

Pool Quotas

LU-11023

Why

- **Need to limit users and groups usage of limited resource (i.e. flash OSTs)**
- **Pools currently only select OSTs for inclusion in layout**
- **But they do define groups of OSTs generally**
- **Re-use the group definitions for quotas**

Confusing point

- **A file does not belong to a pool. An object belongs to an OST that may belong to a pool(s)**
- **Objects have usr/grp associated with them...**
- **...but the user and group quotas are already tracked by each OST**

Pool quotas is therefore just policy

- **Quota master already tracks quota grants for each OST independently**
- **Computes total outstanding grant for any identifier**
- **Decides new grant based on limits in admin_quotafile**
- **Can compute total outstanding for any subset (i.e. pool)**
- **Just need to add new limits file for this pool**
- **Reply to grant request must satisfy all quota limits**
 - Simply the `min()` of all limits: global and each pool

Problems?

- **OSTs may belong to multiple pools**
 - OST objects count toward all pool limits
- **Pool membership may change**
 - Calculation of pool granted quotas changes
 - May be suddenly over quota
- **Flash OSTs may be allocated outside of pools**
 - OST objects are still tracked, still count toward pool quotas
- **Allocator and quotas are independent**
 - Can allocate to OSTs without enough quota to write
 - This is a (different) general problem with quotas
 - Really want space reservations, not just quota grant