Data Management/Visualization on the Grid at PPPL

Scott A. Klasky
Stephane Ethier
Ravi Samtaney

PPPL
PRINCETON PLASMA PHYSICS LABORATORY

The Problem

- Simulations at NERSC generate GB's – TB's of data.
- The transfer time for practical visualization is too long.
  - Using "standard" ftp, we get 600 KB/s.
  - GTC data is over 3TB, transfer time = 2 months.
- Interferes with productivity of our research program.
  - Leaves valuable data unexamined.

PPPL PRINCETON PLASMA PHYSICS LABORATORY

Visualization/Data Management

3 Categories of data and techniques

- Small (less than 1×10^6 mesh points) [1D/2D + Time] <200MB/Run.
  - Often consists of many signals; commonplace in experimental datasets
  - MDS+ is used for this type of data
    - Grab data from a server to local client
    - Since the dataset is small, the transfer time is small.
  - Scalable display walls allow one to present many functions at one time. [NOTE: NO NEED for a cluster to drive wall for <20 screens]. KEEP THE COST DOWN!
  - IDL is de-facto standard for this type of viz. in fusion community.
  - We are developing Elvis as part of the Fusion Grid (http://www.fusiongrid.org). [Feibush/Klasky/McCune]
    - Java.
    - Fully collaborative with a multi-tier architecture; send the minimal amount of information to collaborate.
    - Web-based graphics.
    - Links to IDL math routines.

Visualization/Data Management

- Medium (less than $10x10^6$ mesh points) [Typically 2D/3D + Time] [<50GB /simulation]
  - Can be analyzed on **1** PC
  - Use HDF5.
    - Parallel I/O with MPI2
    - Highly supported by the "general public"
    - FAST! (look at http://w3.pppl.gov/~pletzer/performance/hdf5-netcdf/index.html
  - Data is analyzed many times, so grabbing data for every signal from a network does not make sense; 1 minute transfer medium 10 minute transfer for large signals.
  - Locality of data is CRUCIAL for this type of data!
  - Use AVS/Express 6 ($500/license –onetime cost, with a $2,000 yearly maintenance fee)
  - We continue to try other viz. products, but they do not meet our needs (IDL, OpenDX, SciRun,. VTK, Amira)

PPPL PRINCETON PLASMA PHYSICS LABORATORY

Visualization/Data Management

- Large (> 100×10^6 mesh points);
  - [3D + Time] [n-TBytes/simulation]
  - Key area of research, which must use parallel visualization!
  - Locality of data is even more CRUCIAL for this type of data!
  - Remote Visualization is common for this type of data.
    - Possibly look at streaming-video running from some viz. supercomputer to analyze results.
  - We often give up features, for usability.

PPPL
PRINCETON PLASMA PHYSICS LABORATORY

Our “Proposed” Solution

- *Data Grids*:
  - Support for existing AVS/Express-based data analysis infrastructure at PPPL that depends on having local copies of the data.
  - This entails leveraging on the emerging DOE Science Grid infrastructure to provide efficient automated data migration and mirroring capabilities between PPPL and NERSC.
- *Remote/Grid-Distributed Visualization*:
  - Build upon the tools that are already inside the Cactus framework
    - NERSC/NGI-developed Visapult system
  - Provides support for interactive remote visualization of large fusion datasets.
- *Advanced Display Technologies*:
  - Leverage existing advanced display technologies
    - High-resolution tiled display walls assembled at both PPPL and NERSC
  - Build tools for parallel data analysis.
    - PPPL is working with IBM research fo a parallel MPI version of OpenDX

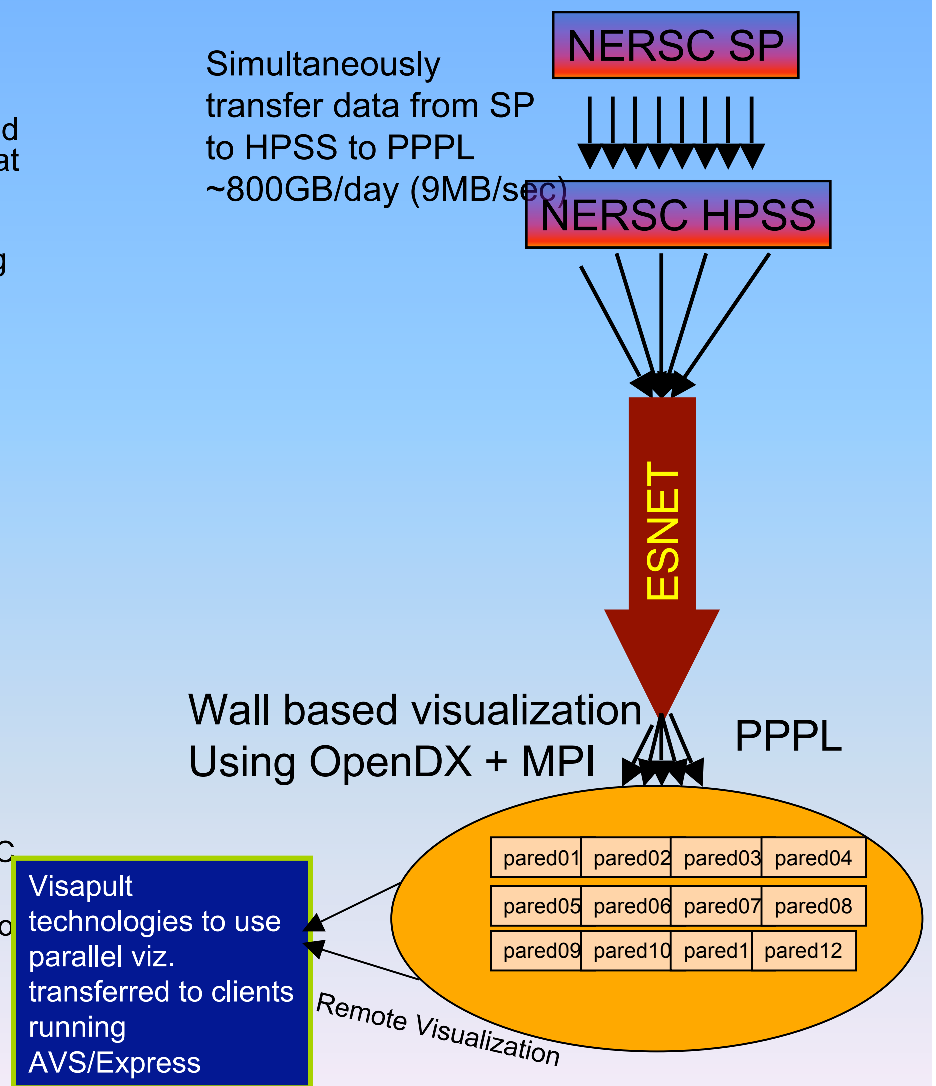

What we need from MICS/DOE

- 2 CS PostDocs — 2 years
  - 1 at PPPL, 1 at NERSC
- 2 Graduate students (Picasso program)
  - 1 at PPPL, 1 at NERSC
- OC12 ($100K)

Our Definition of a Princeton Grid Computer

- A group of “trusted” cluster computers who *WANT* to share resources at certain times.
- Standard Operating System
  - Linux
  - Keep directory structures similar
  - Parallel Virtual File System (separate on each cluster)
  - Keep firewall interaction at a minimum!!
- Part of TCF
  - 19 node (38 AMD processors); gigabit network. PPPL
  - 10 node (20 AMD processors); gigabit network. P.U.
  - 10 node (20 AMD processors); gigabit network. GFDL.

PPPL
PRINCETON PLASMA PHYSICS LABORATORY

Exploratory Analysis Grid Linked Environment (EAGLE)

- Objective: Utilize the Cactus computational framework to create a comprehensive grid-based simulation and analysis environment.
- Provides important capabilities such as
  - Rapid prototyping
  - Parallelization
  - Architecture-independent checkpointing/ remote visualization/ computational monitoring
  - Computing in a distributed heterogeneous environment.
- Future realistic simulations of fusion devices may require more computing power than any single machine can provide
  - Enables the creation of large-scale multi-physics codes that span the spectrum from small-scale microturbulence to device-scale simulations
  - A move to the distributed/metacomputing paradigm can help address these needs.
    - Executes on wide-area distributed heterogeneous computing resources.
    - Complement the existing SciDAC funded national fusion collaboratory.

PPPL
PRINCETON PLASMA PHYSICS LABORATORY