The WiderNet Project

eGranary Digital Library

Final Report for Leland Initiative/AED Grant

This support from the Leland Initiative and AED, from 4/1/2003 to 8/30/2005, has been critical to the development of the eGranary Digital Library and the underlying premise that local caching of educational materials on low cost storage devices has a role to play in the digital strategy of every institution with inadequate Internet access.

In the two years of this program we have grown the collection from 500,000 to 2,000,000 documents, in this way developing a whole new social contract and construct through which a significant number of authors and publishers are able to deliver their digitized knowledge to the seven out of eight people in the world who do not have access to the Internet.

We have demonstrated that a sizable potential subscriber base exists by growing the eGranary Digital Library to over 60 installations with nearly 90% paying for the service.

We have developed a mechanism for collecting content, organizing content, and updating content in existing installations, making the process of expanding the eGranary Digital Library, or creating similar digital libraries, a concrete and repeatable reality.

More importantly, we have drawn the attention of a considerable number of organizations and individuals to the concept of asynchronous content delivery and bandwidth conservation.

Today the eGranary Digital Library is a very real and welcome resource for scores of subscribers and is being seriously considered for global application by the likes of the World Health Organization, UNESCO, the U.S. State Department, national governments, and library associations. It has attracted another 14 months of funding from the Hewlett Foundation, additional grants of $25,000 to build strategic and business plans towards sustainability, and individual gifts and grants totaling $40,000.

Without this initial support from the Leland Initiative, much of this would be a pipe dream – scrawlings on a white board and meager efforts by volunteers and overworked professors. Instead, today the eGranary Digital Library is a polished product with great potential to serve the information poor with a first-class knowledge access experience.

Basic Goals

We started this project with seven goals and 20 indicators. We met or exceeded most, although some of our goals changed as we grew more aware of the needs and capacities of the technology and our subscribers.

We are confident and pleased with the results of this grant period, although the toughest test is yet to come: sustainability.

The goals and outcomes are indicated in the following chart, while more detail on each goal and indicator is contained in the accompanying "Summary Report on Goals."

Goal	Outcome
Bring together multiple stakeholders to plan the project	Met
Produce and distribute an expanded demonstration library (at least two million documents)	Met
Digitize a small set of African scholarly works	Partially Met
Demonstrate that contemporary satellite broadcast techniques can be modified to work at most African universities.	Partially Met Possibly Exceeded
Develop and demonstrate the satellite broadcast protocol and customized OS to provide USENET, email, and proxy services.	Met
Build a coalition of stakeholders to further the development of new content, capacity building, and new installations.	Met
Increase the number of installed libraries to 60 (with half or more paying for their own equipment.).	Met

Alliance Building

This grant period was marked by significant outreach to potential alliance partners. This has not been an easy task. As with many new technologies, people do not generally warm up to the idea on the first hearing. In many cases, our initial contacts were met with confusion, if not derision.

Librarians, in general very protective of their patrons and concerned about a host of qualitative issues surrounding information equity and quality, have proved to be the most skeptical.

On the other end of the spectrum, a bevy of western technology enthusiasts have become champions of the eGranary Digital Library. While this can be a good thing, too much glee over technology can be off-putting for those who are less familiar with IT and make nervous those who are wary of marketing pitches by western shills with overblown solutions.

So we walk a fine line between "selling" the eGranary Digital Library and recruiting partners and making them stakeholders.

One group, the Virtual Hospital, took a year and a half to decide to include its content, even though the group's main aim is to make high quality medical resources available for free. It simply took that long for the members to be consulted and become aware of the need and the niche that the eGranary Digital Library occupies.

In another situation, we found that the enthusiasm of the African Studies Association was dimmed by the prospect that contributing the 25-year archive of its journals would impact their legal and financial agreements with JStor, the U.S.' largest collective academic journal archiving project. JStor

representatives had been unresponsive to requests for clarification and, after being put on the spot by members of the African Studies Librarians group at a recent annual conference, turned to personal intimations and declared they would not participate in eGranary Digital Library efforts.

So, while we have made significant progress in introducing the eGranary Digital Library concept and attracting a healthy amount of attention to the potential of this technology, we still have a lot of work to do to educate and motivate potential partners.

That being said, a considerable amount of progress has been made in our promotional efforts. We have…

- presented the eGranary Digital Library to new Fulbright Scholars heading to Africa for three years in a row
- presented the concept to the staff of the Carnegie Corporation (who tried to replicate the project, only marginally succeeded, and then put out a paper highlighting the potentials and pitfalls of such a project)
  http://www.carnegie.org/sub/pubs/virtlibreport-low.html
- started negotiations with the Nigerian National University Commission to implement a nation-wide program for creating, installing, and supporting eGranary Digital Libraries
- made special presentations to hundreds of librarians at the Standing Conference of Eastern, Central and Southern Africa Library and Information Associations (SCECSAL)
- delivered papers on the eGranary Digital Library at the African Studies Association annual conference for two years in a row
- met in Geneva, Switzerland with program officers of the World Health Organization (who are putting together a $20 million proposal to use the eGranary Digital Library as a platform to deliver training to heath care workers across Africa)
- demonstrated the eGranary Digital Library to education program officers at UNESCO in Paris
- demonstrated the eGranary Digital Library in presentations to over 100 international educators at MIT's Learning International Networks Consortium (LINC)
- had the eGranary Digital Library described in dozens of newspaper and journal articles, as well as numerous Web sites and blogs. (See appendix.)

Our efforts to spread the technology to other areas of Africa have been boosted dramatically by the University of Iowa library's chief Africanist, Dr. Edward Miner. He has become our major champion in East Africa, introducing our program and providing demonstrations to dozens of colleagues in Uganda, Tanzania, and Kenya while opening the door for WiderNet Project staff to make presentations and do installations on half a dozen campuses.

Long-term discussions with Computers4Africa led to their main technologist and master teacher, Herbert Busiku, spending three weeks in technical training in Iowa City. He is now installing eGranary Digital Libraries in Uganda secondary schools, using donated computers to make servers.

We have crafted a strategic partnership with the World Computer Exchange, whose board recently decided that one third of its upcoming shipments to developing countries should include eGranary Digital Libraries.

We have entered into discussions with the U.S. State Department's American Corners group, providing demonstrations to dozens of program participants in Nigeria, as well as program officers in Kenya, Tanzania, and Uganda.

While we still await breakthroughs with these promotional efforts, we feel that we have done the grunt work required to lay a good foundation and build trust.

Finally, we've done significant work with our colleagues and partners to identify reasonable fee structures that we can adopt to give the project financial security. We have determined that we need to develop a steady income stream from all subscribers and propose to do this in the form of alliance memberships / update service subscription fees. In conversations with nearly 100 partners, we've identified that an annual fee of between $100 and $200 would be perceived as fair. We would then charge more to those who would expect more frequent updates and direct technical support from the WiderNet Project.

Project Research into the Bandwidth Conundrum

Nearly all of our potential partners need time and information to consider the nature of the bandwidth conundrum that stymies Internet use in the developing world.

Those who do not have bandwidth sometimes prefer to focus all their energies on buying bandwidth, not understanding that the tiny connection they might be able to afford will deliver limited functionality.

Those who have a miniscule amount of bandwidth, even if only being provided to a narrow group of professors or paying students, can be so proud of this achievement that they dismiss the need for augmentation.

Those in the West, who have in a few short years become habituated to the idea of plentiful bandwidth, have a hard time conceiving of NOT being connected, let alone the futility of sharing a 128k Internet connection with 100 users.

Then, many people still operate under the impression that ample bandwidth will be universally available in a few short years. For someone unfamiliar with the perniciousness of poverty and underdevelopment in the Third World, it is hard to conceive of the generations of technical, economic, civil, and social development that under girds the West's Internet facility.

So we have been doing research into the speed, reliability, cost, and usability of Internet connections.

We have developed a bandwidth speed test which appears as a link on the project's home page (http://www.widernet.org). When this link is selected, a file is downloaded to the requesting computer and the process is timed. The results are stored in a database and patrons are able to see their results over time. This test reveals the download speeds that users experience after the Internet link has been shared to their community, a more realistic assessment than simply measuring how much bandwidth the institution has purchased. We have done thousands of such tests from hundreds of clients in dozens of countries. Typical test speeds for top U.S. universities range between 7Mbps and 17Mbps. Typical speeds at the few connected African universities range from 14Kbps to 112Kbps – 150-1,200 times slower than their U.S. counterparts. (Interestingly, we've timed files opening from the eGranary at 68Mbps in several countries. This means African students are opening documents three to nine times faster than students at Big Ten universities in the U.S.!)

We have developed a program to test the reliability of email and Web servers. We ping the machines every 30 minutes to see whether they respond and store the results in a database. This testing reveals that most schools in Africa have difficulty providing six hours of connectivity a day, while many experience Internet outages lasting from days to weeks.

We have done some network "sniffing" to evaluate what patrons do on Internet connections. We have found that African Internet browsing essentially mirrors American Internet browsing: the bulk of it is recreational. And it may be that the Internet connection fuels Brain Drain: the top category of search engine requests are those for school, work, and emigration opportunities elsewhere. (This is one area in which we do not have enough samples for our assessments to be accurate, but it deserves more attention.)

Finally, we have researched institutional Internet connectivity costs, demonstrating that U.S. institutions enjoy ample bandwidth for the cost of one or two professor salaries; while some of our sister institutions in Africa spend the equivalent of 20+ full-time professor's salaries for a meager 1mbit connection.

We use the results of this research to educate our potential partners and enable them to make better distinctions around issues of Internet connectivity and the use of local resources and services to pursue their core missions more efficiently and effectively.

The eGranary Digital Library Appliance

In our first dozen or so installations, we found we expended a considerable amount of time coaching our new subscribers through the process of setting up and configuring their servers. At the same time, we identified the potential of redesigning our file system so that we could serve files through a customized proxy, making the eGranary Digital Library look and act like the Internet itself. (So a user types in a URL and gets the page just as though it came from the original server.) Finally, we wanted to add a search engine to the eGranary Digital Library, but could find no client-based solutions that would work with such a large amount of content.

We needed a plug-and-play server.

We addressed all three problems by prototyping a server in our labs with a limited set of content. This worked better than imagined, so we set about the six month process of rearranging all our data, re-scraping Web sites, rewriting our content management system, and developing a customized proxy server based on Apache.

We developed both Windows and Linux versions of the server, switched to 250GB hard drives, added the Lucene indexing service, slapped on the new eGranary Update service and delivered the first eGranary Digital Library Appliance in early 2005.

This is a plug-and-play server that provides all the services of the eGranary Digital Library to potentially thousands of clients within the intranet. It includes space for local content (which can be included in the search engine's index.) And it acts as a proxy to the Internet – if a connection exists – so that patron requests are transparently forwarded to the Internet if the requested document does not exist in the eGranary Digital Library.

Added benefits of the eGranary Appliance:

- Links from one site in the eGranary that point to another site in the eGranary work, improving the patron's experience

- Provides a "singing for the same songbook" experience: patrons can exchange URLs and follow email links whether using the eGranary or the Internet
- Links on portals or Web sites designed by patrons work both on the eGranary and the Internet
- It has built-in DHCP, DNS, and other network services, requiring little ICT experience to set up and manage the system
- Runs well on Pentium IIs and IIIs
- Keeps extensive logs of usage and errors

While not conceived in our original proposal to the Leland Initiative, the development of the eGranary Appliance has dramatically improved the eGranary Digital Library. It required a considerable investment of time and skill, but the payoffs have been tremendous.

Shortly after the grant period ended, we reworked the proxy and search engine so that it could be run on a stand-alone client. That gives us the capacity to deliver an external USB hard drive – the size of a paperback novel – to any computer on the planet and within minutes turn it into a knowledge access station with 3 million documents instantly at the patron's fingertips. It looks and acts just like the Internet, only thousands of times faster than a local Internet connection. (A 220MB video file opens in 5 seconds!)

Areas of Ongoing Concern

Building Institutional Capacity

The most critical non-technical challenge we face is our subscriber's institutional capacity to host and nurture the eGranary Digital Library.

With over 60 installations in the field, it is safe to say that only half of the host institutions are implementing the eGranary Digital Library well. In one case, the ICT staff had installed the eGranary Digital Library and reported using it everyday (to the point of addiction) yet further investigation showed they had done nothing to promote the resource to others on campus. We have since made a shift to focusing on librarians as campus champions, but even then the track record has been mixed. At one library, the librarian had restricted the use of the eGranary Digital Library to herself and her staff, inviting staff and students to submit their search requests in writing and then returning to their patrons a printout of the discovered resources. At another institution the eGranary server has been physically moved from the library to the computer center several times as the staff try to work out how to implement it.

With its high speed access, the Granary Digital Library is a powerful tool for teaching information literacy. It gives patrons the capacity to "browse at the speed of thought" through hundreds of documents an hour, comparing and contrasting, and rapidly building bibliographies. The Lucene-powered search engine gives patrons a broad spectrum of options for crafting queries that locate specific information in little time. And the ability to create and store customized local portals makes it possible for instructors and reference librarians to pull disparate resources together into a cohesive collection for their students and patrons. Yet few librarians have been trained to this skill level.

Promotion of technical innovations requires specific skills and aptitudes. Few librarians are trained in this area. Occasionally they put up a poster. In one outstanding effort, the librarians at Ahmadu Bello University in Zaria, Nigeria, invited hundreds of staff and students to a 30-computer training lab for

extensive eGranary Digital Library training. But this kind of effort is rare. As well, it need not fall to librarians to make changes in how instructors teach. In a better, well-funded world, we should focus on building integrated approaches on campuses that involve department administrators, instructors, master teachers, librarians, and technologists to devise new ways of teaching with technology on campus.

Finally, for most librarians receiving the eGranary Digital Library it is 1996 all over again. In 1996 in the U.S., few librarians had expertise with the use of the new-fangled Internet. In fact, many of them vociferously opposed the idea for a whole host of qualitative concerns. Librarians using the eGranary Digital Library for the first time have to deal with the same plethora of personal, professional, and even political challenges that their colleagues faced when the Internet was introduced in to their lives – not the least of which is that their young students seem to grasp the technology much more quickly than they.

Full Partnership Rare

Despite our best efforts to date, we have few subscribers that truly see themselves as full partners.

While we actively solicit feedback and ideas, we're more likely to hear third-hand complaints that a particular subject area is not covered well rather than direct requests for new content.

We install eGranary Digital Libraries and ask patrons to notify us when something isn't working, but few do.

We ask subscribers to provide us with server logs and semi-annual reports; very few do.

In fact, as we conduct field demonstrations and trainings, we emphasize that ours is a social enterprise requiring everyone to do their fair share to build the collection, support the development, and promote it to their colleagues. Yet most people, in the U.S. and abroad, struggle with the concept. They prefer to see the eGranary Digital Library as a product.

This attitude is so pervasive that our executive committee has more than once revisited the idea of turning some portion of the project into a commercial enterprise.

Still, we remain committed to "serving the information poor without making them too much poorer" and are exploring other ways to develop a more robust community around the eGranary Digital Library concept.

Improving Quality Control

An early strategy and assumption was that we could pour new content into the eGranary Digital Library as quickly as possible and then count on our subscribers to let us know when one of the tens of millions of links wasn't working. As mentioned above, we get little such feedback.

For awhile we tried to have volunteers and student staff poke their way through our copied Web sites to detect errors in the mirror. However, this proved too cumbersome as many sites can have tens of thousands of links and may have hundreds of levels.

As this grant was drawing to a close we launched into a quality control initiative, setting up systems with computers to test links and record results. By comparing our error rate against the original Web site, we can establish an acceptable margin of error.

As we start getting more logs from subscribers, we'll have another tool in our quest for better quality. Server logs not only record which resources are being used, but also those requests which generated errors. The logs will also reveal requests to Internet sites not included in the eGranary Digital Library, giving us real feedback on user's interests and desired content.

Conclusions

We have a tiger by the tail.

In demonstrating the potential of distributing digital content without the Internet, we have created a lot of excitement and enthusiasm, changing people's notions of knowledge sharing and opening doors that had previously gone unnoticed.

This idea is growing and will soon become a "movement."

Today's eGranary Digital Library represents the collective efforts and contributions of hundreds of authors, publishers, programmers, librarians, instructors, and students around the globe. We have clearly demonstrated that the eGranary Digital Library can provide lightning-fast access to educational materials -- video, audio, books, journals, Web sites -- even where no Internet access exists. By removing the barriers imposed by inadequate infrastructure and costly connectivity, the eGranary makes it possible to put immense libraries into the hands of the smallest schools and increase access to information everywhere.

Now our goal is to grow the application of this "Internet in a Box" technology to thousands of schools, clinics, libraries, and homes across the developing world. We envision smaller, vertical eGranaries -- focused on education, or health, or civil society -- being distributed on CD-ROMs, flash drives, on the hard drives of laptops, and even cell phones with adequate memory.

We have the Leland Initiative and, now, the Hewlett Foundation to thank for providing the basic sustenance to bring this concept to fruition. While we still require funding to keep our doors open and the ideas growing, we can now identify a clear potential for eventual self-sustainability.

On behalf of the staff of the WiderNet Project and our many stakeholders who have helped to shape and grow the eGranary Digital Library, we extend our warmest thanks to the friends/stakeholders at the Leland Initiative for underwriting this project at this very critical juncture.