

MAX Quarterly Report

January 1, 2010 - March 31, 2010

Overview

The MAX project is building a GENI facility known as the "Mid-Atlantic Crossroads GENI (MAX GENI) Facility". This is an experimental fiber-based regional network in the Washington DC metro area. MAX GENI 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 MAX GENI capabilities by remotely located researchers.

MAX GENI 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. MAX GENI 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 2 (3/31/2010):

MAX.S2.b: NetFPGA as a GENI Resource Plan (this item was originally in Quarter 1 list, moved to Quarter 2)

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

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, Quarter 2

MAX.S2.b: NetFPGA as a GENI Resource Plan (complete)

The MAX GENI Substrate deployment includes research nodes with NetFPGA (www.netfpga.org) cards. The plan for these NetFPGA resources is to make them available as a general MAX GENI Substrate resource for use on an individual basis, or to combine with other host and network resources into an experiment specific topology. This will allow researchers to build an experiment topology that includes NetFPGA hosts and other substrate resources such as PlanetLab slivers and also interconnect them via dedicated network resource provisioning. The specific use of the NetFPGA hosts will be up to the individual researchers. However, we plan to pre-load the OpenFlow software in anticipation of a desire for experimentation in this area. For instance, a researcher could request a topology that included multiple NetFPGA hosts that are interconnected in a ring. This type of topology would lend itself to experimentation with OpenFlow protocols and associated higher-level applications. We also had a meeting

with Rob Sherwood from the OpenFlow project and talked in some detail about options for more extensive use and experimentation with OpenFlow. Our work in this area may expand in the future.

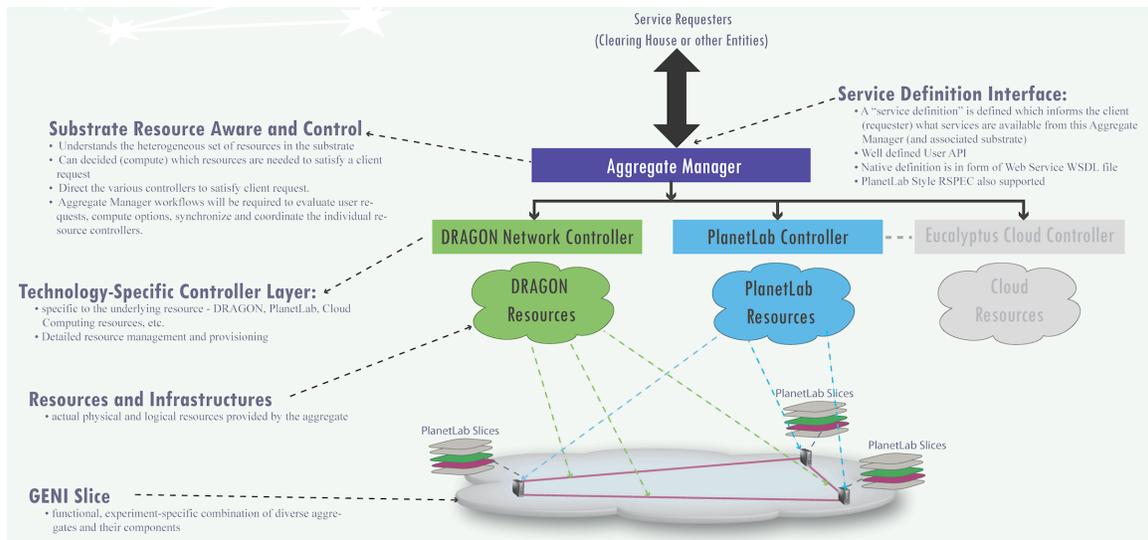
We evaluated the possibility of deploying PlanetLab or other virtualization technologies on the NetFPGA hosts. However, there does not currently appear to be a good method for sharing the NetFPGA resource in the context of a virtual environment. For this reason, we feel that providing the NetFPGA, as a reserveable unit is the most valuable way to treat this resource at this time.

MAX.S2.g: NetFPGA as a GENI Resource Implementation (complete)

The MAX GENI Substrate includes three high-end machines with NetFPGA (www.netfpga.org) cards. Two machines are located at the University of Maryland lab at College Park campus, and 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 connected into the MAX GENI Substrate DRAGON Network components and are now available for use and experimentation. The MAX GENI Aggregate Manager includes the NetFPGA machines as a discoverable resource via the ListCapabilities and ListNodes calls. We have few more items to complete on the Aggregate Manager task to fully incorporate these NetFPGA resources into the dynamic topology instantiation. The details of how these resources are included in the Aggregate Manager is outlined in the status presented in Milestone MAX.S2.h below. The general approach is to include the NetFPGA hosts as a top-level resource in the ListCapabilities and ListNodes service calls of the MAX Aggregate Manager. These resources could then be included in subsequent slice and topology creation commands.

MAX.S2.h: DRAGON Aggregate Manager Enhancement Initial Implementation (complete)

The MAX Aggregate Manager architecture and design is complete. 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:



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 a WSDL file (AggregateGENI.wsdl).

-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, PlanetLab provisioned user slices, and NetFPGA hosts. The DRAGON resource controller is the existing dynamic provisioning system that is deployed on multiple research and education networks throughout the world. The PlanetLab controller is MyPLC PlanetLab Central/SFA system that 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 of this architecture is extensibility in terms of accommodation of new resources (and associated controllers) as they become available. As shown in the above figure, planned future activities include the addition of Eucalyptus resources to this capability set. As a result a slice topology can be instantiated that may span multiple resource types within the substrate.

We have made much progress in the feature enhancements during this reporting period. We have completed the implementation of the features described below. We have a little more work to do to complete the documentation and provided detailed instructions and guidance to facilitate researcher use.

The status and enhanced feature set of the MAX GENI Aggregate Manager is described below.

i) MAX-GENI Aggregate Service Definition

The services have been restructured into four groups.

- **Compute Slice Services** include *CreateSlice*, *DeleteSlice*, *QuerySlice*, *UpdateSlice*, *StartSlice* and *StopSlice*.
- **VLAN Networking Services** include *CreateSliceVlan*, *DeleteSliceVlan* and *QuerySliceVlan*.
- **RSpec-Based Slice Services** include *CreateSliceNetwork*, *DeleteSliceNetwork* and *QuerySliceNetwork*.
- **Resource Information and Topology Services** include *ListCapabilities*, *ListNodes*, *ListSlices* and *GetResourceTopology*.

These web services are roughly equivalent to the SFA services. For example, the *list** services offer information similar to “*sfi list*” and the compute slice services also have corresponding “*sfi*” commands. In addition, we use VLAN networking services to provision point-to-point VLAN connections between two compute nodes/slivers. The *Create/Delete/QuerySliceNetwork* services are the RSpec-based named slice operations available via “*sfi*” but we add the VLAN networking components that bridge the PlanetLab nodes/slivers. The *GetResourceTopology* service by default can provide the physical substrate resource topology description including both compute and network components. It can also provide the overlay topology description of a named slice.

ii) MAX GENI Aggregate RSpec

We have enhanced our MAX specific RSpec definition to better capture the integration of the compute and networking resource components. Our approach is to leverage existing work in dynamic networking deployments and incorporate a new compute resource description tailored to the GENI environment. We have composed a max-RSpec XSD schema that can describe the substrate infrastructure and any overlay slices using a common schema. The compute resources part includes generic compute nodes, Planetlab slivers, NetFPGA nodes, and Eucalyptus slivers. The network part is extended from Open Grid Forum (OGF) Network Markup Language (NML) network topology description. The schema also includes description for how these components are interconnected to form a topology of compute and network resources.

iii) Implementation of MAX GENI Aggregate Manager Services

At the time of our last status report we had implemented the following services: *ListCapabilities*, *ListNodes*. We have now completed the remainder of the services which are: *ListSlices*, *CreateSliceVlan*,

DeleteSliceVlan, QuerySliceVlan, CreateSlice, DeleteSlice, QuerySlice, StartSlice, StopSlice, CreateSliceNetwork, DeleteSliceNetwork, QuerySliceNetwork and *GetResourceTopology*.

Of particular note are the new *CreateSliceNetwork, DeleteSliceNetwork, QuerySliceNetwork* and *GetResourceTopology* services. These now allow us to offer RSpec based named slice operations with coherent compute and network resource provisioning. For example, a researcher can request a three node triangle slice network via *CreateSliceNetwork* by providing the below RSpec description. The MAX Aggregate Manager will take care of the slice creation and management. The researcher can then later use the *QuerySliceNetwork* to check on slice states and *DeleteSliceNetwork* to terminate the slice.

-----MAX GENI Rspec Example-----

```
<rspec id="xi_rspec_slice1">
  <aggregate>geni.maxgigapop.net</aggregate>
  <description>Xi's test rspec</description>
  <lifetime id="time-1271533930-1271563981">
    <start>1271174406</start>
    <end>1271563981</end>
  </lifetime>
  <computeResource id="urn:aggregate=geni.maxgigapop.net:rspec=my-test-max-rspec-slice1">
    <planetlabNodeSliver id="urn:aggregate=geni.maxgigapop.net:rspec=my-test-max-rspec-
slice1:domain=dragon.maxgigapop.net:node=planetlab3">
      <address>206.196.176.55</address>
      <computeCapacity>
        <cpuType>generic</cpuType>
        <cpuSpeed>2.0GHz</cpuSpeed>
        <numCpuCores>1</numCpuCores>
        <memorySize>256MB</memorySize>
        <diskSize>16GB</diskSize>
      </computeCapacity>
      <networkInterface id="urn:aggregate=geni.maxgigapop.net:rspec=my-test-max-rspec-
slice1:domain=dragon.maxgigapop.net:node=planetlab3:interface=eth1.1001">
        <deviceType>Ethernet</deviceType>
        <deviceName>eth1</deviceName>
        <capacity>100Mbps</capacity>
        <ipAddress>10.10.10.2/24</ipAddress>
        <vlanRange>1001</vlanRange>
        <peerNetworkInterface>urn:aggregate=geni.maxgigapop.net:rspec=my-test-max-rspec-
slice1:domain=dragon.maxgigapop.net:node=planetlab2:interface=eth1.1001</peerNetworkInterface>
      </networkInterface>
      <networkInterface id="urn:aggregate=geni.maxgigapop.net:rspec=my-test-max-rspec-
slice1:domain=dragon.maxgigapop.net:node=planetlab3:interface=eth1.107">
        <deviceType>Ethernet</deviceType>
        <deviceName>eth1</deviceName>
        <capacity>100Mbps</capacity>
        <ipAddress>10.10.30.1/24</ipAddress>
        <vlanRange>107</vlanRange>
        <peerNetworkInterface>urn:aggregate=geni.maxgigapop.net:rspec=my-test-max-rspec-
slice1:domain=dragon.maxgigapop.net:node=planetlab5:interface=eth1.107</peerNetworkInterface>
      </networkInterface>
    </planetlabNodeSliver>
    <planetlabNodeSliver id="urn:aggregate=geni.maxgigapop.net:rspec=my-test-max-rspec-
slice1:domain=dragon.maxgigapop.net:node=planetlab5">
      <address>206.196.176.138</address>
      <computeCapacity>
        <cpuType>generic</cpuType>
        <cpuSpeed>2.0GHz</cpuSpeed>
        <numCpuCores>1</numCpuCores>
        <memorySize>256MB</memorySize>
        <diskSize>16GB</diskSize>
      </computeCapacity>
    </planetlabNodeSliver>
  </computeResource>
</rspec>
```

```

    <networkInterface id="urn:aggregate=geni.maxgigapop.net:rspec=my-test-max-rspec-
slice1:domain=dragon.maxgigapop.net:node=planetlab5:interface=eth1.312">
      <deviceType>Ethernet</deviceType>
      <deviceName>eth1</deviceName>
      <capacity>100Mbps</capacity>
      <ipAddress>10.10.20.2/24</ipAddress>
      <vlanRange>312</vlanRange>
      <peerNetworkInterface>urn:aggregate=geni.maxgigapop.net:rspec=my-test-max-rspec-
slice1:domain=dragon.maxgigapop.net:node=planetlab2:interface=eth1.312</peerNetworkInterface>
    </networkInterface>
    <networkInterface id="urn:aggregate=geni.maxgigapop.net:rspec=my-test-max-rspec-
slice1:domain=dragon.maxgigapop.net:node=planetlab5:interface=eth1.107">
      <deviceType>Ethernet</deviceType>
      <deviceName>eth1</deviceName>
      <capacity>100Mbps</capacity>
      <ipAddress>10.10.30.2/24</ipAddress>
      <vlanRange>107</vlanRange>
      <peerNetworkInterface>urn:aggregate=geni.maxgigapop.net:rspec=my-test-max-rspec-
slice1:domain=dragon.maxgigapop.net:node=planetlab3:interface=eth1.107</peerNetworkInterface>
    </networkInterface>
  </planetlabNodeSliver>
  <planetlabNodeSliver id="urn:aggregate=geni.maxgigapop.net:rspec=my-test-max-rspec-
slice1:domain=dragon.maxgigapop.net:node=planetlab2">
    <address>206.196.176.133</address>
    <computeCapacity>
      <cpuType>generic</cpuType>
      <cpuSpeed>2.0GHz</cpuSpeed>
      <numCpuCores>1</numCpuCores>
      <memorySize>256MB</memorySize>
      <diskSize>16GB</diskSize>
    </computeCapacity>
    <networkInterface id="urn:aggregate=geni.maxgigapop.net:rspec=my-test-max-rspec-
slice1:domain=dragon.maxgigapop.net:node=planetlab2:interface=eth1.1001">
      <deviceType>Ethernet</deviceType>
      <deviceName>eth1</deviceName>
      <capacity>100Mbps</capacity>
      <ipAddress>10.10.10.1/24</ipAddress>
      <vlanRange>1001</vlanRange>
      <peerNetworkInterface>urn:aggregate=geni.maxgigapop.net:rspec=my-test-max-rspec-
slice1:domain=dragon.maxgigapop.net:node=planetlab3:interface=eth1.1001</peerNetworkInterface>
    </networkInterface>
    <networkInterface id="urn:aggregate=geni.maxgigapop.net:rspec=my-test-max-rspec-
slice1:domain=dragon.maxgigapop.net:node=planetlab2:interface=eth1.312">
      <deviceType>Ethernet</deviceType>
      <deviceName>eth1</deviceName>
      <capacity>100Mbps</capacity>
      <ipAddress>10.10.20.1/24</ipAddress>
      <vlanRange>312</vlanRange>
      <peerNetworkInterface>urn:aggregate=geni.maxgigapop.net:rspec=my-test-max-rspec-
slice1:domain=dragon.maxgigapop.net:node=planetlab5:interface=eth1.312</peerNetworkInterface>
    </networkInterface>
  </planetlabNodeSliver>
</computeResource>
</rspec>

```

-----MAX GENI Rspec Example-----

Another important service is the *getResourceTopology* service. Calling this service with a RSpec slice name like the one shown above, “*xi_rspec_slice1*”, will provide detailed information on the compute and network resources for that particular slice. If this service is called without a specific slice name, the service will respond with a description of the entire MAX GENI substrate resources including available compute nodes, network topology, and how they are interconnected. This will allow a researcher to discover the

available resources and subsequently make a request for an experiment specific topology. The below is a reduced output for the currently available MAX substrate.

-----MAX GENI *getResourceTopology* Example-----

```
<?xml version='1.0' encoding='UTF-8'?>
<soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/">
  <soapenv:Body>
    <ns1:GetResourceTopologyResponse xmlns:ns1="http://aggregate.geni.net/services/api">
      <ns1:status>normal</ns1:status>
      <ns1:resourceTopology>
        <ns1:statement>
          <computeResource id="+urn:aggregate=geni.maxgigapop.net:rspec=physical_crdb">
            <computeNode id="urn:ogf:geni:domain=dragon.maxgigapop.net:node=planetlab2">
              <plld>2</plld>
              <networkInterface
id="urn:aggregate=geni.maxgigapop.net:rspec=physical_crdb:domain=dragon.maxgigapop.net:node=planetlab2:interface=eth1">
                <deviceType>Ethernet</deviceType>
                <deviceName>eth1</deviceName>
                <capacity>1000Mbps</capacity>
                <ipAddress></ipAddress>
                <vlanRange>1001-2000</vlanRange>
                <attachedLinkUrn>urn:ogf:network:domain=dragon.maxgigapop.net:node=MAX:port=1-0-9:link=*</attachedLinkUrn>
              </networkInterface>
            </computeNode>
            <computeNode id="urn:ogf:geni:domain=dragon.maxgigapop.net:node=planetlab3">
              <plld>5</plld>
              <networkInterface
id="urn:aggregate=geni.maxgigapop.net:rspec=physical_crdb:domain=dragon.maxgigapop.net:node=planetlab3:interface=eth1">
                <deviceType>Ethernet</deviceType>
                <deviceName>eth1</deviceName>
                <capacity>1000Mbps</capacity>
                <ipAddress></ipAddress>
                <vlanRange>1001-2000</vlanRange>
                <attachedLinkUrn>urn:ogf:network:domain=dragon.maxgigapop.net:node=CLPK:port=1-0-4:link=*</attachedLinkUrn>
              </networkInterface>
            </computeNode>
            ##### Skip more compute nodes #####
          </computeResource>
        </ns1:statement>
      </ns1:statement>
      <topology id="dragon.maxgigapop.net-1212079234">
        <domain id="urn:ogf:network:domain=dragon.maxgigapop.net">
          <node id="urn:ogf:network:domain=dragon.maxgigapop.net:node=MAX">
            <address>140.173.2.233</address>
            <port id="urn:ogf:network:domain=dragon.maxgigapop.net:node=MAX:port=1-0-9">
              <capacity>1000000000</capacity>
              <maximumReservableCapacity>1000000000</maximumReservableCapacity>
              <minimumReservableCapacity>1000000</minimumReservableCapacity>
              <granularity>1000000</granularity>
              <link id="urn:ogf:network:domain=dragon.maxgigapop.net:node=MAX:port=1-0-9:link=*">
                <remoteLinkId>urn:ogf:network:domain=*.node=*.port=*.link=*</remoteLinkId>
                <trafficEngineeringMetric>100</trafficEngineeringMetric>
                <capacity>1000000000</capacity>
                <maximumReservableCapacity>1000000000</maximumReservableCapacity>
                <minimumReservableCapacity>1000000</minimumReservableCapacity>
                <granularity>1000000</granularity>
              </link>
            </port>
            <port id="urn:ogf:network:domain=dragon.maxgigapop.net:node=MAX:port=1-2-1">
              <capacity>1000000000</capacity>
            </port>
          </node>
        </domain>
      </topology>
    </ns1:resourceTopology>
  </soapenv:Body>
</soapenv:Envelope>
```

```

        <maximumReservableCapacity>1000000000</maximumReservableCapacity>
        <minimumReservableCapacity>1000000</minimumReservableCapacity>
        <granularity>1000000</granularity>
        <link id="urn:ogf:network:domain=dragon.maxgigapop.net:node=MAX:port=1-2-1:link="*">
          <remoteLinkId>urn:ogf:network:domain=dragon.maxgigapop.net:node=CLPK:port=1-2-
2:link=140.173.97.94</remoteLinkId>
          <trafficEngineeringMetric>100</trafficEngineeringMetric>
          <capacity>1000000000</capacity>
          <maximumReservableCapacity>1000000000</maximumReservableCapacity>
          <minimumReservableCapacity>1000000</minimumReservableCapacity>
          <granularity>1000000</granularity>
        </link>
      </port>
    </node>
  <node id="urn:ogf:network:domain=dragon.maxgigapop.net:node=CLPK">
    <address>140.173.2.233</address>
    <port id="urn:ogf:network:domain=dragon.maxgigapop.net:node=CLPK:port=1-0-4">
      <capacity>1000000000</capacity>
      <maximumReservableCapacity>1000000000</maximumReservableCapacity>
      <minimumReservableCapacity>1000000</minimumReservableCapacity>
      <granularity>1000000</granularity>
      <link id="urn:ogf:network:domain=dragon.maxgigapop.net:node=CLPK:port=1-0-4:link="*">
        <remoteLinkId>urn:ogf:network:domain=*.node=*.port=*.link=*</remoteLinkId>
        <trafficEngineeringMetric>100</trafficEngineeringMetric>
        <capacity>1000000000</capacity>
        <maximumReservableCapacity>1000000000</maximumReservableCapacity>
        <minimumReservableCapacity>1000000</minimumReservableCapacity>
        <granularity>1000000</granularity>
      </link>
    </port>
    <port id="urn:ogf:network:domain=dragon.maxgigapop.net:node=CLPK:port=1-2-2">
      <capacity>1000000000</capacity>
      <maximumReservableCapacity>1000000000</maximumReservableCapacity>
      <minimumReservableCapacity>1000000</minimumReservableCapacity>
      <granularity>1000000</granularity>
      <link id="urn:ogf:network:domain=dragon.maxgigapop.net:node=MAX:port=1-2-1:link="*">
        <remoteLinkId>urn:ogf:network:domain=dragon.maxgigapop.net:node=MAX:port=1-2-
1:link=140.173.97.93</remoteLinkId>
        <trafficEngineeringMetric>100</trafficEngineeringMetric>
        <capacity>1000000000</capacity>
        <maximumReservableCapacity>1000000000</maximumReservableCapacity>
        <minimumReservableCapacity>1000000</minimumReservableCapacity>
        <granularity>1000000</granularity>
      </link>
    </port>
  </node>
  ##### Skip more network nodes #####
</domain>
</topology>
</ns1:statement>
</ns1:resourceTopology>
</ns1:GetResourceTopologyResponse>
</soapenv:Body>
</soapenv:Envelope>
-----MAX GENI getResourceTopology Example-----

```

The currently deployed capabilities are available via the following mechanisms.

- MAX Aggregate Manager Service Interface
 - Human Readable Service Definition via Web Browser (FireFox) or SOAP Client (ditchnet.org/soapclient)
 - <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.i: Common Control Framework Design Update (complete)

We have continued our common control framework design approach that includes use of the PlanetLab GENI Control Framework and a specifically tailored MAX RSPEC to incorporate the unique features of the MAX Substrate. These unique features revolve around the integration of dynamic network provisioning with host-based resources like PlanetLab slivers. Details of this RSPEC were presented earlier in this status report.

The general approach for our common control framework interoperability is a MAX Aggregate Manager based on an open service interface using standardized web service mechanisms. We feel that this will allow us to conform to any common control framework that may evolve from the multiple GENI Aggregate Managers under development. As described above in the milestone MAX.S2.h (DRAGON Aggregate Manager Enhancement Initial Implementation) status, we now have the base implementation of a fully featured service interface for the MAX GENI Aggregate Manager that includes the following services:

- ListCapabilities
- ListNodes
- ListSlices
- CreateSlice
- DeleteSlice
- UpdateSlice
- StartSlice
- StopSlice
- QuerySlice
- CreateSliceVlan
- DeleteSliceVlan
- QuerySliceVlan
- CreateSliceNetwork
- DeleteSliceNetwork
- QuerySliceNetwork
- GetResourceTopology

It appears that currently a single GENI Common Control Framework is still under discussion and development within the community. The start of an official API being defined here:

- http://groups.geni.net/geni/wiki/GAPI_AM_API

The key services defined here are:

- ListResources
- CreateSliver
- DeleteSliver

- SliverStatus
- RenewSliver
- Shutdown

While we have slightly different names, we feel we do include all the functionality defined above in our aggregate manager service interface. And as with the other aggregate manager development efforts, we have an extended set of capabilities in the form of more advanced and substrate specific services. In general each aggregate is developing its own API that generally includes a more extensive list of services than those defined above.

Our design approach for a common control framework is to develop our open service interface, publish the API, and demonstrate interoperability with other Aggregates. In some cases these interoperability demonstrations will require translation between aggregate services interfaces. We believe this will maximize the chance of an interoperable control framework evolving from the multiple aggregate service development efforts. These interoperability demonstrations will be a focus of our efforts during the next reporting period.

MAX.S2.j: Common Control Framework Implementation (complete)

We have implemented our standard and open web service Aggregate Manager Interface. This is the basis of common control framework implementation. If official GENI common API, schema, and message formats are developed, we will modify our current service interface definition to match this. The current GENI specifications define a set of capabilities but not necessarily specific message formats.

MAX.S2.k: Common Control Framework Documentation (complete)

Our common control framework is documented in the form of standard web service schema and WSDL message definitions. This documentation is located on the MAX GENI web site at the below location.

- <http://geni.dragon.maxgigapop.net:8080/axis2/services/AggregateGENI?wsdl>

(note: this documentation can be viewed in human readable format via a browser like FireFox, or via downloading the file and opening in any text editor)

This specifies the details of the interface to allow users to make calls, and application developers to build more sophisticated applications around the MAX GENI Substrate capabilities. We will be adding to this documentation to describe more fully the services available, and provide additional help on steps for researchers to access these facilities.

MAX.S2.l: User Tool Evaluation and Development Plan (complete)

Now that we have the base set of MAX Aggregate Manager Services in place, we will focus more on development of tools to facilitate the user experience. Our development plan here is to develop several reference APIs, which users can use directly in their experiment specific operations and software. We have already developed the core of these user tools as part of our development and demonstration work. During the next reporting period we will package these along with the appropriate documentation so that they will be sufficient for experimenter use. In addition we will further evaluate other GENI project user tools like GUSH and RAVEN to see if there is value in incorporation of our service interface into those tools.

MAX.S2.m: User Support (complete)

This task includes outreach and support of external researchers who would like to use the MAX GENI 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 tasks to enable easy user access and use.

In terms of reaching out to users in other substrates and clusters we made good progress during this reporting period in providing access and developing interoperability with other groups and projects. This included working the GpENI and ProtoGENI projects to demonstrate interconnection of our substrates. At

GEC7 we demonstrated the building of a topology which utilized dynamically provisioned network resources to build a private network connecting PlanetLab slivers at GpENI (University of Kansas) and MAX (University of Maryland). This included a segment across ProtoGENI that was statically pre-provisioned. We plan to next work on integrating the dynamic provision of ProtoGENI resources into this capability. Our GEC7 demo also included dynamic network connections to JGN2 in Japan. Our hope is that we can expand that collaboration in the future to the Japan CoreLab project where they have a PlanetLab infrastructure in place.

For all these inter-substrate and cluster interconnects, the main purpose is on the expansion of user access and to foster diverse researcher collaborations moving forward. The objective of this substrate/aggregate manager integration effort is to create a richer set of resources for users and to expand the opportunities for these collaborations.

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

The MAX GENI web site has the Aggregate manager documentation here:

- <http://geni.maxgigapop.net/twiki/bin/view/GENI/Software>

The latest version of the MAX AggregateGENI software distribution will also allow the compile and build process to be completed wholly at the Linux terminal if the use of NetBeans is not desired as currently described on the web site. The ANT build script has been modified and all dependency libraries are now included in the distribution and the configuration structure is streamlined. The svn repository contains README and sample configure files which provides detailed instructions on installation and configuration. The effort has greatly improved deployment experience.

Additional documentation will be added to allow reflect a new packaging system to facilitate researcher access and use of the MAX Aggregate Manager service interface. We will 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 3 tasks and milestones as identified. This includes the following: i) completing the deployment, documentation, and user tool packaging for the new features of the MAX GENI Aggregate Manager, ii) addition of x509 based security features for Aggregate Manager Access, iii) work on multi-aggregate interoperation between MAX and other aggregates (ProtoGENI, GpENI, ORCA), iv) increased outreach and work with potential GENI researchers.

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.