



Contribution ID: 7

Type: **Presentation**

## **iRODS: Data-Centric, Metadata-Driven, Data Management**

*Tuesday, 30 January 2018 10:50 (20 minutes)*

iRODS is Open Source Data Management that can be deployed seamlessly onto your existing infrastructure, creating a unified namespace, and a metadata catalog of all the data objects, storage, and users on your system. iRODS allows access to distributed storage assets under the unified namespace and frees organizations from getting locked into single-vendor storage solutions. iRODS can represent data at rest in object, tape, and POSIX filesystems all within the same logical collection. Within the same catalog, iRODS provides the ability to annotate every data object, logical collection, storage resource, and user in the namespace. Through the use of this metadata these entities become actionable and may be operated upon by the integrated rule engine framework. The rule engine framework allows for any operation within your iRODS zone to be a trigger or hook for your code. This affords the creation of automated data management policy which may prevent the operation, provide context to the operation, log the operation or many other use cases. Additionally, iRODS provides the ability to federate any number of zones, allowing for the sharing of not only data across administrative boundaries, but the sharing of infrastructure as well.

Given these features iRODS provides a number of capabilities: automated data ingest, storage tiering, compliance, indexing, auditing, data integrity, provenance, and publishing.

This talk will cover the four core competencies of iRODS, which afford the implementation of these many capabilities. We will then cover emerging data management design patterns as well as existing use cases deployed in production. Finally, we will cover the iRODS software roadmap as well as an overview of the iRODS Consortium.

**Primary author:** Mr COPOSKY, Jason (iRODS Consortium)

**Presenter:** Mr COPOSKY, Jason (iRODS Consortium)

**Session Classification:** Scalable Storage Backends for Cloud and HPC: Foundations