

Taming the Data Deluge with the New Open Source iRODS Data Grid System

Version 1.0 Offers New Generation of Distributed Data Management Power

February 7, 2008

Paul Tooby

In the Information Age, the freedom to easily generate and share digital forms of information is driving life-changing advances in science and medicine, dramatic expansions in communications, big gains in business productivity, and a new flowering in video, music, and other cultural expressions.

At the same time, the digital data we all love is growing explosively. In 2006, humanity produced 161 exabytes of digital data - that's 161 billion billion bytes, or 12 stacks of books stretching from the Earth to the Sun -- more data than our capacity to store it.

This deluge of data is bringing with it unprecedented challenges in organizing, accessing, sharing, and preserving digital information. To meet these challenges, the Data-Intensive Computing Environments (DICE) group at the San Diego Supercomputer Center (SDSC) at UC San Diego has released version 1.0 of iRODS, the Integrated Rule-Oriented Data System, a powerful new open-source approach to managing digital data.

"iRODS is an innovative data grid system that incorporates and moves beyond ten years of experience in developing the widely used Storage Resource Broker (SRB) technology," said Reagan Moore, director of the DICE group at SDSC. "iRODS equips users to handle the full range of distributed data management needs, from extracting descriptive metadata and managing their data to moving it efficiently, sharing data securely with collaborators, publishing it in digital libraries, and finally archiving data for long-term preservation."

The most powerful new feature, for which the Integrated Rule-Oriented Data System is named, is an innovative "rule engine" that lets users easily accomplish complex data management tasks. Users can automate enforcement, or "virtualize" data management policies by applying rules that control the execution of all data access and manipulation operations. Rather than having to hard code these actions or workflows into the software, the user-friendly rules let any group easily customize the iRODS system for their specific data management needs.

For example, when astronomers take new photographs in a sky survey and enter them into a data collection, the researchers can set up iRODS rules to automatically extract descriptive information and record it in the iRODS Metadata Catalog (iCAT), replicate a copy to another repository for backup, create a thumbnail for a Web-based gallery, and run an analysis program to identify related images.

An organization's archivist can configure iRODS rules to identify and retain a collection of digital records for five years, and then move them to another site or destroy them. And if someone requests these records, the archivist can confirm that the current digital copy is indeed an authentic copy of the original. iRODS rules are being developed that will validate the trustworthiness of digital repositories.

Users can apply the growing set of existing rules or write new ones. Rules can also be developed as community-wide policies to manage data.

"One reason policy-based data management is important is that it lets communities integrate across different types of collection structures," said Moore. "What this means is that iRODS lets one community talk to any other community independent of what data management system the other community is using. No matter which technology you pick you aren't isolated."

iRODS is designed to be flexible, growing seamlessly from small to very large needs.

"You can start using it as a single user who only needs to manage a small stand-alone data collection," said Arcot Rajasekar, who leads the iRODS development team. "The same system lets you grow into a very large federated collaborative system that can span dozens of sites around the world, with hundreds or thousands of users and numerous data collections containing millions of files and petabytes of data - it's a true full-scale distributed data system." A petabyte is one million gigabytes, about the storage capacity of 10,000 of today's PCs.

At SDSC alone iRODS and its predecessor SRB technology are already managing one petabyte of data and two hundred million files for 5,000 users.

"It's an advantage that the new iRODS system is open source," added Rajaseker. "This is bringing in collaborators from the US and as far away as France, the UK, Japan, and Australia who are contributing code, so iRODS will quickly add more features."

"We also find that users like the open source approach and have more confidence in adopting the new technology. Open source software makes it possible to assemble a larger development team and interact with a wider range of user communities. This increases user confidence that the iRODS system will be around in the future."

Currently the iRODS team is working with partners to help a number of projects apply the technology, including the National Archives and Records Administration (NARA), the Ocean Observatories Initiative (OOI), the National Science Digital Library, the Temporal Dynamics of Learning Center (TDLC), the UC Humanities, Arts and Social Sciences (HASS) grid and the Testbed for the Redlining Archives of California's Exclusionary Spaces (T-RACES) project, and numerous others.

Version 1.0 of iRODS is supported on Linux, Solaris, Macintosh, and AIX platforms, with Windows coming soon. The iRODS Metadata Catalog (iCAT) will run on either the open source PostgreSQL database (which can be installed via the iRODS install package) or Oracle. And iRODS is easy to install -- just answer a few questions and the install package automatically sets up the system.

Under the hood, the iRODS architecture stores data on one or more servers, which may be widely separated geographically; keeps track of system and user-defined information describing the data with the iRODS Metadata Catalog (iCAT); and offers users access through clients (currently a command line interface and Web client, with more to come). As directed by iRODS rules, the system can process data where it is stored using applications called "micro-services" executed on the remote server, making possible smaller and more targeted data transfers.

"Because it's a second generation effort, iRODS isn't like a new, untested product since we have the knowledge from years of experience with dozens of projects using the SRB," said iRODS software architect Mike Wan. "iRODS includes the familiar functions from the SRB, so people can jump in and easily start using the new system."

Added iRODS senior software engineer Wayne Schroeder, "For a 1.0 release it has a large number of features -- we already knew where we were going as we developed it, and this has made it cleaner and faster."

To help users get started with iRODS, the DICE group is offering several tutorials and workshops in the US and internationally. Following on the very popular Society of American Archivists (SAA) workshop at SDSC last summer, there will be two SAA sessions this summer, with additional tutorials in the US, Europe, and Asia.

The DICE team plans to continue supporting the widely used SRB system well into the future. But as SRB users decide to upgrade, the team is developing a seamless migration path to the more capable and faster iRODS system. As part of this, for a digital data collection at NARA the iRODS team has already migrated one million files from an SRB data grid to an iRODS data grid.

"We migrated not only the data files but also the metadata, access controls, and directory structure," said Moore. "This is an important demonstration that users can migrate collections to different choices of data grid technology without any problem."

In addition to Moore, Rajasekar, Wan, and Schroeder, group members who contributed to the iRODS system include: Sheau-Yen Chen, Lucas Gilbert, Chien-Yi Hou, Arun Jagatheesan, George Kremenek, Sifang Lu, Richard Marciano, Dave Nadeau, Antoine de Torcy, and Bing Zhu. Other collaborators in the iRODS project include the French Institut National de Physique Nucléaire et de Physique des Particules (IN2P3), the UK e-Science Data Management Group at Rutherford Appleton Laboratory, and the High Energy Accelerator Research Organization, KEK, in Japan.

iRODS is funded by NARA and the National Science Foundation (NSF). More information, the iRODS software download, and documentation are available at <http://irods.sdsc.edu>.

Media Contacts: Warren R. Froelich, SDSC Communications, 858 822-3622 Jan Zverina, SDSC Communications, 858 534-5111