

Digital Object Registry
Corporation for National Research Initiatives
GENI Quarterly Report – Jan 1, 2009 through Mar 31, 2009

1. Major Accomplishments

The scope of work on this project is to adapt the Handle System and/or the CNRI Digital Object Registry to create a clearinghouse registry for principals, slices, and/or components in at least one GENI Spiral 1 control framework, capable of supporting limited operations in Year 1. We will also analyze ways in which the Handle System and/or a Digital Object Registry could be used to identify and register GENI software, including experimenter's tools, test images and configurations, and test results. Finally, we will define the operational, scaling, security, and management requirements, plus recommended design approaches, for implementing GENI clearinghouse and software registry services.

During this quarter, we began collaboration with the ProtoGENI group to federate the ProtoGENI clearinghouse records into a digital object registry, proposed a high level design of an experimental control tool, shared a few ideas in the control framework mailing list, and attended the GEC4 in Miami.

1.a. Milestones achieved: We proposed an experimental control tool that uses our Digital Object Architecture, following up on a conversation with Vicraj Thomas and Harry Mussman. The proposed tool, once implemented, would automate the building of a slice using resource combination models that evaluate whether or not resources requested by an experimenter are compatible with each other, and, if compatible, uniquely identifies the combination of those resources for reuse by experimenters, perhaps to repeat the same kind of experiment. Our GPO system engineer, Vic Thomas, concluded that there was insufficient experimental data available at this time to make any further progress on this proposal and it has been tabled for the time being.

1.b. Deliverables made: There were no scheduled deliverables during this quarter.

2. Description of Work Performed During Last Quarter

2.a. Activities and findings

As mentioned in the last quarterly report, we continue to evaluate the high-level designs and requirements for the GENI principals, components, and slice registries to determine how to leverage existing technologies, such as the Handle System and the digital object repository/registry system, to best create a common GENI clearinghouse registry.

While a clearinghouse for a single cluster is one potentially useful way to use our technologies in GENI, a clearinghouse that federates records from multiple clearinghouses across the GENI clusters would potentially be even more useful. Starting down this path, we were able to start collaboration with the ProtoGENI group, who made their clearinghouse data available to us. We are also trying to collaborate with other cluster members to have two or more participants to be able to define a normative GENI-wide accepted clearinghouse model. The fundamental goal of this federated clearinghouse is to serve all the GENI clusters, and, later on, all the GENI suites.

We are presently studying the existing setup of the ProtoGENI clearinghouse, adapting and refactoring its model to incorporate GENI requirements. We have downloaded the clearinghouse modules used to communicate with the ProtoGENI clearinghouse, and acquired relevant credentials and certificates. We have analyzed the security architecture, the metadata schemas for the registered items, and the resolution mechanisms used by the ProtoGENI clearinghouse. While the overall design is a very good one and effective for a cluster, there are a few features that we think could be added. We will design our federated clearinghouse to add those features, such that they will apply to the data that will be retrieved from the ProtoGENI clearinghouse. In our study of the ProtoGENI clearinghouse, we looked for areas we felt could be improved so that we could add value, in addition to potential interoperability, in the construction of a federated clearinghouse. The following observations were made at the end of this quarter.

The security architecture is a standard one that uses Public Key Infrastructure (PKI) methodologies, but the way the certificates are managed in the system lacks scalability. Additionally, the clearinghouse allows registering metadata for the items that are being registered, but does not allow for searching that metadata. For example, while the clearinghouse allows for registering a description about a professor, the clearinghouse does not provide a way to search on that description to retrieve that record. Furthermore, the UUIDs used for resolving records stored in a clearinghouse are not designed for distributed management or administration. That is, the present model does not allow individual organizations to create UUIDs on their own while still using those UUIDs across the organizations in a compatible way. This point goes back to the identifier topic we proposed in the control framework mailing list that addresses distributed administration requirements. While our study on ProtoGENI clearinghouse is still ongoing, we will address these issues when we design and prototype our federated clearinghouse.

In addition to the federation effort, we proposed an experimental control tool design that uses the Digital Object Architecture. The proposed tool enables combining resources requested to conduct experiments, and reuse of that combination of resources, when needed, by recalling the combination key. This tool optimizes the stitching process performed by slice authorities to ensure resource compatibility when experimenters request resources for conducting experiments by computing compatibility of those resources with the help of resource combination models.

Furthermore, once a combination is verified to be internally compatible, a unique key is assigned to that combination, in the form of a handle, to enable reusability of the combination by recalling the combination key.

This idea was well received by Vicraj Thomas. However, he suggested that we not pursue this further until GENI oversees some experiments, in order to leverage experience from those experiments and so decide which resource combination models would be relevant and what compatibility checks would have to be introduced.

2.b. Project Participants

CNRI has discussed its project with a number of other GENI participants, but all work done this quarter was done by CNRI alone or with the cooperation of ProtoGENI. Names and email addresses of CNRI and ProtoGENI participants are available on the GENI wiki page for the project.

2.c. Publications

No publications were produced this quarter. CNRI produced a presentation of the experimental control tool proposal. Those slides are available on the GENI wiki page for the project.

2.d. Outreach Activities

CNRI, specifically project PI Laurence Lannom and Giridhar Manepalli, attended the March/April GEC. We engaged in a number of technical discussions with other projects on the topic of future collaboration. Among other things, we also participated in RSpec discussion held by the ProtoGENI group.

CNRI participated in a control framework teleconference and contributed identifier and interoperability insights leveraging our experience in digital library and information management domains.