



Hi, I'm Ken

I am an operator of the Nebraska US CMS Tier-2 site.

I am the program manager for the seven US CMS Tier-2 sites.

I am also one of two liaisons for CMS-wide Tier-2 sites.

I consider myself to be pretty dumb about networking.

➡ Excuse me if I misuse terms of art!

I think I've answered the questions addressed to the panel already, so I will mention some other issues, specifically about T2.

Opinions are by and large my own; feel free to disagree.



T2-T2 is good news

I agree that we will probably see an increase in T2-T2 transfers.

- ➔ Some T2's are at least as beefy as some T1's! Can handle the load.

This is potentially good news for the OPN:

- ➔ The seven US CMS T2 sites are pretty well connected to each other. If multiple sites want a dataset that's in Europe, then it only has to be brought to the US once, and then it can be moved around within the US.
- ➔ If not for T2-T2 transfers, then we'd be making multiple transfers across the Atlantic.
- ➔ May not work as well in the other direction -- many T2 sites seem to have better connections to FNAL than to their nearest T1, never mind their T2's.



T2 is different from T1

T0→T1 transfers will be fairly continuous, and will *not* grow dramatically with instantaneous luminosity.

- ➔ We're always going to max out the trigger rate, will just cut harder in trigger as collision rate increases. Event size will increase, but not by orders of magnitude.

Transfers to/from T2 will be bursty, and will grow with integrated luminosity.

- ➔ Biggest bursts will be when a reprocessing pass is complete and new datasets are suddenly available.
- ➔ These datasets will keep growing in size as the LHC keeps running, so the bursts will be bigger (longer?) each time.

Given that, a model of scheduling extra bandwidth on demand would seem to be a good fit.



It's the endpoint, stupid

I believe that our transfers are limited not by bandwidth but by site issues.

Example: CMS's Debugging Data Transfers (DDT) effort

- ➔ Trying to commission all T2-T2 channels ($\sim 50 \times 50 \times 2 = 5000$)
- ➔ Had been trying to do this one channel at a time
- ➔ Turns out to be more efficient to do many-to-one or one-to-many at the same time!
 - Look for patterns of problems that arise with any one site
 - One fix can solve problems for many channels

Straightforward assistance to sites could perhaps make a big difference



User analysis stageout

In CMS, jobs go to data. User might have small-ish output from a job that they need to copy back (stageout) to their “home” T2 site.

Failed stageout is currently a major contributor to CMS analysis job failures.

- ➔ Waste of compute time, waste of batch slots when stageout doesn't fail quickly.

These are small transfers, so it's not about the bandwidth -- I think it once again gets back to having working SE's on either end of the transfer.