

Digital Object Registry
Corporation for National Research Initiatives
GENI Quarterly Report - Oct 14, 2008 through Dec 31, 2008

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 foreshortened quarter, CNRI developed an initial technical approach to creating a clearinghouse registry, attended and presented at the GEC meeting in Palo Alto in October, and made initial contacts with other projects.

During this foreshortened quarter, CNRI developed an initial technical approach to creating a clearinghouse registry, attended and presented at the GEC meeting in Palo Alto in October, and made initial contacts with other projects.

1.a. Milestones achieved: There were no scheduled milestones in this quarter.

1.b. Deliverables made: There were no scheduled deliverables during this quarter. CNRI created and delivered two presentations and one poster for the GEC meeting in October. These are available on the GENI wiki page for the project:

<http://groups.geni.net/geni/wiki/DigitalObjectRegistry>

2. Description of Work Performed During Last Quarter

2.a. Activities and findings

CNRI evaluated the high-level designs and requirements for the GENI principals, components, and slice registries to determine how to leverage its existing technologies, such as the Handle System and the digital object repository/registry system, to best create a common GENI clearinghouse registry. Those basic technologies are described below.

The Handle System provides a unified, distributed, and secure identification system that can be directly used to identify all discreet resources and entities within GENI, e.g., principals, slivers, aggregates, and components. One of the key features of the handle system, of particular interest in this context, is the ability to associate any data, or references to data, with each handle identifier, either directly within the handle record or through a pointer to external data. This allows the rapid and reliable resolution of any of a large set of identifiers to the current state data for the identified entity. Another critical functional aspect of the Handle System, beyond its ability to provide distributed storage and retrieval, is its secure administration mechanisms, which guarantee that only authorized entities can appropriately create, modify or delete handle records. This is of special importance within the GENI environment, as much of the information will have to be locally managed but globally accessible and trusted. For relatively small data items, such as the latest version number, the Handle System could be used directly to serve part of the GENI clearinghouse function. In other cases it will provide one or more pointers to external data, in registries or not.

The CNRI Digital Object Repository provides an object storage and retrieval protocol that can be used to provide a common overlay across multiple back-end storage systems. This is a natural fit for GENI and can be used to provide storage capacity beyond that of the handle system while providing the same level of distributed access and security. This technology could provide a standard distributed storage mechanism for experimental data sets, code modules, documentations and any other resources.

Finally, CNRI's digital object generic registry system could be used for registering, indexing and providing search capability over any XML data. More specifically, this system could be used to validate, register, index and provide search capability over Rspecs, principals, and code within the GENI framework. A more dynamic application of this registry system could be extended to registering slices with the intent to record or share slice information with other researchers. And in all cases, multiple instances of the digital object registry can be federated to provide a single search and access point across multiple registries.

CNRI designed an approach to taking the three specific technologies listed above and applying them within the specific context of the GENI Clearinghouse and the Experiment Control Tool systems. The technologies could also be used as standalone systems within the GENI framework as independent data repositories, registries or as a standalone identification system.

In the context of the clearinghouse, CNRI would be able to leverage the registry technology and the handle system technology to provide a principal, component and slice registry. This approach is shown graphically in the poster that was created for the GEC meeting in October, which is included in this report as an attachment.

The principal registry will be the least dynamic of the three but will require the highest level of security to make sure that each principal's identification,

authentication and authorization are kept up to date within the GENI framework. This registry will primarily leverage the handle system but could also use the registry for user discovery. The handle system would be used to identify, consistently and authoritatively define policies and permissions within the entire GENI framework

The component registry would leverage all three technologies, as it needs to securely identify all GENI components, make them searchable and potentially associate stored code, data sets and other resources with each component. The registry would be fairly dynamic, as it would have to be able to provide real time information about the status, including availability, of each component.

Finally, the Slice registry would be built in a similar manner to the component registry but it would be used to record the life cycles of slices as well as to provide real time information about slice status.

CNRI further designed a way to use the CNRI service composition framework, which itself leverages the handle system and the registry, to set up slices and to validate and coordinate components and resource allocation. Key to this framework is the ability to pair resources to matching rules such that a set of resources that compose a given slice can be dynamically evaluated for validity of slice composition. Is the given resource valid within the slice, is it available, and do you have rights to it? This approach would provide an extensible framework for establishing and recording experiment and system configuration rules. The ability of the composition framework to provide this service will require sufficient level of resource description in each relevant Rspec.

2.b. Project Participants

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

2.c. Publications

No publications were produced this quarter. CNRI produced two presentations and one poster for the GEC meeting in October. Those 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 October GEC. CNRI's President and CEO, Robert Kahn, also attended, although his attendance was not covered by this contract. We gave two presentations and created a poster for the poster session and engaged in a number of technical discussions with other projects on the topic of future collaboration.

Finally, we have participated in a number of conference calls with Harry Mussman, Vicraj Thomas of the GPO. These calls were devoted to exploring CNRI's technology base and planning on its use in the current and future GENI environments.