THE IRODS CONSORTIUM

CHARTER

TERMINOLOGY

iRODS – integrated Rule-Oriented Data System, data management software that associates data with metadata and implements automated workflows, secure collaboration, and data virtualization.

iRODS-based data middleware – Any data middleware implementation that is based on the iRODS technology and interoperable with iRODS protocols for authentication, transport, rule language, server interaction, and client access.

UNC-Chapel Hill – University of North Carolina at Chapel Hill

RENCI - The Renaissance Computing Institute – RENCI develops and deploys advanced technologies to enable research discoveries and practical innovations. RENCI was launched in 2004 as a collaborative effort involving the UNC Chapel Hill, Duke University and North Carolina State University. RENCI is a center within UNC-Chapel Hill.

DICE – The Data Intensive Cyber Environment (DICE) Center at the University of North Carolina at Chapel Hill. This group is the administrative home of the original iRODS development team.

CONTENTS

EXECUTIVE SUMMARY

The integrated Rule-Oriented Data System (iRODS) is a leading open-source data management solution that has been adopted worldwide by the academic research community, research organizations, government agencies, and private interests. The iRODS Consortium is tasked with promoting the sustainability of the iRODS technology and furthering its adoption and continued evolution. To this purpose, the Consortium works synergistically with Consortium members and the existing iRODS community to promote the iRODS data middleware technology, to evangelize iRODS among potential users, to promote new advances in iRODS, and to expand the adoption of iRODS-based data middleware technologies. The establishment of the iRODS Consortium provides the mechanism for bringing together members with vested interest in the success and sustainability of iRODS.

WHY THE CONSORTIUM IS NEEDED

iRODS is a community-driven, open-source, software environment. The iRODS effort is focused on providing research communities with distributed, policy-based data management technologies while concurrently engaging in research to identify new approaches for addressing challenges in data management and preservation. iRODS was originally funded primarily through research grant awards from the National Science Foundation (NSF). iRODS is now a leading technology for managing data within research communities and across the global research community.

The sustainability of the iRODS technology—exemplified by financial solvency, stability of the data management solution across software releases, and the ability to meet new demands, to adapt to new technologies, and to support new approaches to data management—is of paramount concern to potential adopters of iRODS. Many organizations need to see a comprehensive strategy for sustainability before they can adopt iRODS as a key component of their data management infrastructure. However, no single organization has the necessary resources, market visibility, or mission to ensure the sustainability and growth of iRODS as an open-source technology.

Ensuring that iRODS is sustainable and broadening the adoption of iRODS represent significant challenges. Toward this end, RENCI and UNC-Chapel Hill have worked with the iRODS development team and the iRODS community to evolve iRODS into a technology appropriate for use in enterprise-level deployments, which is easily adopted by users and developers.

Key features of this effort include the packaging and certification of binary distributions, implementing a greater level of system-level testing, modularizing of the code base to allow for pluggable components and easier integration with legacy systems, and hardening the code to resolve bugs and security concerns. The iRODS team integrates new features and bug fixes after the software code has been thoroughly tested and brought to conformance with the best software engineering practices.

The iRODS Consortium is a consortium of universities, research organizations, companies, and government agencies that together work to provide the guidance and obtain the funding needed to ensure that iRODS evolves in a manner that protects the investments of the Consortium members.

The Consortium funds and guides, in a neutral, non-competitive environment, the development of iRODS for use by Consortium members and non-members. Consortium members receive a variety of benefits, including prioritized access to support, training, and consulting, as well as the opportunity to influence the development roadmap for future software releases, thus providing both adopters and resellers of iRODS and other iRODS-based data middleware technologies a way to protect their investments.

VISION AND MISSION STATEMENTS

VISION STATEMENT

The iRODS Consortium develops the iRODS technology into a widely-adopted, sustained, middleware technology that enables effective and efficient data management, sharing, and integration within and across academic, federal, and industry enterprises.

MISSION STATEMENT

The mission of the iRODS Consortium is to ensure the sustainability of the integrated Rule-Oriented Data System (iRODS) and to further its adoption and continued evolution. To this end, the Consortium works to standardize the definition, development, and release of iRODS-based data middleware technologies, evangelize iRODS among potential users, promote new advances in iRODS, and expand the adoption of iRODS-based data middleware technologies through the development and release of iRODS as open-source, mission-critical, production-level software.

CONSORTIUM ORGANIZATION

The iRODS Consortium operates under RENCI, a research institute of the University of North Carolina at Chapel Hill. The Consortium consists of several working groups composed of both Consortium staff and member representatives, and a team of senior officers and staff located at RENCI.

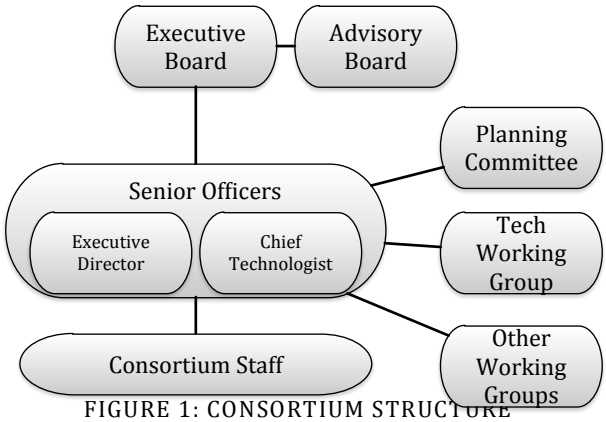

FIGURE 1: CONSORTIUM STRUCTURE

EXECUTIVE BOARD

Governance of the Consortium is provided through an Executive Board that consists of representatives from RENCI, DICE, and other Members with appropriate rights.

ADVISORY BOARD

The Consortium maintains an advisory board of officials with expertise in data technologies, iRODS, and academic, federal, and industry sectors that represent existing and potential users of iRODS. The advisory board serves as a forum for the governance of the Consortium, for evangelism of Consortium activities and iRODS in general, and for discussions of technical information related to Consortium activities.

PLANNING COMMITTEE

The Consortium includes a Planning Committee that serves to develop strategic plans and oversee implementation of plans for Consortium Activities. A key function of the Planning Committee is to approve information about Release Roadmaps for iRODS. Representation from voting and participating member organizations provides critical input and oversight for the success of this Committee.

TECHNOLOGY WORKING GROUP

The Consortium includes a Technology Working Group that serves to organize technical activities, develop technical plans, and oversee implementation of technical activities engaged in by the Consortium. A key function of the Technology Working Group is to produce draft software standards for approval by members, to produce draft release plans for Consortium developed software, and to develop Consortium documentation and software, including iRODS. Inclusion of technical staff allocated from select member organizations in the Technology Working Group aid in ensuring the iRODS middleware technology and a breadth of technical know-how and best practices informs Consortium practices.

EXECUTIVE DIRECTOR

The Consortium's Executive Director is responsible for ensuring the Consortium is responsive to its vision, mission, and the needs of member organizations. The Director provides for management of administrative and infrastructure operations and staff, oversight of the financial operations of the Consortium, management of relationships with Consortium members, recruitment of new Consortium members, development and presentation of Consortium status reports, and organization of events conducted by the Consortium.

CHIEF TECHNOLOGIST

The Consortium's Chief Technologist ensures that all products delivered by the Consortium, including iRODS, other related software products, documentation, standards, and technical reports are produced on time with a level of quality expected by member institutions and the broader IT community. The Chief Technologist works with technical experts in member organizations and other technical experts to ensure the Consortium products incorporate innovative technological advances and best practices for the management of data. The Chief Technologist also ensures that the best new features and capabilities generated by the broader iRODS community are incorporated into iRODS.

CONSORTIUM STAFF

In addition to the Executive Director and Chief Technologist, the Consortium employs a core team focused on the development of the iRODS software, including software developers, software testers, technical writers, and staff that provide technical support in the use of iRODS. The core team consists of staff from RENCI.

From time to time, the Consortium works directly with staff allocated by other member organizations to further aid in software development, testing, and support. Staff from member organizations work closely with Consortium staff to provide cross-fertilization of knowledge and best practices, thus creating a virtual development team that furthers the Consortium mission.

CONSORTIUM MEMBERSHIP

Consortium memberships at the following levels are open to all public and private entities, including research organizations, government agencies, and companies.

All Consortium members are granted rights to: participate in the Consortium Planning Committee and other working groups and committees formed as needed to advance Consortium activities; helpdesk support, training, and consulting (at levels outlined in the Membership Benefit Table); attend Consortium events; access Consortium documentation materials; and use the Consortium to disseminate appropriate member software, documentation, and marketing material. Summaries on the benefits of each membership level can be found in the iRODS Consortium Membership Benefit Table.

Premier Members are entities that contribute membership fees of $150,000 or more per year. All Premier Members are granted voting rights on the Executive Board and Planning Committee, thus providing Premier Members the right to influence the Governance of the Consortium and to influence Consortium activities, such as the software release roadmap for iRODS. Premier members receive the highest level of support from the Consortium and have the highest priority for accessing Consortium resources. In addition, Premier Members receive the highest priority for requests submitted to the Consortium for iRODS development and support, and for collaborating on new funding opportunities. Premier Members are permitted to include their hardware (up to 10 rack units, or "10U") in the iRODS continuous integration process for no fee.

Sustaining Members are entities that contribute membership fees of $75,000 or more per year. Like Premier Members, Sustaining Members are granted voting rights on the Executive Board and Planning Committee, thus providing Sustaining Members the right to influence the Governance of the Consortium and to influence Consortium activities. Requests from Sustaining Members receive high priority in accessing Consortium resources and support from the Consortium.

Professional Memberships are geared toward entities that desire voting rights on the software release roadmaps or entities that would benefit from the greater access the Consortium affords for networking and exposure to groups that have adopted or are adopting iRODS. Professional members are granted voting rights on the Planning Committee, and thus have direct influence on Consortium activities, such as the release roadmap for iRODS and the development of data grid standards. Professional membership fees are set at $35,000 per year.

General Memberships are geared toward entities that wish to be involved in Consortium activities and benefits without voting privileges. General members may participate in Consortium committees and working groups, including the Planning Committee. In addition, general members receive comprehensive access to Consortium technologies, documentation, and events. General membership fees are set at $10,000 per year.

CONSORTIUM ACTIVITIES

SOFTWARE DEVELOPMENT AND DOCUMENTATION

The Consortium performs the necessary software development and documentation activities to develop iRODS technology that is highly extensible, well tested, and production-quality. A primary objective of the Consortium is an open-source implementation that can be used by Consortium members and non-members 1) to deploy iRODS data middleware systems, 2) to support applications that depend on iRODS technologies, and 3) to extend the iRODS technology. The Consortium will continue to provide core iRODS technology without licensing fees to Consortium and non-Consortium members.

The Consortium guides the integration of new functionality and source code from iRODS after the source code has been reviewed, developed to Consortium standards, and determined to fit within the requirements for data-middleware technologies. The Consortium develops extensions to iRODS-based technologies based upon directions set forth by Consortium members. The Consortium's core team also coordinates and integrates software development and testing efforts, including the production of software documentation.

The Consortium maintains one or more web sites that provide access to current software implementation releases, software documentation, mailing list management and archive, a list of frequently asked questions, an issue tracking system, and the source code repository.

PRODUCTIZATION, TESTING, AND CERTIFICATION

A key function of the Consortium is to deliver packaged binary and source code distributions of the iRODS software and extensions to the software. The iRODS software development team develops and implements roadmaps for the development and deployment of packaged binary distributions, with the goal of making iRODS broadly available on a selection of operating systems, guided by the needs of Consortium members and non-members.

All released distributions are tested and certified by the Consortium. The Consortium, as part of its normal activities, manages and approves the requirements for testing and certification.

The Consortium guides the development and distribution of certified and non-certified extensions to the data middleware technologies. To encourage the development of extensions that are reliable and well engineered, the Consortium develops test suites to determine whether extensions properly implement standardized protocols and Application Programming Interfaces (APIs). The Consortium maintains a web site where members and non-members can download extensions provided by the Consortium, as well as extensions and links to open-source and proprietary extensions provided by the community and vetted by the Consortium.

PROMOTING THE USE AND ADOPTION OF IRODS-BASED TECHNOLOGIES

The Consortium pursues a number of dissemination and marketing efforts to encourage the adoption of iRODS-based data middleware technologies in industry, federal agencies, and academic institutions. In addition, the Consortium pursues efforts to encourage the further development of data-middleware extensions, e.g., storage resource drivers, in order to expand the range of technologies that iRODS-based data systems can interface with.

These efforts include the development and maintenance of liaison relationships with other organizations, vendors, and individuals that are responsible for:

- Making decisions regarding the technologies adopted for data management and archiving.
- Developing standards for interfacing to data storage resources.
- Developing standards for accessing, querying, and retrieving data from data systems.
- Developing data technologies that augment or can be incorporated into data-grid technologies.

The Consortium also works closely with other organizations to develop iRODS-based data middleware products and enhancements that promote the further adoption of iRODS-based technologies. The Consortium promotes the iRODS-based data middleware technologies through several avenues, including:

- The development and publication of case studies in different scientific, federal, and industry domains.
- The development and posting of test suites for performance and functionality.
- Professional conference seminars and the publication of conference proceedings and posters.
- Publication of papers in journals and in magazines and books.
- The development and publication of best practices for data grid usage, administration, and extension development.
- The organization and implementation of annual conferences for administrators, developers, and Consortium members and non-members. The conferences include tutorial sessions and presentations from Consortium members. The annual conference also includes status reports on Consortium plans and provides a forum for users and developers to discuss needs, successes, relevant technologies, and ongoing projects.

STANDARDIZING IRODS-BASED TECHNOLOGIES

The iRODS Consortium defines a set of expected behaviors of iRODS-based data middleware to:

- Provide automated testing capabilities that certify compliance with expected behavior.
- Provide a stable definition of the core iRODS data middleware technology for inclusion in all releases of the implementation software, facilitating upgrades at sites with mission-critical installations.
- Provide a standardized definition of the interfaces between the core data middleware technology and other technologies for inclusion in all releases of the implementation software.
- Provide robust, productized software that implements the core data middleware technology and interfaces, while supporting extensions that are independent of core functionality.

The Consortium staff works with Consortium members and the iRODS user community to define and document the core iRODS data middleware technology, standard interfaces and protocols for interacting with the technology, and standardized testing procedures. In addition, the Consortium works with community groups to make recommendations and to adopt updates to the core definitions and standards for iRODS-based data middleware. Finally, the Consortium staff drives iRODS development and operates the testing facilities to meet the established standards.

As part of these efforts, the Consortium investigates the development and publication of standards for data-middleware technologies that go beyond those developed for the iRODS implementations.

SUPPORTING ADOPTION AND USE THROUGH CONSULTATION AND TRAINING

The Consortium recognizes that professional consultation, support, and training are necessary for the success of many iRODS data grid users. The Consortium works to ensure that users who adopt iRODS have access to experts who can meet their needs through two approaches: First, the Consortium organizes, prioritizes, and manages the use of Consortium staff for support, consultation, and training. Second, the Consortium actively promotes third party vendors who can provide appropriate support, consultation, and training to Consortium members.

CONSORTIUM SUPPORT AND FUNDING

The Consortium operating expenses fall into three broad categories: staffing, technical infrastructure, and operating expenses. Staffing consists of the core team at RENCI, and any additional personnel provided by member organizations. The core Consortium team includes:

- Senior Officers: Executive Director and Chief Technologist
- System Development Team: developers and software testers
- Support Team: iRODS specialists to provide training, support, and administrative capabilities to Consortium members and support staff to provide technical writing, web development, and event planning.

Technical infrastructure includes costs for maintaining and running the development environment, testing environment, and web servers. The testing environment includes hardware and software that are capable of rapidly deploying a wide range of distributed iRODS deployments using virtualization technology and automated testing software.

Operating expenses include costs for travel, marketing, and event hosting. These costs will scale as membership grows, in order to cover a broader geographic area and membership base. Typical business expenses, e.g., personal computers, printing, mailing, office space, conference calls, etc., are covered in-kind by RENCI as part of its commitment to the Consortium.