt, M. Beck, T. Moore, B. Dempsey, D. Weiss, Improving the Availability of Internet2 Applications and Services, (in submission).

[3]. I2-DSI Applications Workshop, University of North Carolina at Chapel Hill, March 4-6, 1999, http://dsi.internet2.edu/apps99.html.

[4] Akamai, http://www.akamai.com/

[5] M. Beck, B. Dempsey, and T. Moore, The Internet2 Distributed Storage Infrastructure (I2-DSI) Project, homepage at http://dsi.internet2.edu/.

[6] K. Moore, J. Cox, and M. Swany, Sonar Project, University of Tennessee Department of Computer Science, http://icl.cs.utk.edu/projects/sonar/overview.html.

[7] Cisco DistributedDirector, http://www.cisco.com/warp/public/cc/cisco/mkt/scale/distr/index.shtml

[8] Beck, M., T. Moore, B. Dempsey, and R. Chawla, "Portable Representation of Internet Content Channels in I2-DSI", in 4th International Web Caching Workshop, San Diego, CA, April 1999.

[9] Sandpiper Networks, http://www.sandpiper.com/

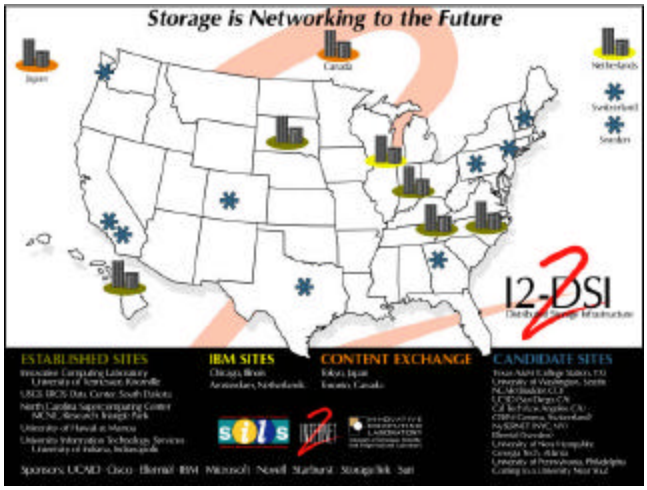

Figure 1: I2-DSI Infrastructure Map (May 1999)