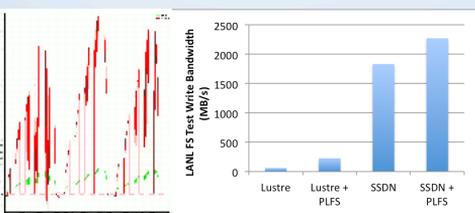
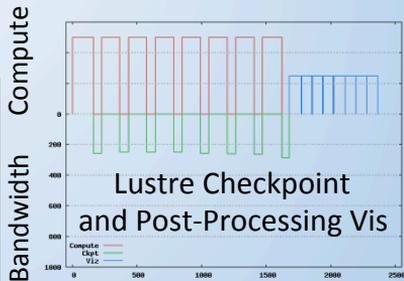
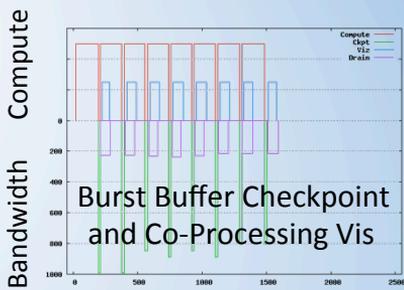
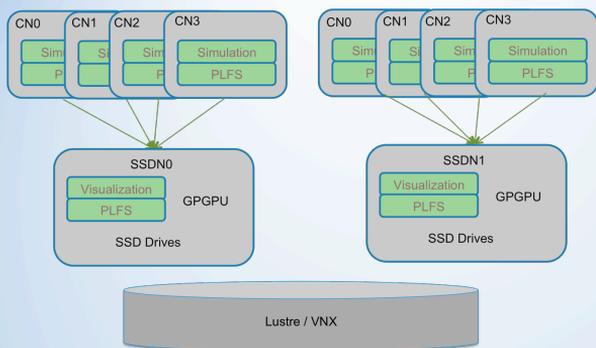


Co-processing Simulation and Visualization on a Prototype Exascale Burst Buffer Storage System



Petascale storage designs will not scale up to Exascale

- The cost for the number of mechanical disks required for fast bandwidth would be too expensive
- The cost for the number of SSDs required for storage capacity would be too expensive
- Massively parallel concurrency would overwhelm file systems – N to N writes will be intractable for thousands and millions of processing elements

LANL and EMC's Burst buffer prototype and demonstration

- Fast checkpointing via EMC burst buffer nodes with SSDs is cheaper than rotational-only bandwidth
- Asynchronous drain to Lustrre does not impede the simulation
- Jitter-free co-processing analysis and visualization tasks on burst buffer nodes
- Application checkpointing and I/O is transparent; PLFS ensures data is visible regardless of location

Burst Buffer checkpointing: An Integrated Solution

LANL PLFS file system	Fast write capability
LLNL SCR checkpoint restart	Data staging
EMC Burst Buffer nodes	SSDs for bandwidth
Lustrre/VNX disk storage	Storage capacity

EMC²
where information lives™

LA-UR 11-06271
Los Alamos
NATIONAL LABORATORY
EST. 1945

NNSA
National Nuclear Security Administration

EMC: John Bent, Sorin Faibish, Uday Gupta, Gary Ma, Jim Pedone, Percy Tzelnic, Beijing FastData Team

LANL: Jim Ahrens, HB Chen, Gary Grider, John Patchett, Jon Woodring