Technical Report for the project of "Integrating a CRON (Cyberinfrastructure of Reconfigurable Optical Network) Testbed into ProtoGENI" (GENI Proposal 1794)

Period: April 2011 – June 2011

PI: Seung-Jong Park
Computer Science and Center for Computation & Technology
Louisiana State University
289 Coates Hall
Baton Rouge, LA 70803
sjpark@csc.lsu.edu

Co-PI: Rajgopal Kannan
Computer Science
Louisiana State University
289 Coates Hall
Baton Rouge, LA 70803
rkannan@csc.lsu.edu

I. Major Accomplishments

Scope

This effort provides a reconfigurable optical network emulator aggregate connected to the GENI backbone over Louisiana Optical Network Initiative (LONI). The role of optical network emulation in GENI is to provide a predictable environment for repeatable experiments, and to perform early tests of network research experimentation prior to acquiring real network resources. The tools and services developed by this project will integrate with the ProtoGENI suite of tools. The aggregate manager and network connections between LONI and GENI for this project will also allow other LONI sites to participate in GENI.
For the second year from 2010 to 2011, the scope of work includes

(1) implementation of GENI AM API,
(2) trusting all four GENI Clearinghouses,
(3) finding experimenters and providing supports for experimenters,
(4) performing demos with experimenters at GECs,
(5) exporting the status of CRON to the GMOC, and
(6) hosting tutorials at GECs.

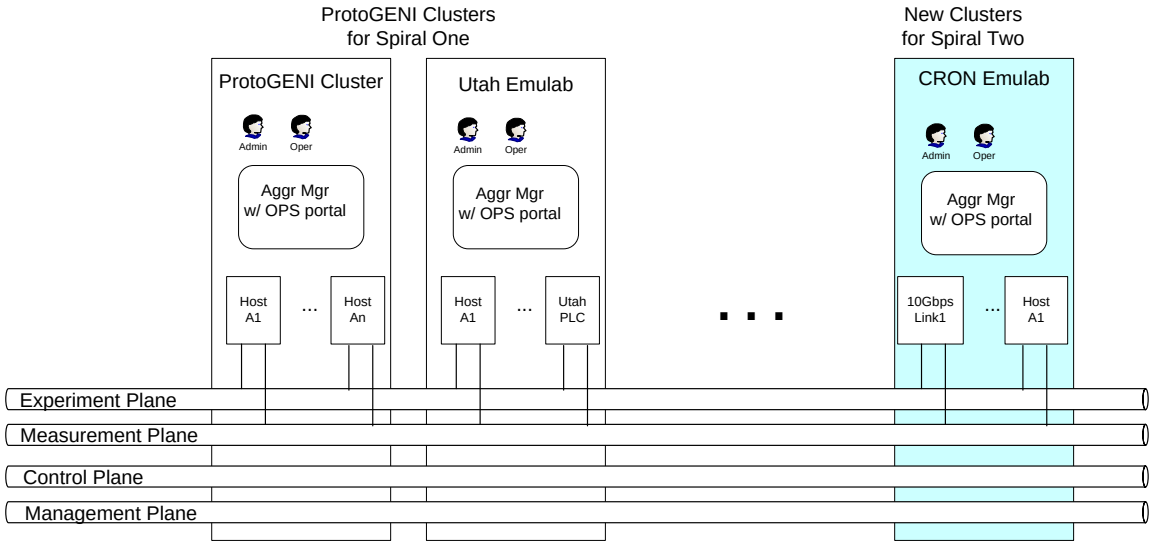

Figure 1. Implementation of Federation of CRON into ProtoGENI Cluster C at the Spiral 3

A. Milestones Achieved Between April 2011-June 2011

Note: (6mo) means activity that completes by the 6th month after start of project. Also, (7-12mo) means activity that starts in the 7th month and continues through the 12th month.

Milestone 1: Outreach to current CRON user community. Due 5/27/11

- Web pages that include sample experiments involving GENI and CRON resources to show the current CRON user community the benefit of becoming GENI experimenters
- Advertize this web page to the current CRON user community
- Updated experimenter guide for GENI experimenters using CRON resources
  :

Status: Accomplished

For the purpose of tutorial, video and example files are linked at our CRON website (http://cron.loni.org/crondemo.php)

Milestone2: Implementation of GENI API over Internet2. Due 6/30/11

- Continuously operation aggregate manager that implements the GENI API

Status: Accomplished

To accomplish those two milestones, we first needed to upgrade CRON testbed based on recent release from ProtoGENI which includes recent GENI API. At the same time, we needed to move CRON testbed

from LSU networks (www.cron.cct.lsu.edu) to LONI (Louisiana Optical Network Initiative) networks (www.cron.loni.org). Due to restricted security administration, users could not have full access to our resources. For example, outside GENI collaborators and experimenters could not have a direct access to public IP addresses because of LSU network security policy using network firewalls.

To resolve these issues, CRON testbed has been moved from domain "www.cron.cct.lsu.edu" to "cron.loni.org", with CNAMEs like "www.cron.loni.org" and "boss.cron.loni.org". From "loni.org", we have requested a whole class C IP space for CRON testbed open to public. Then, the CRON testbed control network has been re-designed to have multiple sub-network spaces by virtual LAN technology in order to accommodate security and extension. Currently, CRON testbed is designed to have control network, experiment network and monitor network, since all nodes inside are remote manageable high performance servers. The communication features of CRON are outstanding private optical 10Gbps experimental network and public 1Gbps control network on each nodes.

After CRON had been updated based on the recent version of ProtoGENI (release version of April 2011), CRON has continuously operated an aggregate manager that implements the recent GENI API.

B. Deliverable Made

Not available during this quarter

II. Description of work performed during last quarter

A. Activities and Findings

1) Federation CRON with PlanetLab clusters through Internet2 in GENI framework

CRON testbed has been upgraded to latest Emulab API which includes recen GENI API. And we also upgraded "ops.cron.loni.org" with latest Emulab API as users access server, public mailing server and file system, and upgraded "boss.cron.loni.org" with lastest Emulab API plus ProtoGENI API to have direct federation to ProtoGENI and have support for map interface "Flack". Authorized official SSL certificates have been approved from "lsu.edu". And we have customized CRON testbed API to provide functionalities for 10Gbps optical network.

After upgrading CRON, we have collaborated with collaborators at Utah group and MAX (Mid-Atlantic Crossroads) project to aggregate CRON with other ProtoGeni/Planetlab sites through GENI framework. Federation between CRON and BBN GPOLab within ProtoGeni has already been demonstrated at GEC 9. Federation between CRON and PlanetLab at MAX site is ongoing.

For GENI aggregation, at first, a GENI *clearinghouse* authenticates experimenters and issues them credentials needed to obtain GENI resources for experimentation. Then GENI *aggregates* provide

resources to experimenters with GENI credentials. A GENI *slice* holds a collection of computing and communications resources capable of running an experiment or a wide area service. And *RSpec* is the mechanism for advertising, requesting, and describing the resources used by slice

CRON have used the GENI Aggregate Manager API, including Flack and Omni. The GENI Aggregate Manager API provides a common interface to Aggregate Managers, including PlanetLab, ProtoGENI, and OpenFlow.

Also, network stitching will be provided for CRON to connect into GENI through Internet2. GENI network stitching operation is to construct a topology of substrates as represented by their Aggregate Managers. Each Aggregate Manager has a unique Rspec which defines its Substrate resources. Rspecs is a topology description of the individual substrate.

Experimentors will get slivers including several machines from CRON at cron.loni.org(ProtoGeni) and several machines from MAX(PlanetLab). Using Flack/Omni to get slivers at each site with allocated resources and exchanged Rspec. Then we will get ION circuit from LONI/MAX to connect the resources. Finally, we will start test on the resources through the connection in Internet2.

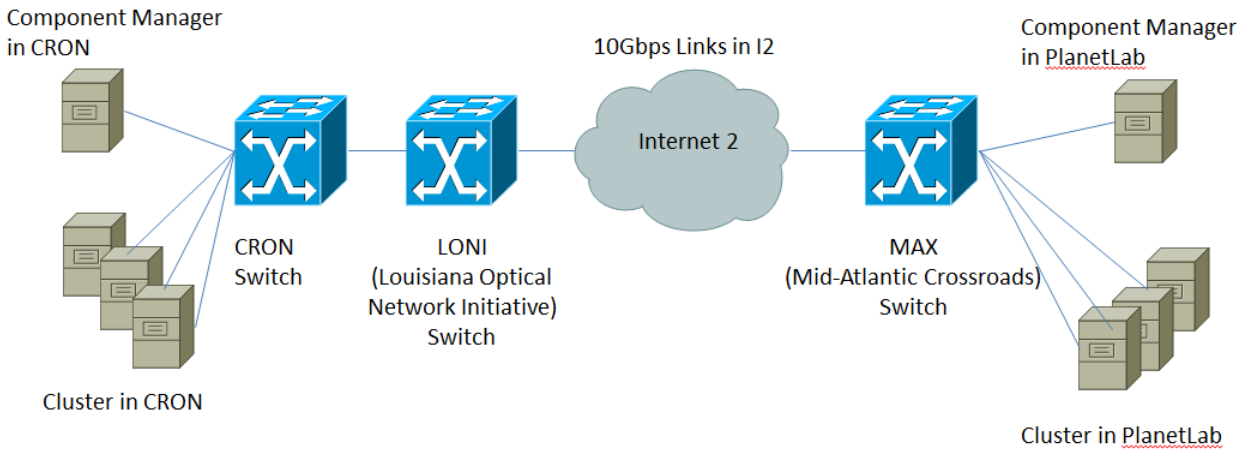

2) Demo Preparation Showing Multiple Virtual Clusters over Virtual High Speed Networks

To prepare the demonstration over a slice connecting resources at CRON and other testbeds, we use a software, called Eucalyptus, which is a computing and storage resources management system. It manages a set of machines and dynamically allocates part of them to users on demand. The system consists of 2 major components, the Hypervisor, the virtual machine management software, on each physical machine and the virtual machine allocation system among all physical machines.

The fundamental unit for resource allocation is a Virtual Machine (VM). One physical machine is divided into multiple VMs and serves multiple users in parallel so that the utilization for the physical hardware is better than it is with the traditional one physical machine to one user manner.

The Eucalyptus supports two types of Hypervisors, Xen and Kernel-based Virtual Machine (KVM). Both of them can run on the machines on the CRON Testbed. After considering the hardware compatibility and average performance, we decide to use the Xen. The following 2 advantages make Xen the winner.

(1) Xen does not need any special CPU Virtualization support. The KVM requires the CPU for its hosting machine to have the Intel VT-x or AMD AMD-V hardware extension to accommodate virtualization while the Xen not only supports the equivalent hardware virturalization but also provides a pure software alternative. Old machines can also be virtualized with the Xen.
(2) Xen has better average performance on the different networks. According to some previous evaluations [1], Xen outperforms KVM on both 10G Ethernet and 10G InfiniBand networks. It helps the VMs on CRON better utilize the 10GE links. It may also help us to develop some new services with the 10G InfiniBand on CRON2 in the future.

References

[1] Recommendations for Virtualization in HPC, Nathan Regola & JC Ducom.

B. Project Participants

PI: Seung-Jong Park

co-PI: Rajgopal Kannan

Graduate Students: Chase Pierson, Mohammed Azad, Cheng Cui, Lin Xue

C. Publications

N/A

D. Outreach Activity

N/A

E. Collaborations

ProtoGENI Project Group at University of Utah

OnTimeMeasure Project Group at Ohio State University

GMOC Project Group at Indiana University

MAX project at PlanetLab Control Framework

Openflow Switch Project Group at Stanford University

F. Other Contributions