Environmental Data Exchange Server for Wildfire Simulation and Visualization

Principal Investigator: John Graham*
Co-Investigators: Ilkay Altintas^ and Jessica Block*

*California Institute for Telecommunications and Information Technology
Qualcomm Institute and ^San Diego Supercomputer Center
University of California, San Diego

Overview

Wildfire models require many disparate datasets to initialize and update to compute. The core data components include elevation, vegetation, and weather, yet none of these data are readily available at fine spatial scale at relevant temporal scales to study and respond to real fire events. Two Cyberinfrastructure (CI) efforts are using these available datasets and combining them with computational tools that assimilate available real-time data into physical wildfire models often comprised of computational fluid dynamics (CFD) and weather modeling tools. One of these efforts, called WIFIRE, was awarded by NSF to University of California at San Diego (UCSD) to build an end-to-end cyberinfrastructure (CI) for real-time and data-driven simulation, prediction and visualization of wildfire behavior. WIFIRE partners with in this efforts with TelaScience, a joint cyberinfrastructure collaboration between SDSU, UCSD's Calit2, and the global open source GIS community. The effort to provide real-time global environmental data for disaster relief has been pioneered by TelaScience, lead by John Graham, but has hit its limit of data ingestion to global environmental models. With data the US Government has made available in the last year, computing requirements have surpassed the TelaScience server capabilities. Additionally, through this partnership, after a review of wildfire data and behaviour modeling state-of-the-art, we have discover that this research community would greatly benefit from high-end computing platforms. Therefore, we are requesting cycles on Gordon and Comet to design a system that can deliver the required cyberinfrastructure and tools to the wildfire research and disaster response communities.

Research Objectives

Our goal is to improve wildfire spread simulations by improving the accuracy of the data inputs to the models. WIFIRE hosted a 2-day workshop at UC San Diego January 12-13, 2015 by some of the leading wildfire modelers in the world. One glaring need described by the community is access to environmental data for their models. There has been tremendous progress in data being made available measuring the fire environment on earth. For instance, a startup satellite company called Planet Labs has deployed 28 micro satellites that will collect enough imagery equivalent to a 10-terapixel image each day. In the near future, we will be using the Planet Labs Flock to do change detection to find burn scars and to create simplified landcover products (due to its limited number of spectral bands). We were recently approved in the first group of the Planet Explorers program that gives us unlimited access to their API and content.

Thanks to the recent such advances in access to data, the great new computing challenge now is ingestion, integration and standardization of these large datasets and then be able to run existing models on high resolution inputs. The reward, however, will be a greater understanding of fire behaviour and more effective emergency response for vulnerable cities worldwide.

Data and Models

We have ambitious plans to integrate many large datasets. Satellite imagery from NOAA's GOES Satellite and NASA's MODIS and NPP Satellites are useful for fire event detection using the thermal signature. High resolution radar and local weather station data are used to improve the accuracy of atmospheric and fire models. Multi-spectral satellites like MODIS and Landsat create products such as combustible fuel load index and enhanced vegetation index to provide greater detail about the combustibility of the fuel on the ground across the area being modeled. Landcover datasets, currently static and updated every 5-10 years can be updated daily with these satellites as well.

Atmospheric weather forecasts are used to initialize fire propagation models. The Weather Research and Forecasting model (WRF) is the standard model used Nationally that has many product extensions including a fire propagation model called WRF-SFire. The code has already been designed to be run on MPI machines. The code developed by the TelaScience group uses Portable Batch System (PBS) to achieve the scaling of the underlying libraries. Other models we use are Farsite (US Forest Service) and Firefly (University of Maryland) for fire propagation, and Wind Ninja (wind assimilation). A new satellite, called Soil Moisture Active Passive (SMAP), is launching at the end of January and will provide soil moisture in the top 2 inches of the ground. A San Francisco company called Planet Labs has launched Flock-1. At this point 71 Dove satellites have been launched and have started producing a global 5 meter mosaic once a day then combine them into a cloud free mosaic that is constantly updating. Having access to this data will allow us to create daily updated burn scar and landcover datasets which will improve the accuracy of the models.

UCAR, the consortium of universities focused on atmospheric and earth systems, hosts a data portal called Unidata via what they call the Environmental Data Exchange Server (EDEX). The EDEX needs to run in a virtual machine because of all the daemon processes required to run 24/7. Also Unidata provides the code in a Centos 6 repository making the installation and updating of the code seamless. This is the first step to producing a Science Gateway for Wildfire Modeling, which does not currently exist in any format. There is no central repository for fire data, which is why this proposal is so important.

Another speedup in processing can be obtained by processing parts of the model on GPGPU's . The new K80 Nvidia GPUs can be accessed by our pipeline as part of the WRF model code which can be compiled with MPI and OpenMPI using PGI compilers

- http://dx.doi.org/10.1175/JTECH-D-12-00218.1
- http://www2.mmm.ucar.edu/wrf/users/workshops/WS2014/ppts/1.4.pdf

All of these pipelines need to run 24 hours a day to enable a rapid response to a real wildfire event.

Computational Methodologies

We are requesting

- Gordon - 4T SSD scratch space will be needed for some of the larger CFD models
- Comet - Virtual Machines, Tesla K80 GPU's with 4992 Cuda Cores will greatly reduce the CPU load and decrease the time it takes to run models.

We intend to export the Data Oasis storage allocation over the Prism 100GE network via NFS to the TelaScience XenServer stack and WIFIRE server (12 cores 128GB RAM 4TB of RAID 10) to create other downstream products. The 100GE Prism network enables research on other Prism connected campuses.

In this first year we will concentrate on producing the required datasets for initializing the wildfire modeling. In the second year, we will expand our allocation request to support larger domain model runs that are executed directly on compute nodes, keeping the need for virtual machines at a minimum. We will also provide experimental EDEX VMs that can be used to test new advanced processing pipelines. These VMs can be cloned and spun up to run short term experiments that can later be included in the master EDEX without interfering with the primary pipeline. These VMs can also be used as a CAVE client if users want to use the AWIPS II standard weather forecasting display and analysis. Users would be able to access the VMs using VNC remote desktop from any personal computer.

Integration to Existing Cyberinfrastructure

This request will be supplementing existing high performance computing that we will integrate seamlessly. Everything we produce will be exported and mounted as a file system to our other clusters. We will continue to use the TelaScience XenServer stack at Calit2 for running small models during the testing phase and continue to expand and upgrade the older systems with new hardware and storage. We also want to extend the Luster file system on the 35TB Data Oasis allocation request over the Prism network using an NFS intermediary device enabling Virtual Machines running on the TelaScience XenServer to NFS mount the data being produced by the EDEX pipeline.

Currently the EDEX server is available for anyone with a .edu domain via virtual machines on Amazon and on Azure but neither have the radar processing enabled because of the costs of bandwidth. There is interest at Unidata in switching from Amazon and Azure over to Comet so more downstream users and a full radar feed could be enabled. Having the EDEX on the 100GE Prism network will improve access to downstream users of the .edu community.

An EDEX server with all products enabled requires 16 cores with 32GB RAM and 500GB SSD constantly running 24 hours a day. The EDEX has its own scrubbing system to remove data that has completed all the steps in the downstream product pipeline. Currently we are running an EDEX VM on the TelaScience XenServers in the Calit2 Terascale data center. The Xenserver consists of two quad 6 core servers each with 128GB of ram and 48 cores combined. Both connected with a Cisco Nexus 5000 CNA fabric switch. The servers boot from a 6TB iSCSI Dell Equallogic storage device and the data is stored on two 48TB Sun Thors and two 24TB Sun Thumpers. All devices are connected with 10GE CNA's but we continue to be I/O bound. To solve some of the poor performance due to virtualization overhead and motherboard limitations, we use RAMDISK to ingest the data from the Unidata Internet Data Distribution system (IDD). All the processing is done in RAM and only the output data is written to the slow storage. Recently we attempted to add SSDs to the XenServers but found we gained no I/O performance due to several factors. The primary one being the motherboard backplane is just too slow as it is now 7 years old.

The TelaScience pipeline uses the same frontend to ingest the data. The IDD sends and the Local Data Manager (LDM) receives and forwards to the processing pipeline. We plan on merging our pipeline code in the EDEX system and contribute it back to the Unidata community. The custom code we have written for visualization of the global weather data from the Unidata IDD via LDM consist of three libraries: libNDIS for NEXRAD radar, libGINI for the GOES weather satellite imagery and grib_countour for vectorizing GRIB data into KML animations for playback on google earth. Previously we had an allocation on Dash (Pre-Gordon) that used to execute wrapper code we write to run the WRF-SFire Wildfire model using PBS job scheduling. This code will allow us to run the actual model runs on a compute node outside of a Virtual Machine that spawned the job.

The ultimate goal is to have an automated system that can launch a simulation based on satellite detections, 911 reports and manually input ignition events.

Appendix

Currently the TelaScience system is processing less than 1/10th of the full feed from the IDD feed. The virtual machine consumes 8 cores and 16GB of ram.

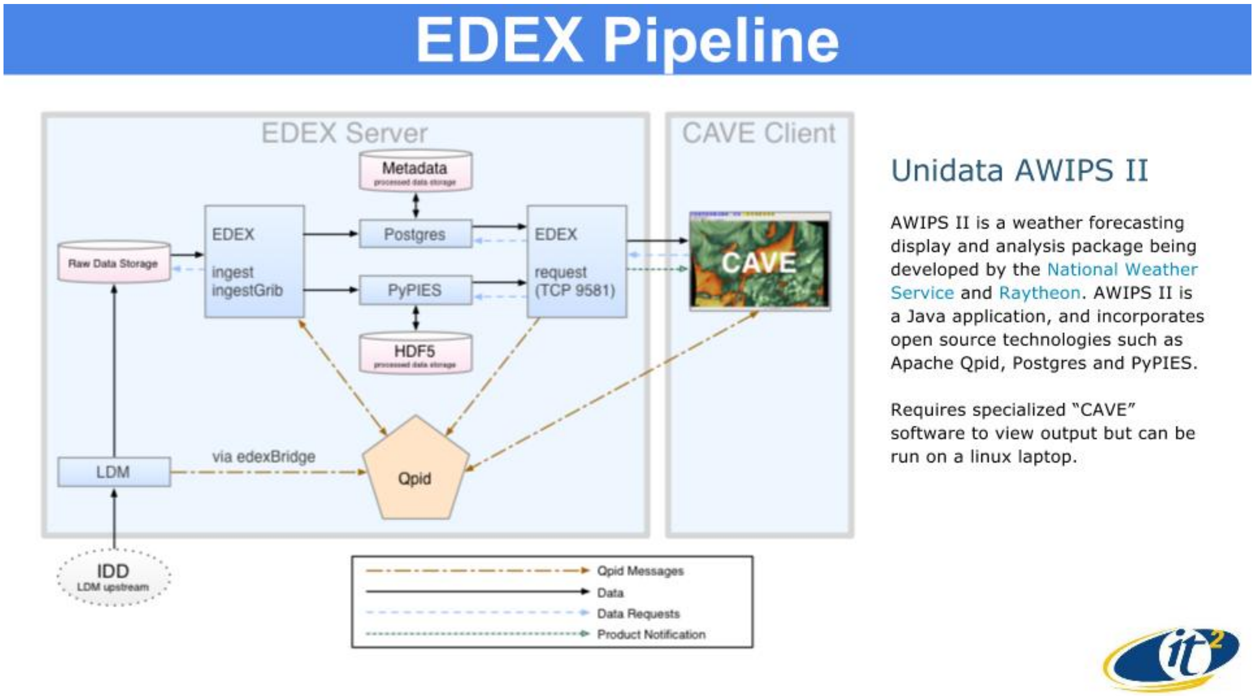

Figure 1. EDEX

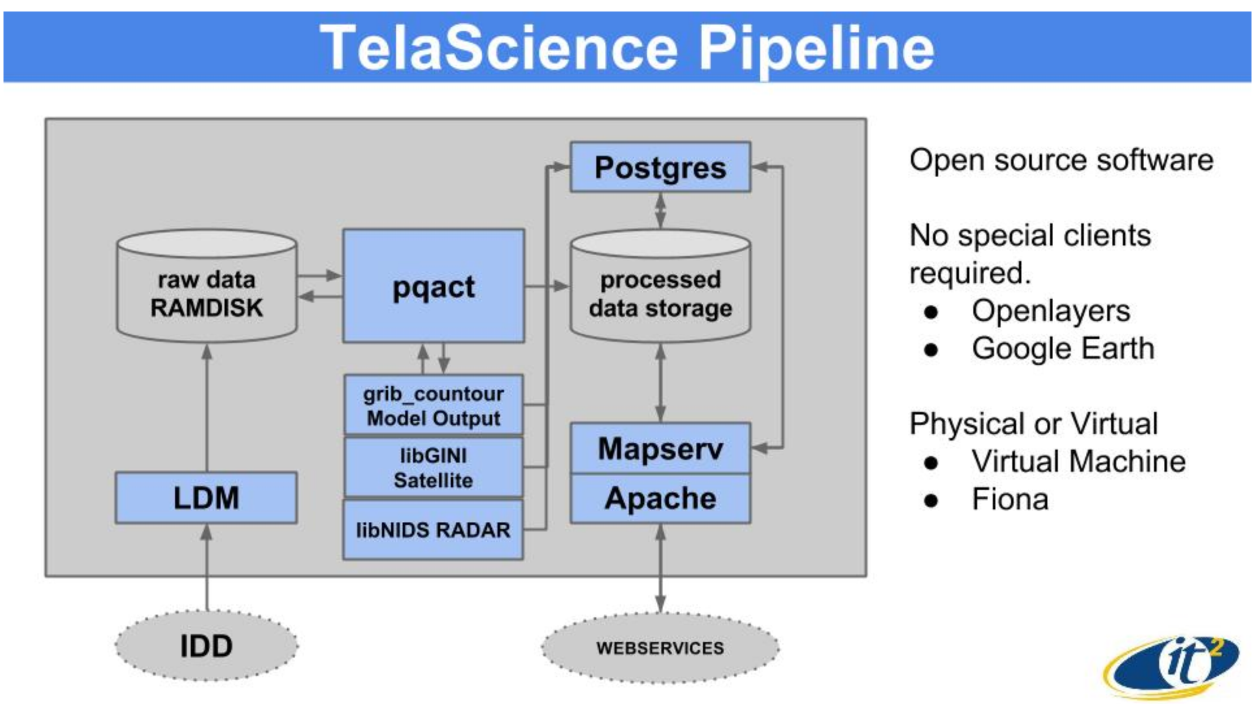

Figure 2. TelaScience

Current Feed	Average (M byte/hour)		Maximum (M byte/hour)	Products number/hour
CONDUIT (NCEP)	3745.433	[38.138%]	14655.975	76891.865
NGRID (NOAA High Res Models)	3046.015	[31.016%]	12497.987	22951.295
NEXRAD3 (radar)	1395.999	[14.215%]	5040.291	83128.71
FNMOC (Oceanography)	1123.077	[11.436%]	6834.16	3064.269
HDS (High Resolution Data Service)	357.259	[3.638%]	727.038	18090.006
NIMAGE (NOAA Satellites)	153	[1.558%]	598.423	198.825
LIGHTNING	0.01	[0.000%]	0.301	23.495

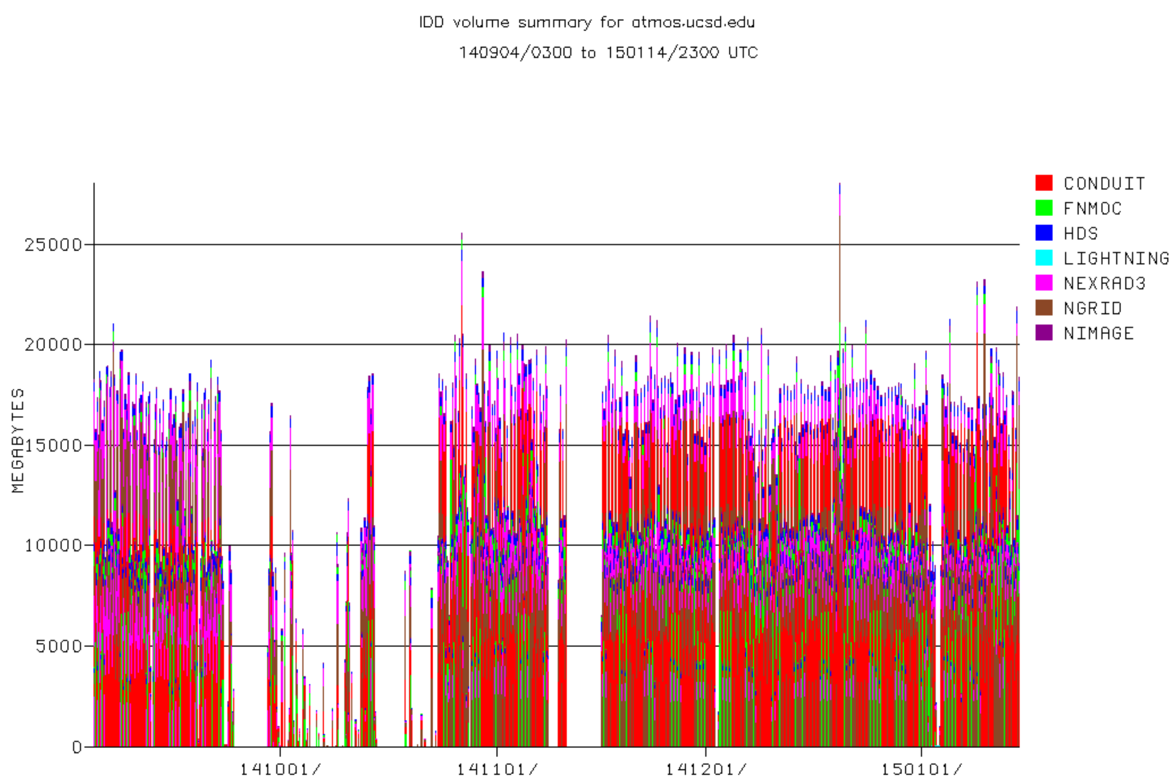

Figure 3. Unidata IDD feed
http://rtstats.unidata.ucar.edu/cgi-bin/rtstats/rtstats_summary_volume1?atmos.ucsd.edu+GRAPH