

Guest Editorial

Peter Deutsch Bunyip Information Systems

*T*here's something happening on the Internet. The way we use it and the way we think about it seem to be changing in a fundamental way.

Perhaps this statement is a little strong, but over the past two years there has certainly been a number of changes in the way we perceive and make use of the net. Of course, the number of users continues to climb at a dizzying pace and the sheer volume of traffic continues to climb along with the population, but there's more to it than that. As someone who has been involved in developing new services and educating people on how to use them, it seems to me that there has been a basic shift in the way the net is perceived by its users and advocates.

I'm not sure that we can point to any one event or activity as the catalyst for this change, but certainly it is influenced in part by the huge numbers of new users who are coming on-line and the differing demands they are making of the net.

The Internet is no longer just an interesting experimental testbed for a few researchers developing basic communication protocols, nor is it any longer solely a playground for a few lucky individuals at selected research and educational institutions. The net is now a daily tool for hundreds of thousands of people who couldn't tell an IP packet from a burst of tty line noise. These people are using the Internet for a whole range of new and different tasks.

Changes are happening not just at the level of individual users. We are seeing more research and development activity aimed at distributed network applications and the technologies that make them possible.

We are seeing new and different techniques, tools and projects taking hold that try to make use of this unique new computing environment.

Some of the tools and techniques are fairly traditional, migrated onto a new medium for connectivity, but there are also a range of newer applications and approaches that point the way to good things to come. To recall the old proverb, "we are living in interesting times."

Given all this activity, the time seemed right for a sampling of the state of the art for readers of this journal. I was pleased to have been asked by Mike O'Dell to take on the task of coordinating submissions for this special issue on tools and techniques for Internet user services.

When Mike asked me to take on this project, we spent some time discussing the type of papers we hoped to attract. There are any number of approaches we could have taken. For example, we could have concentrated on algorithms and basic research that might influence what will be deployed in the near future.

Alternatively we could have tried for a "catalogue of projects" approach, in which authors of successfully deployed projects were asked to describe their work. There are certainly any number of interesting projects out there now, each taking its own particular bite out of the pie and each certainly deserving greater exposure in a community such as this.

Neither approach seemed quite right. The excitement I detect as a player in this particular sandbox right now extends across a broad spectrum of research and development activities. Sure, there are people who are demonstrating the benefits of new approaches in distributed database maintenance and updating algorithms, but there are also people building and deploying truly useful projects that show the way for the next generation of Internet directory tools. I wanted to capture the range and depth of this activity (and if possible, the wonderful sense of community) we find in this area right now.

After some discussion with a number of people working in this field, we have instead elected to opt for a "spread spectrum" approach in which the selection of papers span the range of theory and practice encountered in bringing new services to the net. We've even thrown in a survey paper that we hope will serve as a signpost for those wishing to find out where to go for more information about the various projects in the field right now.

Our aim has not been to try to provide a paper on every possible project under development on the Internet right now. That would obviously take up far more space than we have available. It is enough if we can provide enough evidence that good things are happening that we persuade you to check it out.

One of our major objectives has been to provide some indication of the types of technologies that we expect to make a contribution to future Internet-wide services. In making our choices, we had to rule out some good papers that might, for example, illustrate techniques that functioned well in a LAN environment but that collapse when deployed in an environment of hundreds of thousands or even millions of machines.

It turns out that when building tools for the Internet, some things that people thought would be hard turn out to be relatively easy (like indexing all of the currently available anonymous FTP archives). On the other hand, some things that look like they should not be all that difficult to do are providing quite a challenge (just try finding somebody's email address, or maybe the name and location of a new service you might have heard about. It's still a non-trivial exercise, despite any number of attempts at the problem.)

What we have sought to do is to bring you some of the techniques that show promise in addressing the kinds of problems we are encountering as we work to bring up services in this new medium.

Our first paper is by Richard Golding of U.C. Santa Cruz. Richard presents an architecture for maintaining replicated collections of information in a large scale distributed computing environment. Richard's work with weak-consistency algorithms holds great promise for those wishing to maintain multiple replicas of information in our new environment.

The second paper, by Clifford Neuman of U.S.C.'s Information Sciences Institute, covers the design of the Prospero distributed file system. Prospero is a successful attempt to design a file system that can address the problems of scaling, naming and performance encountered when used across extremely large distributed networks. An extremely flexible model, Prospero has found use in a number of other projects and is intended to provide one component of the Virtual System Model for distributed computing.

The third paper, by Peter Danzig, Shih-Hao Li and Katia Obraczka of U.S.C., describes the architecture for "Indie," a distributed indexing system intended to provide support for future resource discovery and access tools. Indie is designed to provide automated creation and maintenance of distributed "information brokers" which can provide an Internet directory of services. To date, the resource discovery step has been perhaps the least understood, and least successfully addressed aspect of providing Internet services. Such work shows that we are now starting to find ways to approach this problem.

We wrap up with a survey paper prepared by Mike Schwartz, Alan Emtage, Cliff Neuman, and Brewster Kahle that examines the current state of the art in Internet information discovery. This paper provides a brief overview of many of the currently deployed information systems, and observations about the implications of their underlying mechanisms. I hope this paper will serve as a useful primer for newcomers to the field, as well as a point of departure for further explorations into relationships between information systems.

As I mentioned, this issue of *Computer Systems* is intended only to provide a snapshot of work in progress. Given the current level of activity, and the continued growth in demand for users, we can expect that the goal of practical, useful Internet services will provide a source of inspiration for some time to come. We are truly working in interesting times.