NASA Interests in File System and I/O R&D in Support of HPC

Ellen.M.Salmon@nasa.gov

HEC-IWG File Systems and I/O R&D Workshop

August 16-17, 2005

Examples of NASA Applications Driving HPC File Systems and I/O

- Earth Science
  - Example: Weather modeling, 2010
    - 10**11 observations per day
    - Data ingested into cleaning/gridding process
    - Input: 1TB/day
    - Archive: 1 PB/day (12 GB/s sustained average)
  - http://esto.nasa.gov/files/ctreqreport.pdf
- Visualization
  - Example: Hyperwall-2, as early as 2006
    - 250 Mpixels, up to 350 Gb/s output
    - From live simulations and analysis

NASA Recent HPC Filesystem Experiences

- Shared filesystems (Fibre Channel plumbing)
  - "Local" sharing with fairly homogeneous members
    - Smaller number (~dozen) of larger cluster clients (64p-512p)
    - **Metadata scaling** and *interoperating with maturing Linux clients* are current major issues
  - Long-distance WAN sharing desired, but
    - Stability concerns: participating sites cannot afford to have a glitch at one site bring down the entire WAN cluster
    - Security concerns: near-term would be overcome by only sharing filesystems "outside the walls"

NASA HPC Application Needs

- Mixed-size I/O, not necessarily well-formed
- Applications have very long development and modification times

- **Strong I/O middleware** needed so innovations can plug in via API and libraries with as little disruption as possible

NASA HPC Facility Needs

- Great need for better **integration of filesystems with *higher-level* metadata** (via APIs or libraries) to allow for data management
  - At present users have difficulty managing the data because they have few tools to do so
  - Higher-level metadata of interest
    - "Management" metadata, e.g., creation/expiration, quality of service
    - Domain-specific content metadata
- **Tracking**: desire to understand better who is using the data and how, and how well the system is serving

Summary: Areas for R&D

- Low/mid-level filesystem metadata
  - Scaling, management, stability in the shared cluster environment
  - Long-distance cluster complications (stability, security)
- I/O middleware
  - Libraries and APIs crafted so innovations are accessible with minimal disruption to applications
- Interface and/or integration of filesystems with higher-level metadata (file management, content-specific, etc.)