

dCache Storage Cluster at BNL

Tuesday, 24 March 2009 08:00 (20 minutes)

Over the last (2) years, the USATLAS Computing Facility at BNL has managed a highly performant, reliable, and cost effective dCache storage cluster using SunFire x4500/4540 (Thumper/Thor) storage servers. The design of a discreet storage cluster signaled a departure from a model where storage resides locally on a disk-heavy compute farm. The consequent alteration of data flow mandated a dramatic re-construction of the network fabric.

This work will cover all components of our dCache storage cluster (from door to pool) including OS/ZFS file-system configuration, 10GE network tuning, monitoring, and environmental. Performance metrics will be surveyed within the context of our Solaris 10 production system as well as those rendered during evaluations of OpenSolaris and Linux. Failure modes, bottlenecks, and deficiencies will be examined.

Lastly, we discuss competing architectures under evaluation, scaling limits in our current model, and future technologies that warrant close surveillance.

Presentation type (oral | poster)

Oral

Author: PETKUS, Robert (Brookhaven National Laboratory)

Co-authors: Dr SMITH, Jason (Brookhaven National Laboratory); Mr MCCARTHY, John (Brookhaven National Laboratory); Ms KARASAWA, Mizuki (Brookhaven National Laboratory)

Presenter: PETKUS, Robert (Brookhaven National Laboratory)

Session Classification: Poster session

Track Classification: Hardware and Computing Fabrics