

Experimental Needs for High Performance Network Research

Martin Swany

1 Introduction

My group's work generally focuses on issues around, and techniques for, high-performance data movement. We have worked on many public experimental facilities as well as constructing our own testbeds. We have had many successful endeavors, but have also had our share of frustration. Striking the balance between packaging a common service and providing the level of customization that we require is clearly challenging. This position paper describes some of our needs for an effective experimental environment.

2 Bare Metal Access

Our work is generally focused on high-performance data transfer and virtualization generally adds unacceptable overhead. Even in cases in which we work to quantify and mitigate I/O related overhead from virtual environments, it is necessary to have control over the hypervisor.

Even in the case of bare-metal access, we often require full control over the operating system version, packages installed, and configuration.

3 Layer-2 Topologies

We have done various experiments with protocols other than IP. To do this outside of our lab, we need to be able to assemble Layer-2 network topologies. We have done this in various ways in the current GENI testbed and over Internet2 ION and AL2S and with ESnet's OSCARS.

Despite a great deal of effort by a number of talented people, this remains problematic. There are often manual steps and the services are fragile. There is a consistent gap between what we understand networks to be capable of doing and what is straightforward to offer as a service.

4 High-performance specialized hardware

We have worked with various pieces of specialized network hardware. In the case of FutureGrid, we were the primary user of the network impairment device. When we first started, we had a great deal of dialog with the operations group explaining that their control model made it impossible for us to use for experimentation.

At the same time, there are other devices in e.g. GENI that we would like to use but which seems extremely difficult. There are e.g. some programmable network resource that are ostensibly “in GENI” but which closer inspection seems to indicate are impossible to really use by anyone other than the resource owner. Again, this is not a situation with an easy answer.

5 Long-running services

One of the considerations for the NSF Cloud is the model under which it will support long-running network services. We have been working with the USGS to distribute satellite data using our NSF-funded network storage system, the Data Logistics Toolkit. Using the current GENI infrastructure has proven to be difficult if not impossible.

6 Measurement Capabilities

Despite personally working on instrumentation and measurement components in GENI for some time, there are still many situations in which the data available to us is insufficient for publication. The type of systems work that we engage in requires deep instrumentation, and even then, needs careful experiment construction to lead to any real understanding. This is an issue that we see across various infrastructures.