On-Demand Research Networks: End-to-End Virtualization in GENI

Michael Zink
Department of Electrical and Computer Engineering
University of Massachusetts Amherst
Amherst, MA, USA
zink@ecs.umass.edu

I. INTRODUCTION

Innovation is defined as a process that transforms new ideas into products, services, or processes. In order to demonstrate innovation in computer networks one has to go beyond the traditional approaches of theoretical analysis and simulation. It is necessary to demonstrate the transformation of new ideas through the implementation of these ideas in an actual network. A successful demonstration with real components, real data, and, especially, real users will reveal the innovative potential of new ideas. As in many other cases, this last step is always the hardest, especially, if the innovation is concerned about new mechanisms for the future Internet. Until recently experiments that demonstrate new ideas had to be executed in the actual Internet or in smaller, isolated "lab environments". Unfortunately, performing experiments in today's Internet is not always possible due to the following reasons. First, unintentional misconfiguration or misbehavior could have a dramatic effect impacting the Internet's performance. Second, certain new ideas, e.g., a new routing protocol or an alternative to IP, can simply not be tested in the existing infrastructure. The goal of the Global Environment for Network Innovations (GENI) of the National Science Foundation (NSF) is to create an infrastructure that will allow the demonstration of new ideas in large-scale networks.

II. EXAMPLE

The following example demonstrates how GENI can support network innovation for research in the area of closed-loop sensor networks. In recent years, we have seen sensor network become more and more an integral part of the Internet. These cyber physical systems tightly connect their computation elements with the control of their physical elements by establishing a closed-loop between the two. They are called "closed-loop" because data generated by the sensors is subsequently analyzed and then used to determine future actuation of the sensors. Often applications that run on top of closed-loop sensor networks are time critical. E.g., they are used to track severe storms with weather radars or objects with PTZ cameras. The existing public Internet infrastructure is not suited well for such networks since guarantees on network throughput and delay cannot be given. Having access to networks that would provide low-layer and performance isolated services between the sensor nodes and computation elements (which do not necessarily have to be located close to the sensors) would allow the evaluation of new mechanisms that satisfy the sensor network requirements. With these services new transport mechanisms (on top of layer 2) can be implemented, tested, and demonstrated. Next to the networking services mentioned above compute services to analyze sensor data are also required. Natural candidates for such compute services would be compute clouds since they allow for temporary allocation of computational and storage resources. A requirement that cannot be met by existing compute clouds so far is the provision of guarantees.

In GENI, we are in the process of creating a closed-loop sensor network that is open to the research community for experimentation. In the Virtual Sensing Environments (ViSE) project [1] we have created a sensor network that currently consists of

three nodes, each equipped with a weather radar, PTZ camera, and weather station. Researchers can gain access to this test bed via a portal[1]. At GEC7 we have shown, in a combined demonstration with other research groups in our cluster, the ability to transport data from our test bed via a dynamic VLAN to a compute cloud at Duke/RENCI, where the data analysis was performed. The resources required to execute this demonstration were dynamically allocated through our cluster's control framework.

In a second GENI project, Data-Intensive Cloud Control for GENI (DiCloud) [2], we aim to develop a complete environment for researchers to conduct data-intensive experiments in GENI from start (the data collection point) to finish (processing and archiving). Data-intensive requirements of high bandwidth sensor network platforms can benefit from tight integration with both the compute and storage resources offered by cloud computing platforms, and the emerging "big data" cloud software platforms.

With the combination of both projects a researcher can allocate all necessary resources required to establish a closed-loop sensor network for environmental sensing. Not only can a researcher allocate these resources but the virtualized nature of the resources offers a variety of options to evaluate new mechanisms and techniques in the test bed. The isolation between slices in the test bed prevents that an experiment in one slice influences an experiment in another slice.

With the presented approach, experimenters have the opportunity to obtain access to a closed-loop sensor network test bed on an on-demand basis, allowing them to evaluate new mechanisms without having to create their own test bed.

III. Future Infrastructure

After almost 2 years of creating the GENI infrastructure through test beds, control frameworks, and new applications initial but very basic experiments have demonstrated the capability of GENI to support network innovation. To expand experiments and demonstrate innovations on a larger, even global, scale the following challenges have to be addressed in the coming years.

- **Users**: The existing and still growing GENI infrastructure is designed to be available for the larger research community. It should be used by a large group of experimenters allowing them to evaluate new ideas and approaches in real world scenarios. This group should not be limited to networking researchers but should also include scientist that use sensors or other infrastructure for their research.
- **Expansion:** The existing GENI infrastructure should be constantly extended to allow for large-scale experiments. The planned "GENI Racks" initiative will make an important contribution to this goal in the area of nationwide, wired infrastructure. But besides an expansion of the backbone the leaves should also grow in number. One potential expansion would be the inclusion of further closed-loop sensor networks.
- **International federation**: For very large-scale experiments it would be desirable to allocate resources on a global scale. E.g., an atmospheric scientist in Germany might be interested in analyzing data generated by weather radars located in the US. This would potentially require the federation of several control frameworks to allow the scientist to allocate the required sensing, networking, and computing resources.

IV. Bibliography

[1] D. Irwin, N. Sharma, P. Shenoy, and M. Zink. Towards a Virtualized Sensing Environment. In *Proceedings of 6th International ICST Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities*, May 2010.

[2] D. Irwin, P. Shenoy, E. Cecchet, and M. Zink. Resource Management in Data-Intensive Clouds: Opportunities and Challenges (invited paper). In *Proceeding of the 17th IEEE Workshop on Local and Metropolitan Area Networks*, May 2010.

[1] http://geni.cs.umass.edu/vise/