SAND2013-4282P

NNSA New Mexico Alliance for Computing at Extreme Scale

Lustre Experience on Cielo

Joel Stevenson (ACES, SNL)

Cielo Lustre L2 Milestone Review, March 11, 2013

cielo
NNSA New Mexico Alliance for Computing at Extreme Scale

File System Performance with CTH

- Compared write times with Lustre and Panasas for Jason Wilke's CTH simulations (32K core CCC-1, 64K core CCC-2, and 16K core CCC-3)
  - Lustre writes were 1 min; Panasas writes were 1-30 min
  - Simulation progress was faster with Lustre (larger is better in tables below)
    - CCC-1 data set: lustre 310 usec/24 hr (12.9 usec/hr); panasas 251 usec/24hr (10.5 usec/hr)
    - CCC-2 data set: lustre 203 usec/24 hr (8.5 usec/hr); panasas 182 usec/24hr (7.6 usec/hr)
    - CCC-3 data set: lustre 137 usec/24 hr (5.7 usec/hr); panasas 129 usec/24hr (5.4 usec/hr; restart adjust)

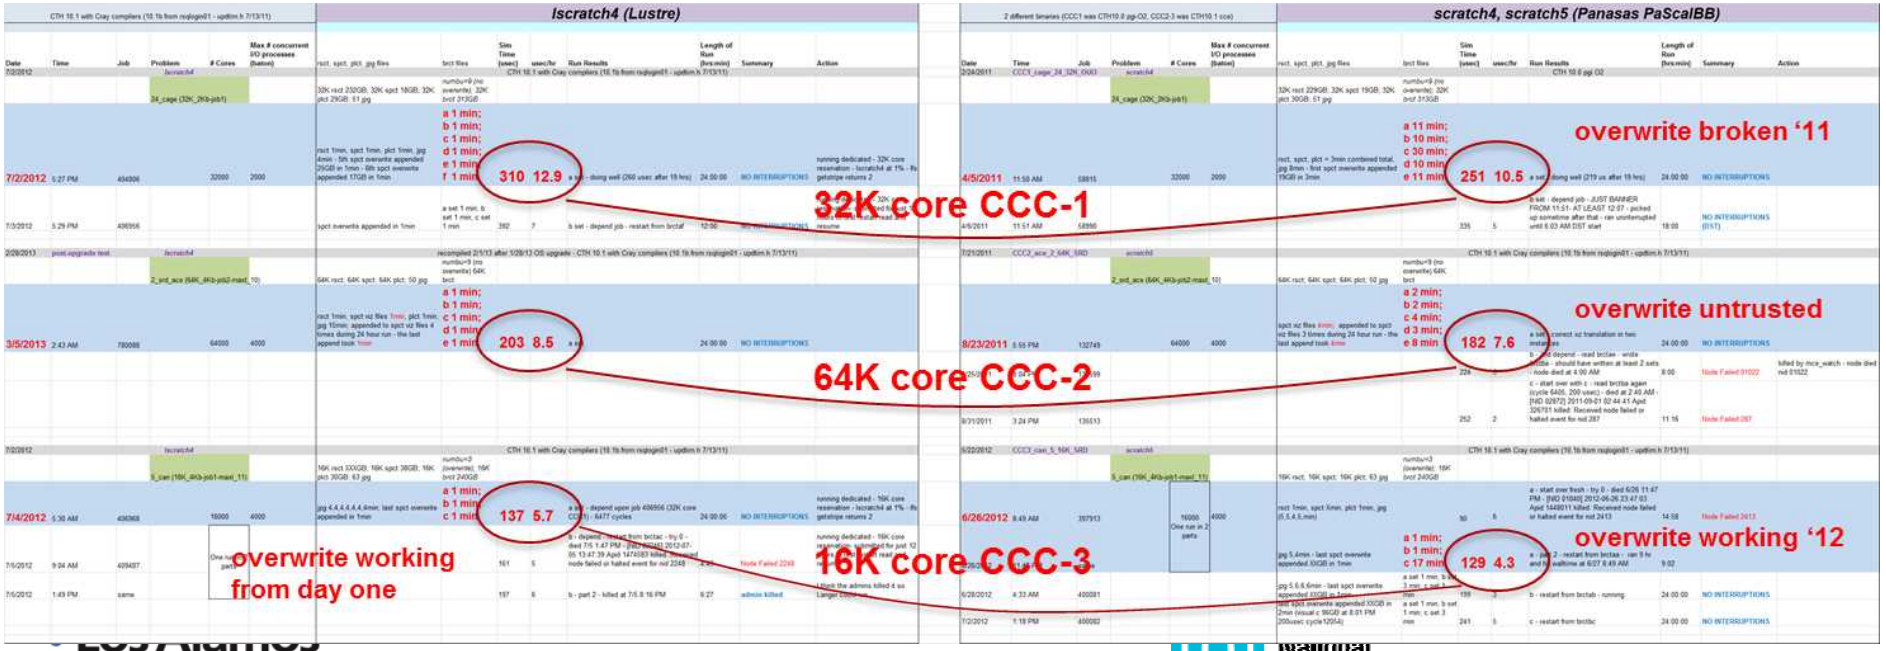

New Mexico Alliance for Computing at Extreme Scale

Write Times for CTH Restart Files

- **CCC-2 Panasas**
  - 64K core CCC-2 data set (baton = 4K files)
  - Not overwriting restart files
  - Exceed 1.4M files in "accessed" directories and paging occurs, reducing application performance; stay below threshold and write times = 2-10 minutes

- **CCC-3 Panasas and Lustre**
  - 16K core CCC-3 data set (baton = 4K files)
  - Overwriting restart files
  - Switched to Lustre on 7/17/12
  - Panasas:Avg = 4.1 minutes, Max = 38 minutes
  - Lustre: Avg = 0.8 minutes, Max = 14 minutes

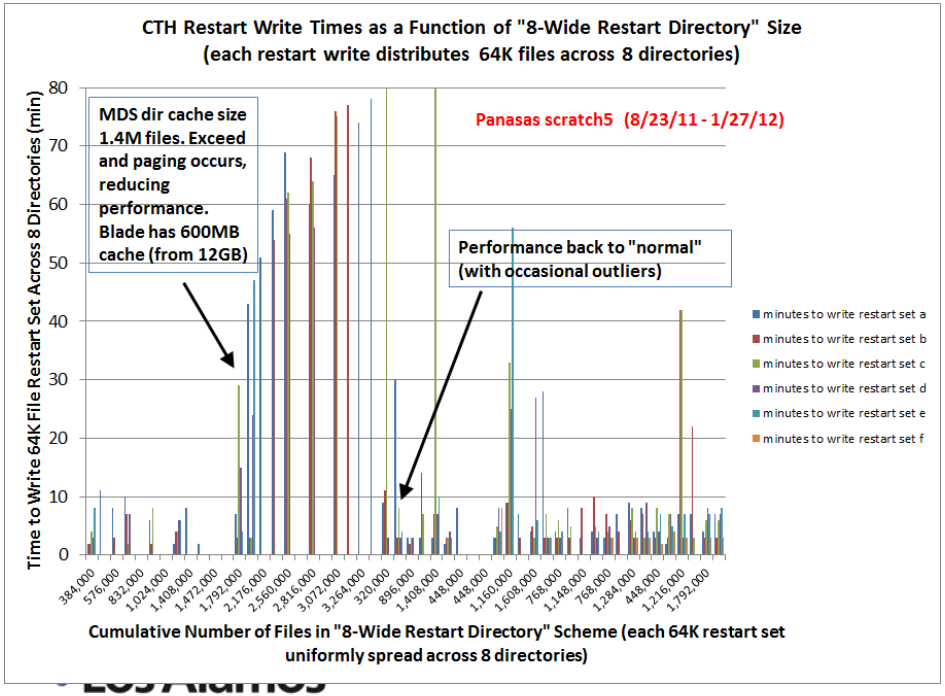

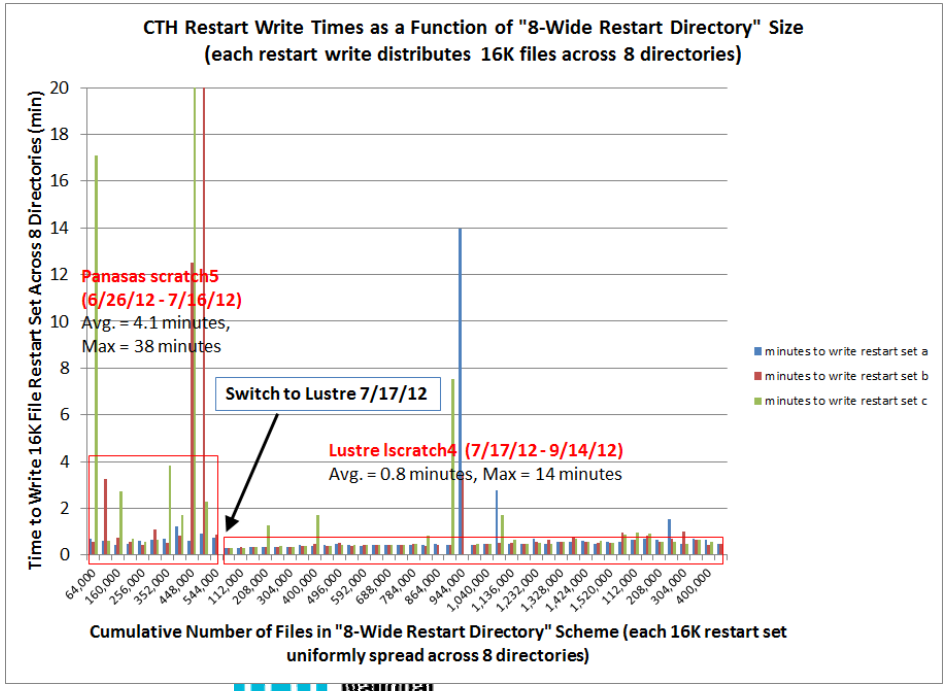

cielo
NNSA New Mexico Alliance for Computing at Extreme Scale

CCC-3 Simulation Progress

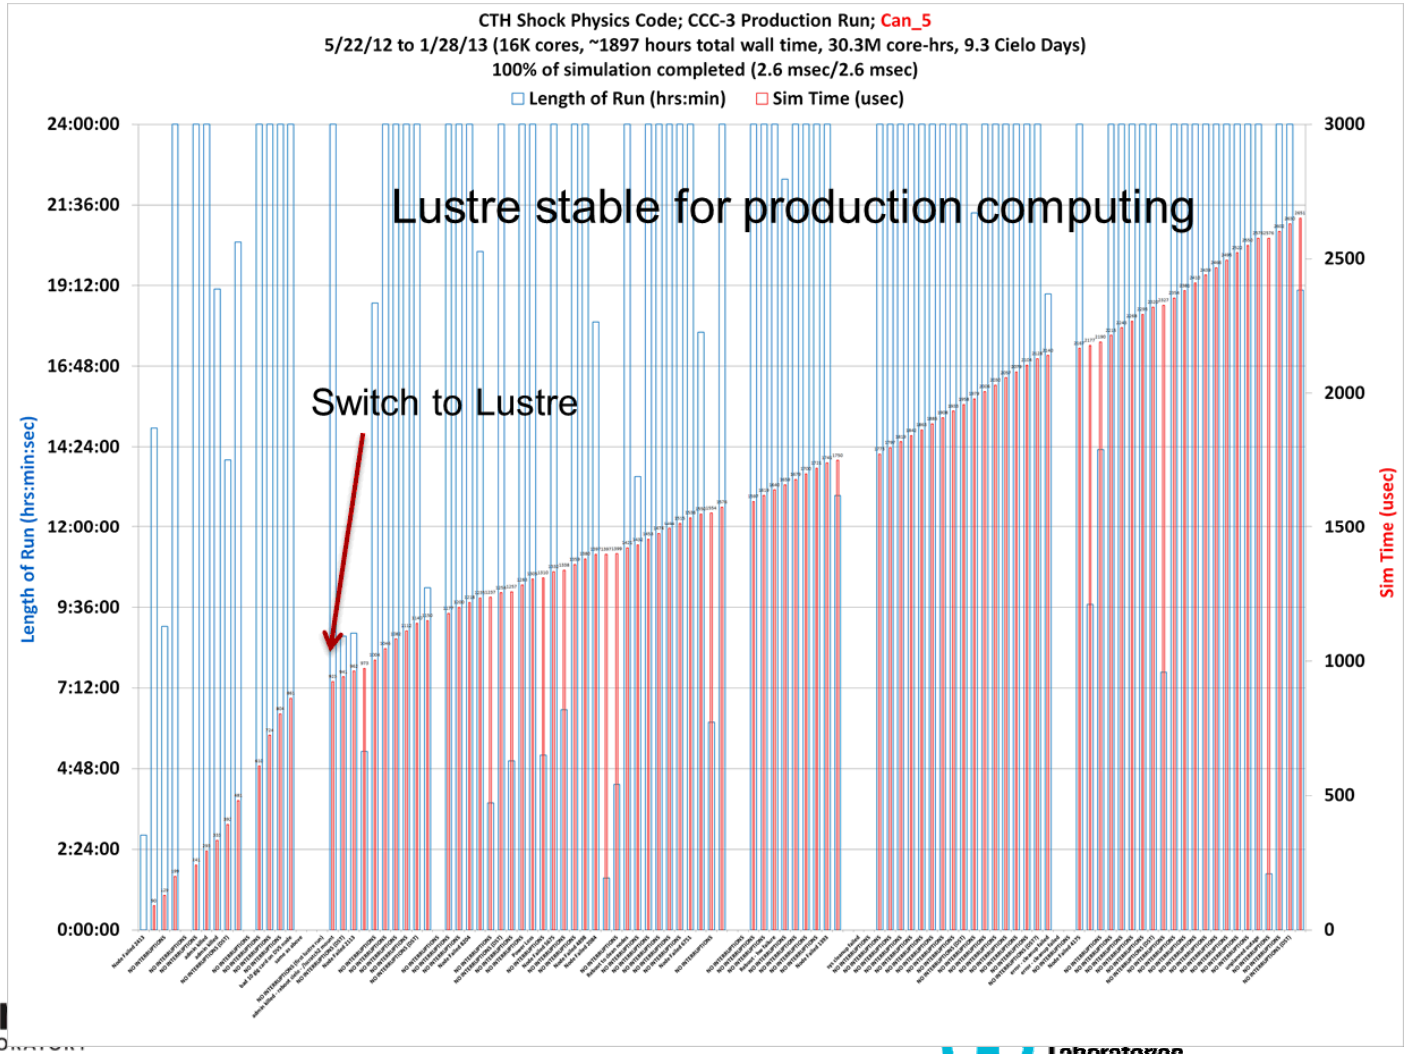

Los Alamos NATIONAL LABORATORY EST.1943

Laboratories

NNSA

cielo
NNSA New Mexico Alliance for Computing at Extreme Scale

Conclusion

- The Cielo file system transition to Lustre has been completed, providing increased performance and reliability needed for full production of the Cielo platform
- The Cielo file system has demonstrated an acceptable production user environment for production capability-class simulations - application users can get their work done on Cielo without major file system issues