DuraSpace's Solution Communities:
Marshalling the Resources for Open-source Development

In his opening plenary address to the Open Repositories Conference, James Hilton made the statement that "open-source software is free like a puppy." This statement succinctly summarizes the need for institutions that benefit from freely available software to get involved in its ongoing development, that an investment of resources is always required. Anyone who has worked with information technology in libraries, museums and archives knows that compared with the total cost of buying vendor software and making it actually do the desired job, there can be a great deal of room to save money while making a significant investment of resources in the process.

The popular mythology that has developed around the open-source movement, fed both by wishful thinking and media hype, would have it that the software is developed by a bands of fervently engaged software developers, working in their spare time, magically making it all happen. Nothing could be further from the truth. Interesting prototypes, experimental systems and new directions often do emerge as a result of ad hoc and creative collaborations. Mature, well-tested production-ready open source software and best practices, however, are methodically developed and tracked by professionals who are also stakeholders in its ongoing success.

Viable open-source software depends on communities of developers who are deeply engaged, whose time is usually provided by companies and institutions that need the software and value their stake in guiding its design and development. Successful open source project personnel have demonstrated that they can start-up and organize processes while developing governance models and collaborative infrastructures for managing the software and their communications. However, keeping them going usually depends how well the developers can convince their managers of the value.

DuraSpace's Solution Communities program is aimed at directly engaging the technically aware managers and technical practitioners who can make resources available for the software development process. The goal is to bring communities of interested people together around specific use cases that are important to their organizations, to share information and get to know each other, in order to develop the level of trust that creates the conditions that makes collaboration possible.

Rather than take a top-down approach to organizing, DuraSpace has positioned itself as the catalyst for self-organizing communities. Taking lessons from others who are applying research about emergence in complex systems and "small world" theory (AKA, "the six degrees of Kevin Bacon"), the effort builds off the

basic idea that communities of interest develop around knowledge bases and good communication infrastructure.

Building a community around a particular use case depends on starting out with a few people who are willing to get the knowledgebase started and to start hosting events where people who are interested in the subject can connect with each other. This organizing group must be made up of people who see the effort as central to the needs of the organization they represent, and be able to devote some time to getting it off the ground.

Three roles have been identified that need to be taken on by members of the organizing group. These roles may be initially served by one person, but can just as easily be held down by a group. Note that these roles are as much about taking on the responsibility for getting something done as about doing it.

The first, and most important, role is the “steward.” More of a moderator than a director, the steward is responsible establishing and maintaining the vision for the community, and for taking responsibility for making sure that it has what it needs to operate. This can mean setting up times for the organizing group to have phone calls, organizing birds-of-feather sessions at relevant conferences, etc.

The second role is the “knowledgebase gardener.” This is the person or group who takes care of the knowledgebase by initially setting up the organizing structure and starting to add information to it. The metaphor of gardening is an apt one for this role. Encouraging the growth of the site, weeding and transplanting content to keep it healthy and useful are all involved.

The third role is the ”evangelist” whose job it to find and recruit participants who would be useful to the community. If the beginning this is a very important role, helping to expand the organizing group, recruit contributions to the knowledgebase, etc. Over time, as the community reaches a useful size and starts to take off, this role may become less important.

There are currently five active solution communities:

- Data Curation
- Media Archives
- Preservation and Archiving
- Scholar’s Workbench
- Small Archives

Each of these efforts started off by creating a vision for the community that sets the goals and defines its purpose. They then started wikis on the Confluence site provided by DuraSpace, and each of them set up a Google group that provides the mailing list for the growing community. DuraSpace has also begun a Crowdvine that also

provides a way for people to connect. The solutions community wiki site can be found at http://www.fedora-commons.org/confluence/display/FCCWG/Home.

This presentation will discuss the general idea of solution communities and their value to the open-source software development process. It will also discuss the experiences has so far with the active communities and demonstrate some of each of their knowledgebases.