NowCasting: UMass/CASA Weather Radar Demonstration

Michael Zink
Department of Electrical and Computer Engineering
University of Massachusetts Amherst
Amherst, MA, USA

David Irwin, Emmanuel Cecchet, Prashant Shenoy
Department of Computer Science
University of Massachusetts Amherst
Amherst, MA, USA

I. Introduction

The NSF's Global Environment for Network Innovations (GENI) project's goal is to create a virtual laboratory for the exploration of future Internets at scale.

At presents researchers from all across the United States are in the process of creating testbeds, and new network and distributed systems technologies that build the basis of the virtual GENI laboratory.

After three years of development users are now starting to use this infrastructure for the exploration of new ideas in the area of computer networking and distributed systems. During GENI Engineering Conference 9 in November 2010 some examples of how the GENI infrastructure can be used for research were given through a series of demonstrations of a variety of applications.

In our presentation, we will give an overview on our demonstration of a short-term weather forecasting application that runs in the cloud. The demonstration is an end-to-end demonstration that stitches together sensing, networking, and computing resources. Data needed as input for the short-term forecasts (also called Nowcasts) are generated by weather sensors (in this specific case CASA weather radars). The data are then transmitted over dedicated layer 2 links to remote compute resources, and the final Nowcasts are made available to end users via a web portal. The heterogeneous resources required for this application are reserved by using the GENI/Orca control framework.

II. Demonstration Overview

The primary purpose of our GENI alpha demonstration was to exploit GENI's sliceable heterogeneous computing and networking infrastructure to improve weather "nowcasting," i.e., short-term weather forecasts. A key goal of this demonstration was to emphasize the benefits of the Orca control framework architecture's support for heterogeneous resources, which includes high-bandwidth sensors, network switches, and computational elements, amongst others.

Nowcasting differs from forecasting in its focus on highly-accurate short-term forecasts starting at 1 minute in the future, and are particularly critical in severe weather situations to provide advance warning as soon as possible. Nowcasting is both bandwidth- and computationally-intensive, since it must simulate future storm movements using real-time data feeds from active radars as quickly as possible. For example, for a 1 minute Nowcast to be useful, the prediction must be available in almost real-time: a 1 minute Nowcast that takes 2 minutes to compute has little or no value. The hypothesis of our Nowcast experiment is that GENI enables a more efficient short-term severe weather forecasting system by triggering the creation of on-demand computational and networking resources to sense, transmit, and process radar data as severe weather approaches.

Our demonstration proposal builds on the successful Cluster D plenary demonstration at GEC7, which utilized both networking resources (Orca/BEN/iGENI), as well as sensing and computational resources (ViSE/DiCloud). In that

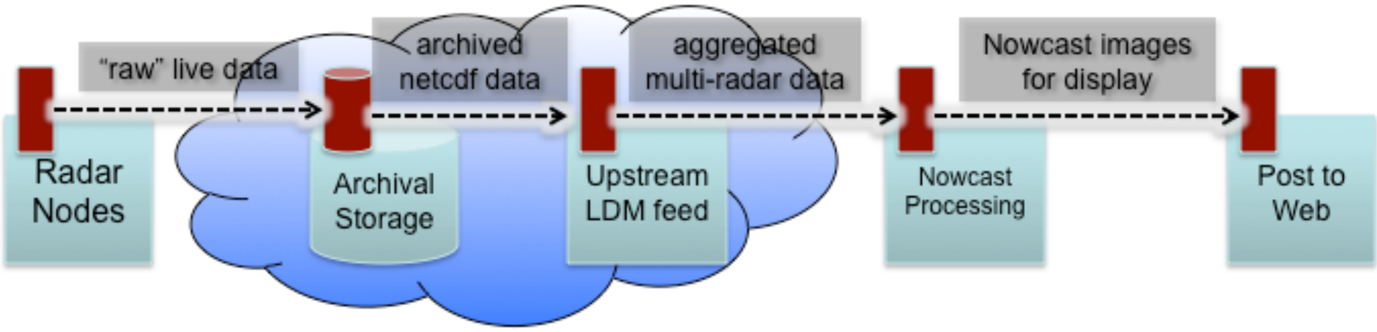

Figure 1 Mapping Nowcast workflows onto GENI

demonstration, GENI Cluster D showed the coordinated allocation and stitching of an isolated VLAN across multiple points-of-presence, including UMass-Amherst, RENCI, and Duke University. However, our use of the VLAN was simply to display a visualization of radar data from a ViSE radar. In the GEC9 alpha demo we built on this demonstration by integrating computation from real-time Nowcasting of high-bandwidth data fed over the VLAN. Since severe weather is generally a rare occurrence, GENI's ability to provide on-demand network and computing capacity in the event of a storm is important: dedicating such expensive resources for relatively rare events is not cost-effective. With GENI, as severe weather forms, both network and computational resources may be shifted from lower-priority tasks to high-priority forecasting tasks. Importantly, radar systems require both a mix of high-bandwidth networks and substantial computation. As noted above, the more network resources that are available the lower the latency for starting Nowcasting, and the more computational resources that are available the more parallel Nowcasts, e.g., 1, 5, 10, 15-minute Nowcasts, are able to execute. **Figure 1** illustrates how we mapped the Nowcast workflow onto GENI. An overview on the network resources used for the demonstration and the geographic locations of each single component are shown in **Figure 2**.

Below we briefly outline a few additional aspects of the demonstration:

- **Amazon Resources for Overflow.** Amazon EC2 was used to archive the data generated by the radars. The DiCloud software was used to allocate the EC2 resources and monitor the occurring cost caused by the resource usage.
- **Nowcast Visualization.** Nowcast data was visualized on a Google Maps overlay, and highlighted both the forecast data as well as the observational data that matches with the forecast time, to point out severe weather events. As a result, viewers should be able to discern the accuracy and speed of the Nowcasts.

III. Acknowledgements

This demonstration would not have been possible without the support from Jim Kurose, Eric Lyons and Jorge Trabal (UMass); V.Chandrasekar, Evan Ruzanski, and Yanting Wang (CSU); Ilia Baldine, Jeff Chase, Anirban Mandel (Duke/RENCI); Sandra Cruz-Pol, Jose Colom, Gianni Pablos (UPRM); Mark Berman (GPO).

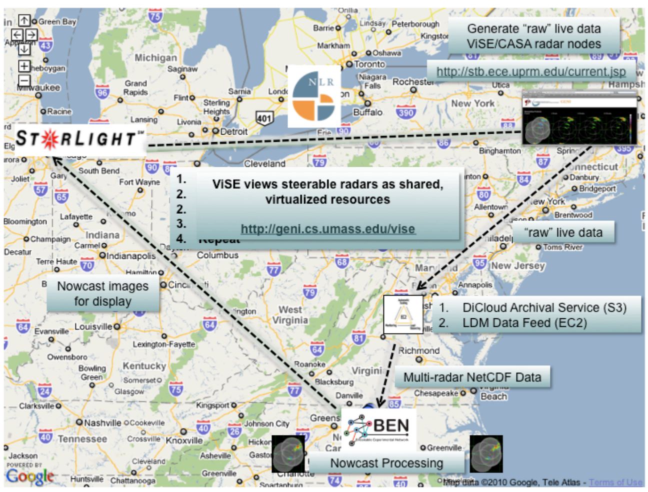

Figure 2 Networking topology for demonstration