Breakout 1B
Security (Device, File System, on the SAN, Transactional, Access Control, Need to Know, etc.)

Session Coordinators: Grider

Session Scribes: Miner

Session Presenter: Grider

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

- *Cluster File Systems that support GSS, ACL's, need to know etc. Kerberos and Private Key*
- *NFSv4 security, GSS, ACL's, need to know*
- *In place encryption (1694)*
- *Block level encryption*
- *End to end encryption/ user level with key mgmt scheme*
- *Key management schemes that live on*
- *Network based security like ISEC*
- *Extended attributes in Linux api's*
- *Role based models for administration*
- *Multi level database security*
- *Hardware assistance for security*

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

1. Security at scaled up workloads including transactional security on the SAN
2. Standardize on one set of ACL’s (?) (including multi-realm) including scaling up ACL’s, and make them really usable
3. Standard API for end to end encryption?
4. Kernel support for tickets/keys etc. / pags (nearly done)
5. Key management that lives on
6. Data stewardship that is encrypted over different time scales (5, 10, 50, million)
7. Multi-realm user mapping, including scale up of number of realms
8. Usability of security infrastructure/mgmt including integration of various layers
9. Tracking of access and changes, tracking provenance, signatures, over different time scales (5, 10, 50, million)
10. How do we store signatures with files
11. Security exception recovery in a robust way
12. How do you know you achieved digital destruction (destroy the key, backups, archives)
13. Denial of service for file/storage system (should this be in the reliability section?)
14. Practical security policies
15. Security benchmarking/verification, something better than tiger teams
16. How do you know that the security code is pure and controlled
17. Can you secure data independent of the protocols used to access it, should it be a property of the data, how do you expose keys to protocol

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

- Network distributed denial of service
- Antivirus

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

1. Usability of security infrastructure/mgmt including integration of various layers and usability of ACL's
    1. Total 45 Government 13 med-long term
2. Key management that lives on forever
    1. Total 38 Government 11 med-long term
3. Multi-realm security including user mapping, with scale up of # of realms
    1. Total 25 Government 9 med term
4. Security at scaled up workloads including transactional security on the SAN
    1. Total 19 Government 10 med term
5. Security benchmarking/verification, something better than tiger teams
    1. Total 20 Government 2 <- disagreement
    2. (.edu and .com like this) med-long term
    3. At least we need a list of desired items or best practices, at least publish what the Tiger teams would check on. (file system threats need to be focused on)
6. Data stewardship that is encrypted over different time scales (5, 10, 50, 1000 years)
    1. Total 13 Government 1 <- disagreement
    2. (.edu likes this) med-long term
    3. Government interested in tactical solution and this looks like an elegant problem for .edu's