



Architecture of the Regional Datastore in Trinity College Dublin

Brian Coghlan (coghlan@cs.tcd.ie), John Walsh, Stephen Childs, Geoff Quigley (geoff.quigley@cs.tcd.ie), David O'Callaghan, Gabriele Pierantoni, John Ryan, Neil Simon - Trinity College Dublin and Keith Rochford (rochfordk@cp.dias.ie) - Dublin Institute of Advanced Studies

e-INIS and the National Datastore

e-INIS is a federation of core electronic infrastructure providers dedicated to the provision of a sustainable national e-infrastructure supporting advanced academic research activities in Ireland. It aims to provide a cohesive e-infrastructure of a scale that enables internationally competitive research.

The project is funded under the Irish Higher Education Authority's Programme for Research in Third-Level Institutions (PRTL), a component of the National Development Plan.

It offers four main classes of capability: access to advanced computing facilities, specialist expert user support and training, secure network and grid services and pilot data management services - the last of these comprising the National Datastore, a federation of Regional Datastores. The first three of Regional Datastores are at Trinity College Dublin, University College Cork and Dublin Institute of Advanced Studies and will be online by the end of 2009.

Federation of storage allows data to be managed under a single namespace whilst retaining the ability to replicate across sites and to place data near compute facilities or other producers and consumers.

Classes of Data and Storage

Data access can be characterized via the properties of mutability and frequency of access.

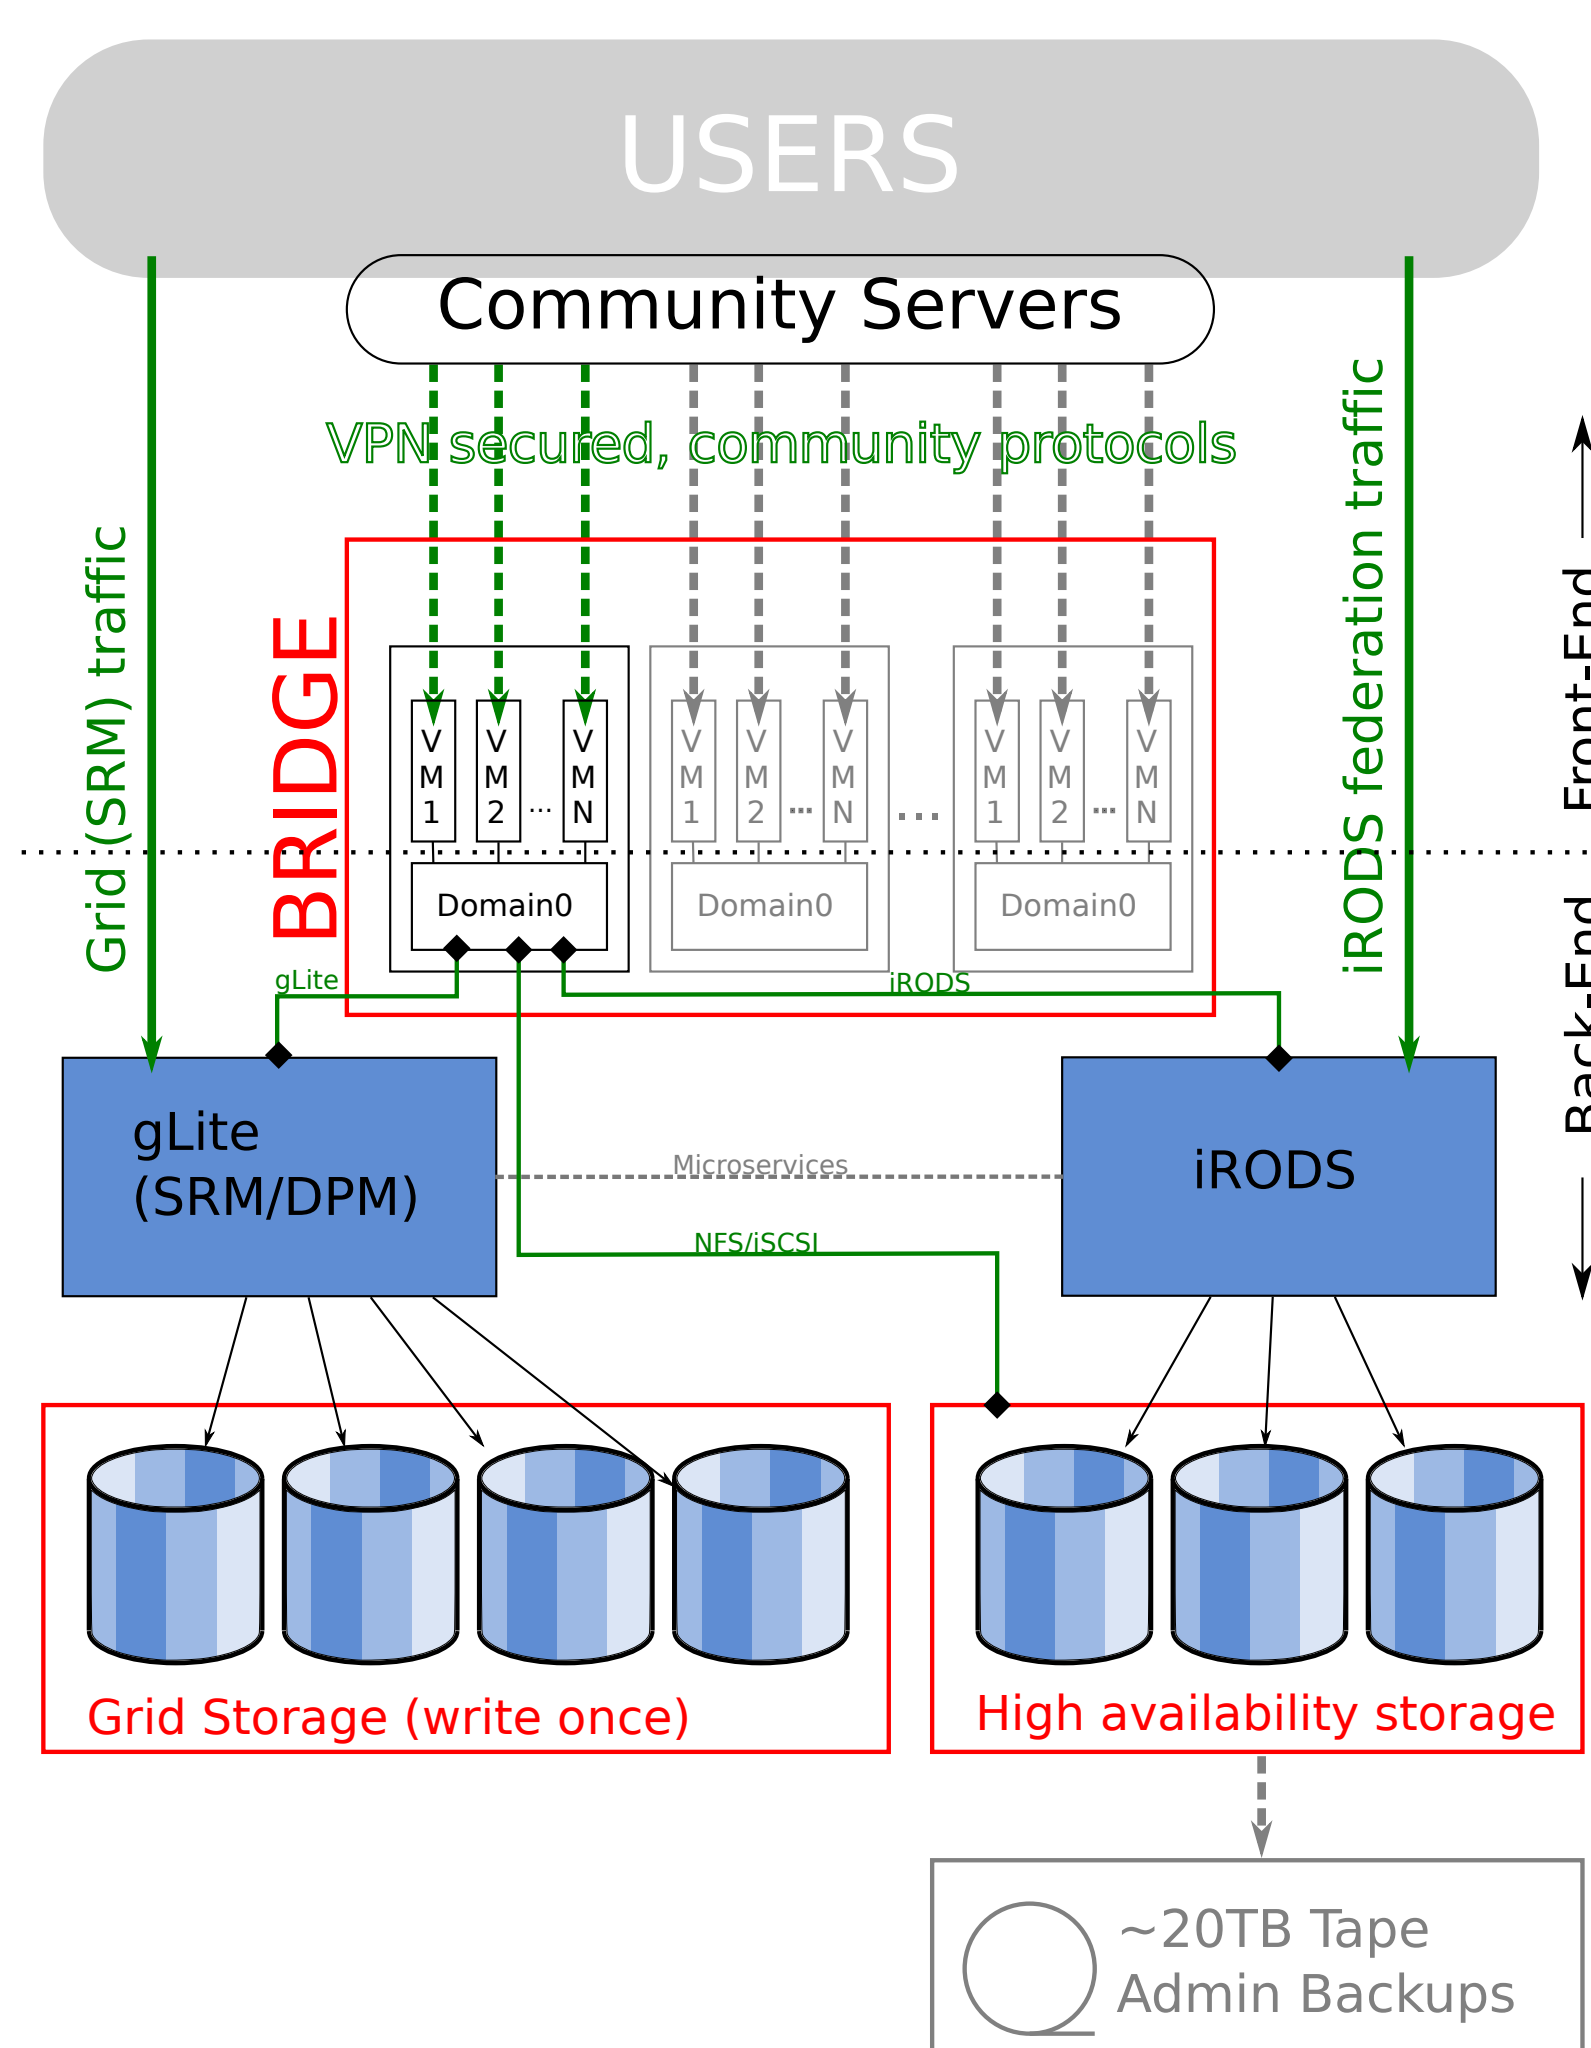
- Immutable data is written once (and read multiple times)
- Mutable data is changed in-place
- Frequently accessed data needs to be stored in an 'online' facility
- Infrequently accessed data can be stored in nearline facilities and brought online automatically, with some extra latency
- Rarely accessed data can be stored offline in facilities that require administrator intervention to bring online

We use gLite Grid technology with inexpensive hardware to support the immutable online accesses. Performance and resilience are achieved through replication and aggressive caching.

For mutable data, maintaining consistency of replicas becomes more of a problem. More expensive, high availability, hardware is used.

Nearline solutions utilising MAID technology are under investigation for future expansion. The limited tape capability is being reserved for administrator-initiated offline backups. Long term curation of data is outside the remit of the pilot.

Bridge Servers




By decoupling the front end from the back end, the bridge facilitates:

- Insulation of user communities from each other's activities
- Implementation of front-ends by communities
- Freedom of communities in how front-ends are implemented
- Freedom of communities to define their own security policies
- Scaling of the back end

Grid-Ireland Operations Centre

In Ireland, the universities and other institutions of advanced education and research are represented by the Higher Education Authority, and computing systems at these Institutions are interconnected by HEAnet.

Grid-Ireland is a managed layer above HEAnet providing grid services. The aim is to enable communities of users, for example, astrophysics, geneticists or linguists to construct virtual organisations above Grid-Ireland.



The national Grid is closely monitored and maintained by the operations team, based in Trinity College Dublin. The Grid-Ireland OpsCentre is the EGEE Regional Operations Centre (ROC) for Ireland

Regional Datastore Hardware in TCD

The Ops Centre in TCD hosts the first site to be online in the federated national datastore. This builds upon conventional grid storage, which will continue to be available in the current way, with a new architecture that uses technology such as iRODS and the notion of 'bridge servers' as a translation layer to provide complete flexibility to user communities in how they access their data while ensuring all write accesses are secured using PKI based on grid certificates. So far in excess of 0.5 PB of storage is running at TCD (with plans to extend that) with sites at UCC and DIAS due to come online by the end of the year.

