Lost in a Labyrinth of Workstations

Joann J. Ordille
joann@cs.wisc.edu

Barton P. Miller
bart@cs.wisc.edu

Computer Sciences Department
University of Wisconsin-Madison
1210 W. Dayton Street
Madison, Wisconsin 53706
(608) 262-1204

Millions of users will be lost in tomorrow's labyrinth of workstations. In fact, today's users are already lost. An internetwork of workstations makes myriad applications, services, and computing environments available to its users. Users, however, can access only resources that they can find, and finding a resource is a difficult, if not impossible, task.

To find a resource, you must name it. Today's naming labyrinth is comprised of islands of uniform naming in a maze of heterogeneity. First, it has a proliferation of name spaces, and no single way to name the variety of resources available to users. Is the resource available through the Domain Name System [10], X.500 [4], the CSNET Name Service [8], the Resource Location Protocol [1], Archie [13], Netfind [15], Profile [12], the Wide Area Information Servers [7], Univers [2], the Knowbot Information Service [5], the Network Library System [17], a company personnel database, or some other place? Second, the largest name services, like the Domain Name Service and X.500, are hierarchical services. Since most of us have trouble finding files in our own directory trees, it is no surprise that users cannot find things in other users' naming trees, trees that they did not create. Heterogeneity and confusion result from both the proliferation of services (and service interfaces), and the nonuniform tree structure in hierarchical services.

To the extent that users can deal with heterogeneity, they now roam here and there looking for what they need. The sheer size of the name space daunts all but a brave few. Both the number of workstations and the information organized by those workstations is growing at an alarming rate. In the last ten years, we have seen exponential growth in the number of hosts on the Internet [9]. This growth is accompanied by an exponential decline in the cost of disk space, and a corresponding increase in the amount of information available through hosts on the network [6]. Systems are proposed, like Melampus [3] and the Information Mesh [16], that would integrate most of the information available on our workstations into one world-wide database. How are we as users to survive this information glut? How are we to avoid being forever lost in a labyrinth of global proportions?

Only we, as designers of workstations and networks, can remove this labyrinth for

This work was supported in part by an AT&T Ph.D. Scholarship, National Science Foundation grants CCR-8815928 and CCR-9100968, Office of Naval Research grant N00014-89-J-1222, and a Digital Equipment Corporation External Research Grant.

users. Our first step is a commitment to provide descriptive name services. Descriptive name services allow users to discover resources by describing some attributes of the resources. The labyrinth disappears for users, because the descriptive, i.e. relational, query language is non-procedural. Users do not navigate the name space to find resources; the descriptive name service does. Users are freed from the need to understand the heterogeneous name services and complex structures in the name space.

By removing the labyrinth from the user's world, we cast it more deeply into our world of operating systems and networks. The second step in our solution is providing high performance descriptive name services that unravel the complexities of the name space. Many descriptive name services fail in this task. They require users to direct the search for resources by specifying target name services [5], target name space servers [12], or lower-level branches of the name space tree [4]. These services must ask the users for guidance, because they are cowed by today's naming labyrinth. It is difficult to process queries quickly in the distributed, heterogeneous and global relational database that our name space has become, but queries must be fast if users are ever to find anything in our internetworks. When confronted with the need for speed, these services abandon the user by saying: "If you can tell me where that name is, I can get it for you quickly."

When users tell descriptive name services where to look for resources, the lost lead the lost through the labyrinth. Another case is found in services that automate the steps a user would take in finding a resource. Archie locates files available on a network by making periodic exhaustive listings of files on hosts with anonymous `ftp` directories [13]. Netfind locates people by using information from USENET messages to direct `finger` queries to the most likely hosts for particular users on the Internet [15]. These services act like scouts in the labyrinth. They use hints in the environment to direct them to places where the needed resources are likely to be found. Scouts may get lucky and find the needed resource, or they may remain lost, like the users before them.

Communities tend to be friendly or hostile to scouts: friendly if they think the scout will only discover what they want to have publicized, and hostile if they think the scout with reveal that which is better left unnoticed. Typically each community has some things to reveal and some to keep hidden, so we need an alternative to scouts. We do not want our name service to depend on scouts trying all the doors in our systems.

We cannot rely on users or scouts to save the day. We must cooperate with the owners and administrators of naming data to unravel the labyrinth. *Those who control naming data must provide us with hints and instruction on how to find the data.* The national highway labyrinth is effectively negotiated by road signs and maps. Local service providers, like restaurants, post signs to help themselves be found. Local communities provide maps and guides to their resources that are incorporated into summaries for larger geographic areas. This is the kind of information we need from the controllers of naming data, and they, like service providers and communities on highways, are the most motivated to provide this kind of information. Our descriptive name service will succeed if we elicit, organize, and encourage the cooperation of the owners of naming data.

Two efforts are currently pursuing this kind of descriptive name service. The Networked Resource Discovery Project proposes to use signposts to locate resources [14]. This approach works when you are looking for one of many instances of a service. For example, it works when you are traveling down the highway looking for any restaurant and see a restaurant sign. Following signs does not work when you are looking for an individual item that is not nearby. For example, since there are no signs for Minneapolis on the interstate system in Chicago, signs do not help you travel from

Chicago to Minneapolis. You are reduced to recursive scouting of nearby places (Indiana, Wisconsin, Rockford, St. Louis) in the hopes of finding an appropriate sign. If you go to Wisconsin, you might find Minneapolis eventually. If you get to St. Paul, don't give up! If you go to Indiana, you might find Miami first.

Nomenclator is a system that combines maps with signs to get us from Chicago to Minneapolis without going through Miami [11]. Maps tell us where we are going and signposts direct us along the way. Early Nomenclator results encourage us that there are significant performance advantages to be had by using a combination of signs and maps to locate resources. If we pursue this approach to descriptive name services, we and our users will be experienced travelers in a interconnected system and never lost in a labyrinth.

REFERENCES

[1] M. Accetta, ''Resource Location Protocol,'' Request for Comments 887, DDN Network Information Center, SRI International, Menlo Park, CA (December 1983).

[2] M. Bowman, L. Peterson, and A. Yeatts, ''Univers: An Attribute-Based Name Server,'' *Software – Practice Experience* **20**(4), pp. 403-424 (April 1990).

[3] L. -F. Cabrera, L. Haas, J. Richardson, P. Schwartz, and J. Stamos, ''The Melampus Project: Toward An Omniscient Computing System,'' Research Report RJ 7515 (70030), IBM Research Division, San Jose (June 1990).

[4] International Telegraph and Telephone Consultative Committee (CCITT), ''The Directory,'' Recommendations X.500, X.501, X.509, X.511, X.518-X.521 (1988).

[5] R. E. Droms, ''Access to Heterogeneous Directory Services,'' *Ninth Joint Conference of IEEE Computer and Communications Societies (INFOCOMM),* San Francisco, pp. 1054-1061 (June 1990).

[6] J. L. Hennessy and D. A. Patterson, *Computer Architecture: A Quantitative Approach,* Morgan Kaufmann Publishers, Palo Alto (1990).

[7] Brewster Kahle, ''Wide Area Information Server Concepts,'' Technical Report TMC-202, Thinking Machines Corporation, Cambridge, MA (August, 1990).

[8] L. Landweber, M. Litzkow, D. Neuhengen, and M. Solomon, ''Architecture of the CSNET Name Server,'' *SIGCOMM Symposium on Communications Architectures and Protocols*, Austin, pp. 146-149 (March, 1983).

[9] M. Lottor, ''Internet Growth (1981-1991),'' Request for Comments 1296, DDN Network Information Center, SRI International, Menlo Park, CA (January 1992).

[10] P. V. Mockapetris, ''Domain Names - Concepts and Facilities,'' Request for Comments 1034, DDN Network Information Center, SRI International, Menlo Park, CA (November 1987).

[11] J. J. Ordille and B. P. Miller, ''Nomenclator Descriptive Query Optimization in Large X.500 Environments,'' *ACM SIGCOMM Symposium on Communications Architectures and Protocols,* Zurich, pp. 185-196 (September, 1991).

[12] L. L. Peterson, ''The Profile Naming Service,'' *ACM Transactions on Computer Systems* **6**(4), pp. 341-364 (November 1988).

[13] R. P. C. Rodgers, N. H. F. Beebe, and A. Emtage, ‘‘Archie – Internet Archie Server Listing Service,’’ Manual Page, UCSF School of Pharmacy, San Francisco, CA (1991).

[14] M. F. Schwartz, ‘‘The Networked Resource Discovery Project,’’ *IFIP XI World Congress*, San Francisco, pp. 827-832 (August 1989).

[15] M. F. Schwartz and P. G. Tsirigotis, ‘‘Experience with a Semantically Cognizant Internet White Pages Directory Tool,’’ *Internetworking: Research and Experience* **2**(1), pp. 23-50 (March 1991).

[16] K. R. Sollins, ‘‘Supporting the Information Mesh,’’ *Workshop on Workstation Operating Systems*, Miami(April, 1992).

[17] University of Wisconsin, ‘‘Network Library System (NLS),’’ *Internet Resource Guide*, BBN Systems and Technologies Corporation, Cambridge, MA, p. 4.2.14 (August 1989).