How will Petascale systems change what we have been doing? -or- What will I (hopefully) be doing on my summer vacation?

Phil Andrews

Patricia Kovatch

San Diego Supercomputer Center

How long will it take to write memory to disk on a Petascale system?

- **1 million cores with 1 GB of memory/core -> 1 PB of total memory**
  - We generally allocate 2 GB of memory/CPU now so what will apps do? Idle cores due to memory size and BW
- **1 PB probably too expensive, assume 0.5 PB**
- **Performance changes with Petascale?**
  - Expect parallel file system performance ~1 TB/s
  - ~10 minutes to write to disk at this speed, no significant changes

- **Memory will be more important than cores!**

How much will it cost to save Petascale system's memory to archival storage?

- **1 TB cartridge costs ~$100 -> 0.5 PB costs ~$50K!**
  - Assuming 0.5 PB total system memory
  - Write memory 4X/week -> **$10M/year for tapes?!!**
- **Or consider actual archival usage at SDSC**
  - DataStar: ~5 TB/memory -> $1M/year in tapes
  - Petascale: 100X memory -> **$100M/year in tapes?!!!**
  - **Much more expensive than power costs!**
- **Storage size?**
  - Data parked on parallel file system for 3 months at a time -> **24 PB file system, 50 PB more likely**
  - Assuming full system memory written 4X/wk * 0.5 PB * 12 weeks
- **Storage more expensive than flops!**

New Data Flow Model Needed for Petascale systems

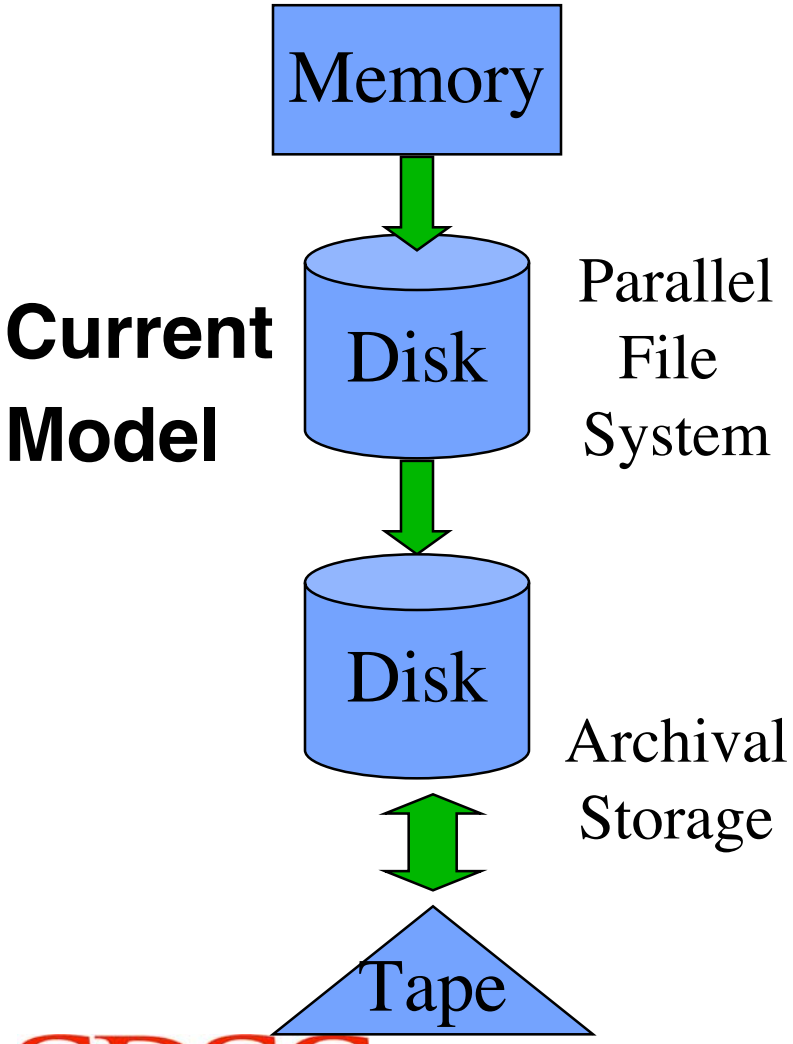

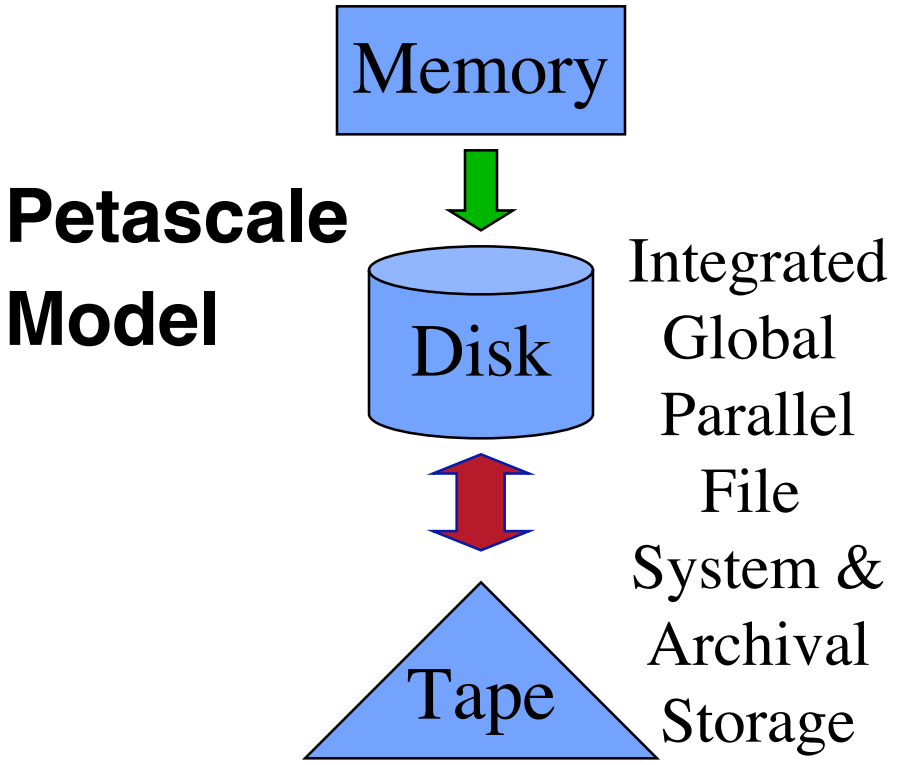

Most data has a short-term life and never stored on tape

How will Petascale systems change what we have been doing?

- **Memory-driven computing**
  - Allocations and queuing based on memory, not cores
  - Memory is the scare resource, cores sit unused
- **Disk-driven storage**
  - Allocations based on storage
  - Tape costs are prohibitive
  - RAID 6 and other schemes essential for highly reliable file systems
  - Integrated global parallel file system and archival storage
  - Users perform real-time, concurrent analysis and visualization
  - Faster/cheaper to rerun job than to restage data from archive