Trust in the Process

Another lesson I was fortunate to learn early in my career is that it is possible to both be successful and have fun at the same time. The Archie project demanded a tremendous effort over a long period of time, but we lightened the load with inflatable plastic baseball bats, stuffed Tasmanian Devil dolls, and massive amounts of chocolate. Software design and development is a creative process and as such benefits from a little structured freedom.

Perhaps the most important lesson I learned from my experience with Archie was that I had to trust the Darwinian, iterative development process, even when I could not see where it was taking us. In a world where you do not always have the time for traditional engineering and development processes, it is perhaps encouraging to realize that others have made similar voyages and not foundered upon the shoals.

As I recall saying a number of times during the project, "We don't know where we're going, but we're making great time. . ."—perhaps a fitting epitaph for any Internet-based development project. ■

Peter Deutsch is a technical consultant specializing in Internet applications. He was founding co-chair of the URI Working Group at the IETF, served on the Advisory Board of the Canadian Institute for Scientific and Technical Information at the National Research Council of Canada, and was a developer of WHOIS++. Deutsch has an MSc in computer science and a BSc in mathematics and computer science from McGill University, as well as associate degrees in French and electronic technologies from Chaffee College.



Larry Smarr, *NCSA, University of Illinois*

Infrastructure for Science Portals

A major driver of Internet growth has been the United States' decades-long national policy of funding basic research and infrastructure prototyping in its universities. This unique national system has enabled universities to experiment with capabilities years ahead of their mass commercial distribution. From Arpanet to the NSFnet, this far-reaching partnership between government, universities, and industry planted the seeds that grew into the modern cyberworld. More recently, NCSA Mosaic and NCSA's Web server software catalyzed the development of the World Wide Web.

This "explore the future" methodology has been institutionalized in the NSF's Partnerships for Advanced Computational Infrastructure program. Each PACI partnership—NPACI (http://www.npaci.edu), led by the University of California, San Diego, and the Alliance, led by the University of Illinois at Urbana-Champaign (http://www.ncsa.uiuc.edu)—brings together dozens of leading researchers at some 75 institutions to create the services necessary to prototype a seamless, integrated computational and collaborative environment known as the Grid.[1]

Collaboration through the PACI Grid

Grid functions can be bisected into two logical grids: the *computational grid* and the *access grid*. Through the computational grid, scientists will be able to access virtually unlimited computing and distributed data resources, while the access grid will provide a group collaboration environment.

PACI researchers are developing sophisticated middleware (such as Globus and Legion) designed to render complex Grid environments manageable while enabling application-level control of performance-critical features. Such middleware will make it possible for Grid-enabled applications to discover available network, computer, and storage resources; to obtain advance reservations for those resources required for a particular purpose; and to handle authentication, policy, resource allocation, and communication issues. It will also maintain a desired performance level across a multisite complex and monitor application performance during execution.

Through a Web browser, users will be able to view all the Grid resources and services in the PACI Virtual Machine Room (VMR), a simulated single physical site. The VMR requires development and deployment of common schedulers, security mech-

anisms, and so on, to all sites, as well as tuning of individual sites to achieve a uniform interoperability. This allows PACI to support the computing needs of national academic researchers by providing peer-reviewed access to a diverse collection of high-performance computing technologies.

Eventually, strong market forces will push scalable computing toward a commodity building-block approach. Already we have seen supercomputer processors move from expensive custom vector processors to RISC workstation processors to Intel PC processors. Operating systems also moved from supercomputer-only to Unix and now are heading toward NT and Linux. Interconnects will have to follow a similar logic, merging with Internet and telecom standards like Gigabit Ethernet and OC-192.

Eventually, strong market forces will push scalable computing toward a commodity building-block approach.

As a result, computing will become a distributed function with thousands of processors in national centers, hundreds at universities, dozens in laboratories, and one to four on the desktop. Early visionary experiments like the SETI@home project (http://setiathome.ssl.berkeley.edu), prime number (http://www.entropia.com), and public key breaking (http://www.distributed.net) projects are pointing toward a future in which the Net itself, with its tens of millions of computers, can be electronically reconfigured by end users into a scalable computing fabric.

The access grid is the ensemble of resources used to support group-to-group human interaction across the Grid. The access grid expands the concept of remote group access to a digital framework that includes supercomputers, scientific instruments, and remote visualization. It is designed to support large-scale distributed meetings, collaborative sessions, seminars, lectures, tutorials, and training, enabling both formal and informal group interactions. Entry points may be desktops or large-format multimedia displays integrated with intelligent or active meeting rooms.

Grid users may require one or both grid types for a given application. For example, virtual-engineering prototyping requires access grid virtual-reality capabilities for distributed collaboration on a given design, and also requires significant computing and data resources to perform complex simulations with real-time updates. The computational and access grids rely heavily on the availability of high-speed networks, such as NSF's vBNS and Internet2's Abilene.

Collaborative Workbenches

The Web has enabled a fundamental shift in how science is conducted, allowing researchers to communicate their results to a worldwide audience, and students of all levels to access the latest tools and information, collaborate with peer groups, and interact with experts. PACI scientists have extended the Web from a multimedia document retrieval function to a tightly coupled framework tying together a suite of computational applications with the underlying distributed data. This software architectural approach creates Web-based science workbenches, Internet sites that give users access to advanced tools for data mining, analysis and visualization, tutorials, and problem solving using remote supercomputers. The Biology WorkBench (http://biology.ncsa.uiuc.edu) is used by more than 10,000 researchers in 25 countries in over 150,000 computing sessions a month. Similar prototypes are being developed for chemistry, chemical engineering, nanomaterials and nanodevices, and cosmology.

Looking forward to the widespread use of the Grid, the PACI program is designing a second-generation workbench architecture called *Science Portals*, which will avoid the limitations of the current workbench designs. The idea is to create a distributed object architecture for Grid-based, scientific problem solving.

The portal concept, which has emerged as a powerful idea in the ".com" private sector, is the starting point for our vision for the next generation of workbenches. A key feature of portals is that users can customize not only the content, but also the presentation (page layout, level of detail) and a set of tools and services (e-mail, chat, stock quotes, and so on).

The Alliance's Science Portals are to be based on emerging standards for distributed computing component systems, such as CORBA, and Java technologies like Jini and Enterprise JavaBeans. In addition, they will exploit the recent revolution in XML metadata standards. The major departure from the emerging commercial portal frameworks will be in our integration of advanced Grid services, very large scale data analysis, distributed components that run on high-end parallel systems, and advanced visualization and collaboration services.

One of researchers' fastest growing needs is an enhanced ability to discover, analyze, mine, and publish data and metadata from catalogs, archives, and instrument sources. To explore arbitrary archives of distributed data from their desktops, users need a new generation of analysis tools that exploit Grid resources to combine remote filtering with wide-area query searches. PACI has worked extensively with industry to develop object-oriented toolkits for knowledge discovery and visual analysis of information gathered from multiple databases and data warehouses.

Electronic Communities

From desktop, laptop, or palmtop, users should be able to draw upon the Grid's computational resources. The user's only task should be to specify the computational problem by using a portal-based, visual-programming tool in which the icons represent remote supercomputing computations and data sources. A collaborative, electronic scientific notebook would automatically journal the research experiments and results.

Science has always been a collaborative activity, with the scientific literature and disciplinary meetings serving as the "glue" that provided for asynchronous communication of ideas. Early technological improvements such as telephone and fax greatly shortened the timescale for communication and accelerated the rate of scientific discovery. In the past 15 years, the Internet added e-mail and then the Web to these collaborative mediums. With the Internet's next generation will come the services of the Grid and the formation of electronic communities through Science Portals. These will add the real-time tool access—to computing, visualization, distributed data, scientific instruments, and synchronous collaboration—to the asynchronous technologies of e-mail and the Web. ■

I would like to thank the members of the Alliance for help in writing this article.

REFERENCE

1. I. Foster and C. Kesselman, eds., *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann, San Francisco, 1998.

Larry Smarr is director of the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign, and director of the National Computational Science Alliance. Smarr has a PhD from the University of Texas at Austin. He is a member of the National Academy of Engineering and a Fellow of the American Physical Society and the American Academy of Arts and Sciences.

kc claffy, *UC San Diego Supercomputer Center*

Measuring the Internet

> The so-called science of poll-taking is not a science at all but a mere necromancy. People are unpredictable by nature, and although you can take a nation's pulse, you can't be sure that the nation hasn't just run up a flight of stairs.
>
> —E.B. White, *The New Yorker*, Nov. 1948

Internet traffic behavior has been resistant to modeling. The reasons derive from the Internet's evolution as a composition of independently developed and deployed (and by no means synergistic) protocols, technologies, and core applications. Moreover, this evolution, though "punctuated" by new technologies, has experienced no equilibrium thus far.

The state of the art, or lack thereof, in high-speed measurement is neither surprising nor profound. It is a natural consequence of the economic imperatives in the current industry, where empirically grounded research in wide-area Internet modeling has been an obvious casualty. Specifically, the engineering know-how required to develop advanced measurement technologies, whether in software or hardware, is essentially the same skill set required to develop advanced routing and switching capabilities. Since the latter draw far greater interest, and profit, from the marketplace, it is where the industry allocates engineering talent.

A common complaint about traffic measurement studies is that they do not sustain relevance in this environment where traffic, technology, and topology change faster than we can measure them. Moreover, the proliferation of media and protocols make the acquisition of traffic data almost prohib-