Summary

- Phil asked me to provide one slide of what I would like to see from networking that we do not have today.
  - For the most part, I don't worry about any networking - it works reliably.
  - When it doesn't work, it's usually a catastrophe & lots of attention & effort go into getting it fixed rapidly
  - My biggest worry is that we can't afford the ideas people are proposing. 10GE is still expensive to deploy today

- I decided to concentrate on network operational items
- All are achievable with 3-6 months of effort for each item
- None of the items really involves any research, just effort.

- (1) Expose network switch information & offsite WAN testing results to CMS monitoring to CMS for integrated view
- (2) Build weathermap plots for major CMS components such as tape, dCache, BlueArc, interactive nodes, workers, etc
- (3) Combine network information with dCache/xrootd code to improve overall robustness & performance

Monitoring

- CMS has 2 Nexus 7000,~10 6509s & recently several 4948Es.
- The networking team provides a lot of details on each port - mrtg graphs of rates, error counters, etc.
    - We have alarms in place to find large-scale problems. For example we know when we've lost connectivity to entire rack of workers. Once alarmed, it is straightforward to determine if it is a network problem or a power problem by looking at this detailed network information.
- However, we are mostly blind to single port issues, excessive error counts, low IO rates. There are some tests for these items on the nodes themselves, but we do not have any tests on the switch information itself.
- There have been plenty of cases, most of them intermittent, where network connectivity between pairs of nodes is lost. This is very hard to debug without a ready-to-deploy network test. When the node is a critical node, this connectivity loss is very serious.   Most of the problems have been traced to duplicate IPs, misconfigurations, duplex issues, & bad patch cords.
- As the data model evolves, experiments need to be more mindful of network issues and have their data transfer software react appropriately

- We need to expose networking monitoring & switch information to the CMS monitoring framework
- We need a proactive, quick & rolling check of every pair-wise connection in CMS subnet,
    - skipping those ports that are already clearly working
    - "force" mode to check specific problematic connections, or everything we own
- Individual port information need to be continuously monitored & a summary of values outside nominal ranges need to reported (and investigated) when found. This includes error counters & IO rates, and especially bonded IO rates.
- Statistics like the top 10 dcache pool "talkers" would help us balance files on the pools to smooth out traffic
- Test failures of external WAN environments should be monitored under the CMS monitoring framework as well, especially if we are going to a more network-centric data model. Right now I count on emails to find this out.

- For example - here is offsite LPCOPN monitoring showing some issues between FNAL & our LHC peers
- (also note that the worldwide network doesn't look perfect as one naively imagines)

To FNAL

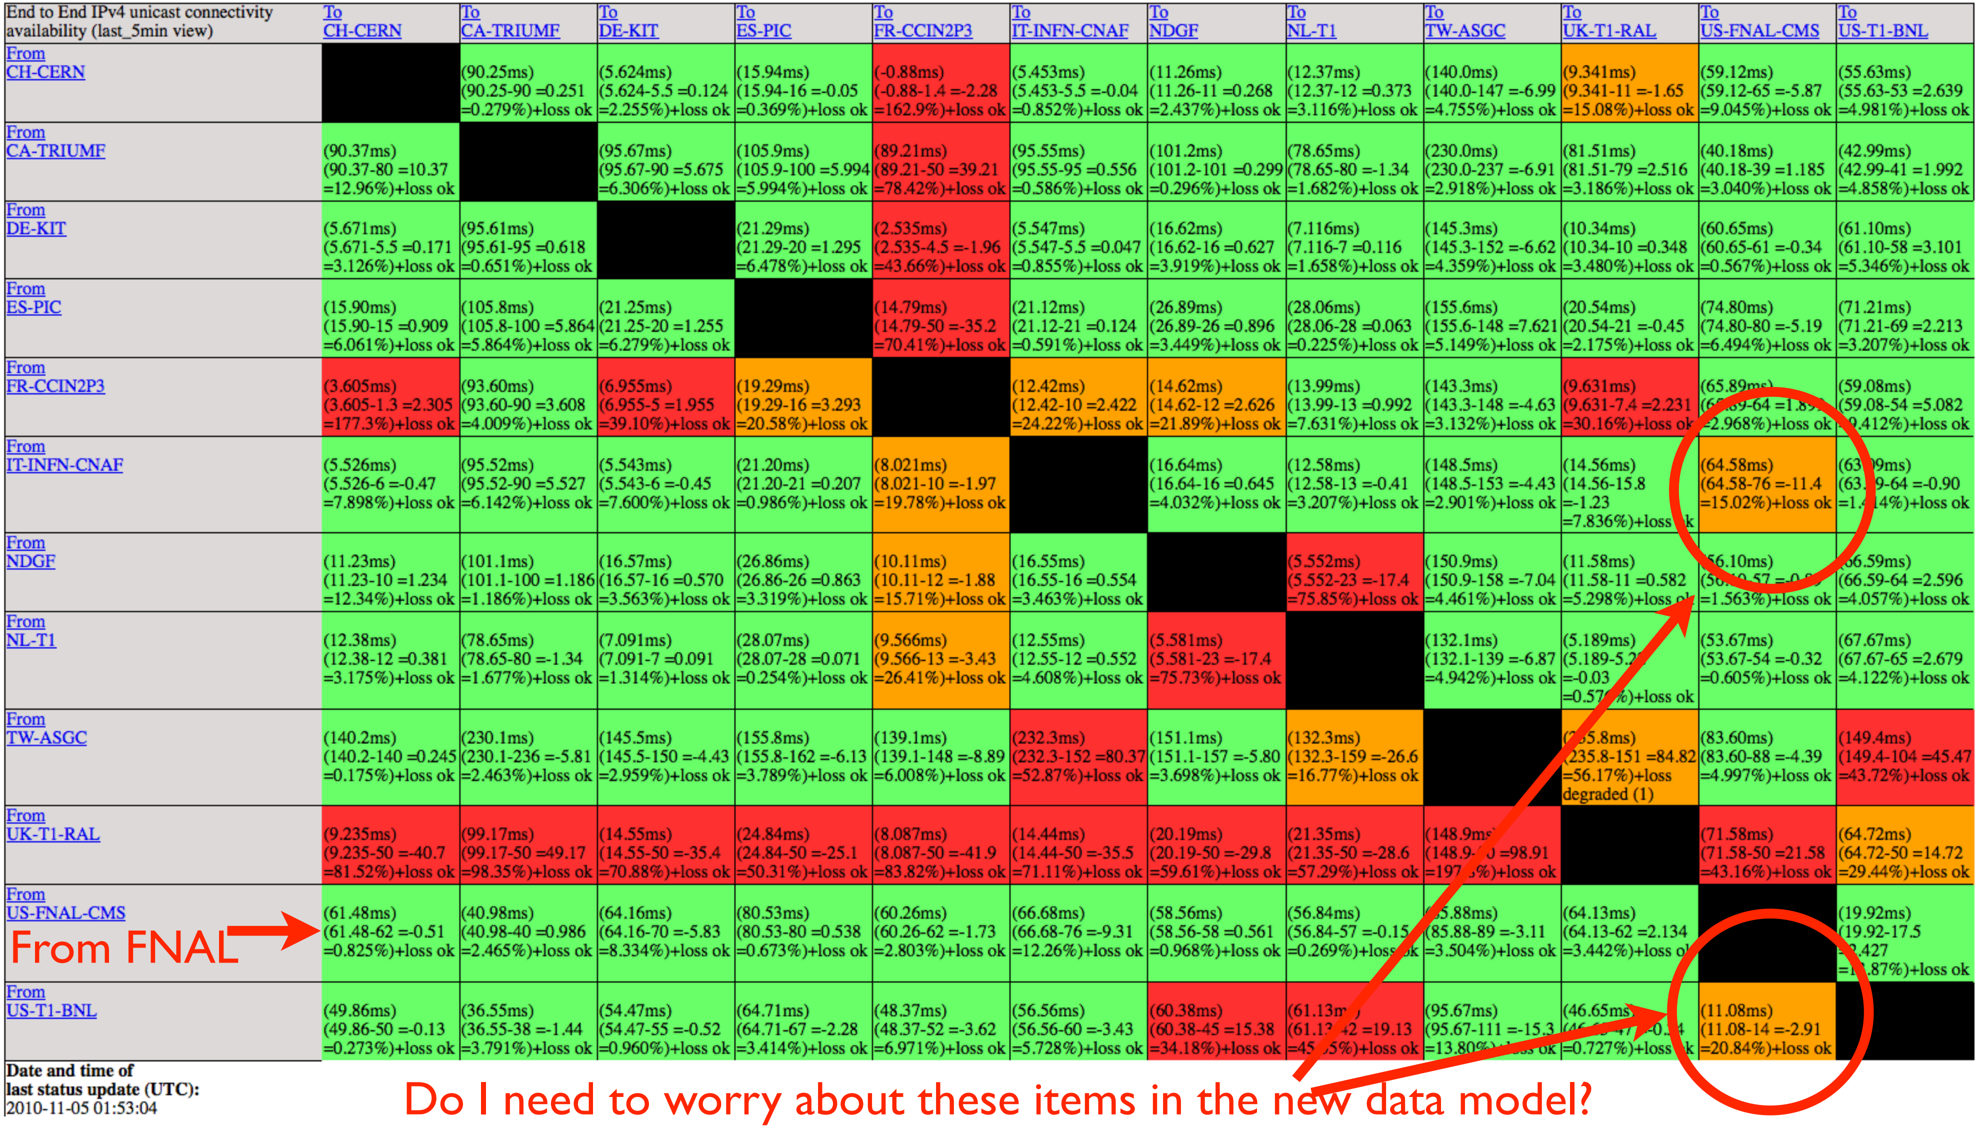

End to End IPv4 unicast connectivity availability (last_5min view)	To CH-CERN	To CA-TRIUMF	To DE-KIT	To ES-PIC	To FR-CCIN2P3	To IT-INFN-CNAF	To NDGF	To NL-T1	To TW-ASGC	To UK-T1-RAL	To US-FNAL-CMS	To US-T1-BNL
From CH-CERN		(90.25ms) (90.25-90 =0.251 =0.279%)+loss ok	(5.624ms) (5.624-5.5 =0.124 =2.255%)+loss ok	(15.94ms) (15.94-16 =-0.05 =0.369%)+loss ok	(-0.88ms) (-0.88-1.4 =-2.28 =162.9%)+loss ok	(5.453ms) (5.453-5.5 =-0.04 =0.852%)+loss ok	(11.26ms) (11.26-11 =0.268 =2.437%)+loss ok	(12.37ms) (12.37-12 =0.373 =3.116%)+loss ok	(140.0ms) (140.0-147 =-6.99 =4.755%)+loss ok	(9.341ms) (9.341-11 =-1.65 =15.08%)+loss ok	(59.12ms) (59.12-65 =-5.87 =9.045%)+loss ok	(55.63ms) (55.63-53 =2.639 =4.981%)+loss ok
From CA-TRIUMF	(90.37ms) (90.37-80 =10.37 =12.96%)+loss ok		(95.67ms) (95.67-90 =5.675 =6.306%)+loss ok	(105.9ms) (105.9-100 =5.994 =5.994%)+loss ok	(89.21ms) (89.21-50 =39.21 =78.42%)+loss ok	(95.55ms) (95.55-95 =0.556 =0.586%)+loss ok	(101.2ms) (101.2-101 =0.299 =0.296%)+loss ok	(78.65ms) (78.65-80 =-1.34 =1.682%)+loss ok	(230.0ms) (230.0-237 =-6.91 =2.918%)+loss ok	(81.51ms) (81.51-79 =2.516 =3.186%)+loss ok	(40.18ms) (40.18-39 =1.185 =3.040%)+loss ok	(42.99ms) (42.99-41 =1.992 =4.858%)+loss ok
From DE-KIT	(5.671ms) (5.671-5.5 =0.171 =3.126%)+loss ok	(95.61ms) (95.61-95 =0.618 =0.651%)+loss ok		(21.29ms) (21.29-20 =1.295 =6.478%)+loss ok	(2.535ms) (2.535-4.5 =-1.96 =43.66%)+loss ok	(5.547ms) (5.547-5.5 =0.047 =0.855%)+loss ok	(16.62ms) (16.62-16 =0.627 =3.919%)+loss ok	(7.116ms) (7.116-7 =0.116 =1.658%)+loss ok	(145.3ms) (145.3-152 =-6.62 =4.359%)+loss ok	(10.34ms) (10.34-10 =0.348 =3.480%)+loss ok	(60.65ms) (60.65-61 =-0.34 =0.567%)+loss ok	(61.10ms) (61.10-58 =3.101 =5.346%)+loss ok
From ES-PIC	(15.90ms) (15.90-15 =0.909 =6.061%)+loss ok	(105.8ms) (105.8-100 =5.864 =5.864%)+loss ok	(21.25ms) (21.25-20 =1.255 =6.279%)+loss ok		(14.79ms) (14.79-50 =-35.2 =70.41%)+loss ok	(21.12ms) (21.12-21 =0.124 =0.591%)+loss ok	(26.89ms) (26.89-26 =0.896 =3.449%)+loss ok	(28.06ms) (28.06-28 =0.063 =0.225%)+loss ok	(155.6ms) (155.6-148 =7.621 =5.149%)+loss ok	(20.54ms) (20.54-21 =-0.45 =2.175%)+loss ok	(74.80ms) (74.80-80 =-5.19 =6.494%)+loss ok	(71.21ms) (71.21-69 =2.213 =3.207%)+loss ok
From FR-CCIN2P3	(3.605ms) (3.605-1.3 =2.305 =177.3%)+loss ok	(93.60ms) (93.60-90 =3.608 =4.009%)+loss ok	(6.955ms) (6.955-5 =1.955 =39.10%)+loss ok	(19.29ms) (19.29-16 =3.293 =20.58%)+loss ok		(12.42ms) (12.42-10 =2.422 =24.22%)+loss ok	(14.62ms) (14.62-12 =2.626 =21.89%)+loss ok	(13.99ms) (13.99-13 =0.992 =7.631%)+loss ok	(143.3ms) (143.3-148 =-4.63 =3.132%)+loss ok	(9.631ms) (9.631-7.4 =2.231 =30.16%)+loss ok	(65.89ms) (65.89-64 =1.89 =2.968%)+loss ok	(59.08ms) (59.08-54 =5.082 =9.412%)+loss ok
From IT-INFN-CNAF	(5.526ms) (5.526-6 =-0.47 =7.898%)+loss ok	(95.52ms) (95.52-90 =5.527 =6.142%)+loss ok	(5.543ms) (5.543-6 =-0.45 =7.600%)+loss ok	(21.20ms) (21.20-21 =0.207 =0.986%)+loss ok	(8.021ms) (8.021-10 =-1.97 =19.78%)+loss ok		(16.64ms) (16.64-16 =0.645 =4.032%)+loss ok	(12.58ms) (12.58-13 =-0.41 =3.207%)+loss ok	(148.5ms) (148.5-153 =-4.43 =2.901%)+loss ok	(14.56ms) (14.56-15.8 =-1.23 =7.836%)+loss ok	(64.58ms) (64.58-76 =-11.4 =15.02%)+loss ok	(63.09ms) (63.09-64 =-0.90 =1.414%)+loss ok
From NDGF	(11.23ms) (11.23-10 =1.234 =12.34%)+loss ok	(101.1ms) (101.1-100 =1.186 =1.186%)+loss ok	(16.57ms) (16.57-16 =0.570 =3.563%)+loss ok	(26.86ms) (26.86-26 =0.863 =3.319%)+loss ok	(10.11ms) (10.11-12 =-1.88 =15.71%)+loss ok	(16.55ms) (16.55-16 =0.554 =3.463%)+loss ok		(5.552ms) (5.552-23 =-17.4 =75.85%)+loss ok	(150.9ms) (150.9-158 =-7.04 =4.461%)+loss ok	(11.58ms) (11.58-11 =0.582 =5.298%)+loss ok	(56.10ms) (56.10-57 =-0.89 =1.563%)+loss ok	(66.59ms) (66.59-64 =2.596 =4.057%)+loss ok
From NL-T1	(12.38ms) (12.38-12 =0.381 =3.175%)+loss ok	(78.65ms) (78.65-80 =-1.34 =1.677%)+loss ok	(7.091ms) (7.091-7 =0.091 =1.314%)+loss ok	(28.07ms) (28.07-28 =0.071 =0.254%)+loss ok	(9.566ms) (9.566-13 =-3.43 =26.41%)+loss ok	(12.55ms) (12.55-12 =0.552 =4.608%)+loss ok	(5.581ms) (5.581-23 =-17.4 =75.73%)+loss ok		(132.1ms) (132.1-139 =-6.87 =4.942%)+loss ok	(5.189ms) (5.189-5.22 =-0.03 =0.579%)+loss ok	(53.67ms) (53.67-54 =-0.32 =0.605%)+loss ok	(67.67ms) (67.67-65 =2.679 =4.122%)+loss ok
From TW-ASGC	(140.2ms) (140.2-140 =0.245 =0.175%)+loss ok	(230.1ms) (230.1-236 =-5.81 =2.463%)+loss ok	(145.5ms) (145.5-150 =-4.43 =2.959%)+loss ok	(155.8ms) (155.8-162 =-6.13 =3.789%)+loss ok	(139.1ms) (139.1-148 =-8.89 =6.008%)+loss ok	(232.3ms) (232.3-152 =80.37 =52.87%)+loss ok	(151.1ms) (151.1-157 =-5.80 =3.698%)+loss ok	(132.3ms) (132.3-159 =-26.6 =16.77%)+loss ok		(235.8ms) (235.8-151 =84.82 =56.17%)+loss degraded (1)	(83.60ms) (83.60-88 =-4.39 =4.997%)+loss ok	(149.4ms) (149.4-104 =45.47 =43.72%)+loss ok
From UK-T1-RAL	(9.235ms) (9.235-50 =-40.7 =81.52%)+loss ok	(99.17ms) (99.17-50 =49.17 =98.35%)+loss ok	(14.55ms) (14.55-50 =-35.4 =70.88%)+loss ok	(24.84ms) (24.84-50 =-25.1 =50.31%)+loss ok	(8.087ms) (8.087-50 =-41.9 =83.82%)+loss ok	(14.44ms) (14.44-50 =-35.5 =71.11%)+loss ok	(20.19ms) (20.19-50 =-29.8 =59.61%)+loss ok	(21.35ms) (21.35-50 =-28.6 =57.29%)+loss ok	(148.9ms) (148.9-50 =98.91 =197.8%)+loss ok		(71.58ms) (71.58-50 =21.58 =43.16%)+loss ok	(64.72ms) (64.72-50 =14.72 =29.44%)+loss ok
From US-FNAL-CMS	(61.48ms) (61.48-62 =-0.51 =0.825%)+loss ok	(40.98ms) (40.98-40 =0.986 =2.465%)+loss ok	(64.16ms) (64.16-70 =-5.83 =8.334%)+loss ok	(80.53ms) (80.53-80 =0.538 =0.673%)+loss ok	(60.26ms) (60.26-62 =-1.73 =2.803%)+loss ok	(66.68ms) (66.68-76 =-9.31 =12.26%)+loss ok	(58.56ms) (58.56-58 =0.561 =0.968%)+loss ok	(56.84ms) (56.84-57 =-0.15 =0.269%)+loss ok	(85.88ms) (85.88-89 =-3.11 =3.504%)+loss ok	(64.13ms) (64.13-62 =2.134 =3.442%)+loss ok		(19.92ms) (19.92-17.5 =2.427 =13.87%)+loss ok
From US-T1-BNL	(49.86ms) (49.86-50 =-0.13 =0.273%)+loss ok	(36.55ms) (36.55-38 =-1.44 =3.791%)+loss ok	(54.47ms) (54.47-55 =-0.52 =0.960%)+loss ok	(64.71ms) (64.71-67 =-2.28 =3.414%)+loss ok	(48.37ms) (48.37-52 =-3.62 =6.971%)+loss ok	(56.56ms) (56.56-60 =-3.43 =5.728%)+loss ok	(60.38ms) (60.38-45 =15.38 =34.18%)+loss ok	(61.13ms) (61.13-42 =19.13 =45.55%)+loss ok	(95.67ms) (95.67-111 =-15.3 =13.80%)+loss ok	(46.65ms) (46.65-47 =-0.34 =0.727%)+loss ok	(11.08ms) (11.08-14 =-2.91 =20.84%)+loss ok	

Date and time of last status update (UTC):
2010-11-05 01:53:04

From FNAL

Do I need to worry about these items in the new data model?

Experiment Component IO Rate

- We have measures of IO rates by our applications.   For example, we know dCache rates and daily tape rates

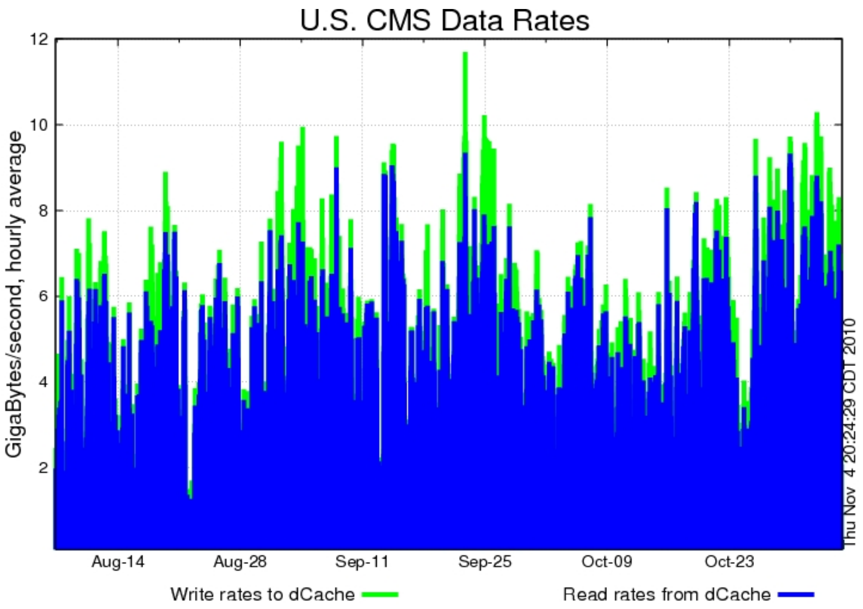

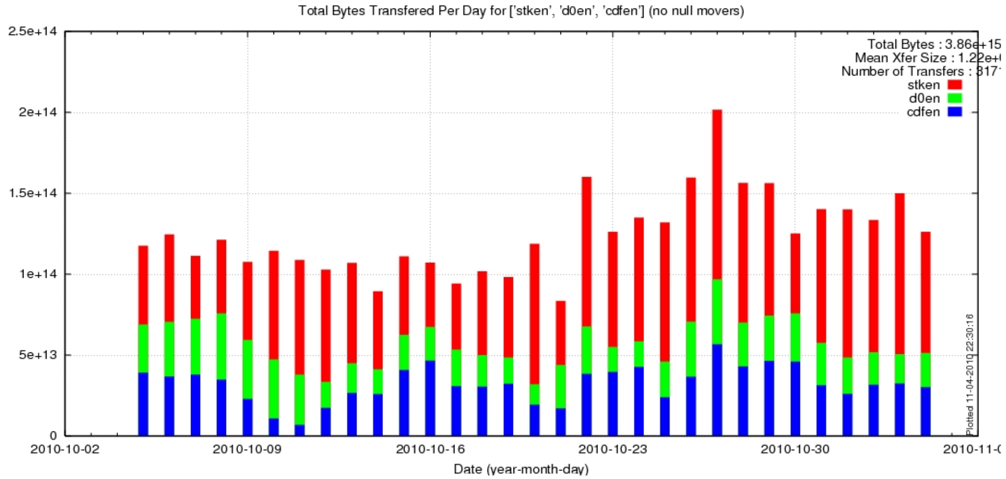

- It would be nice to have this information supplied from the network for each major component in CMS - such as dCache, tape, BlueArc, etc.
- Besides the source component (dCache, tape), rates to destination components should be also tracked. For example, a plot of dCache rates to workers versus interactive nodes would be very useful
- A weather map grouped by experimental component rather than network component is needed.
  - You need to expose which portion of the traffic was standard flows & which were QOS flows.

- Efforts like this CMS weather map are a fantastic start, and it would be great if they could be expanded further as suggested on the previous slide

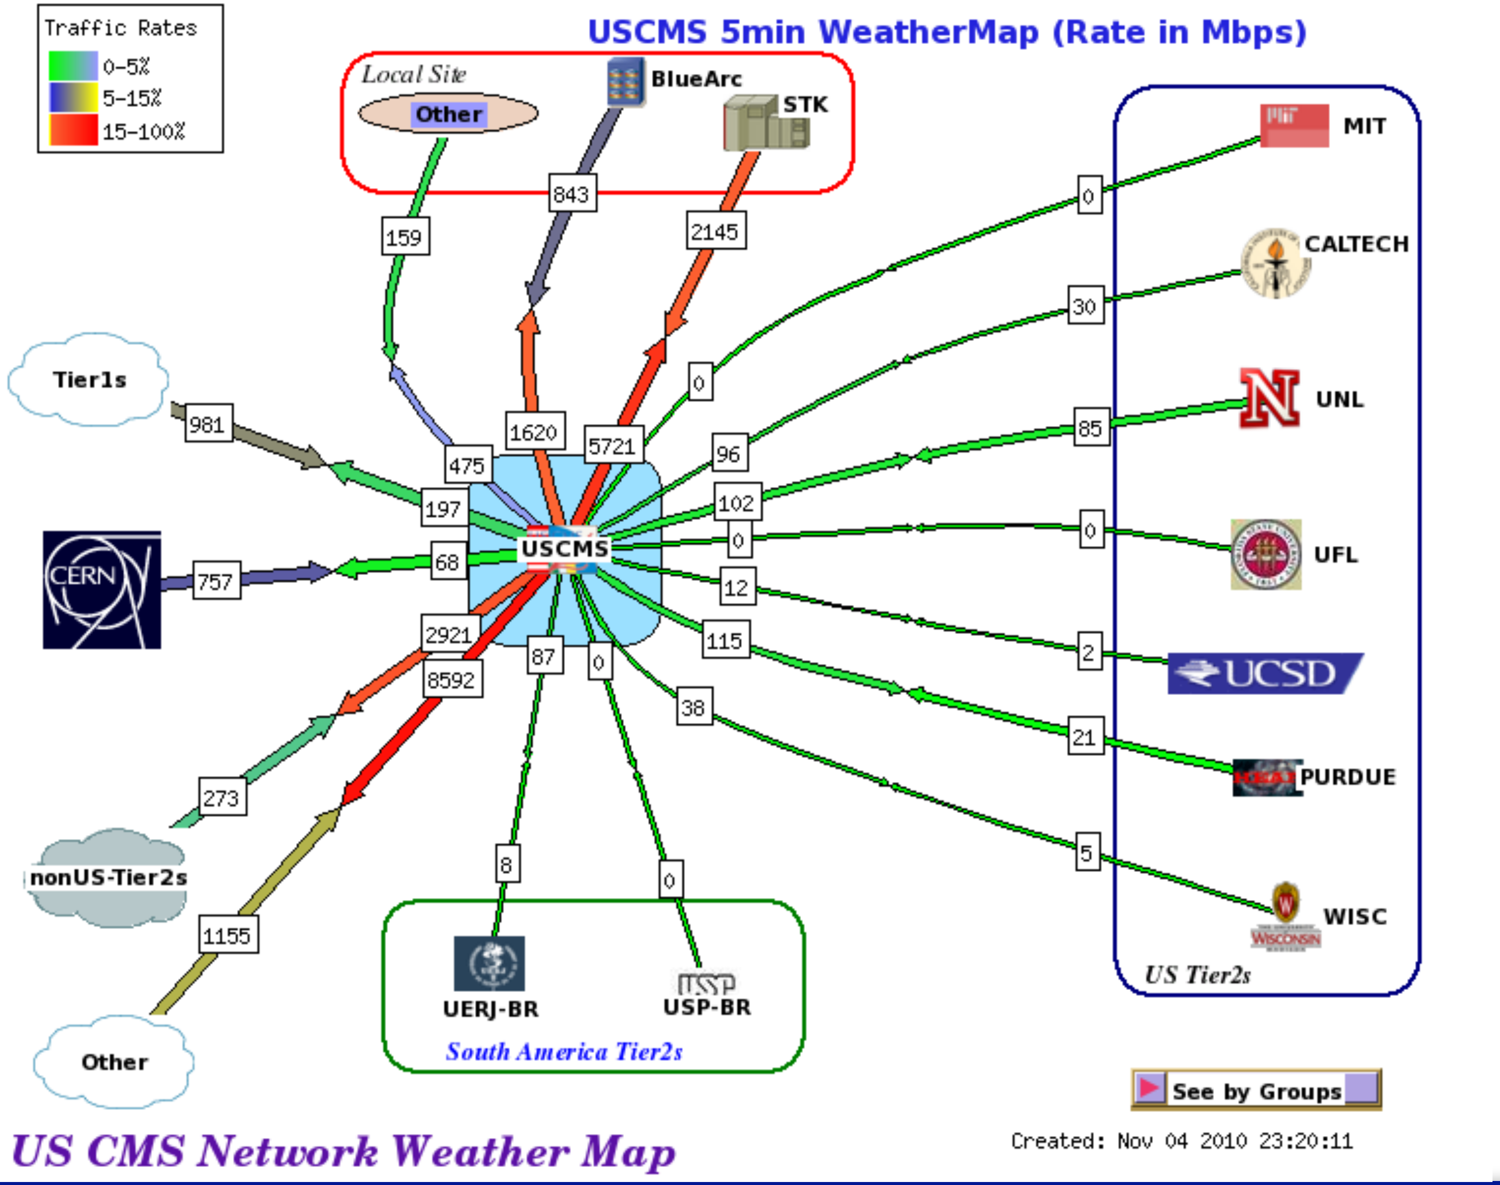

Use network info in data model

- U.S. CMS uses a federated model of data storage, with data currently being delivered by dCache and xrootd. This was a deliberate choice & it provides operational robustness & required performance at acceptable costs.

- This federated solution is not without its problems. We have O(200) storage arrays deployed. When user requests are uniformly spread across all units, the federated dCache/xrootd applications work well. When requests are clumped to a few storage arrays, dCache/xrootd both suffer from resource exhaustion which causes many problems, such as very poor IO performance & re-fetching files from tape when the files are already on disk, etc.

-

- Attempts have been made over many years to fix this in the dCache software, such as counting number of transfers, manipulating cost functions, and dynamically balancing files on each pool. Each attempt has had moderate but incomplete success.

- The goal should not be to set artificial limits that prevent the problem from happening - the goal should be to use the full capabilities of the hardware the experiments purchased. The real problem is there isn't any throttling that evaluates and predicts network performance before assigning transfers to a pool. Detailed immediate past IO rate trends from the network switch should become part of the information evaluated in assigning transfers to data movers.
    - This has to be abstract enough to work in many environments, not just cisco specific ones. The code should accept a generic network info stream, capable of being generated for any vendor. It can be null since it is only one part of the mover assignment process

- The Castor model deploys a back-end network, separate from the user network, to move data between pools. This should be looked at as well. P2P copies require lots of network bandwidth & affect users dramatically.