



Cornell University  
Center for Advanced Computing

# Introduction to Visualization on Stampede

Aaron Birkland

Cornell CAC

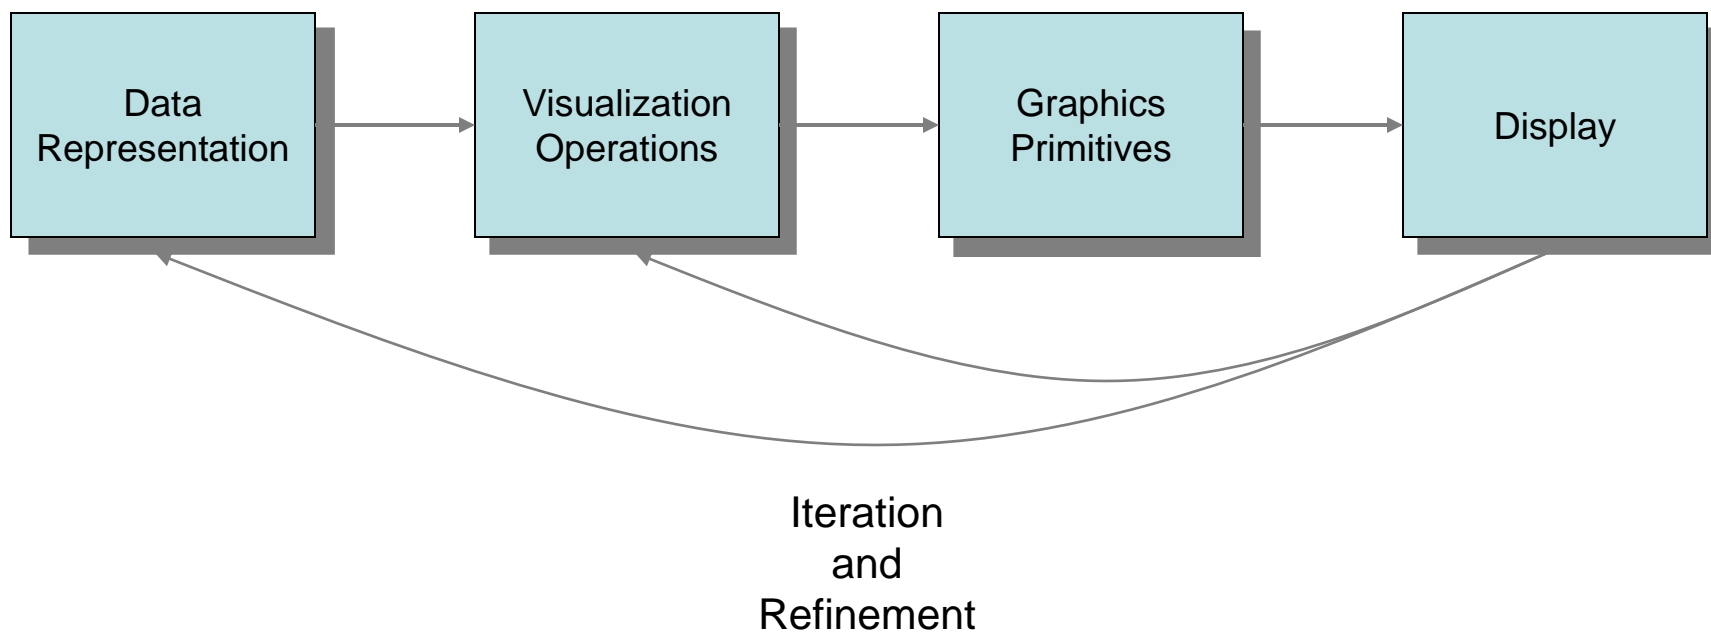
With contributions from TACC visualization training materials

*High Performance Computing on Stampede*

*January 15, 2015*




# Interactive Visualization





## Large data, Remote Systems



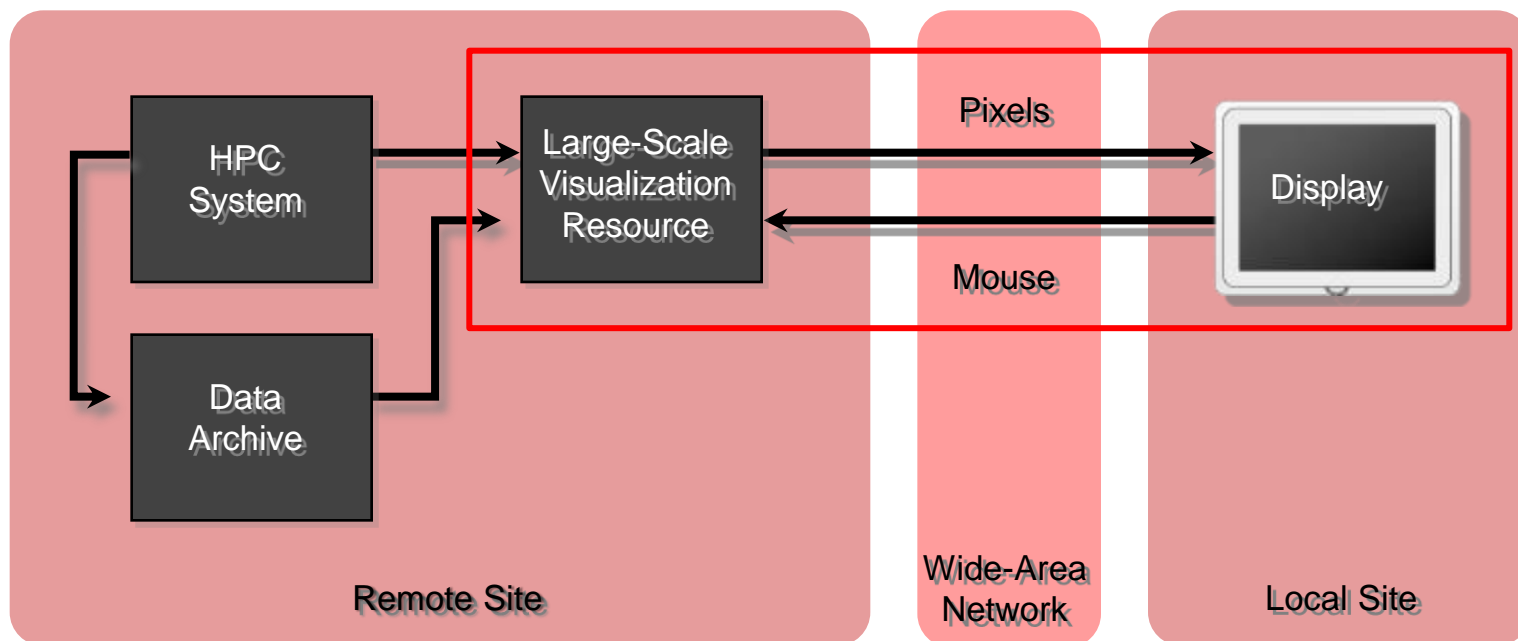


## Visualization nodes

- Stampede:
  - 128 nodes
  - Nvidia Tesla K20 GPU in each node
  - viz, gpu queues
  - 32 GB RAM, 16 cores
  - Share Stampede's lustre filesystems
- Maverick
  - 132 nodes
  - 256GB RAM, 20 cores per node
  - Nvidia Tesla K40 GPU in each node,
  - Lustre filesystems separate from Stampede, but \$WORK accessible



## Remote Visualization Model





## VNC


- Desktop process runs on remote *server*. vnc session
  - Windows, applications, mouse position
- Rendering occurs on server
  - Render on remote GPU. Send pixels to client
- Collaboration
  - Many can join vnc session, share control of mouse.
- VNC password to protect *\*session\** (use `vncpasswd`)
  - Share passwd with collaborators! Don't use login passwd!!





## Visualization Session

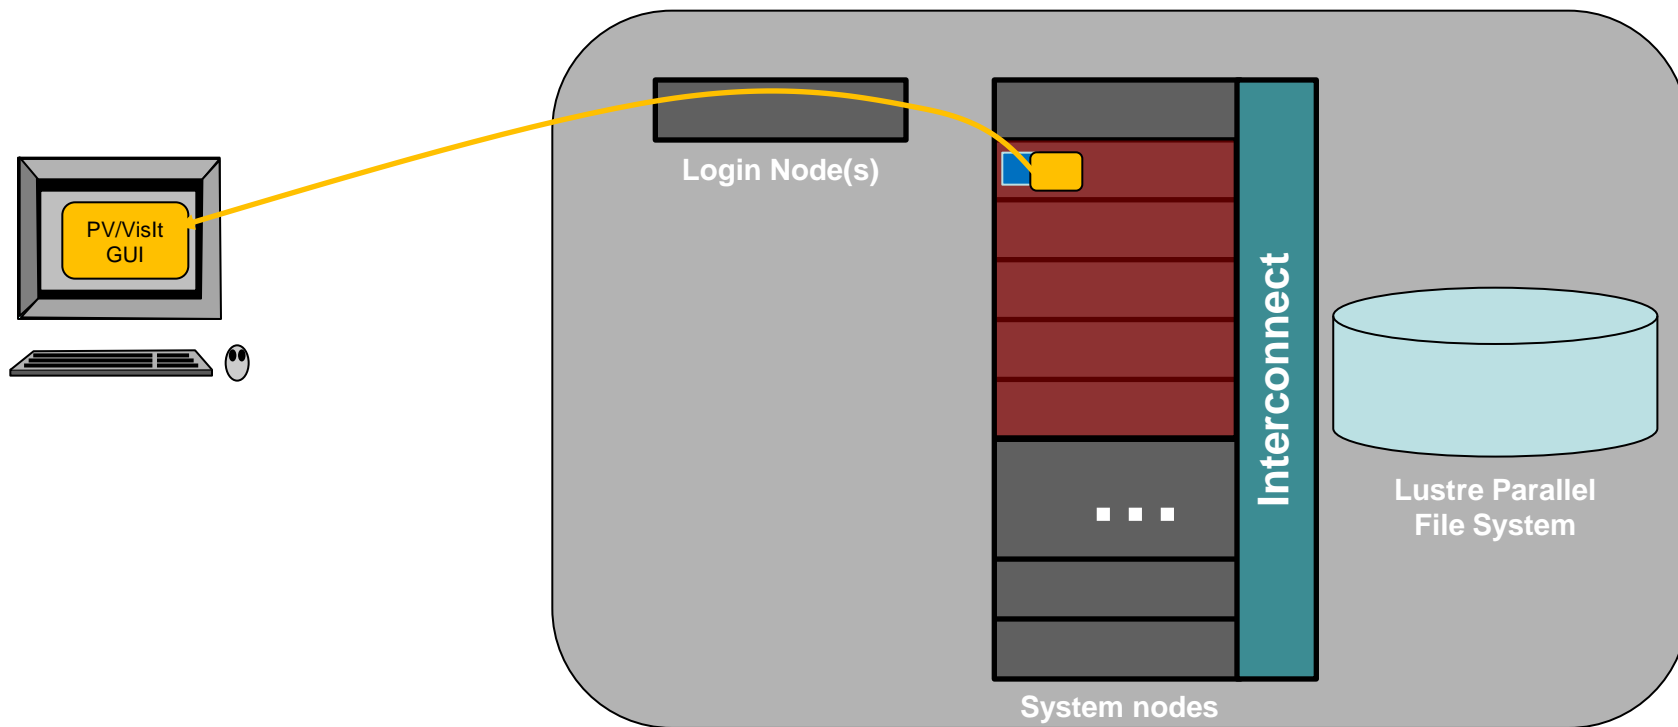
1. Allocate set of nodes on visualization system. This will start a VNC server on one node, to which you will connect





# Visualization Session

2. From that desktop, launch the graphical application

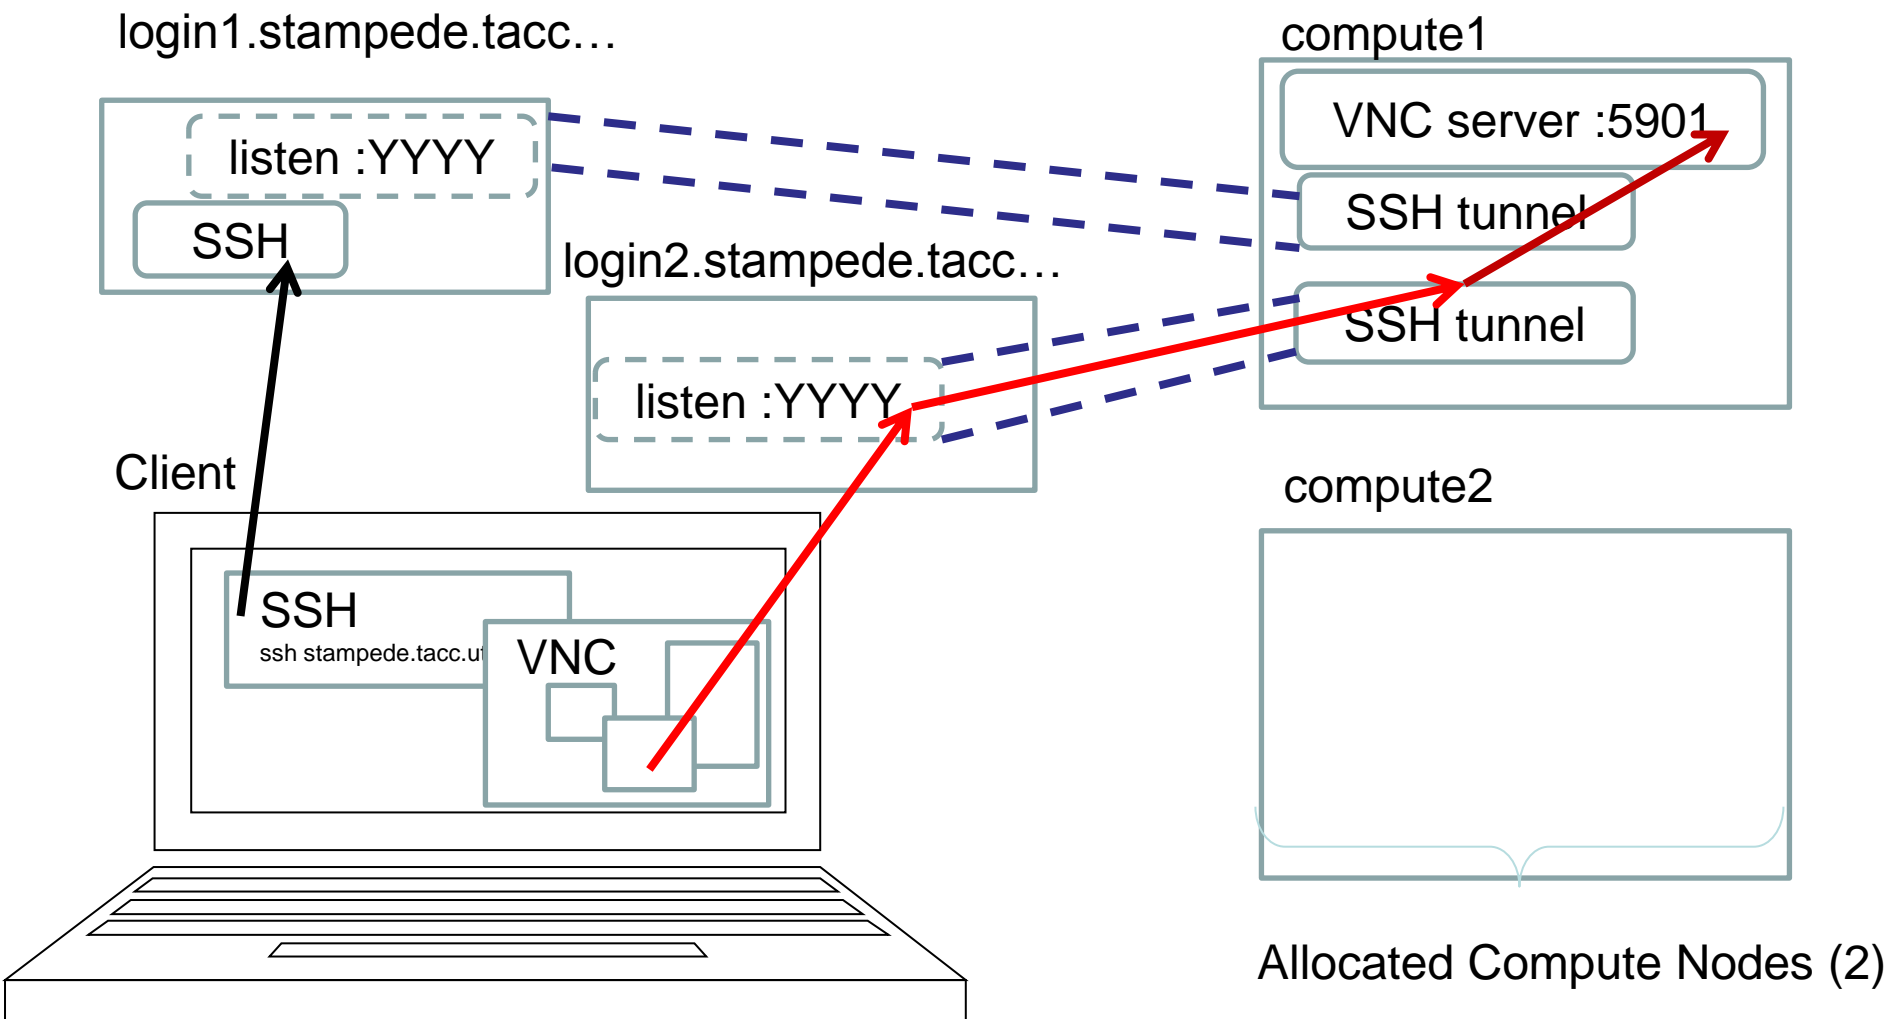






## Start VNC session on Stampede (the hard way)

- Run `vncpasswd` at least once to set initial vnc password
  - No need to do this again, unless you wish to change password.
- Stampede provides convenient job script
  - `sbatch /share/doc/slurm/job.vnc`
- Control node allocation via SLURM params
  - `sbatch -N 8 -A MY_ACCOUNT /share/doc/slurm/job.vnc`
- VNC desktop runs on compute node, private vnc port opened on login nodes just for you. Look for `vncserver.out`  
Created reverse ports on Stampede logins  
Your VNC server is now running!  
To connect via VNC client:  
SSH tunnel port 15754 to `stampede.tacc.utexas.edu:15754`  
Then connect to `localhost::15754`





## Start VNC Session on Stampede (the Easy way)

- <http://vis.tacc.utexas.edu>
- Hides commandline complexity
- Built-in web-based VNC client
- Jobs last as long as you are logged in to portal (unless timeouts reached)
- View viz node availability and queue status

The screenshot shows the TACC Visualization Portal interface. The browser address bar displays <https://vis.tacc.utexas.edu/#>. The page title is "TACC Visualization Portal". Below the title are navigation tabs for "Home", "Jobs", and "Help". The main content area is titled "Start a Job" and contains the following fields and options:

- Resource:** A dropdown menu with "Maverick" and "Stampede" (selected).
- Project:** A dropdown menu with "TG-StA110019S".
- Session type:** Radio buttons for "VNC" (selected), "iPython Notebook", and "R Studio".
- Desktop resolution:** A dropdown menu with "1280x1024".
- Number of nodes:** A dropdown menu with "4".
- Wayness (processes per node):** A dropdown menu with "16".

A light blue informational box contains the following text: "Note: increasing the number of nodes will only increase performance for parallel applications (e.g. ParaView or VisIt). The wayness parameter is only relevant to parallel applications, and determines how many processes are spawned per node when the parallel application is executed."

At the bottom of the form are two buttons: "Start Job ▶" and "Set VNC Password 🔑".



## Visualization applications

- Support for many input formats
  - Some may be better than others for certain tasks
- Aim for realtime point and click image manipulation
  - Data exploration
- Pipeline of data refinement or visualization operations
- Lots of tweakable parameters
- Parallel rendering (more on this later)
- Ease of use and suitability for certain tasks can vary
  - Definitely not one-size-fits-all



## ParaView

- <http://www.paraview.org/>
- Open-source, multi-platform parallel data analysis and visualization application
- Mature, feature-rich interface
- Good for general-purpose, rapid visualization
- Built upon the Visualization ToolKit (VTK) library
- Primary contributors:
  - Kitware, Inc.
  - Sandia National Laboratory
  - Los Alamos National Laboratory
  - Army Research Laboratory



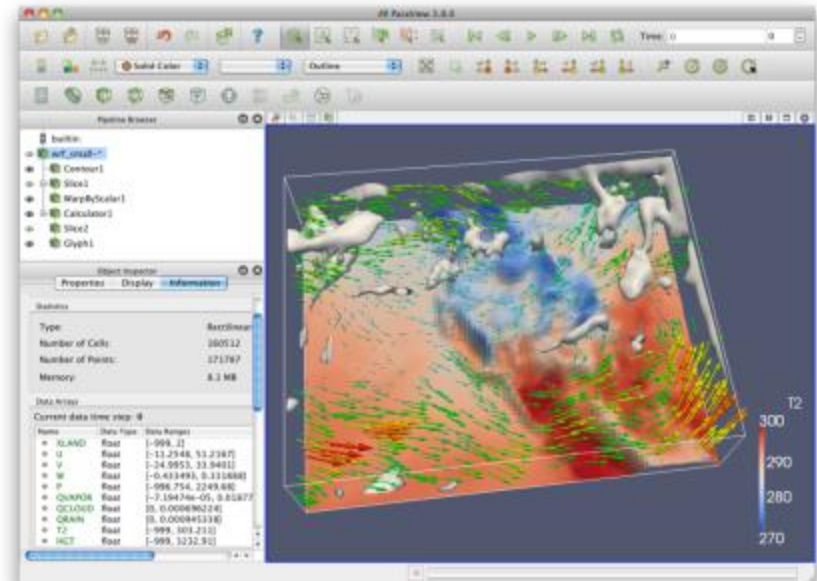
## ParaView

- Supports derived variables
  - New scalar / vector variables that are functions of existing variables in your data set
- Scriptable via Python
- Saves animations
- Can run in parallel / distributed mode for large data visualization



## ParaView

- All processing operations (filters) produce data sets
- Can further process the result of every operation to build complex visualizations
  - e.g. can extract a cutting plane, and apply glyphs (i.e. vector arrows) to the result
    - Gives a plane of glyphs through your 3D volume





## VisIt


- <https://wci.llnl.gov/codes/visit/>
- Open Source, Multiplatform, interactive parallel visualization and graphical analysis tool
- Developed by the Department of Energy (DOE) Advanced Simulation and Computing Initiative (ASCI)
- Although VisIt was developed for visualizing terascale data, it is also well suited typical desktop simulations
- Can run in parallel/distributed mode for large-scale visualization






## VisIt

- VisIt's visualization capabilities are grouped into two categories:
  - Plots are used to visualize data and include boundary, contour, label, mesh, pseudocolor,
  - Operators consist of operations that can be performed on the data prior to visualization. (Examples include slice, isosurface, threshold among others)






# Points, Meshes, and Coordinates




Medical scan

(a) Structured Points




(b) Rectilinear Grid




(c) Structured Grid

Engineering Model


N-body simulation



(d) Unstructured Points



(e) Polygonal Data



(f) Unstructured Grid


Extracted Surface

From *The Visualization Toolkit* by Schroeder et al.




# Data


- Values at each point
- Type and nature will determine applicable techniques
  - Scalar, Vector, Tensor?
  - Discrete? Continuous?
  - Nominal, Ordinal, Interval, Ratio?
- Now what do want to show about your data?



(a) Structured Points



(d) Unstructured Points




(f) Unstructured Grid



## Surface Shading (Pseudocolor)

Given a scalar value at a point on the surface and a color map, find the corresponding color (and/or opacity) and apply it to the surface point.




User: smc  
Sat Sep 20 13:10:41 2008



## Isosurfaces (Contours)

- Surface that represents points of constant value with a volume
- Plot the surface for a given scalar value.
- Good for showing known values of interest
- Good for sampling through a data range






## Volume Rendering

Expresses how light travels through a volume

Color and opacity controlled by transfer function

Smother transitions than isosurfaces








## Particle Traces (Streamlines)

Given a vector field, extract a trace that follows that trajectory defined by the vector.

$$P_{\text{new}} = P_{\text{current}} + V_P \Delta t$$

Streamlines – trace in space



Pathlines – trace in time





## Graphics Primitives

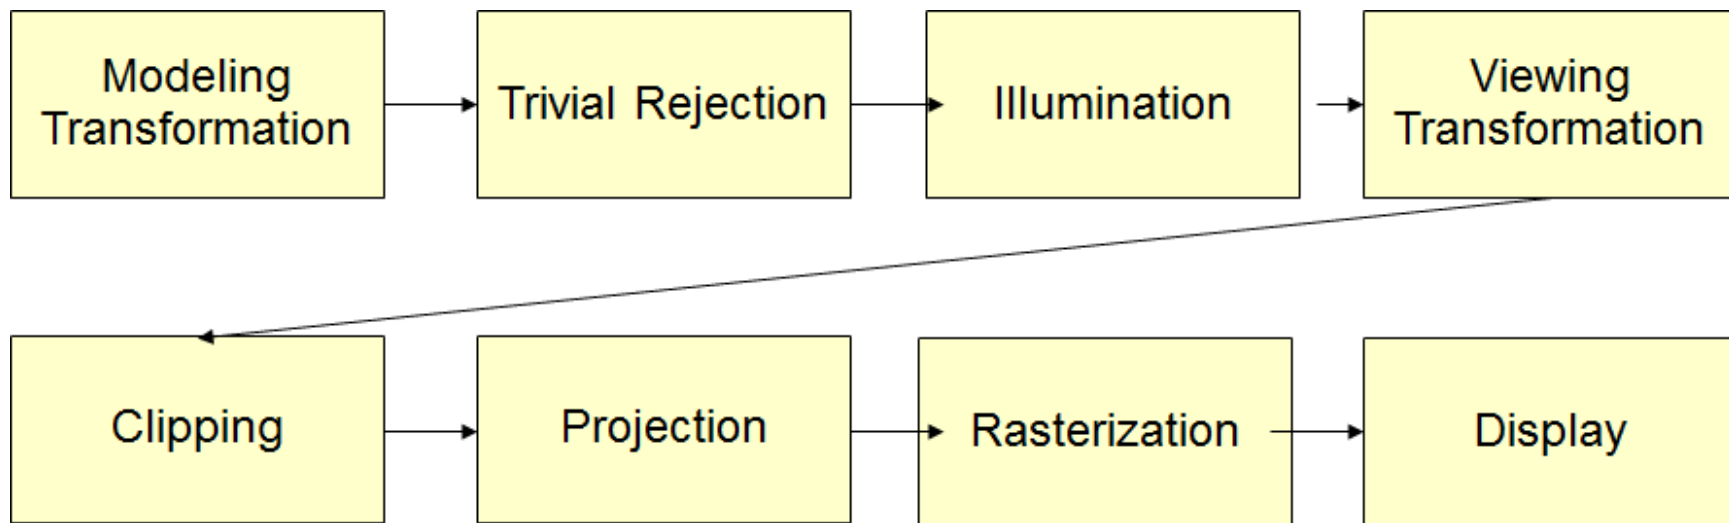
- Basic unit: Polygons, Colors, Textures, Opacity
  - Flat surface formed between points
  - This surface may have an associated color or texture, or opacity
- Complex surfaces composed of several polygons
- A dataset in and of itself!







# Graphics Pipeline





## Parallel Visualization




- Why? Performance
  - Processing may be too slow on one CPU
    - Interactive visualization requires real-time frame rates
    - **Use lots of CPUs**
    - Shared-memory/multicore *or* distributed
  - Data may be too big for available node
    - Virtual memory works, but paging is slow
    - **Use lots of nodes to increase physical memory size**
    - Big shared-memory/multicore scaling is costly (\$/CPU)

Increase interactivity or feasibility



## Memory Utilization


- Some visualization techniques cause memory use to skyrocket!





## Memory Utilization: Regular Grids


- Specified by:
  - (x,y,z) origin
  - (nx, ny, nz) counts
  - Data array
- Requires very little memory





## Memory Utilization: Regular Grids


- Chop off corner -> need an unstructured grid to represent data points
- Specified by
  - Explicit list of vertices
  - Explicit list of triangles
- Memory use can go up *many* times





## Memory Utilization: examples


- Mummy.vtk:
  - Structured Grid
  - (128x128x128)
  - 2MB raw data
- Contour: 7MB
  - Polygonal Mesh
- Slice of Contour: .1MB
- Tetrahedralize: 520MB!!
  - Unstructured Grid
  - Data points -> Tetrahedrons





## Visualization scales with HPC

- Large data produced by large simulations require large visualization machines and produce large visualization results
- Data and all derivations in memory, cumulative!





## Parallel Algorithms: Data Parallelism

### Data parallelism

Data set is partitioned among the processes and all processes execute same operations on the data.

Scales well **as long as the data and operations can be decomposed.**

|           |   | Timesteps        |                        |                    |
|-----------|---|------------------|------------------------|--------------------|
|           |   | 1                | 2                      | 3                  |
| Processes | 1 | Read partition 1 | Isosurface partition 1 | Render partition 1 |
|           | 2 | Read partition 2 | Isosurface partition 2 | Render partition 2 |
|           | 3 | Read partition 3 | Isosurface partition 2 | Render partition 3 |





## Parallel algorithms: What doesn't work

- Streamlines!
  - Not data-parallel
  - Partial streamlines must be passed from processor to processor as the streamline moves from partition to partition
  - No more parallelism available than the number of streamlines!
  - If  $>1$  streamlines pass through the same partition, you may not even get that



## Rendering

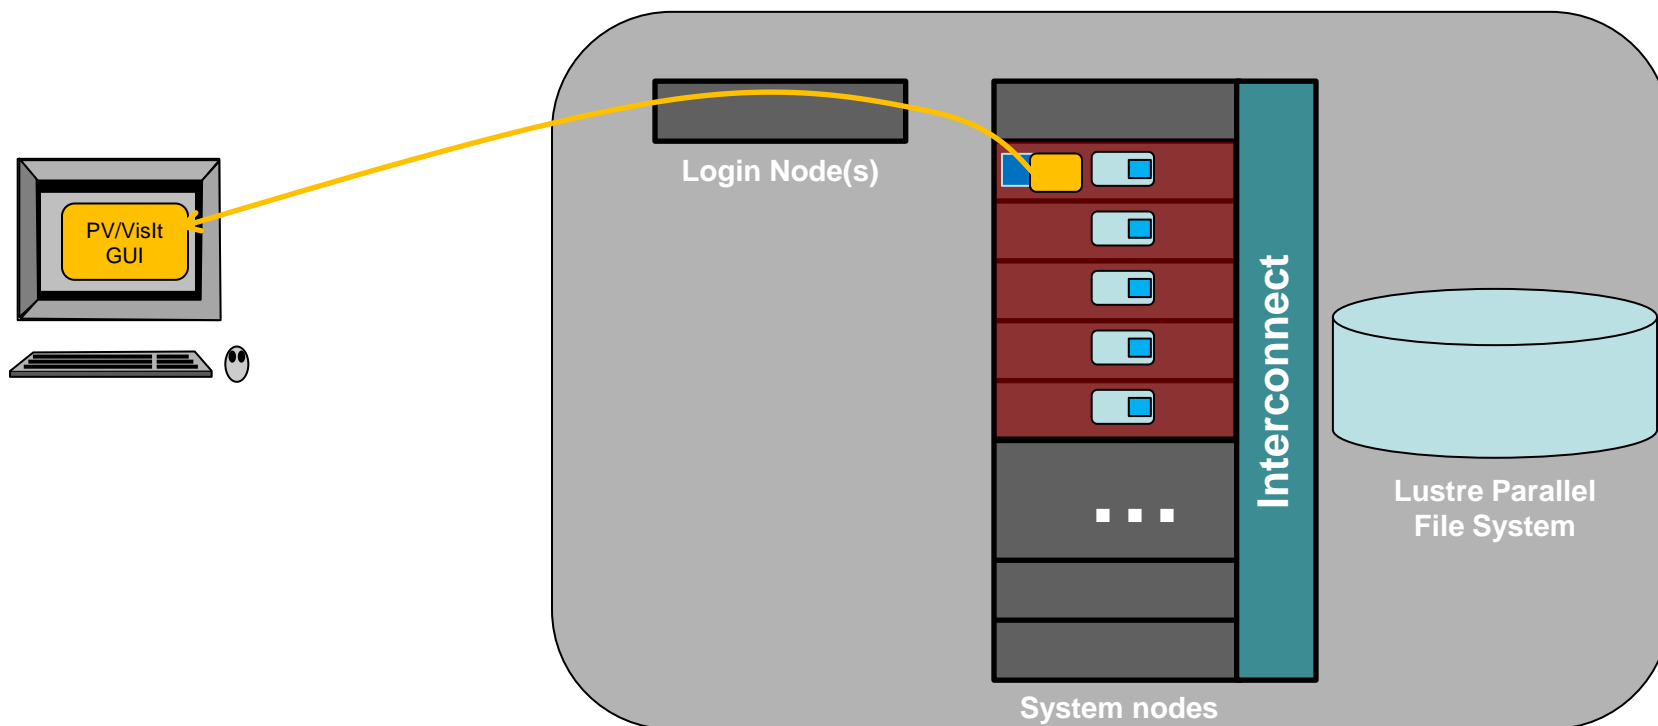
- Many graphics primitives spread out over nodes
- Rendering solutions
  - 1. Gather triangles onto one node, render there
    - Best when there's not a lot of data to render
  - 2. Render triangles in place, gather and Z-composite the results
    - Best when there *is* a lot of data to render
    - Overhead is *almost* independent of data size
- VisIt and ParaView both do it both ways
  - User controls threshold, but both apps aim for reasonable defaults
- Now how do we get rendered graphics to the user?



## Visualization Session

### 3. Start Paraview or VisIt Server Processes

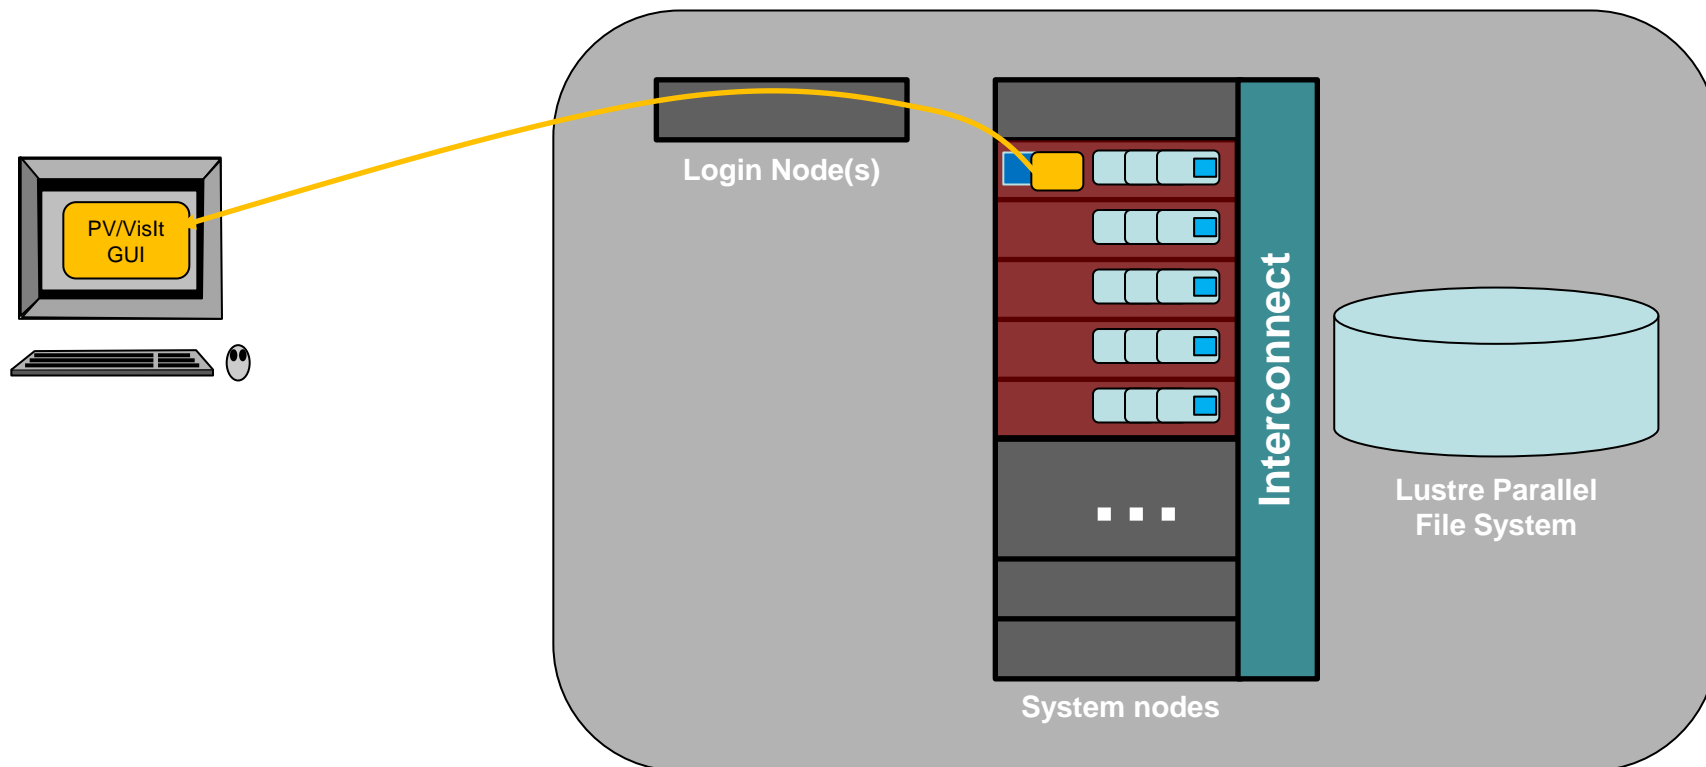
(Paraview and VisIt can do this automatically, with the right configuration)





## Visualization Session

- Multiple processes/node to take advantage of multiple cores/node -- wayness





## Launching Vis Applications

- Applications using OpenGL (i.e. all visualization apps) need to be wrapped with `vglrun <app>`
  - This is a workaround for the fact that vnc servers do not support OpenGL natively
- This starts the visualization GUI only.
  - Parallel backends are launched on-demand by visualization app, or manually by user
  - If not using parallel mode, then you're done!
- VisIt simply asks if you want parallel or serial mode
  - Params automatically determined by session params
- Paraview needs to be told to run backend processes via `ibrun`



## Parallel Data Management

- *Data must be distributed* across parallel processes to take advantage of resources
- *Explicit* Parallel formats use separate files for partitions
- *Implicit* parallel formats have a structure where data partitions can be deduced from file structure
  - .vtk legacy, silo, raw
- *Non-parallel* formats need to be read serially and distributed in order to be used in parallel
  - Overhead!
  - Vtk xml formats (.vtu, .vti, etc)



## Parallel Data Management

- Read the manual!
  - Vis software has varying support for file formats
  - True parallel I/O may not be implemented for some formats
  - Vis software will try to “hide” it’s failings
- Example: ParaView (from FAQ)
  - *Currently there are only a few readers that truly work in parallel: VTK files (not legacy), partitioned legacy VTK files, ParaView data files, HDF5 files, EnSight master server files, and raw (binary) files can be read in parallel. For demonstration purposes, ParaView will distribute pieces of a data set when the reader cannot. Unfortunatley, this is an inefficient process.*