

More Getting Started with Prospero

Sanjay Joshi

Abstract

This document is intended as a guide to the Prospero documentation, for those interested in knowing more about Prospero. It presents a brief survey of the literature on Prospero and briefly describes the organization of Prospero.

Introduction

Prospero is a useful tool for organizing internet resources. It is a distributed file system, which allows the users to organize objects located on different nodes in the Internet, *i.e.* several users may construct multiple, customized views of the system. One of the best papers which explains this concept is [1]. This should be the first paper to read if you are interested in learning about Prospero.¹

Implementation Details

The initial paper [1] gives the principles behind Prospero. However, the best document to learn about the implementation details is Clifford Neuman's Ph.D. dissertation [3]. This document [3] is available by anonymous FTP from PROSPERO.ISI.EDU as `/pub/prospero/papers/prospero-neuman-thesis.ps`. The appendices A and D & E, respectively of [3] have since been revised and are available as a User's Manual [2] and a Protocol Manual [4]. The User's Manual [2] should be read after reading [1]. It describes how to use the basic features of Prospero.

¹All the references mentioned in this document are available via anonymous FTP from `/pub/papers` directory at PROSPERO.ISI.EDU. They are also available as part of the Prospero distribution. A good way to obtain the Prospero distribution, if you're not running Prospero yet, is to anonymously FTP `/pub/prospero/prospero.tar.Z` from PROSPERO.ISI.EDU. Before setting up Prospero, the documents in `getting-started.txt`, `INSTALLATION.m`, and `INSTALLATION.s` should be read.

After getting familiar with Prospero through [2], the next document to read is [3]. Specially, chapter 6 and appendices B and C are very useful. Chapter 6 lays out the design of Prospero. The design is based on the client-server model. The clients handle the name resolution and directory servers on the hosts storing the directories look up the requested names in the directories and return the result to the client.

After the distribution has been untarred, the directory server source files are located in a directory called **server**. The server uses the functions defined in the library **lib/psrv**. The client source files are located in the directory **user** and use the functions from two libraries, **lib/pfslib** and **lib/pcompat** libraries. These two libraries, **pfslib** and **pcompat**, are described in detail in appendix B of [3]. The appendix C describes the structures (directory information) maintained by the Prospero file system. In addition to reading appendix C, a few of the related files in **include** directory should be consulted, specifically **include/pfs.h**. The Prospero clients and directory servers use the Reliable Datagram Protocol (RDP) to communicate with each other. The detailed description of this Protocol is given in [4], which is available by anonymous FTP as **/pub/prospero/doc/prospero-protocol-v5.ps.Z**.

References

1. Neuman, B. C., Prospero: A Tool for Organizing Internet Resources, Electronic Networking: Research, Applications, and Policy, 2(1), Spring 1992.
2. Neuman, B. C., The Prospero User's Manual, University of Washington, June, 1991.
3. Neuman, B. C., The Virtual System Model: A Scalable Approach to Organizing Large Systems, Department of Computer Science and Engineering, Tech. Rep. 92-06-04, University of Washington, June, 1992.
4. Neuman, B. C. and Augart, S., The Prospero Protocol, Version 5, ISI, USC, Aug., 1992.