

December 2011

Data Analysis Services Group - *December 2011*

News and Accomplishments

VAPOR Project

Project information is available at: <http://www.vapor.ucar.edu>

• TG GIG PY6 Award:

Yannick and John continued to try to resolve performance issues with PIOVDC. After much troubleshooting and development of simple benchmark codes it was determined that the underlying netCDF4 library was the culprit: the library simply does not perform well with even the most trivial of cases. The PIOVDC code was migrated to pnetCDF (pnetCDF is an alternative to netCDF4, produced by Argonne Natl. Lab), and performance of raw netCDF test cases (test cases write data directly to a netCDF file, bypassing translation to a VAPOR VDC) improved by an order of magnitude. When data is written to VDC performance is still suboptimal, but this is believed to be due to a simple buffering problem that we hope to resolve shortly.

• XD Vis Award:

Yannick continued work VAPOR's ShaderMgr to provide support for GLSL compilation pre-processing (e.g. #ifdefs).

• KISTI Award:

Year one work was wrapped up in November. The KISTI project will resume in the spring. We have received a verbal confirmation from KISTI that they will fund year two of the project.

• Development:

With the year one KISTI work behind us we resumed efforts to prepare the 2.1 release of VAPOR:

- All known bugs in the 2.1 release candidate designated "must fix" (over 30 bugs) were fixed. Notable defects included:
 - There was a problem with our extensibility API: When authors of VAPOR extensions produced new releases, these could not easily be backwards compatible with the previous release. We modified the API so that such extensions can easily become backwards compatible.
 - The global and U.S. geo-referenced map images produced by NCL were not correctly registered. A workaround to address NCL's lack of a direct raster output had to be devised, and new images generated. NCL example scripts demonstrating the workaround were produced.
- Installers for 2.1 were built for Linux and Mac OSX platforms
- Regression testing commenced

We continue to respond to queries on the VAPOR mailing list, and to post and fix bugs.

• Administrative:

The ARRA quarterly report was submitted for the XD Vis award

• Education and Outreach:

John organized and chaired a 1/2 day workshop on visualization tools and technologies at the Fall AGU meeting in San Francisco. The workshop was attended by over 40 people. Alan gave a talk on VAPOR at the workshop, and both John and Alan attended the AGU conference.

• Consulting:

Alan met with Yuan Ho (Unidata IDV), discussed how Vapor performs flow integration. Yuan indicated that they will be providing similar flow integration capabilities in IDV, but not for some time.

Software Research Projects

Feature Tracking: xxx

Climate data compression: xxx

Data Analysis & Visualization Lab Projects

File System Space Management Project

- Continued to work sporadically on the FMU design and documentation.

Accounting & Statistics Project

- xxx

Security & Administration Projects

- xxx

System Monitoring Project

- xxx

CISL Projects

GLADE Project

- Reviewed GPFS 3.4 documentation set for upcoming SPXXL meeting. Focused on 3.4 features and new "GPFS native RAID administration": pdisk, RecGroup, DeclusteredArray, Vdisk, mmexportfs, mmimportfs
- Reviewed DCS3700 document: 4U unit with two 6Gbps 4x SAS host interface + single 6Gbps 4x SAS expansion with optional HIC(Host Interface Card): four-port 8Gbps FC or two-port 6Gbps SAS
 - Features: FlashCopy, VolumeCopy, Enhanced remote Mirroring.
 - Each DCS3700 can support two DCS3700 expansion enclosure (up to 180 drives)
 - Minimum 20 drive (4 each on 5 drive tray) is due to uniform air flow requirement
 - ESMs(Environmental Service Modules) on SBB (Storage Bridge Bay) A and B
 - Firmware determination: view->sybsystem profile (NVS RAM version, Firmware version, drive firmware, ESM card firmware) controller firmware: controller icon -> physical view -> to repeat for each controller

Lustre Project

- xxx

Data Transfer Services Project

- Still awaiting establishment of a production UCAS authentication based MyProxy server for use with the GLADE GridFTP service. I have not heard of any plans to deploy this yet.

Lynx Project

- xxx

Batch Systems & Scheduler Project

- xxx

NWSC Planning

- xxx

Production Visualization Services & Consulting

- xxx

Publications, Papers & Presentations

- xxx

System Support

Data Analysis & Visualization Clusters

- Installed the latest version of the ANSYS Fluent software and licenses.
- Made NCL 6.0.0 the default version in /fs/local/bin.
- Installed the scipy, matplotlib, Basemap, Nio, and numpy packages for Python-2.6.5.
- Troubleshoot for Gary Strand's issue on NCO command jobs (EV 70268). Traced of ncks command behavior: the reported poor performance was due to redundant reading of the same blocks under GPFS. On local hard drives with 8k block size it quickly goes through the filtering operation (ncks -x -v TH), but on GPFS it revisits overlapping regions with 2MB blocks repeatedly. In terms of file system performance, we noticed nothing wrong. The code itself seems to assume the old 8k block while the command automatically inherits the GPFS blocksize from the system library. (read and seek calls with 2MB stride). We suggested that CSG contact the developer Charlie Zender.
- Made backups of the GDS server and its config than runs on **twister1**. **Twister1** and **twister2** will be decommissioned along with **frost** in early 2012.

GLADE Storage Cluster

- Adjusted the incremental backups of /glade/home to start before the full backups at the suggestion of EMC to avoid possible incomplete backup sets. With the backups partitioned into four subsets, the larger subsets are requiring more than 12 hours to complete a full backup. I plan to repartition the backup groups from four to six or seven subsets in January.
- Increased the quota for the data02 dsszone fileset at the request of DSS.
- Removed 119584 empty, old directories from /glade/scratch.
- Made backups of the /var/mmfs directories on the oasis servers.
- Replaced the Power Supply for G channel (DDN9900), F channel showed error but reseating it removed the error lights
- Fan module for controller C showed left fan fail alarm set, but fans are spinning and no immediate overheat sign. Opened a case for FRU swap (DDN #39108), the replacement Fan module came the next day and was replaced.
- Rebuilt disk 44F (oasisd) after failure.
- Responded to SSG request to umount the gpfs_bhome and removed it from mounttab (using mmremotefs delete).

Data Transfer Cluster

- Found and reported a bug in the GridFTP server that caused data mover backends to deadlock and hang around when a striped transfer was aborted. Provided debug information to Globus about the problem.

Other

- xxx