Breakout 4A
Management, reliability, managed performance, and availability at scale (file system management, SAN management, device management, autonomics, RAID, etc.

Session Coordinators: Grider

Session Scribes: Macaluso

Session Presenter:

Session Writeup:

Each Breakout session will provide

1) Current high level topics of File Systems and I/O Research in this area

2) Areas that need to have more research focus

3) Areas that have or will have too much research focus

4) Some rough consensus ranking of areas that need more focus,
less focus,
and overall recommendations including
Short term research needs
Long term research needs

There will be a presentation of this material for each session done by the session leader and a write up for inclusion in the workshop documentation.

Current high level topics of File Systems and I/O Research in this area

- *Autonomics*
- *Scalable rebuilding techniques*
- *Multiple concurrent RAIDs*
- *Multi-pathing*
- *Server fail over techniques*

Areas that need to have more research focus (designate short and long term)

1. *Research into Autonomics (adaptive/self healing/predictive)*
2. *How do you operate where component failure is a normal case (scale)*
    1. *Degrades with scale*
3. *Performance/determinism during recovery at scale (raid, network, etc.)*
4. *Performance robustness (dealing well with variance in workload, configuration, etc.)*
5. *Multi-path networking at scale for I/O footprints, determinism capabilities.*
6. *Server fail over techniques*
7. *Client process migration*
8. *Large scale storage system configuration and tuning*
9. *Identify/research interfaces to collect information from layers and to use that info to effect the system (possibly related to modeling – could be input into modeling (1a), measure effect of the effect*
10. *Reasonable ways to ensure integrity at scale*
11. *Multidimensional virtualization areas management of all this*
12. *Scaling of management/meta management, management of the management tools*
13. Collect and analyze real data on failure, non-determinism, bad performance, recovery times, tuning stuff

Areas that have or will have enough or too much research focus (designate short and long term)

- None

Some rough consensus ranking of areas that need more focus, less focus and overall recommendations including Short /Long term research needs

1. How do you operate where component failure is a normal case (scale)
    1. Total 57 Government 14 med-long term
2. Research into Autonomics (adaptive/self healing/predictive)
    1. Total 58 Government 6 med-long term
3. Large scale storage system configuration and tuning
    1. Total 34 Government 8 med-long term
4. Performance robustness (dealing well with variance in workload, configuration, etc.)
    1. Total 32 Government 3 med-long term
5. Collect and analyze real data on failure, non-determinism, bad performance, recovery times, tuning stuff
    1. Total 28 Government 3 short-med term
6. Identify/research interfaces to collect information from layers and to use that info to effect the system (possibly related to modeling – could be input into modeling (1a), measure effect of the effect
    1. Total 24 Government 0 <-disagreement med-long term
    2. Interesting
    3. Could be combined with 2 and related to 3 and 4