

Silicon Graphics, Inc.

XFS Overview & Internals

13 - Monitoring

November 2006

Monitoring

- TODO: Incorporate Dave Wright's stuff here

Inodes

- `watch grep inode /proc/slabinfo`
- `sar -v`
 - TODO: Need to install
- `slabtop -s c`
 - Monitor the slab

PCP Inode Metrics

- `ig_attempts` # of looks for inode in memory
- `ig_found` # of times inode found in memory
- `ig_freecycle` # of fails because inode was busy
- `ig_misses` # of times had to go to disk for inode
- `ig_dup` # times missed but found from another
- `ig_reclaims` # of times recycled the inode's memory
- `ig_attrchng` # of times changed inode attributes
- `iflush_count` Number of times inodes are being flushed
- `icluster_flushcnt` # of times inodes clustered
- `icluster_flushinode` # of times not able to inode cluster

PCP Directory Metrics

- lookup Number of file name directory lookups
- create Number of directory entry creations
- remove Number of directory entry removes
- getdents Number of “getdent” operations

PCP File Attribute Metrics

- attr.get Number of “get” operations
- attr.set Number of “set” operations
- attr.remove Number of “remove” operations
- attr.list Number of “list” operations

PCP Allocation Metrics

- `allocs.alloc_extent` Number of extents allocated
- `allocs.alloc_block` Number of blocks allocated
- `allocs.free_extent` Number of extents freed
- `allocs.free_block` Number of blocks freed
- `alloc_btree.lookup` Number of allocation btree lookups
- `alloc_btree.compare` # of compares in alloc btree lookup
- `alloc_btree.insrec` Number of extent record inserts in btree
- `alloc_btree.delrec` Number of extent record deletes in btree

PCP Block Map Metrics

- `block_map.read_ops` Number of block map reads
- `block_map.write_ops` Number of block map writes
- `block_map.unmap` Number of block delete operations
- `block_map.add_exlist` # of extent list insertions for files
- `block_map.del_exlist` # of extent list deletions
- `block_map.look_exlist` # of extent list lookups
- `block_map.cmp_exlist` # of extent list compares
- `bmap_btree.lookup` # of block map btree lookups
- `bmap_btree.compare` # of block map compares
- `bmap_btree.insrec` # of block map insertions
- `bmap_btree.delrec` # of block map deletions

PCP Journaling Metrics

- `transaction.sync` Number of transactions waiting to be committed
- `transaction.async` Number of async transactions waiting
- `transaction.empty` Transactions that did not do anything
- `log.writes` Number of log buffer writes
- `log.blocks` Number of log blocks written
- `log.noiclogs` Log entry attempt during memory flush
- `log.force` Calls to `xfs_log_force`
- `log.force_sleep` Calls to `xs_log_force_sleep`

PCP AIL Metrics

- `push_ail.pushes` Number of times the AIL is moved forward
- `push_ail.success` Number of times successful
- `push_ail.pushbuf` Number of times inode locked – pushbuf called
- `push_ail.pinned` Number of times pinned
- `push_ail.locked` Number of times locked
- `push_ail.flushing` Number of times ail was flushing
- `push_ail.restarts` Number of log write restarts
- `push_ail.flush` Number of times a log was forced

PCP Quota Metrics

- `reclaims` Number of disk quota reclaims
- `reclaims_misses` Number of disk quota reclaim misses
- `dquot_dups` Number of duplicates
- `cachemisses` Number of times disk quota misses cache
- `cachehits` Number of times disk quota in cache
- `wants` Number of times dqwants called
- `shake_reclaims` Number of shaked reclaims
- `inact_reclaims` Number of inactive data reclaims

xfsstats

- xfsmisc/xfsstats.pl -f

sgti[®]