

MAX Quarterly Report

October 1, 2009 - December 31, 2009

Overview

The MAX project is building a GENI facility known as the "Mid-Atlantic Network Facility for Research, Experimentation, and Development (MANFRED)". This is an experimental fiber-based regional network in the Washington DC metro area. MANFRED will provide the GENI community with access to a regional optical network consisting of wavelength-selectable switches, 10 Gbps Ethernet switches, and virtual machines. The network infrastructure is building upon the NSF-funded DRAGON network which provides end-to-end dynamic circuit provisioning via a standardized Web Services interface through the use of a distributed GMPLS control plane - ensuring deterministic, high-speed performance over dedicated network resources. DRAGON technologies and software are also deployed on multiple research and education networks throughout the world. This wide deployment will be utilized to facilitate access to the MANFRED capabilities by remotely located researchers.

MANFRED has leveraged the DRAGON network infrastructure (and related technologies) by adding server virtualization capabilities (PlanetLab nodes) at the edges of the network and programmable network hardware (NetFPGA hosts) at two core switching nodes. MANFRED also has the capability to connect researchers in the Mid-Atlantic region to the rest of the GENI community via its connection to the private, high-speed Layer 2 backbone provided by Internet2 and the ProtoGENI project. The project web site is located here: <http://geni.maxgigapop.net>

A detailed listing and description of the MAX Spiral 2 project tasks and milestones is provided in Appendix A. The task and milestones identified for this reporting period are as follows:

Quarter 1 (12/31/2009):

MAX.S2.a: NetFPGA Integration into DRAGON

MAX.S2.b: NetFPGA as a GENI Resource Plan

MAX.S2.c: DRAGON Aggregate Manager Enhancement Design

MAX.S2.d: Common Control Framework Design

MAX.S2.e: User Support

MAX.S2.f: DRAGON Aggregate Manager Documentation

This document provides a summary of the activities accomplished during this reporting period, and also summarizes the focus area and work plan moving forward.

Milestones Status Report for Spiral 2

MAX.S2.a: NetFPGA Integration into DRAGON (complete)

The MAX GENI Substrate includes two high-end machines with NetFPGA (www.netfpga.org) cards. One machine is located at the University of Maryland lab at College Park campus, the other is at the USC/ISI East Facility in Arlington, Virginia. These host machines have Quad 3.0 GHZ Dual Core processors, 4 port Gigabit Ethernet NIC cards, and NetFPGA cards with another 4 Ethernet ports. They are now connected into the MAX GENI Substrate DRAGON Network components and are available for use and experimentation.

MAX.S2.b: NetFPGA as a GENI Resource Plan (underway - expected completion Jan 31, 2010)

We have begun to evaluate options for more fully integrating the NetFPGA host capabilities into the MAX GENI Aggregate. The main issues here revolve around the following areas:

- What virtualization environments to support? (PlanetLab, OpenFlow, others)
- What are the options for slicing NetFPGA resources and linking them to PlanetLab and/or OpenFlow slices? Initial implementations may be limited to a single user (or slice) per NetFPGA device at any one time.
- What are the unique issues associated with NetFPGA resources as it relates to the MAX Aggregate Manager Development?

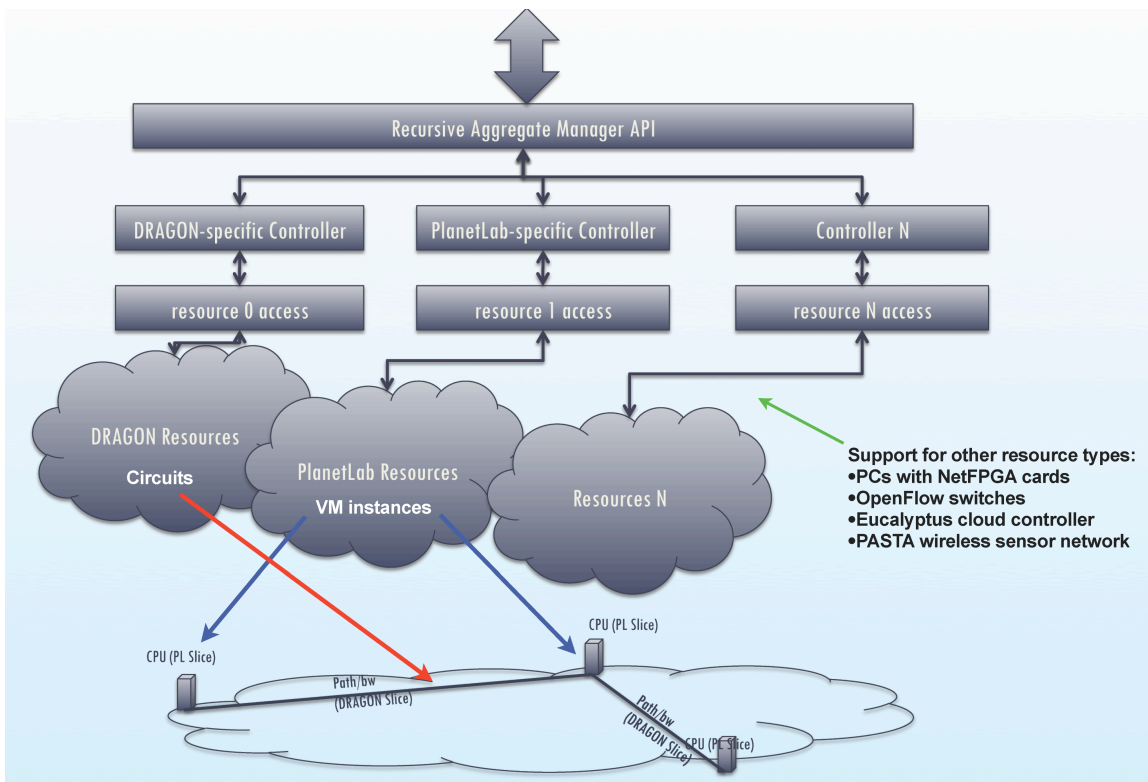
We are making progress in evaluating these issues and developing a more detailed plan moving forward. However we have more work to do before this milestone is complete. As we complete this task we will generate a summary of our plan moving forward for the NetFPGA resources. The ProtoGENI cluster is also using these same NetFPGA devices. We will also seek to obtain more information on the ProtGENI projects plans in this area, with the intent to leverage the work between our two efforts as much as possible.

MAX.S2.c: DRAGON Aggregate Manager Enhancement Design (complete)

The basic MAX Aggregate Manager Design is complete. In addition, we have a good concept for the enhancements we need moving forward to support Spiral 2. The following provides a summary of the MAX Aggregate Manager Design and the general enhancement plan.

The MAX Aggregate Manager software is a Java-based implementation that provides a Web Services API (WSDL) to clients. It is intended to be deployed in Apache Tomcat as an Axis2 service. The basic architecture of this Aggregate Manager is as shown below:

MAX Aggregate Manager External Interface
(Clearing House, Users)



The key architectural features to note regarding this Aggregate Manager are:

-Standard External Interface

The external interface is based on web services with a service definition in the form of WSDL file (AggregateGENI.wsdl). This WSDL based service definition is based on the Slice-based Facility Architecture documents produced by the PlanetLab group. The WSDL definition defines the following main services: ListCapabilities, ListNodes, CreateSlice, DeleteSlice, UpdateSlice, StartSlice, StopSlice, ResetSlice, QuerySlice.

This external interface also utilizes the PlanetLab GENI Control Framework, GENIWrapper. A rspec (max.xml) is included in the PlanetLab Slice Federation Architecture (SFA) repository which defines how to specify topologies which consist of PlanetLab slices interconnected by dynamically provisioned DRAGON dedicated network paths.

-Modularization

Since the resources in any given substrate are expected to be variable and change over time, this architecture accommodates this via a modular approach to resource controllers. For the initial MAX Substrate, the primary resources are DRAGON provisioned network paths, and PlanetLab provisioned user slices. The DRAGON resource controller is the existing dynamic provisioning system which is deployed on multiple research and education networks throughout the world. The PlanetLab controller is MyPLC PlanetLab Central/SFA system which allows slices to be created via human web interface or via application access. The MAX aggregate manager will be responsible for integrating the functions across these multiple resource controllers in response to requests presented at the common external interface.

-Extensibility

Another natural benefit if this architecture is extensibility in terms of accommodation of new resources (and associated controllers) as they become available. As shown in the figure above, additional resources such as cloud computing nodes/OpenFlow nodes/NETFPGA nodes/others can be added and integrated into the overall GENI substrate. As a result a slice topology can be instantiated which may span multiple resource types within the substrate.

The only Aggregate Manager services currently implemented are ListCapabilities and ListNodes. The plan moving forward is to complete the implementation of the rest of the Services such that users can instantiate an experiment topology which combines the provisioning of both PlanetLab slices and associated DRAGON dynamic network paths. This will provide for an experiment topology which is sliced at the both the host and network levels.

The currently deployed capabilities are available via the following mechanisms.

- MAX Aggregate Manager Service Interface
 - Human Readable Service Definition via Web Browser (FireFox)
<http://geni.dragon.maxgigapop.net:8080/axis2/services/AggregateGENI?wsdl>
 - SOAP API Interface
<http://geni.dragon.maxgigapop.net:8080/axis2/services/AggregateGENI>
- MAX MyPLC Service Interface
 - Human Web Browser Interface
<https://max-myplc.dragon.maxgigapop.net/>
 - XML-RPC/SOAP Interface
<https://max-myplc.dragon.maxgigapop.net/PLCAPI/>
- DRAGON Network Provisioning Interface
 - Web Browser Provisioning Interface:
<https://idc.dragon.maxgigapop.net:8443/OSCARS/>
 - Human Readable Service Definition via Web Browser (FireFox)
<https://idc.dragon.maxgigapop.net:8443/axis2/services/OSCARS?wsdl>
 - SOAP API Interface
<https://idc.dragon.maxgigapop.net:8443/axis2/services/OSCARS>

MAX.S2.d: Common Control Framework Design (complete)


We have made progress on the common control framework. This includes an agreement to set of rspecs based on the PlanetLab Control framework. This includes the MAX Aggregate Manager interface which is based on the PlanetLab GENI Control Framework and the GENIWrapper. As part of this we utilize the rspec (max.xml) which is included in the PlanetLab Slice Federation Architecture (SFA) repository. This defines how topologies consisting of PlanetLab slices interconnected by dynamically provisioned

DRAGON dedicated network paths are specified. The PlanetLab and ProtoGeni clusters are both working with PlanetLab based control frameworks. It is expected this will result in a good amount of interoperation between the associated projects. There are still many details that have to be evaluated and considered to realize true interoperability between different substrates and aggregate managers.

Below is an example of a MAX PlanetLab based rspec:

```
<?xml version="1.0" encoding="UTF-8"?>
<rspec xmlns="http://www.maxgigapop.net/sfa/07/09">
  <capacity><netspec name="predefined_physical_topology">
    <nodespec name="planetlab2">
      <node>planetlab2.dragon.maxgigapop.net</node>
      <ifspec name="pl23" linkid="2745"></ifspec>
      <ifspec name="pl24" linkid="2005"></ifspec>
      <ifspec name="pl25" linkid="2271"></ifspec>
    </nodespec>
    <nodespec name="planetlab3">
      <node>planetlab3.dragon.maxgigapop.net</node>
      <ifspec name="pl32" linkid="2745"></ifspec>
      <ifspec name="pl34" linkid="3157"></ifspec>
      <ifspec name="pl35" linkid="2975"></ifspec>
    </nodespec>
    <nodespec name="planetlab4">
      <node>planetlab4.dragon.maxgigapop.net</node>
      <ifspec name="pl42" linkid="2005"></ifspec>
      <ifspec name="pl43" linkid="3157"></ifspec>
      <ifspec name="pl45" linkid="1123"></ifspec>
    </nodespec>
    <nodespec name="planetlab5">
      <node>planetlab5.dragon.maxgigapop.net</node>
      <ifspec name="pl52" linkid="2271"></ifspec>
      <ifspec name="pl53" linkid="2975"></ifspec>
      <ifspec name="pl54" linkid="1123"></ifspec>
    </nodespec>
  </netspec></capacity><request></request>
</rspec>
```

This describes an experiment topology as shown below:



We consider this milestone complete in terms of the objectives for the reporting period. We had originally planned on developing an initial document describing an interoperation plan between cluster substrates. However the larger inter-cluster control framework meetings and discussions did not evolve as originally thought. In addition, the next reporting period includes continued work in this area and an update to the associated designs and documentations. So we plan to move the initial documentation release to the next reporting period. The tasks in this area that we plan to focus on for the next reporting period are:

- Further analysis of what is required to allow slices to be instantiated between other substrates/aggregate managers in PlanetLab Cluster, ProtoGeni Cluster, and ORCA Cluster.
- Documentation of the options with respect to the above issues including an approach for any translators or other adaptation mechanisms which may be required.
- Participation in the GENI Control Framework Working Group

MAX.S2.e: User Support (complete)

This task includes outreach and support of external researchers who would like to use the GENI DRAGON facilities. This task is included in all of the reporting periods and is expected to be an increasing important focus area for Quarter 3 and Quarter 4 reporting periods. For this reporting period our main focus is on making the MAX GENI facility available to those in the immediate MAX region and beyond. The other critical factor in supporting users is to complete the Aggregate Manager and other key technology development task to enable easy user access and use.

In terms of reaching out to users in other substrates and clusters we have begun discussion with other groups. The GpENI (Great Plains Environment for Network Innovation) is utilizing a DRAGON based network provisioning system so we will have a natural ability to interoperate with them and allow users in our respective footprints to have access to each others resources. As we get further along on this work, we also hope to be able to foster collaborations between researchers in each others region. We already have physical connectivity between the MAX DRAGON network and the GpENI peering point in Kansas city. This is in place due to the fact that both infrastructures are connected to the Internet2 DCN service. In addition, the MAX DRAGON infrastructure is also connected to ProtGENI switch in McLean, Virginia. We are actively discussing options with the GpENI project about how to complete and demonstrate interoperability between us. This may involve one or both of the options mentioned above. We hope to have this complete during the next reporting period.

Another key item for this reporting period is to utilize the physical connection we have to the ProtoGeni wide area infrastructure to interoperate with that cluster. We need to have further discussions with them to determine the major tasks for this.

We are also planning to evaluate options for interconnection to the ORCA cluster as well, but we do not have connections in place like we do for GpENI and ProtoGENI at this time. However, we have had some initial discussions regarding control plane interoperation. We will pursue this more during the next reporting period.

For all these inter-substrate and cluster interconnects, the main focus is on the expansion of user access and to foster diverse researcher collaborations moving forward.

MAX.S2.f: DRAGON Aggregate Manager Documentation (complete)

We continue to update the main project web site (geni.maxgigapop.net) with the latest software distributions and configuration documents. Below is a listing of the locations of some of the most critical components:

MAX-DRAGON-GENI Aggregate Manager Repository:
<svn://svn.maxgigapop.net/geni-aggregate>

PlanetLab Control Framework MAX RSPEC:
<http://svn.planet-lab.org/svn/sfa/trunk/sfa/rspecs/aggregates/max.xml>

PlanetLab Control Framework SFA:
<http://svn.planet-lab.org/svn/sfa/trunk/>

DRAGON Software Repositories:
<svn://cvs.maxgigapop.net/dragon-sw>
<http://sands225.east.isi.edu/narb-sw>

DRAGON/OSCARS/IDC Software Package:
<https://wiki.internet2.edu/confluence/display/DCNSS>

Project Participants

Peter O'Neil (MAX)
Abdella Battou (MAX)
Tom Lehman (USC/ISI)
Xi Yang (USC/ISI)

Focus For Next Quarter:

The focus for next quarter is to continue with the Quarter 2 tasks and milestones as identified. As described in the above Quarter 1 milestone status report, the main focus areas for the upcoming period are in the area of i) adding more functionality to the aggregate manager, ii) further progress on physical interconnection and control framework interoperation with other projects and clusters, iii) further integration of the NetFPGA resources into the MAX Substrate and Aggregate Manager functions.

Appendix A

MAX GENI Project

Tasks and Milestones

1. Introduction

This appendix defines the tasks and deliverables for Spiral 2 phase of the MAX project which is constructing the Mid-Atlantic Network Facility for Research, Experimentation, and Development (MANFRED) GENI Facility. A listing of the deliverables and milestones organized by due date is provided below. Section 2 provides a more detailed description of the tasks and deliverables.

Quarter 1 (12/31/2009):

MAX.S2.a: NetFPGA Integration into DRAGON
MAX.S2.b: NetFPGA as a GENI Resource Plan
MAX.S2.c: DRAGON Aggregate Manager Enhancement Design
MAX.S2.d: Common Control Framework Design
MAX.S2.e: User Support
MAX.S2.f: DRAGON Aggregate Manager Documentation

Quarter 2 (3/31/2010):

MAX.S2.g: NetFPGA as a GENI Resource Implementation
MAX.S2.h: DRAGON Aggregate Manager Enhancement Initial Implementation.
MAX.S2.i: Common Control Framework Design Update
MAX.S2.j: Common Control Framework Implementation
MAX.S2.k: Common Control Framework Documentation
MAX.S2.l: User Tool Evaluation and Development Plan
MAX.S2.m: User Support
MAX.S2.n: DRAGON Aggregate Manager Documentation

Quarter 3 (6/30/2010):

MAX.S2.o: DRAGON Aggregate Manager Enhancement Updated Implementation.
MAX.S2.p: DRAGON Aggregate Manager Common Control Framework Integration Plan
MAX.S2.q: DRAGON Aggregate Manager Common Control Framework Integration Implementation
MAX.S2.r: DRAGON Aggregate Manager Common Control Framework Integration Testing
MAX.S2.s: User Tool DRAGON Framework Implementation
MAX.S2.t: User Support
MAX.S2.u: DRAGON Aggregate Manager Documentation

Quarter 4 (9/30/2010):

MAX.S2.v: DRAGON Aggregate Manager Enhancement Performance Testing and Evaluation.
MAX.S2.w: DRAGON Aggregate Manager Common Control Framework Updates
MAX.S2.x: User Support
MAX.S2.y: DRAGON Aggregate Manager Documentation

2. Tasks and Deliverables

a) Virtualization Servers Integration into MANFRED

This task will include the integration of the PC/NetFPGA systems into the DRAGON network environment. These systems will be connected to the edge of the DRAGON network and will be available as a resource for GENI researchers. These NetFPGA cards have multiple Gigabit Ethernet interfaces which will be connected to the DRAGON network to facilitate multiple topologies in support of experiment configurations. The PC/NetFPGA systems will be made available for use as a PlanetLab node or as standalone system with access to NetFPGA card and programming interface. Other organizations in the ProtoGENI cluster are working on techniques to virtualize the NetFPGA resource. This capability will be tested and incorporated as possible into the DRAGON environment.

The Quarterly Milestones for this task are defined below.

- Quarter 1 (12/31/2009):
 - i) (S2.a) NetFPGA Integration into DRAGON. Complete the placement and configuration of the NetFPGA systems into the DRAGON network. This will include configuration and integration to allow use of the NetFPGA systems as end-systems connected to DRAGON provisioned network paths. This will not include integration into the general GENI slice/sliver provisioning system. That functionality will be dependent on results from study in item ii below.
 - ii) (S2.b) NetFPGA as a GENI Resource Plan. Develop a plan to enable the NetFPGA systems to be generally available as part of a GENI slice/sliver and topology provisioning. This will include a review of the work by ProtoGeni and others on the development of NetFPGA virtualization technologies.
- Quarter 2 (3/31/2010):
 - i) (S2.g) NetFPGA as a GENI Resource Implementation. Implement the results of the plan from milestone Task a:Quarter 1:ii above. The expected result is for the NetFPGA systems to be available as MANFRED GENI slice/sliver resource.

b) Enhance DRAGON Aggregate Manager

This task is to enhance the current reference implementation of the DRAGON Aggregate Manager. The objective here is to facilitate the provision of the DRAGON network resources in support of GENI researcher and experiment use. This will include enhancement of the current Aggregate Manager to enable user specification of a topology which includes both compute and network resources. The DRAGON Aggregate Manager will be responsible for authorization and resource verification (topology computation) in advance of resource provisioning in accordance with the user specified topology descriptions. This work will extend the previous DRAGON Aggregate Manager to provide the basis for an extensible control and provisioning framework.

The Quarterly Milestones for this task are defined below.

- Quarter 1 (12/31/2009):
 - i) (S2.c) DRAGON Aggregate Manager Enhancement Design. Complete design for the enhancements to the DRAGON Aggregate Manager to include the features described in above task description.
- Quarter 2 (3/31/2010):
 - i) (S2.h) DRAGON Aggregate Manager Enhancement Initial Implementation. Complete initial implementation of enhanced DRAGON Aggregate Manager based on design work in milestone Task b:Quarter 1:i above. This initial implementation will focus on user interface and topology specification which combines compute and network resources into a common user request format.
- Quarter 3 (6/30/2010):
 - i) (S2.o) DRAGON Aggregate Manager Enhancement Updated Implementation. Complete updated implementation of enhanced DRAGON Aggregate Manager based on design work in milestone Task b:Quarter 1:i above. This updated implementation will extend the work completed in Task b:Quarter 2:i above to include features for more robust user authorization and resource verification (topology computation) techniques.
- Quarter 4 (9/30/2010):
 - i) (S2.v) DRAGON Aggregate Manager Enhancement Performance Testing and Evaluation. Conduct performance testing and evaluation of DRAGON Aggregate Manager. Solicit user feedback on use of DRAGON Aggregate Management. Summarize strengths and weakness of DRAGON Aggregate Manager and develop recommendations for future upgrades.

c) Participate in the development of a common GENI Control Framework or set of Interoperable Mechanisms

This task will include working within the larger GENI community to work on development of a common or compatible set of GENI control frameworks. This may include the normalization of the existing GENI control frameworks to a common GENI interface definition, or possibly the development of a common set of mechanisms for individual framework specification. The latter would allow mechanisms such as translators to be developed to enable cross framework interactions. It is anticipated that the details of the technical solution will be the subject of specially focused GENI working groups consisting of PlanetLab,

ProtoGENI, ORCA, ORBIT, and other control frameworks. This task will include participating in these larger discussions to work toward a common GENI control framework.

The Quarterly Milestones for this task are defined below.

- Quarter 1 (12/31/2009):
 - i) (S2.d) Common Control Framework Design. Work with GENI community to develop an approach for a common control framework and/or other mechanisms to allow interoperation between GENI facilities. Develop initial documentation of this plan in sufficient detail for the various clusters to take implementation action for their respective control frameworks.
- Quarter 2 (3/31/2010):
 - i)(S2.i) Common Control Framework Design Update. Continue work with the GENI community to develop an approach for a common control framework and/or other mechanisms to allow interoperation between GENI facilities. Update associated documentation.
 - ii) (S2.j) Common Control Framework Implementation. Develop an initial common framework implementation sufficient for interoperability testing. Conduct interoperability testing with other GENI Facilities to evaluate ability for multiple GENI facilities to interoperate.
 - iii) (S2.k) Common Control Framework Documentation. Generate a document summarizing the results of interoperability testing and general capabilities for GENI Facilities to interoperate.

d) Integrate enhanced DRAGON GENI Control Framework with DRAGON Testbed

The primary objective of this task is to apply the results of task c above, to the actual DRAGON network and facilities. The common GENI framework, or inter-framework mechanisms, that result from task c will require changes to the current DRAGON Framework and Aggregate Manager. This will include adaptation for the unique set of resources located on the DRAGON network, as well as enhancements to the experimenter facing interface.

The Quarterly Milestones for this task are defined below.

- Quarter 3 (6/30/2010):
 - i) (S2.p) DRAGON Aggregate Manager Common Control Framework Integration Plan. Develop a design and implementation plan for the DRAGON Aggregate Manager to incorporate the results common/interoperable control framework approach developed in Task c above.
 - ii) (S2.q) DRAGON Aggregate Manager Common Control Framework Integration Implementation. Incorporate changes to the DRAGON Aggregate Manager based on the results of milestone Task d: Quarter 3: i above.
 - iv) (S2.r) DRAGON Aggregate Manager Common Control Framework Integration Testing. Conduct testing with updated DRAGON Aggregate Manager and common control framework. Document results and changes needed.
- Quarter 4 (9/30/2010):
 - i) (S2.w) DRAGON Aggregate Manager Common Control Framework Updates. Incorporate updates and changes to the DRAGON Aggregate Manager based on the results of milestone Task d: Quarter 3 above and user feedback.

e) Improve user access to prototype GENI via enhanced tools

This task will include development of user focused tools to facilitate and enhance the ability and ease with which users can access the GENI DRAGON environment. This will likely include working with other project and cluster developed APIs such as GUSH and RAVEN to ensure they work seamlessly with the DRAGON control frameworks.

The Quarterly Milestones for this task are defined below.

- Quarter 2 (3/31/2010):
 - i) (S2.l) User Tool Evaluation and Development Plan. Review the available user tools across the GENI facilities and clusters and identify best candidate for integration into the DRAGON facility and control framework.
- Quarter 3 (6/30/2010):
 - i) (S2.s) User Tool DRAGON Framework Implementation. Incorporate user tools identified in milestone Task e: Quarter 2: i above. into the DRAGON facility and control framework

f) Support increased use of testbed by external researchers

This task will include the outreach and support of external researchers who would like to use the GENI DRAGON facilities. The exact support will be tailored to the individual users and researchers as they are identified.

The Quarterly Milestones for this task are defined below.

- Quarter 1 (12/31/2009):
 - i) (S2.e) User Support. Work with and support interested users on use of the DRAGON GENI Facility.
- Quarter 2 (3/31/2010):
 - i) (S2.m) User Support. Work with and support interested users on use of the DRAGON GENI Facility.
- Quarter 3 (3/31/2010):
 - i) (S2.t) User Support. Work with and support interested users on use of the DRAGON GENI Facility.
- Quarter 4 (9/30/2010):
 - i) (S2.x) User Support. Work with and support interested users on use of the DRAGON GENI Facility.

g) Deliver aggregate manager design documentation to GPO

This task will be to deliver design and usage documentation to the GPO for the DRAGON Aggregate Manager. This is expected to include architecture and design documents, configuration and usage documentation, and open source software distributions.

The Quarterly Milestones for this task are defined below.

- Quarter 1 (12/31/2009):
 - i) (S2.f) DRAGON Aggregate Manager Documentation. Update the design and installation documents for the DRAGON Aggregate Manager. Provide updated open source software distribution of GENI DRAGON Aggregate Manager and other related software.
- Quarter 2 (3/31/2010):
 - i) (S2.n) DRAGON Aggregate Manager Documentation. Update the design and installation documents for the DRAGON Aggregate Manager. Provide updated open source software distribution of GENI DRAGON Aggregate Manager and other related software.
- Quarter 3 (6/30/2010):
 - i) (S2.u) DRAGON Aggregate Manager Documentation. Update the design and installation documents for the DRAGON Aggregate Manager. Provide updated open source software distribution of GENI DRAGON Aggregate Manager and other related software.
- Quarter 4 (9/30/2010):
 - i) (S2.y) DRAGON Aggregate Manager Documentation. Update the design and installation documents for the DRAGON Aggregate Manager. Provide updated open source software distribution of GENI DRAGON Aggregate Manager and other related software.