A Case for Proactivity in Directory Services

Fabián E. Bustamante, Patrick Widener and Karsten Schwan
College of Computing
Georgia Institute of Technology
Atlanta, Georgia 30332, USA
{fabianb, pmw, schwan}@cc.gatech.edu

Innovative and widely distributed applications are enabled by infrastructure layers that allow distributed resources and services to be pooled and managed as though they were locally available. Advances in communication technologies and the proliferation of computing devices have made this possible; two such types of infrastructures are pervasive computing environments and computational grids. An important component of both types of infrastructures is a directory service that provides information about different objects in the environment, such as resources and people, to applications and their users. Well-known examples of such services are the Metacomputing Directory Service for Globus-based environments, and the Intentional Naming System (INS) for applications developed in the Oxygen pervasive computing project. Directory services in both types of environments must support sophisticated object descriptions and query patterns, operate in highly dynamic environments, and scale to an increasingly large number of objects and users.

Traditional directory services have been designed for fairly static environments, where updates are rare (DNS, LDAP, and X.500). Recent work has addressed the issues of expressiveness their object descriptions and query languages (through attribute-value hierarchies as in INS) and considered their scalability (through domain-based partitioning or hierarchical organization as in MDS and Czerwinski et al. SDS). However, these directory services rely on traditional 'inactive' interfaces, where clients interested in the values of certain objects' attributes must explicitly request such information from the server. Czajkowski et al. demonstrate that it is feasible to satisfy widely different information service requirements with a single, consistent framework. Their example applications range from traditional service discovery with relatively static mappings to superschedulers and application adaptation monitors, where objects and their attributes change at fast and unpredictable rates and fresh information is crucial to clients' functionality. In such scenarios, clients in need of up-to-date information have no alternative but to query servers at rates that (at least) match those at which changes occur.

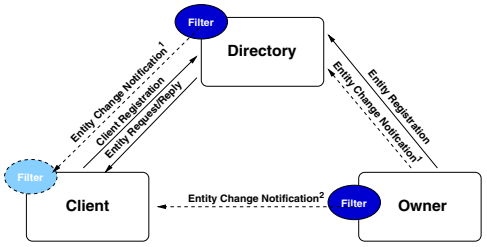

Figure 1. Extending directory services interfaces with customizable proactive mode.

In this paper, we argue that an *exclusively* inactive interface to directory services can hinder server scalability and indirectly restrict the behavior of potential applications [1]. We propose to extend directory services' interfaces with a *proactive* mode by which clients can express their interest in (and be notified of) changes in the environment. These notification channels can be subsequently customized on a per-client basis through client-specified filters. Finally, in order to simplify the handling of client/server failures we adopt a leasing model for client registration to (and customization of) a notification channel. To validate our approach, we have designed and implemented the *Proactive Directory Service (PDS)*.

Our experimental results demonstrate that, contrary to common wisdom, the customization of notification through filter functions executed at directory servers or objects' owners does not necessarily translate into overloaded servers. Indeed, this customization can improve the performance of notification sources as the cost of executing the additional filter code is outweighed by the gains resulting from eliminating unnecessary communications (i.e., executions of protocol stacks).

[1] The problem, in the context of MDS, has been reported by Smith et al. in their *HPDC-9* paper *An evaluation of alternative designs for a grid infromation service*.

Proceedings of the 11 th IEEE International Symposium on High Performance Distributed Computing HPDC-11 2002 (HPDC'02)
1082-8907/02 $17.00 © 2002 **IEEE**