Flash I/O Network Appliance (FIONA) Connected to the 40Gb/s PRISM Network at UC San Diego for Worldwide Access to the IDD

DeFanti/Graham Report to Unidata for UCAR subcontract (Z15-12798)

5-26-2016

The result of this funding to UC San Diego has been to improve collaboration with the Unidata community by installing a Flash I/O Network Appliance ("FIONA")-based AWIPS2 / EDEX server on the 40Gb/s PRISM Network (funded by NSF Award #OCI-1246396) to make the data available to researchers worldwide. John Graham, Senior Development Engineer has been the main builder and integrator of this equipment and software into the ultra-high-speed networks. Tom DeFanti, PI, has supervised the effort.

The EDEX is currently processing:

- Maximum hourly volume: 36,440,360,000 bytes/hour (>36 GB/hr)
- Average hourly volume: 16,892,950,000 bytes/hour (~17 GB/hr)

The EDEX is connected to the Dell S4048-ON switch for use in our SDN experiments, using the Chelsio 40G NIC which supports RDMA offloading. The EDEX server is based on a well-tested multi-year deployment of 10 FIONA dual 40Gbps Haswell CPU servers. These Linux computers have been placed all across California and to Seattle and Chicago on the CENIC/Pacific Wave networks to measure the performance of their 100Gbps network infrastructure. This FIONA configuration is daily tested to be capable of sustained dual stream flows of 37Gbps out of 40Gbps, or 9.7Gbps out of 10, which eliminates bandwidth bottlenecks in the networking and storage system and has allowed us to optimize the data production and dissemination to the research community.

The exact configuration of the EDEX server, as bought is:

- Supermicro 3U Chassis
- X10DAX-O Server Motherboard
- Dual Intel Xeon E5-2687W v3 Haswell 3.1GHz 10 core 160W CPUs
- 256GB DDR4 RAM
- Dual 512GB SSD BOOT Drives
- 9.6TB 10K RPM SAS3 Drives
- 2.4TB PCIe NVMe
- LSI 9300-16i PCIe 3.0 SAS 12Gb/s SAS Host Bus Adapter
- Chelsio T580-LP-CR 40GbE NIC
- Mellanox ConnectX-3 Pro EN 40GbE NIC
- Trusted Platform Module (TPM)

We make the EDEX products available to remote users who want to use the CAVE thin client to visualize the weather products. A CAVE thin client can be accessed using VNC

remote desktop to a Centos 6.6 Virtual Machine in our TelaScience XenServer stack in the Terascale data center at UC San Diego, Qualcomm Institute. We also provide the CAVE client as an OVA file (Open Virtualization Archive). This allows researchers to spin up a Centos 6.6 Virtual Machine fully configured with the AWIPS2 software repository on any virtualization platform, for instance, Xen, VMware KVM, and others.

The FIONA system is directly connected to the PRISM network with dual 40G network connections, with which offer access to the UC San Diego 10Gbps Production Network, XSEDE Resources (https://www.xsede.org/), and off-campus 100Gbps networks like CENIC, PacificWave, ESNET and Internet2. International networks are also reachable.

We have integrated the EDEX into a wildfire simulation pipeline that runs the Farsite/ Wind Ninja ensemble simulations. This supports the NSF-funded WIFIRE program (http://wifire.ucsd.edu/, NSF Award #1331615 under CI, Information Technology Research and SEES Hazards programs). This effort is building an end-to-end cyberinfrastructure (CI) for real-time and data-driven simulation, prediction and visualization of wildfire behavior. See Figure 1 for a screen shot.

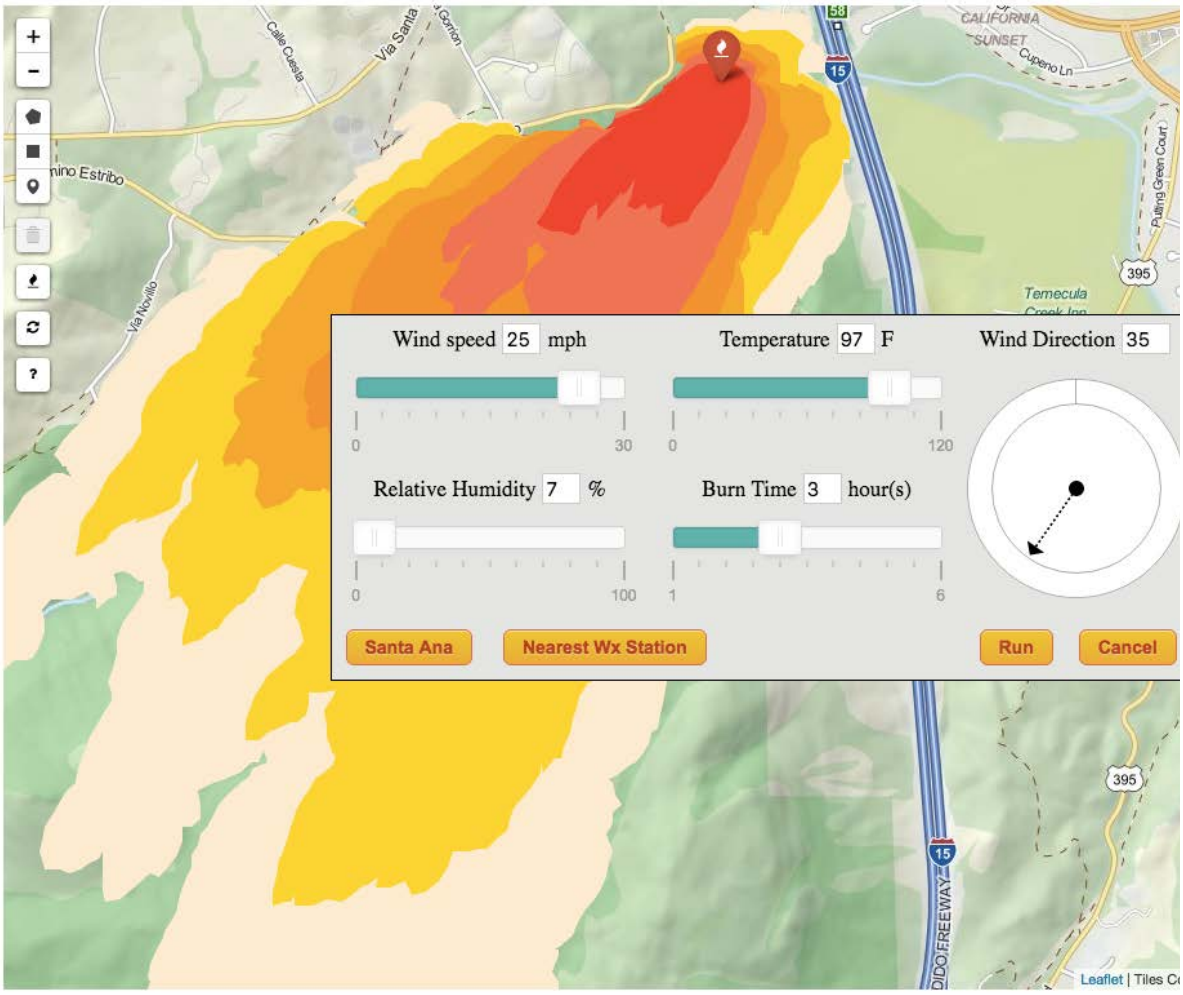

Figure 1. Screen shot of Farsite / Wind Ninja model web interface. The wind speed and direction, temperature and relative humidity are set manually in this interface.

We are finalizing the interface where initialization data is pulled directly from the HRRR forecast model on our EDEX using the Python AWPIS API. We are also working with Michael James at UCAR to implement the 15-minute interval HRRRfsl forecast model which is available, as of May 20, 2016 to our EDEX. This will allow firefighters to run what-if scenarios with greater precision using real weather forecasts during wildfire events and for training. One of the most exciting recent developments is the availability of Jupyter Notebooks on this project. These python-based notebooks are hosted in a web based multi-user platform with a local 5,000-core Nvidia K80 GPGPU for machine vision and deep learning frameworks. This will greatly simplify sharing of research and eliminate the installation of special software. See Figure 2. Classroom use of the ~17 GB/hr coming from the EDEX server is naturally expected, but will not be tracked.

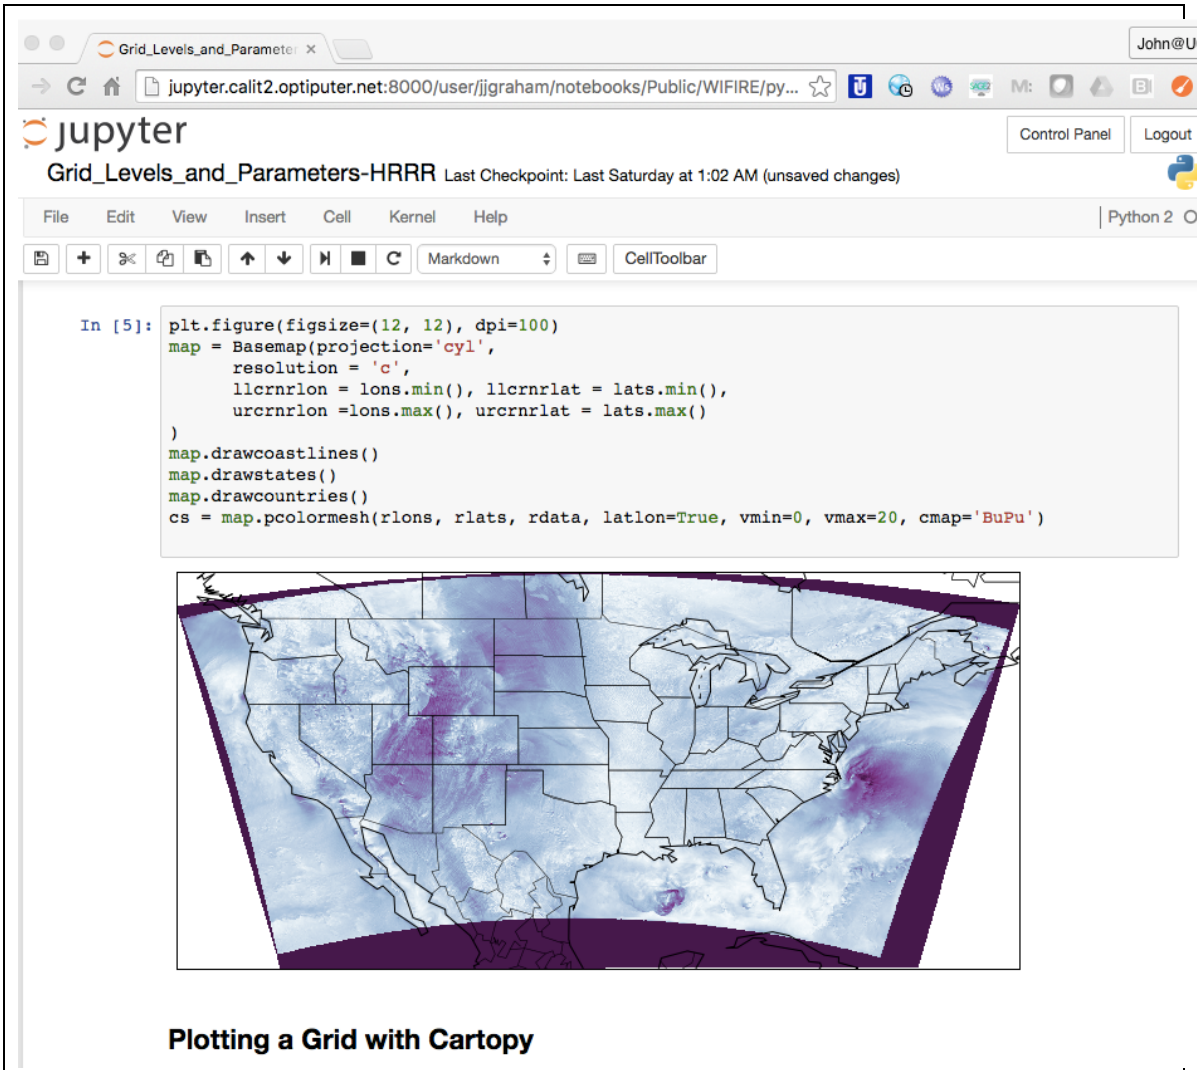

Figure 2. Python-AWIPS Jupyter Notebooks are now available providing new tools to explore and develop even more compelling uses of the EDEX system.