CMS Plans and Site Expectations

Ken Bloom
For the CMS, US CMS and AAA teams
May 13, 2014

Data federations: an arrow in our quiver

- The wide-area data federation allows physicists to access any data, anytime, anywhere (AAA) transparently, with benefits at all scales:
  - Individual users can immediately examine the one event they are interested in, no matter what disk in the world it is on
  - Production-scale data processing is no longer tied to data location, allowing much more flexible use of resources globally
  - Natural solution to the "data problem" for opportunistic resources
- This is working for CMS, but it didn't come for free!
  - Much work up front to optimize WAN reading of files
  - Simple, consistent namespace for all CMS files
  - Reliable network infrastructure with large and growing bandwidth
  - Helpful participation from sites throughout CMS

Status of AAA deployment

- 6 of "8" CMS T1 sites are part of the data federation
  - In: DE, FR, IT, (RU,) UK, US
  - Not in: ES (coming soon), TW ("opportunistic T1")
  - Important caveat: CMS T1 sites are in the midst of disk-tape separation, so that we have greater control over what files are currently on disk. Only the files on disk are actually accessible.
    - In principle this already gives access to a huge amount of CMS data
- 41 of 52 CMS T2 sites are part of the data federation
  - In general, the sites that are not in the federation (without naming names) are smaller and/or less robust
  - ~ 96% of unique datasets resident at T2's are available
- We consider this to be full deployment within CMS!
- Usage level: Files opened at ~few Hz, 10's TB read each day

Applications: fallback

- One of the first applications of AAA was the "fallback mechanism"
  - This is the key to almost every other AAA application….

- Usually, if a job fails to open an input file, it crashes
- The fallback mechanism gives a path for recovery:
  - On file-open failure, CMSSW asks redirector to find file elsewhere
  - Job then reads remote file, user never notices

- More throughput for users, less CPU time wasted on failed jobs
- Makes entire system more robust against single-site storage issues

- A few easy configuration changes needed at sites to do this
  - 47/52 T2 sites have implemented fallback
  - One T1 has not due to firewall issues; discussions/debugging continue on proxy server deployment there

Applications: production with remote data

- "Legacy" reprocessing of 2012 data and associated simulation samples
- Inputs resident at T1 sites
  - T1's ran on data locally
  - T2's ran on simulations read via AAA fallback mechanism
- Flexible use of processing resources is likely to be exploited in Run 2
  - HLT farm already uses this for production reprocessing

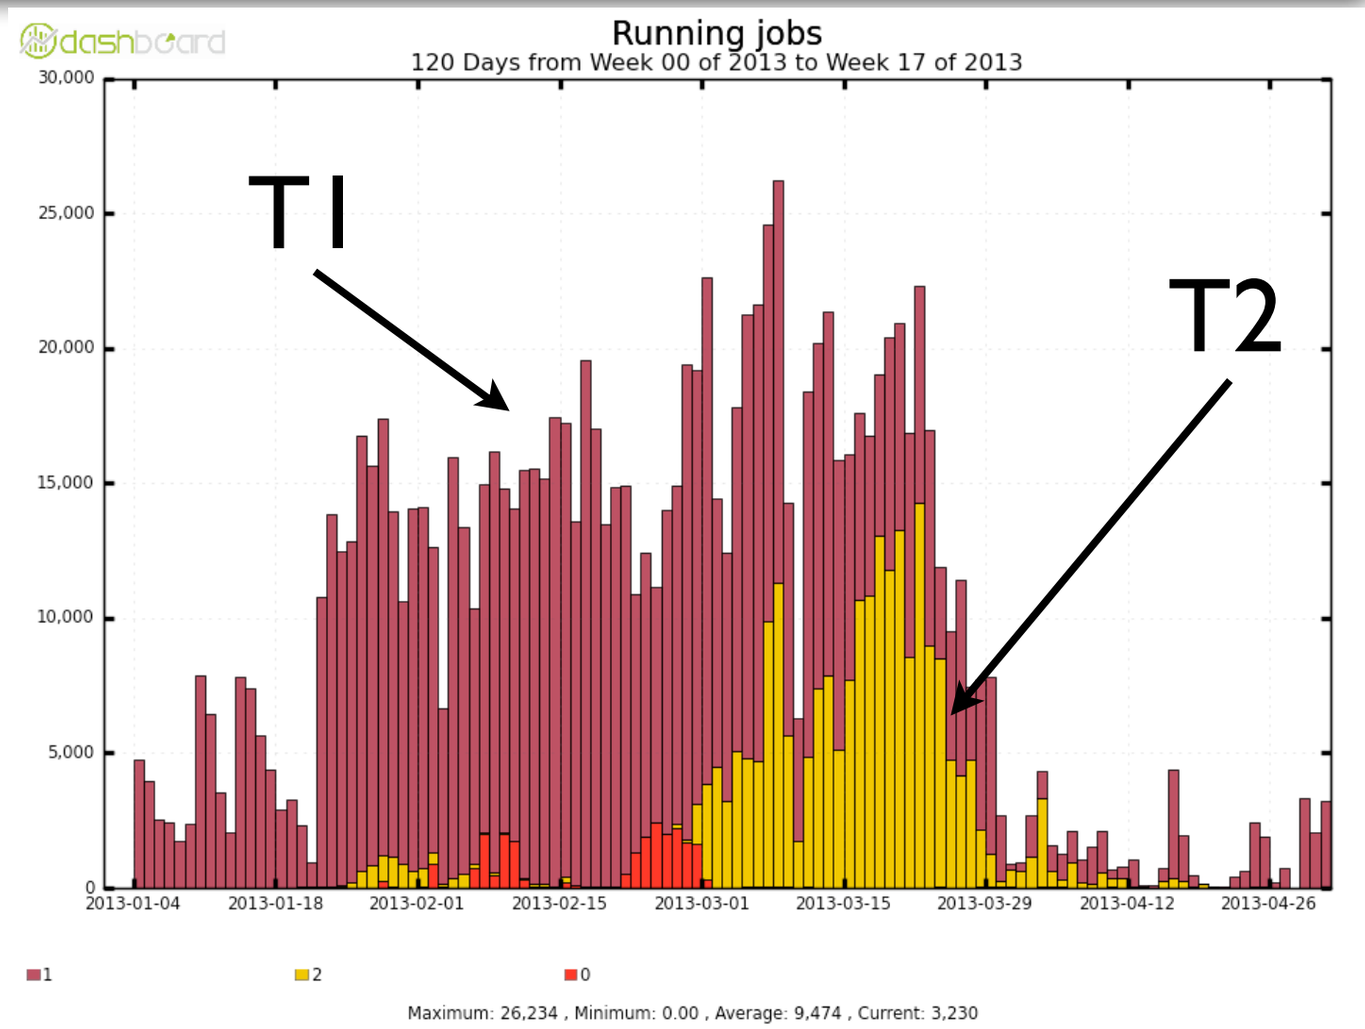

- Sites should be prepared to serve data at a scale similar to that at which they process data!
  - Storage systems, network….
  - We do provide system throttles — use them!

Applications: optimizing resource use

- Sites with popular datasets can have very long batch queues
- Re-direct jobs to another site with free job slots, read data via AAA
  - Smaller CPU efficiency, but jobs can start sooner
  - Achieved by changing scheduling policies in glideinWMS layer, regulate number of jobs to match WAN bandwidth
- So far, only small scale -- overflow amongst four sites in the US, ~O(2K) simultaneous jobs -- but no technical issues block expansion

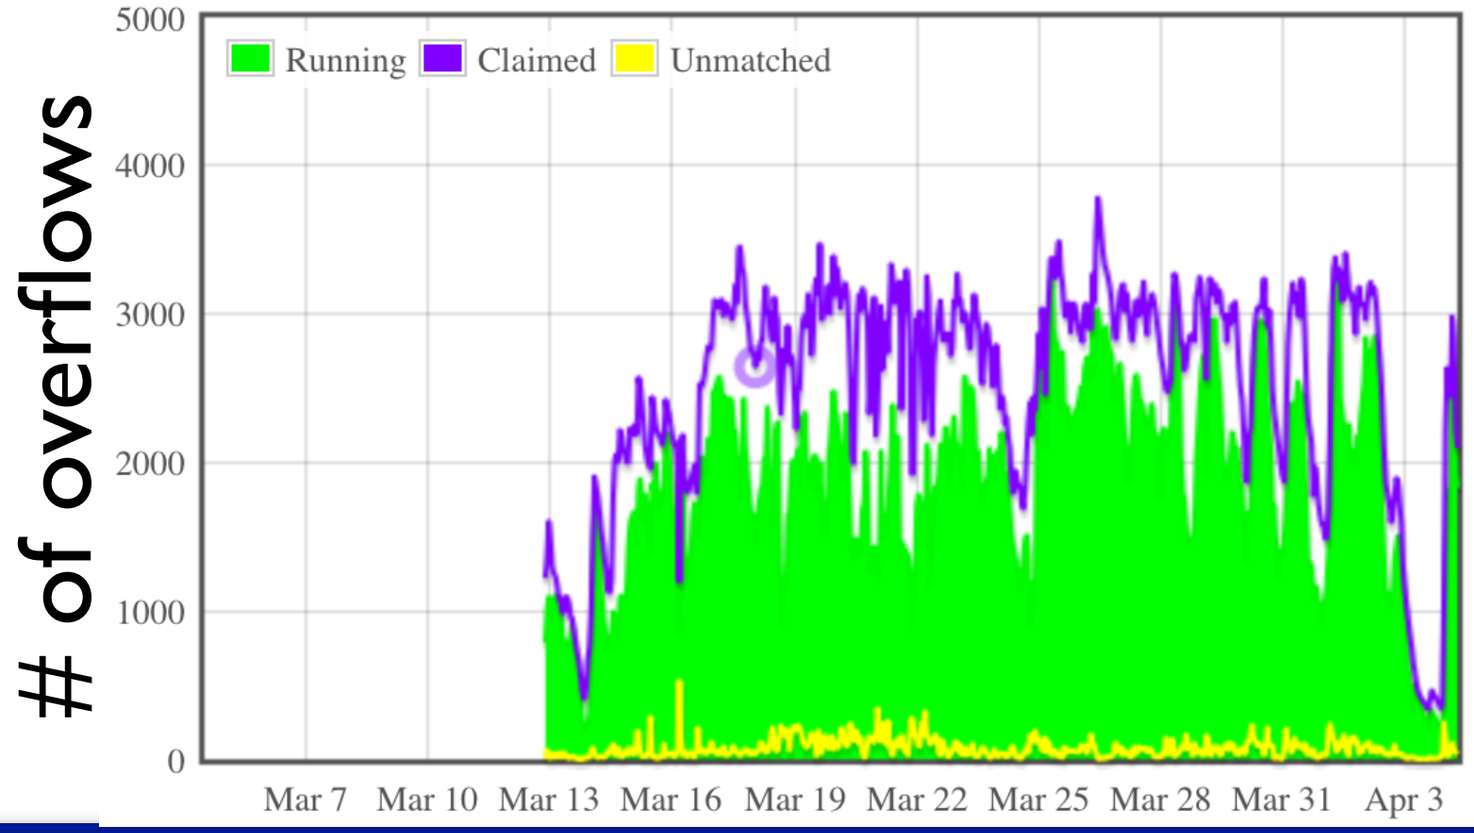

Applications: sites without data

- Some T3 sites are completing entire data analyses through AAA
  - Observed ~800 simultaneous jobs, 2-3 Gb/s WAN input sustained for a week, 99% success rate
  - Much satisfaction with local control over processing resources
  - "At this point, I basically don't pay attention to where the data is and just assume that jobs will find the data and run."
- Exploring possibility of diskless T2 sites at well-networked centers
- Sites that temporarily lose their data due to storage downtime (planned or unplanned) can continue to operate as normal through the fallback mechanism
  - Allows the continuity of processing capacity, system-wide
  - Have seen several successful cases, some planned and some not

Applications: opportunistic usage

- Any data, anywhere means any computer, not just CMS-owned
  - For software, use Parrot and CVMFS for download on demand, brings in 500 MB of files rather than 17 GB
  - Then, read data through AAA fallback mechanism
  - Typical jobs only 2% slower than those running on CMS sites
- Opens the door to any opportunistic resource, e.g. clouds
  - Have run 2K simultaneous jobs across 15 non-CMS OSG sites, including ATLAS sites (thanks)
  - Successful demonstration on Amazon cloud
  - Much CMS development work underway
- We think we are entering an era of constrained computing resources — CMS will need every CPU it can find, and AAA is a key technology for making good use of all of them

Scale tests

- For best operations, want to understand performance of each site
  - Every site will eventually hit some limitation
  - But we want every site to perform as well as possible
- A series of scale tests is now underway:

1) File-open test, how quickly can redirector successfully redirect?
   - Baseline goal 20 Hz, technology goal 100 Hz, testing up to 200 Hz

2) File-serving test, see load that each (and every!) source site can sustain in terms of MB/s, file opens/s etc.
   - Exploring rates up to 800 jobs, each reading 0.25 MB/s

3) Client-hosting test, see load that each (and every!) sink site can sustain, using the same metrics
   - Could a site successfully feed all of its batch slots with AAA?

4) Total-chaos test, do previous tests with many sites at once

Sample file-open test results

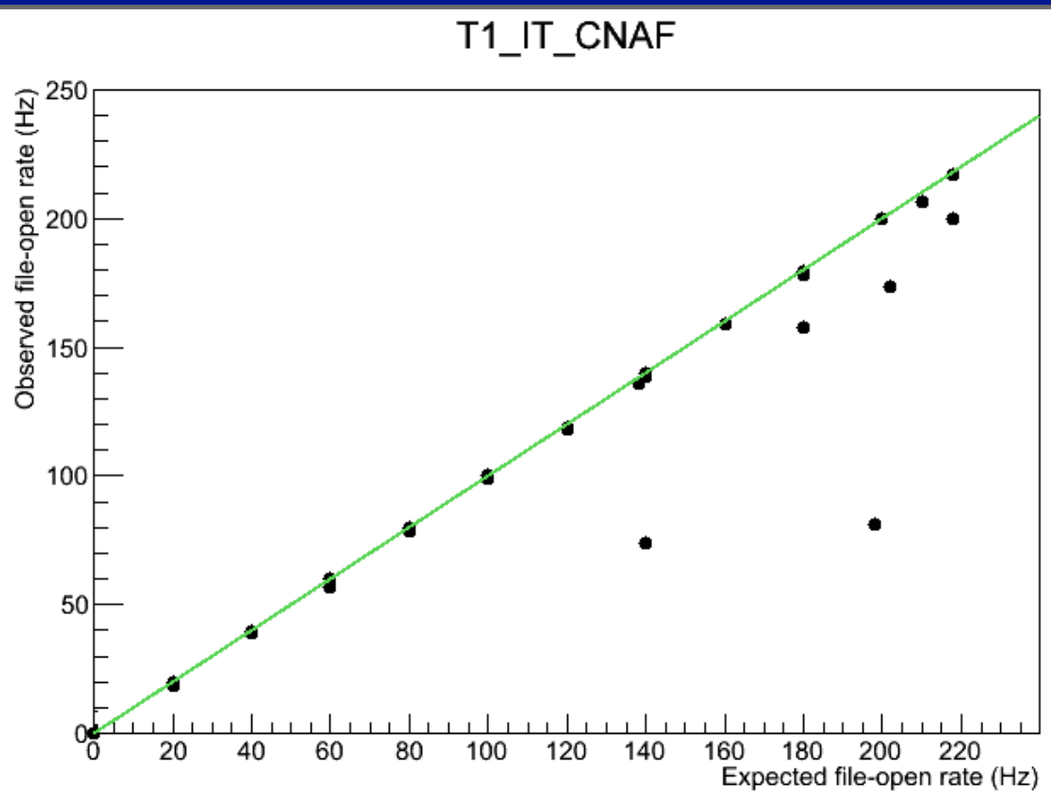

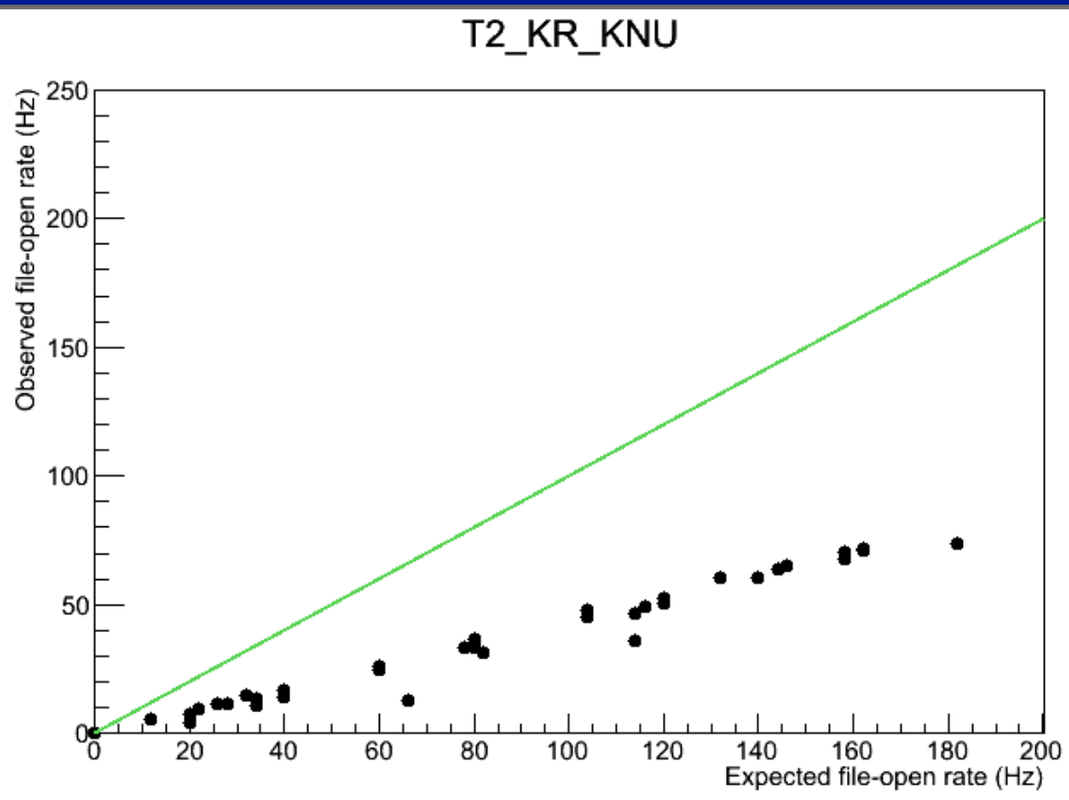

- Results vary by site, for sure
- Possible trend: Hadoop, Lustre, EOS and StoRM systems have performed better than dCache and DPM
  - Probably need to work with developers on this….

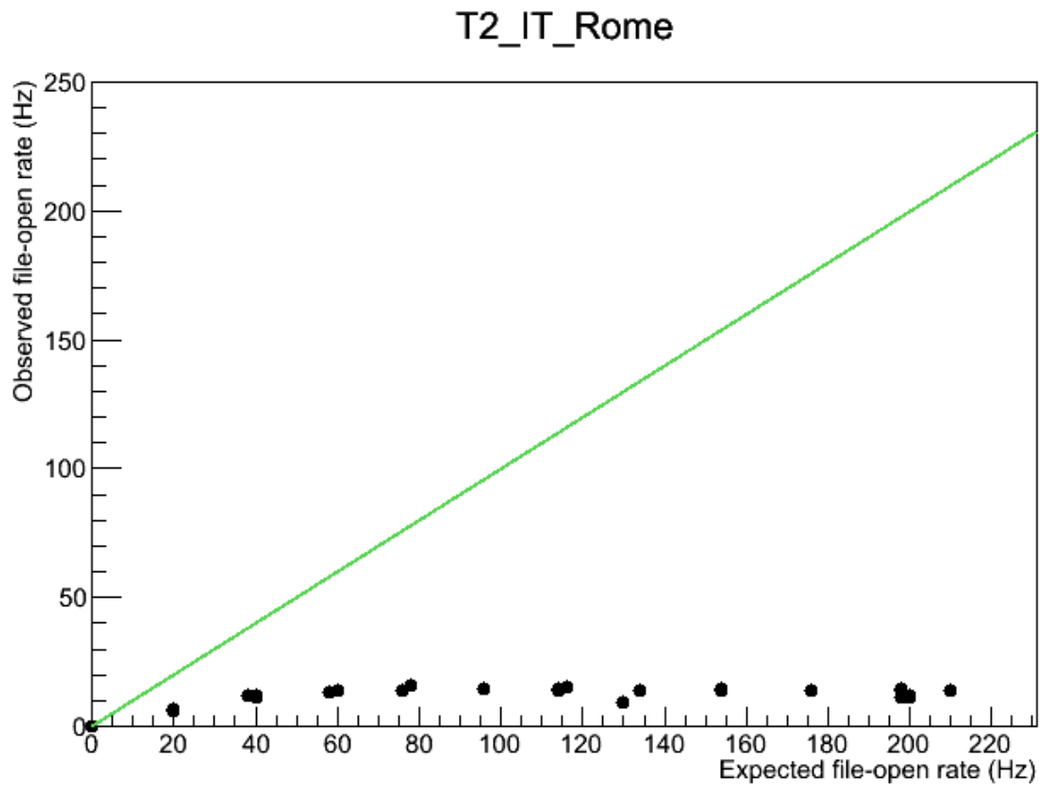

Sample file-read test results

- One can probably learn a lot about storage system behavior from this
- We really hope to systematically get through all sites….
- No results on Test 3 yet, but we hope soon!

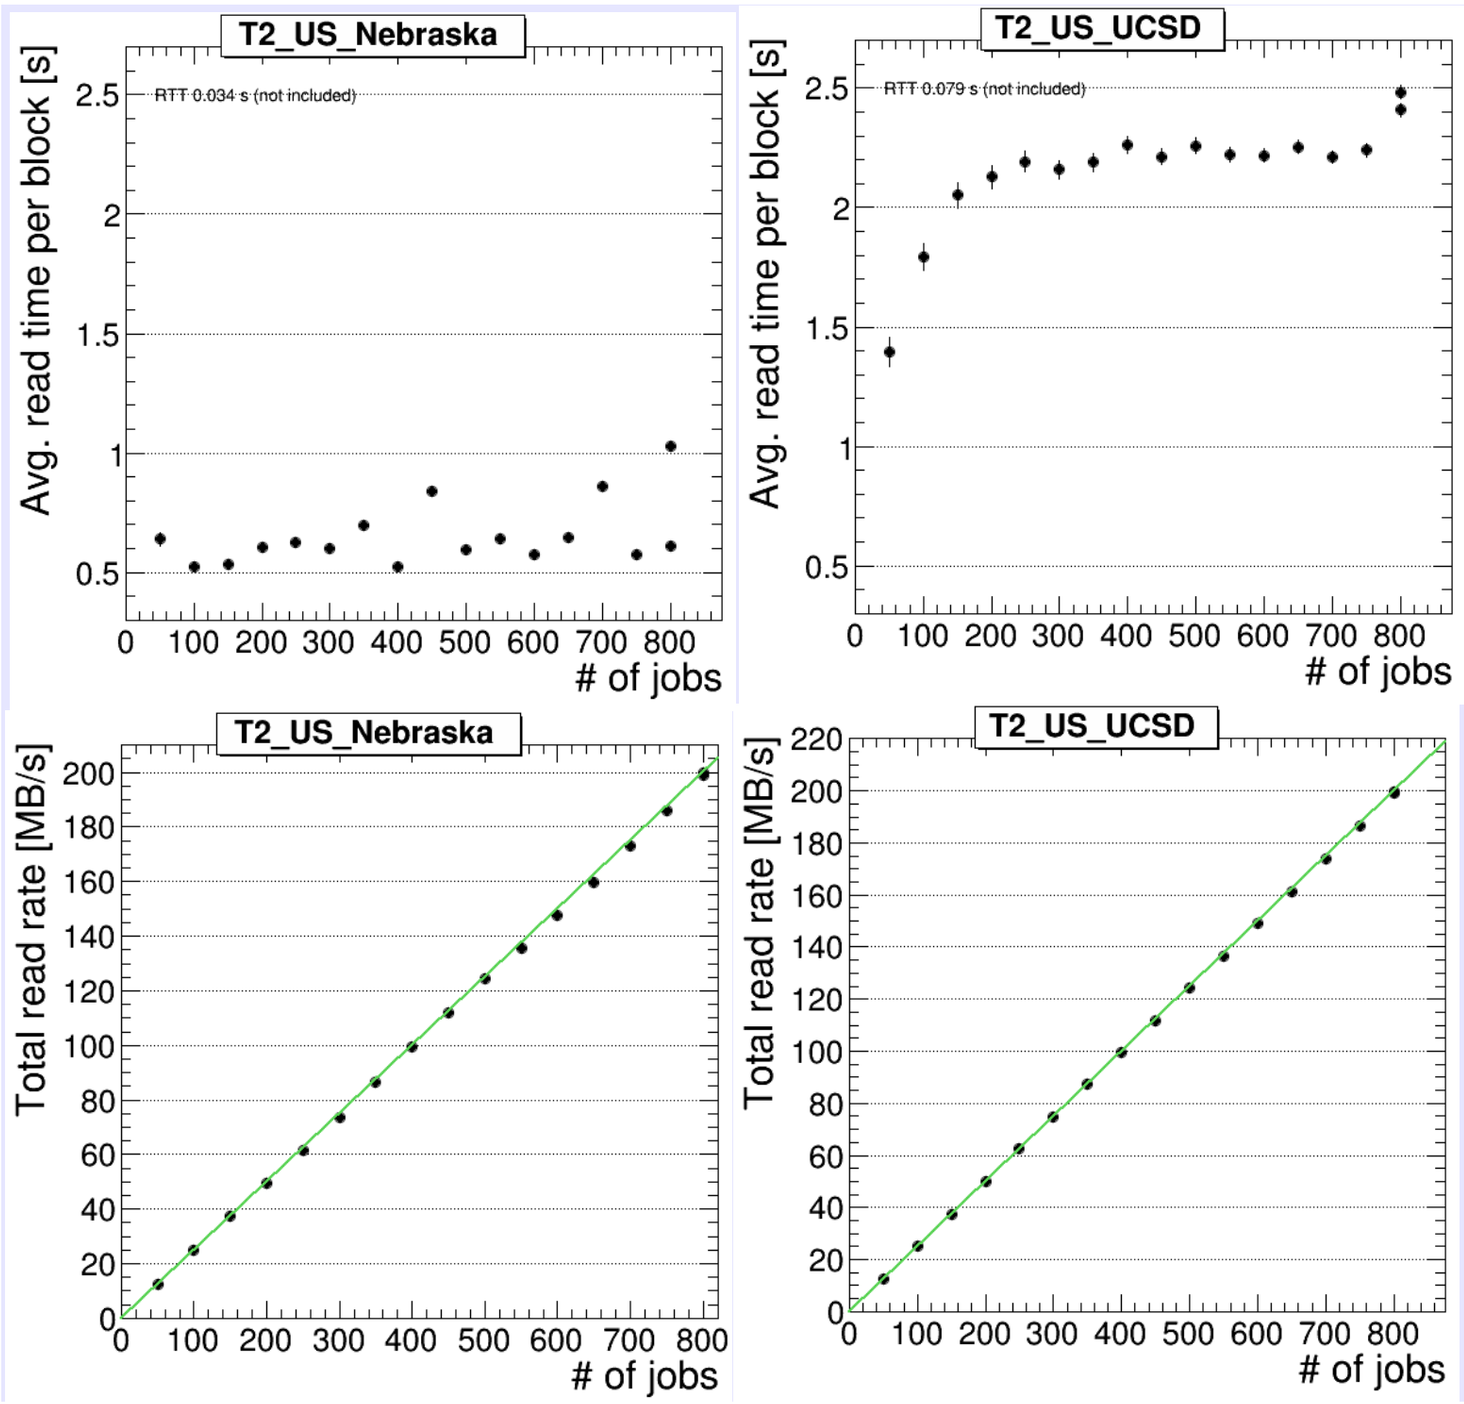

CSA14

- CMS will be holding a readiness challenge in July and August, and AAA testing is among the goals for computing
  - Will try to do "total chaos" test here by centrally forcing jobs to ignore data location and run anywhere
  - Will measure success rates, I/O rates and CPU efficiency for jobs with remote reads, per site, and compare to local reads
  - Usual computing shifters will keep an eye on the health of the system, and AAA team will be available to handle trouble tickets
- If all goes well, users will not notice that they are using AAA….
- We want all sites to be up and running well during the summer to make CSA14 as thorough an exercise as possible

How can sites help?

- Deploy xrootd and join the AAA federation
  - This is largely done; would like to get straggling T2's
- Even easier than that: enable the fallback mechanism
  - As noted, still missing a few sites; too bad, it is helpful for them!
- Enable detailed xrootd monitoring
  - Information about each remote read — user, file, performance, etc.
  - For most systems this is just a few lines in a configuration file
  - For dCache a plugin needs to be installed, only works for >= 2.6
  - Currently miss 3 T1 sites and 21 T2 sites from detailed monitoring
  - We are flying blind without it, hard to evaluate CSA14 metrics at sites that do not have the full monitoring deployed
  - But: last week we realized that there is an issue for multi-VO sites. Will come up with long-term fix, and short-term fix for CSA14.

How can sites help?

- Deploy robust storage systems
  - Sites will be serving data not just to their own batch slots, but (potentially) to many others around the world
  - Good WAN configuration practices also important
- Be responsive to needs of/requests from the experiments
  - We will do our best to alert sites to issues

Outlook

- The data federation, and its implementation through xrootd, has turned out to be a very nice fit with CMS
  - Thanks to robust WAN, straightforward namespace, I/O efforts
  - L. Malgeri, CMS physics coordinator: "It's like a dream come true!"
- We are just starting to understand its implications for the experiment, and for large-scale data management in general
  - LHC Run 2 will be a huge learning experience
  - Much of this learning will take place through working with sites and understanding their experiences
- We're looking forward to future developments in this area and to continued cooperation with the community