CMS Castor2 Experience in SC4 and CSA06

On behalf of CMS Computing

Michael Ernst

December 6, 2006

CMS Dataflow

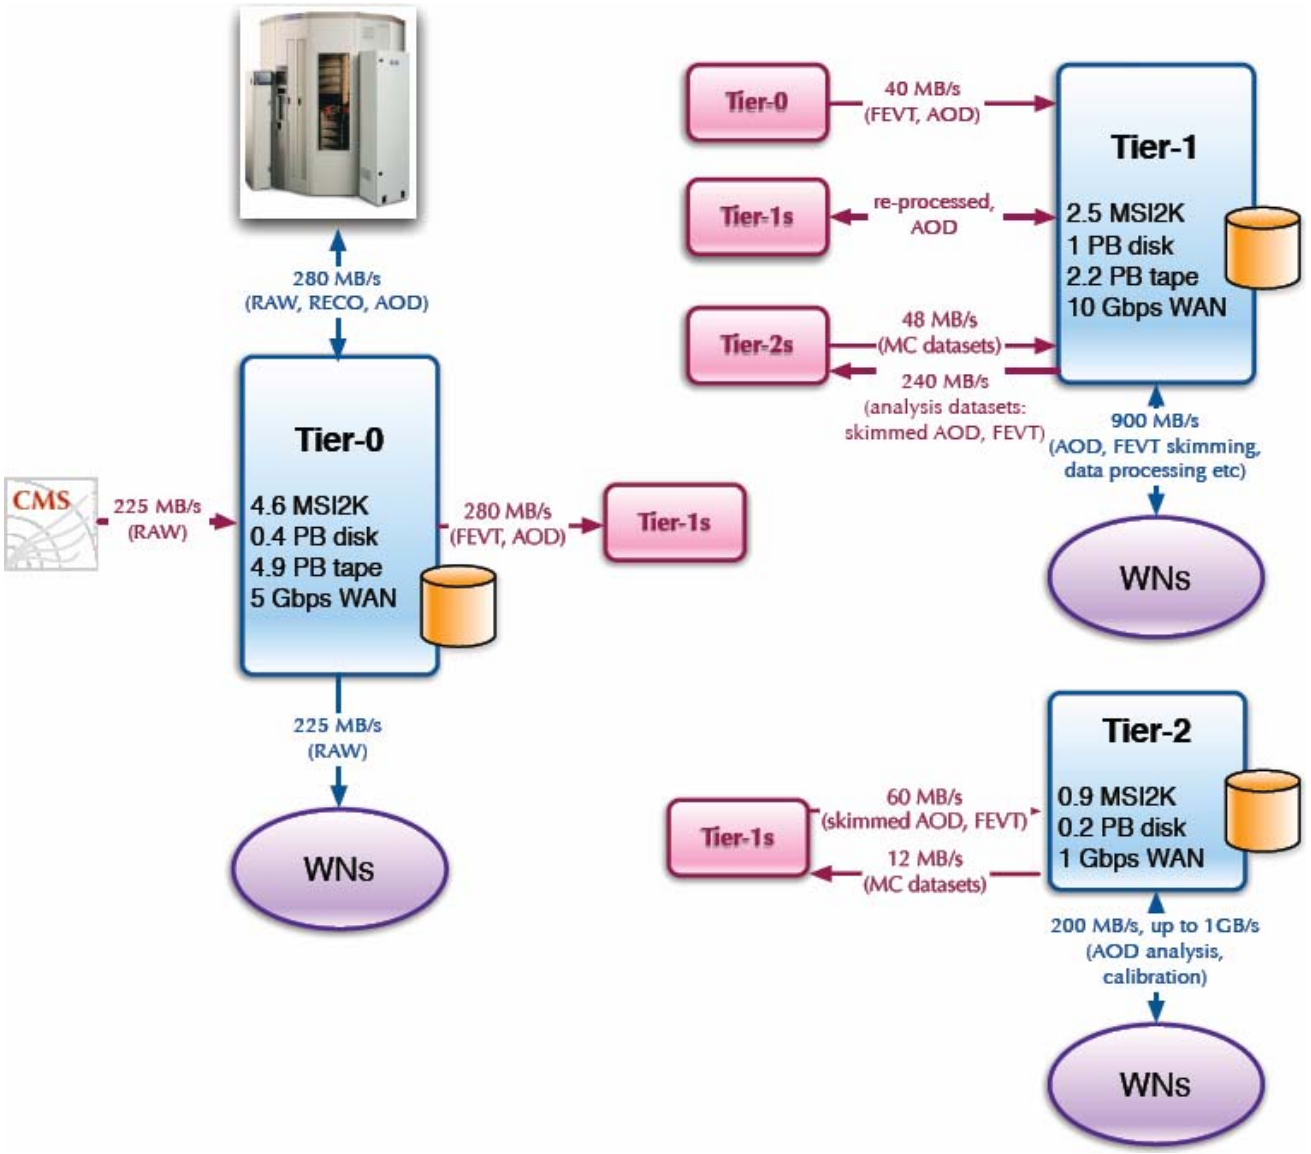

CMS Data Intensive Applications

- Three Application areas are the primary drivers for data access at CERN
  - Magnet Test & Cosmic Data Taking
    - Next Run scheduled for mid October (in parallel to CSA06)
    - Will run for >= 2 weeks
    - O(45)MB/s, 1GB files from the Pit to Castor and from Castor to a few Tier-1 and Tier-2 sites
    - Latency sensitive (Data Quality Monitoring): 10 minutes from close at P5 to availability on WAN pool for application running at CERN
  - Tier-0 Activity
    - 300 MB/s in and out of the t0input buffer (non tape-backed) and t0export (classic tape-backed) pools with file size of ~2GB
    - Total of 400TB of active data in system
      - Factor of 3 more files for CAF (assuming smaller files)
    - Overhead on File Open for disk-resident file should not exceed 2 seconds (on top of network connection time)
    - 40 Hz of reconstruction 20 hours per day in CSA06
      - ~2TB from O(1000) files per day
      - Files prestaged to Castor Disk Buffer
  - Data Replication between CERN and Tier-1 & Tier-2 Centers
    - Data transfer exercises out of CERN to Tier-1 Centers
      - 150 MB/s to tape at destination - PhEDEx driven using FTS Channels
        - ✓ 7 different streams to Tier-1 centers in CSA06
        - ✓ One stream for each Tier-1 center based on relative tape pledges
        - ✓ Pre-challenge tests will start before this weekend
    - Harvesting of Simulated Events for CSA06 Analysis (Tier-1/2 => CERN)
      - O(70)TB
      - Input to Prompt Reconstruction at Tier-0 in CSA06
  - Details on CSA06 on following pages

Pre-CSA06 Exercises: CERN -> Tier-1 Throughput Tests in July

- Goal / Threshold
  - 500 MB/s over one week / 300 MB/s over 3 days
- Preparation
  - Castor WAN Pool: 11 Disk Servers / 5TB each
    - Many thanks to CERN-IT for adding 5 Servers on short notice
  - Reusing datasets that were created as part of CMS Tier-0 test
    - 5 * 4.5 TB, 22.5 TB in total
    - Data placement / distribution organized by Castor
    - Injected in PhEDEx and subscribed to Tier-1 Centers
  - Using PhEDEx over FTS Channels provided by the CERN FTS Server
- Several severe issues discovered
  - Majority was quickly solved by Castor Team

Throughput observed 08/03 – 08/07

Period 08/06 - 08/07

Site	MB/s
FNAL	180
IN2P3	90
FZK	45
RAL	--
ASGC (load test)	35
Aggregate	350

Castor load due to massive/false stage-in requests

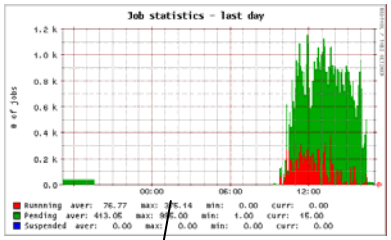

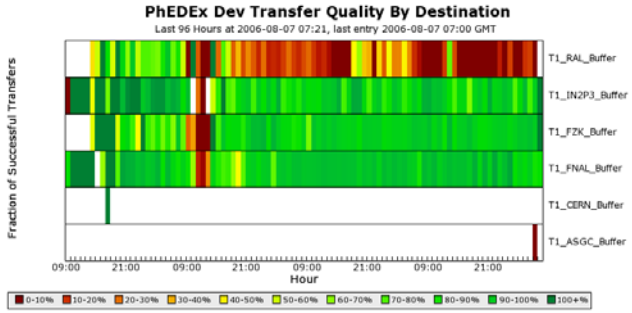

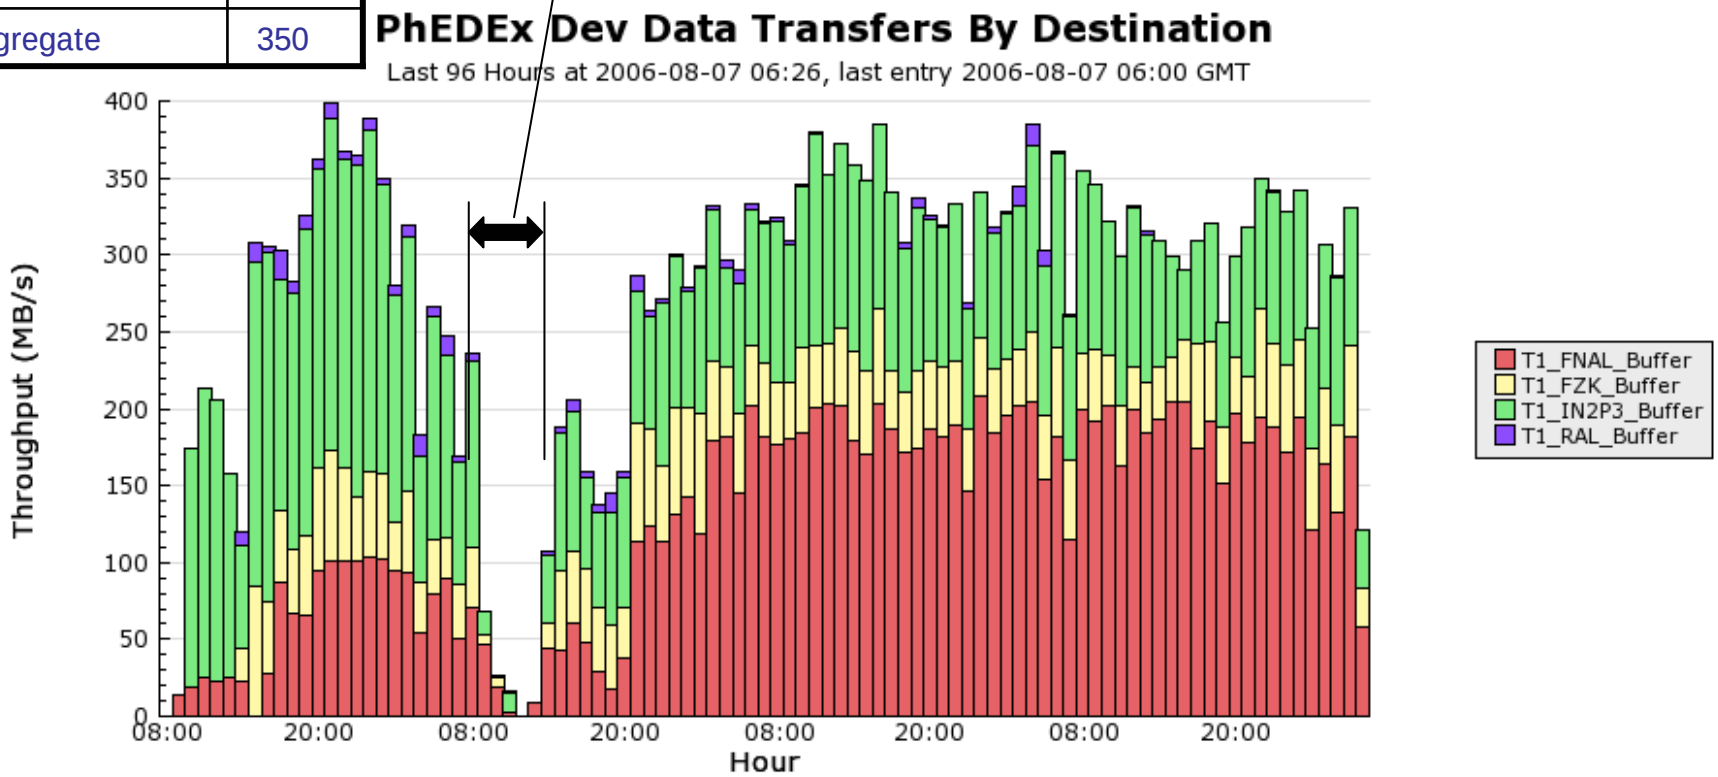

PhEDEx Flooding CMS Castor Job Queue

- User phedex running on lxgate10.cern.ch has submitted 44k prepareToGet requests between 10 am and 1 pm
  - 10am – 11am 15k
  - 11am – 12pm 17k
  - 12pm – 1pm 12k
- 550 Job Slots in "WAN" Queue
  - leaves several servers with no free slots

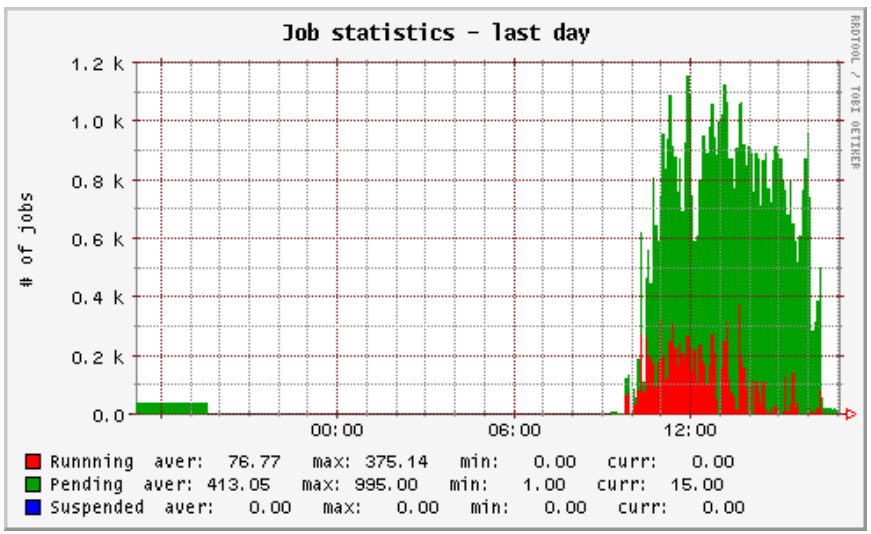

08/04/2006 1 pm

https://lemonweb.cern.ch/lrf-castor/info.php?user=phedex

- PhEDEx Stager Agent has problems w/ stager_qry reply ➡ re-submits stage-in requests at high rate
  - Castor returns diverse (error) strings making the Stager Agent believe the file is not on disk
  - Problem analyzed, temporary fix available since this morning and being tested by CMS
  - Permanent fix expected for mid Sep. 06
- Conclusion
  - all kinds of file access requests launching a mover (i.e. rfio, gridftp, pre-staging/stager_get) require a job slot
  - We can easily create a huge "load" w/o moving a single byte to clients outside CERN
    - Has interrupted DQM during MTCC I

Summary for last 7 Days

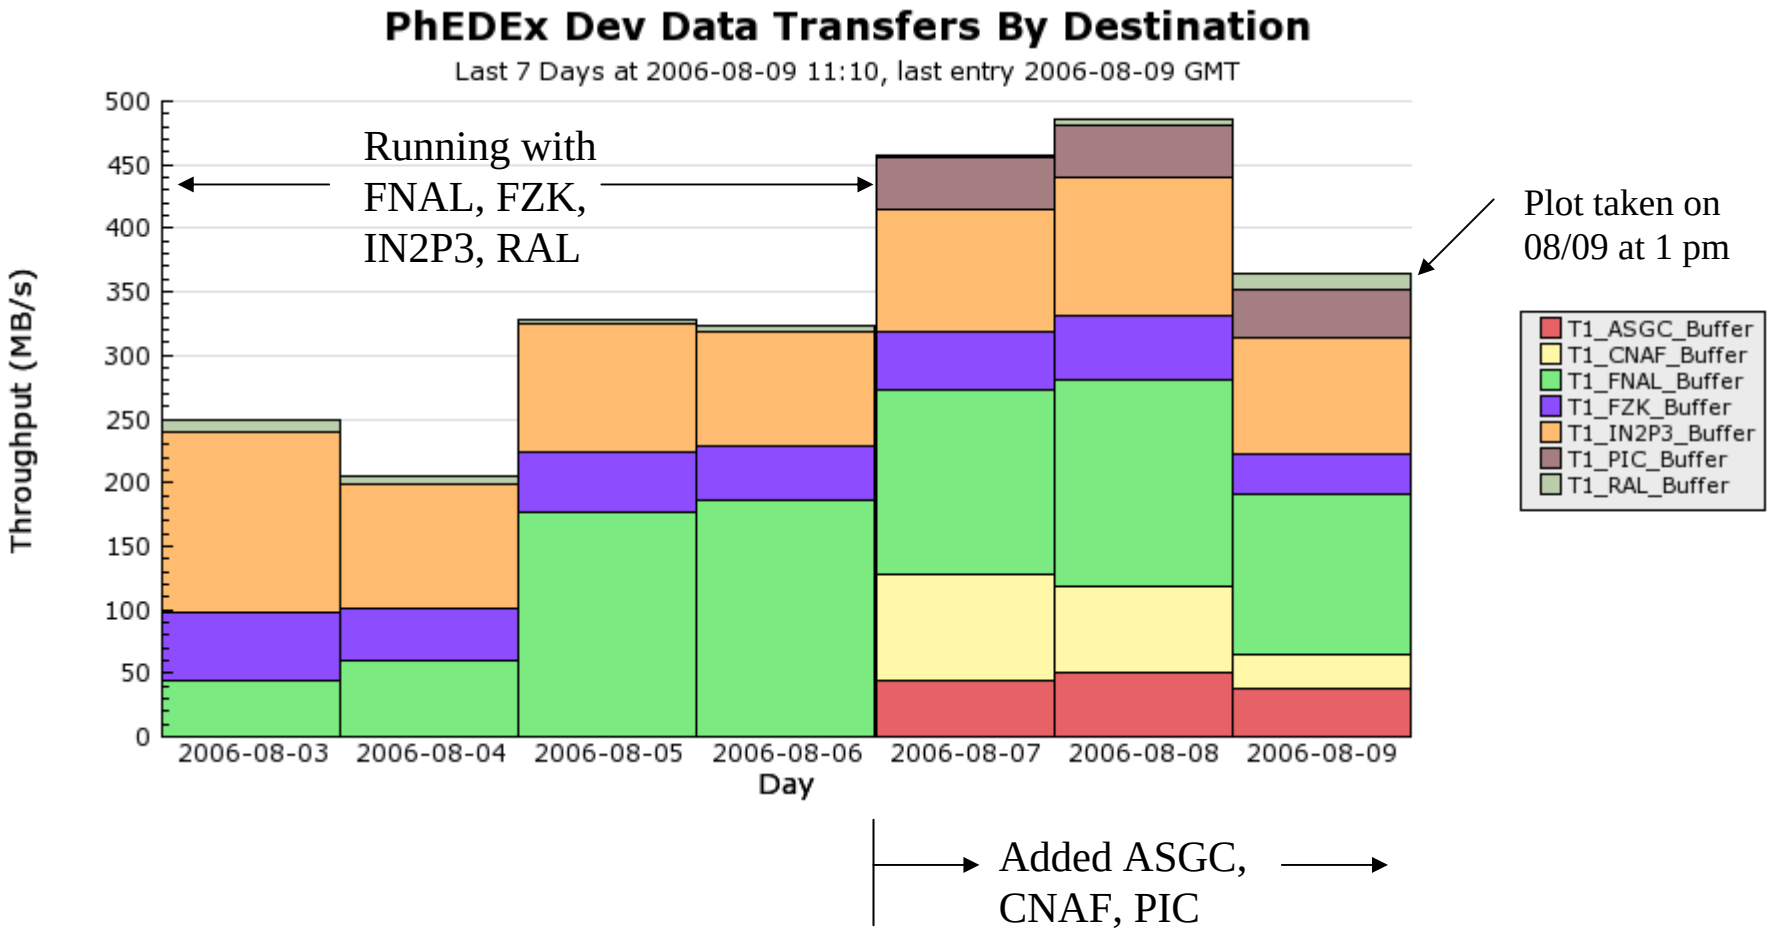

Castor/CMS Collaboration

Though the tests during the summer went reasonably well CMS felt a need for tight integration w/ the Castor Team – the Castor/CMS Collaboration was formed

- Reports to CMS-WLCG Integration Task Force
- Setup a regular (small) meeting between CMS & Castor team
  - Not a technical meeting
- Focusing on understanding each other:
  - What CMS wants to achieve, timescales and the associated metrics
  - Understand what the showstoppers are, what improvements are expected, possibly changes to CMS practices
  - Establish the priorities and timescales for fixing the showstoppers
  - Establish timescales for the expected improvements
  - Agreeing on a process for introducing changes, including communication
  - Iterate over those items
- Participation
  - Tony Cass, Olof Barring, German Cancio, Harry Renshall
  - Michael Ernst + any other relevant CMS person
  - Frédéric Hemmer chairs the meeting
- Frequency
  - Weekly at least at the beginning

CSA06 Computing Metric: Data Transfer

- Tier-0 → Tier-1 (to Tape)
  - Individual goals for each Tier-1, sum to 150 MB/s
    - 25% of 2008 goal
    - 2X the actual sustained rate needed for 40 Hz
      - Exercise clearance of backlog after failure
  - Successful Transfers for 90% of CSA06 interval
    - According to availability metric as defined by WLCG for 2006
  - Not all Tier-1 Centers were prepared to write to Tape
- Tier-1 → Tier-1
  - No such dataflow in CSA06
    - Though exercised to some extent
- Tier-1 →Tier-2
  - Goal: 20MB/s into each Tier-2
  - Threshold: 5MB/s
  - Overall "success" is to have 50% of participants at or above goal and 90% above threshold
  - Successful Transfers for 80% of CSA06 interval

CSA06 Workflows

Execute Analysis Jobs on Tier-1 Resources. Select Data

Transfer Analysis Data to T2s. Validate. Publish Existence Execute Analysis

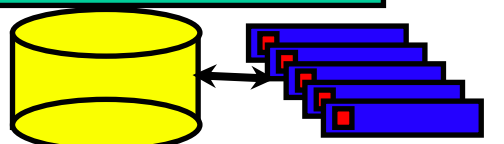

Transfer Analysis Data to T1s. Validate. Publish Existence

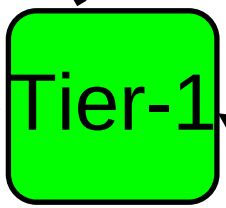

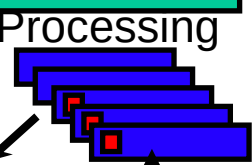

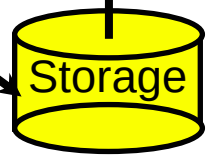

Monitor Success, Rates, Failure Modes, Performance Predefined Metrics

Perform Simulation at T2s Transfer to T1 for Validation

CMS Tier-0 for CSA '06

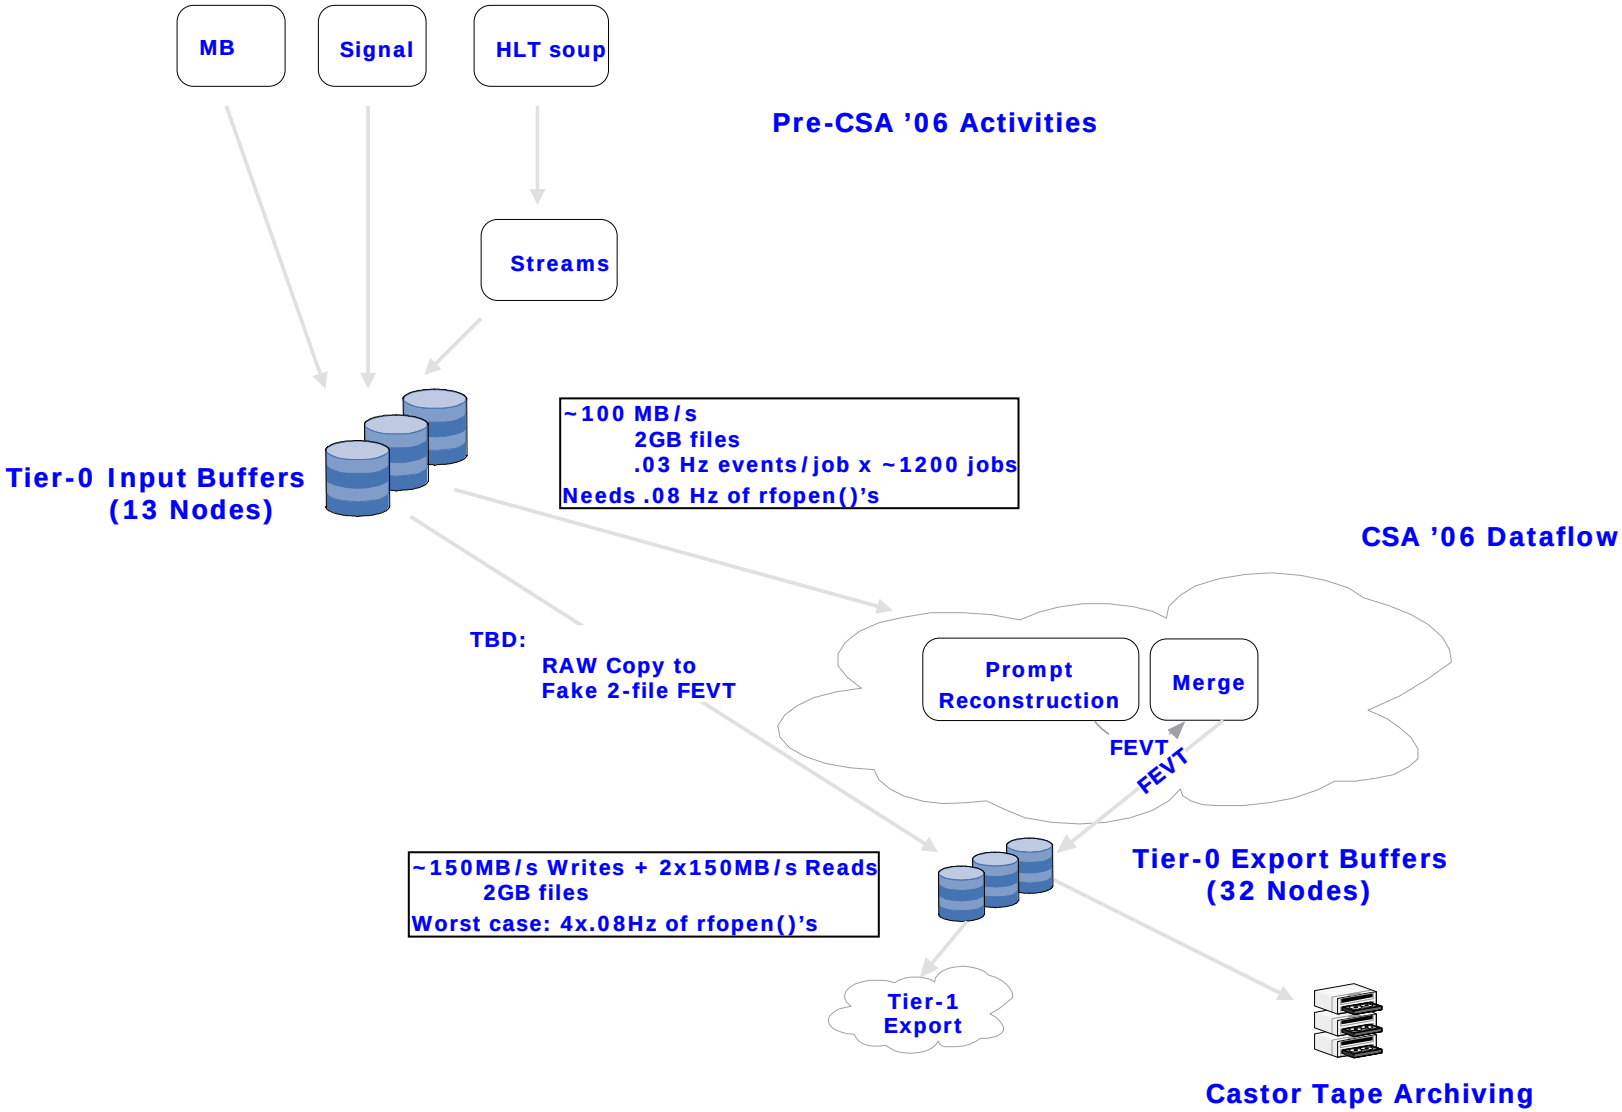

T0 architecture for CSA06: HW

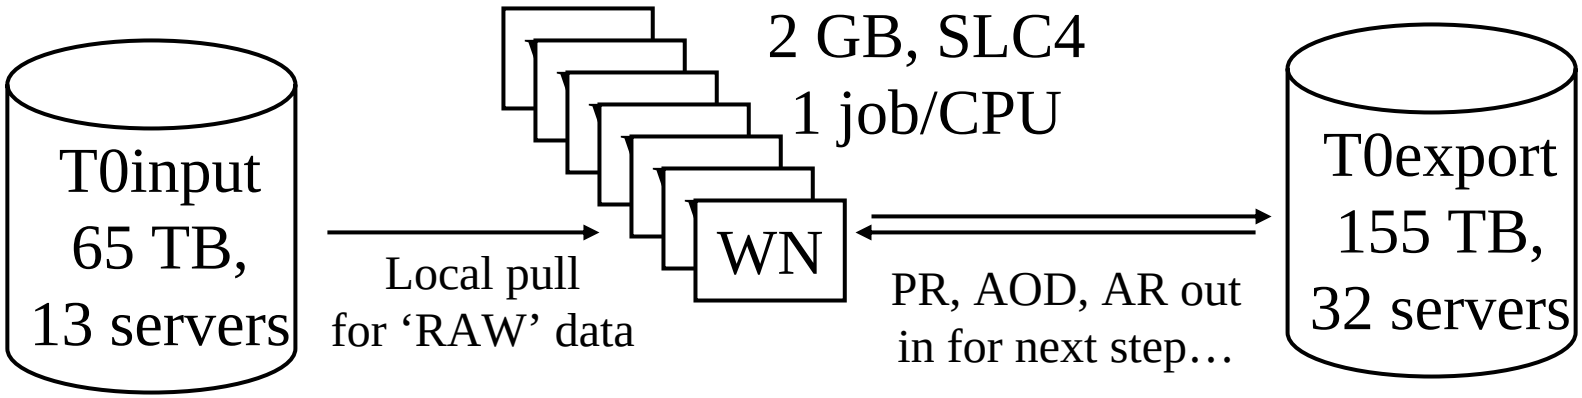

Input data prestaged before start of csa06

lxgate39 (server)

Task steering, logfile collection, job submission

Prompt Reconstruction during CSA06

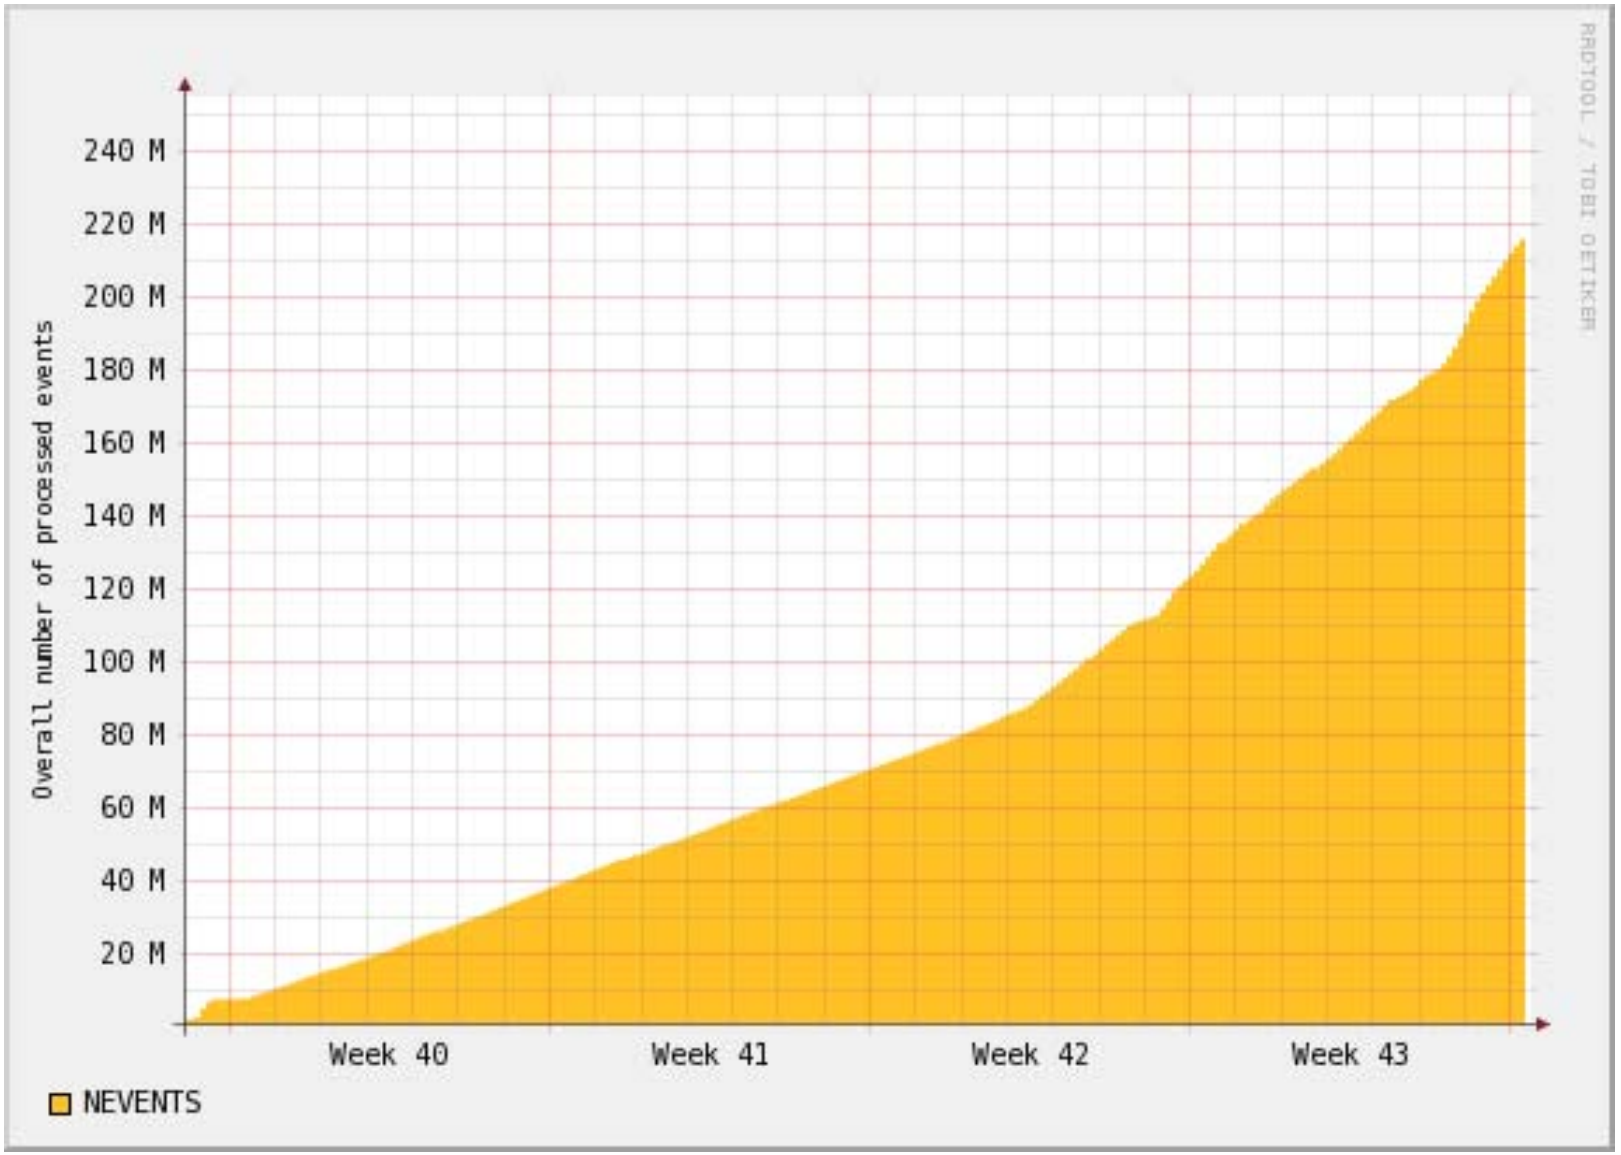

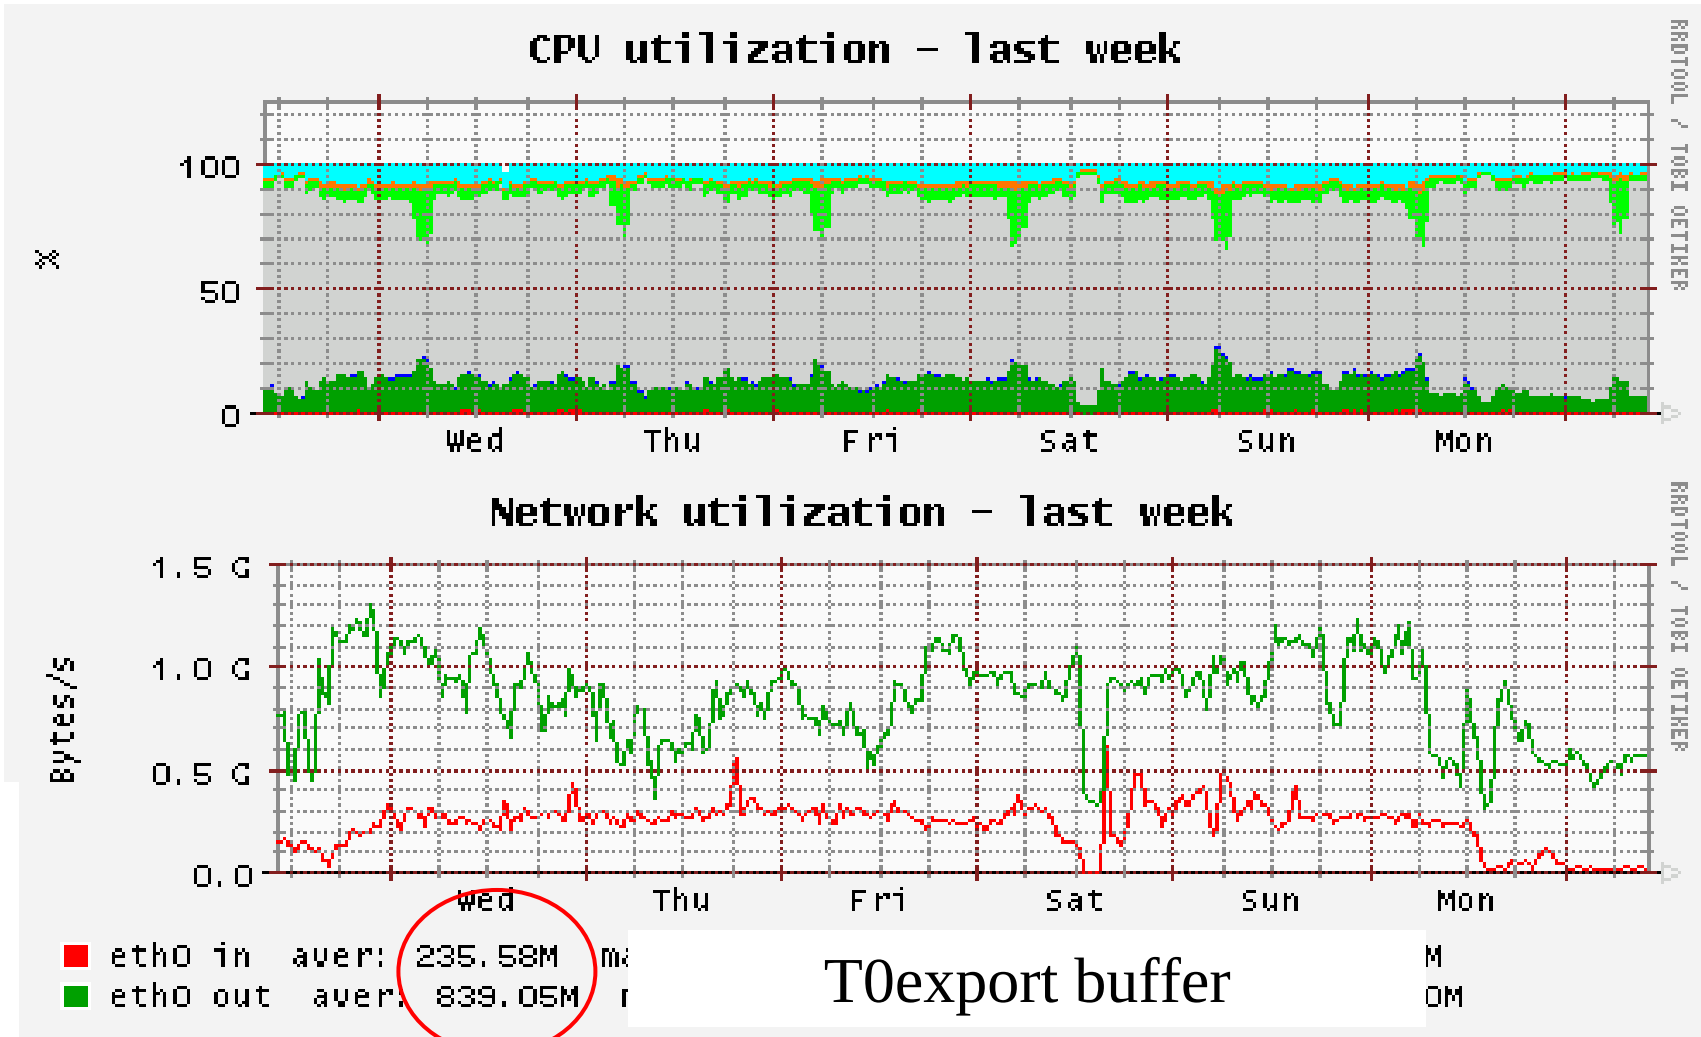
CPU utilization - last week
%
100
50
0
Wed
Thu
Fri
Sat
Sun
Mon
RRDTOOL / TOBI OETIKER
Network utilization - last week
Bytes/s
1.5 G
1.0 G
0.5 G
0.0
Wed
Thu
Fri
Sat
Sun
Mon
RRDTOOL / TOBI OETIKER
eth0 in aver: 235.58M
eth0 out aver: 839.05M
T0export buffer

Issues on Tier-0 Operations

- Processing went very well with an overall job failure rate of well below 1%
  - Small contribution due to rfcp failures

Data Volume moved in CSA06

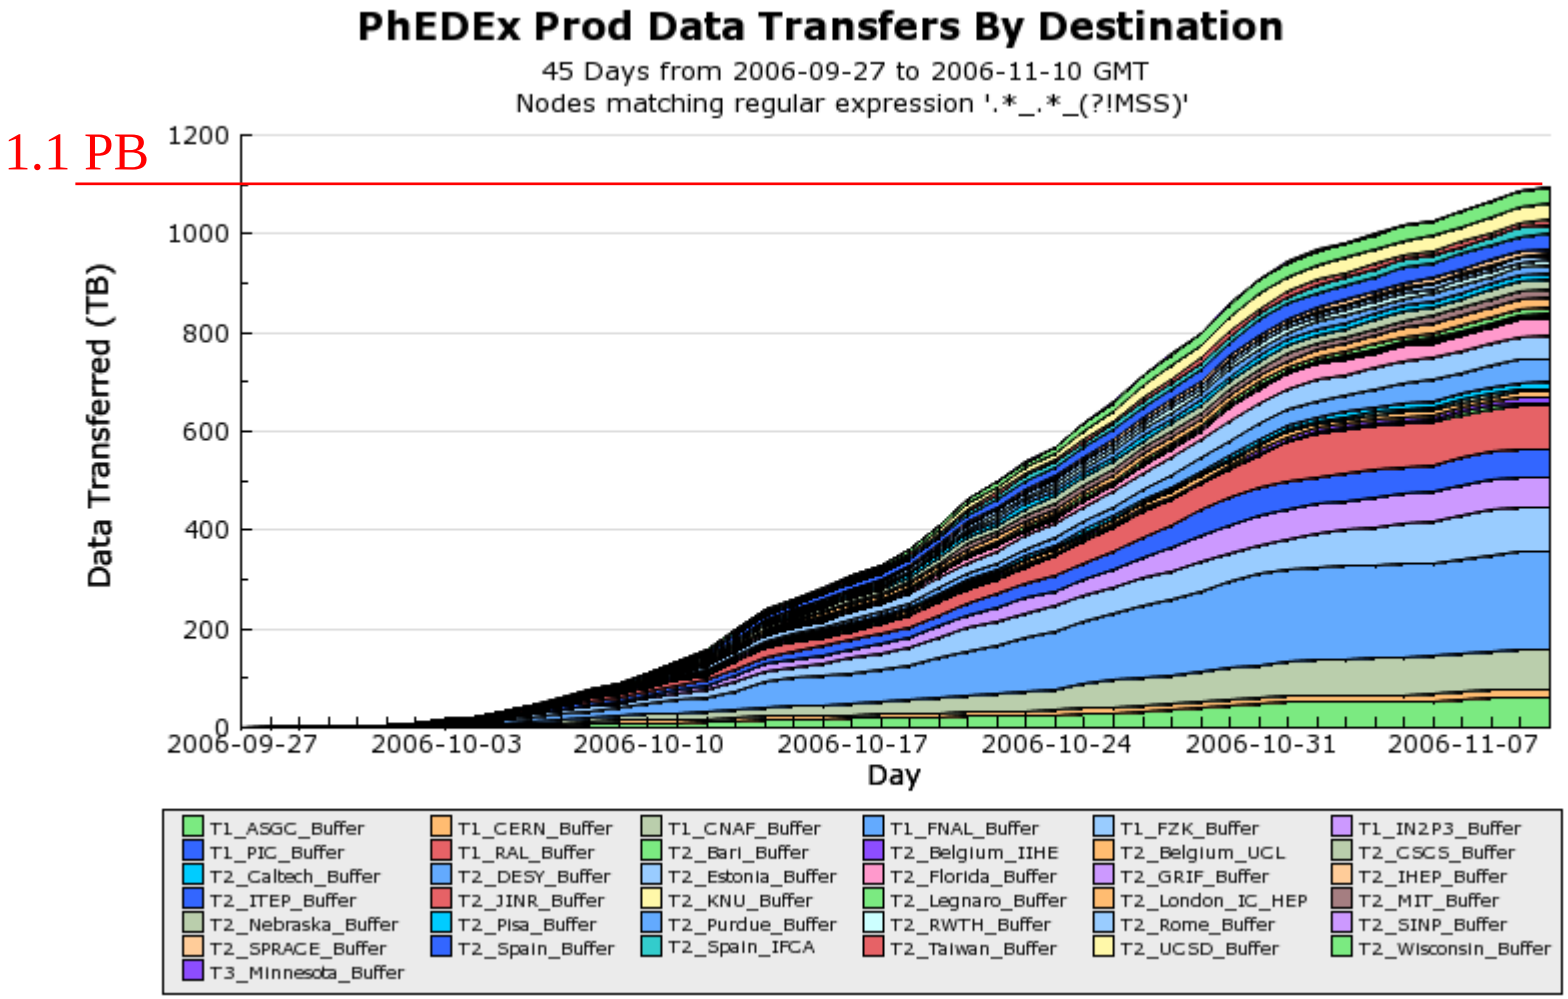

37 Sites, 1.1 PB transferred via the Wide Area Network!

Tier-0 to Tier-1 Throughput for last 30 Days

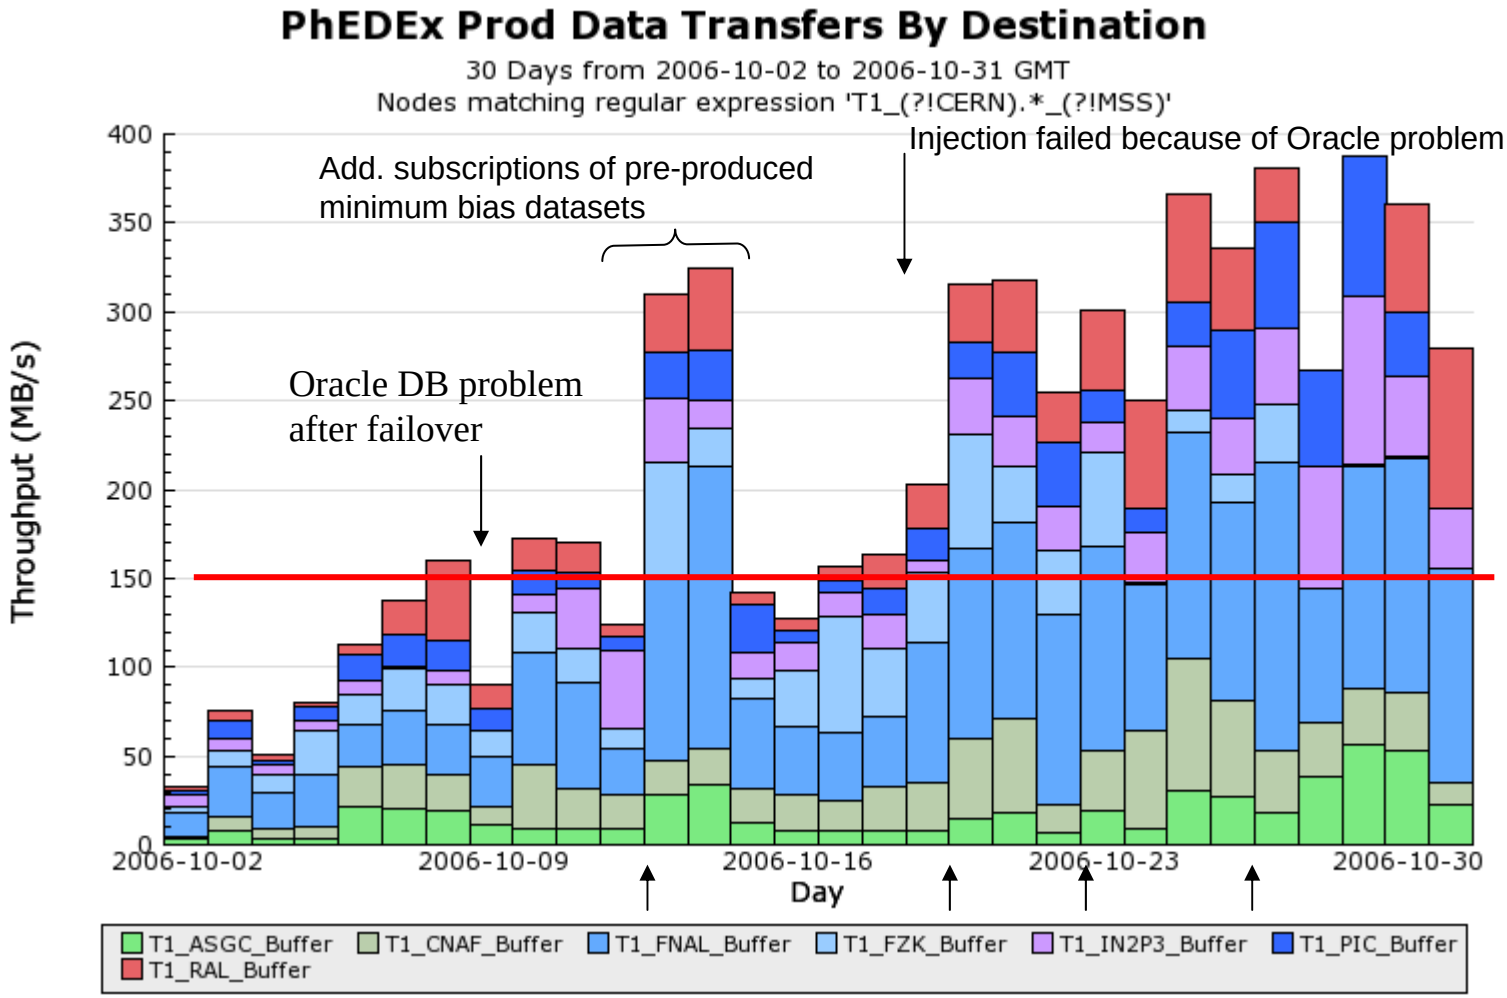

Reconstruction Rate 54Hz 110Hz 170Hz 160Hz

System Response to Subscription of pre-produced Data

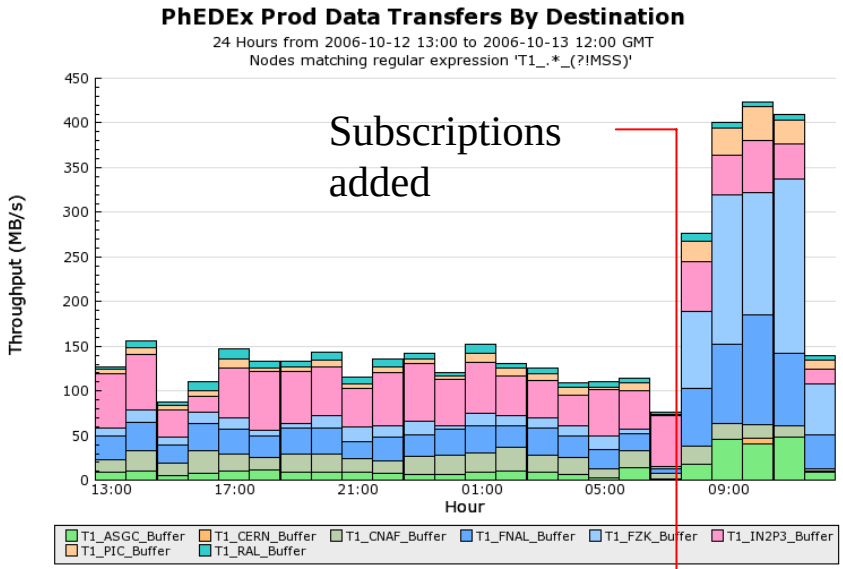

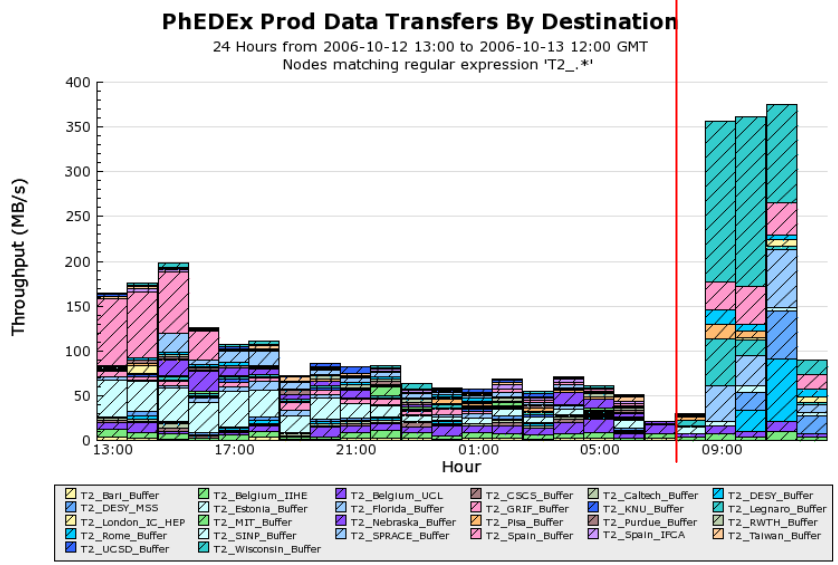

Tier-0 => Tier-1

Tier-1 => Tier-2

Tier-0 to Tier-1 Transfer Quality

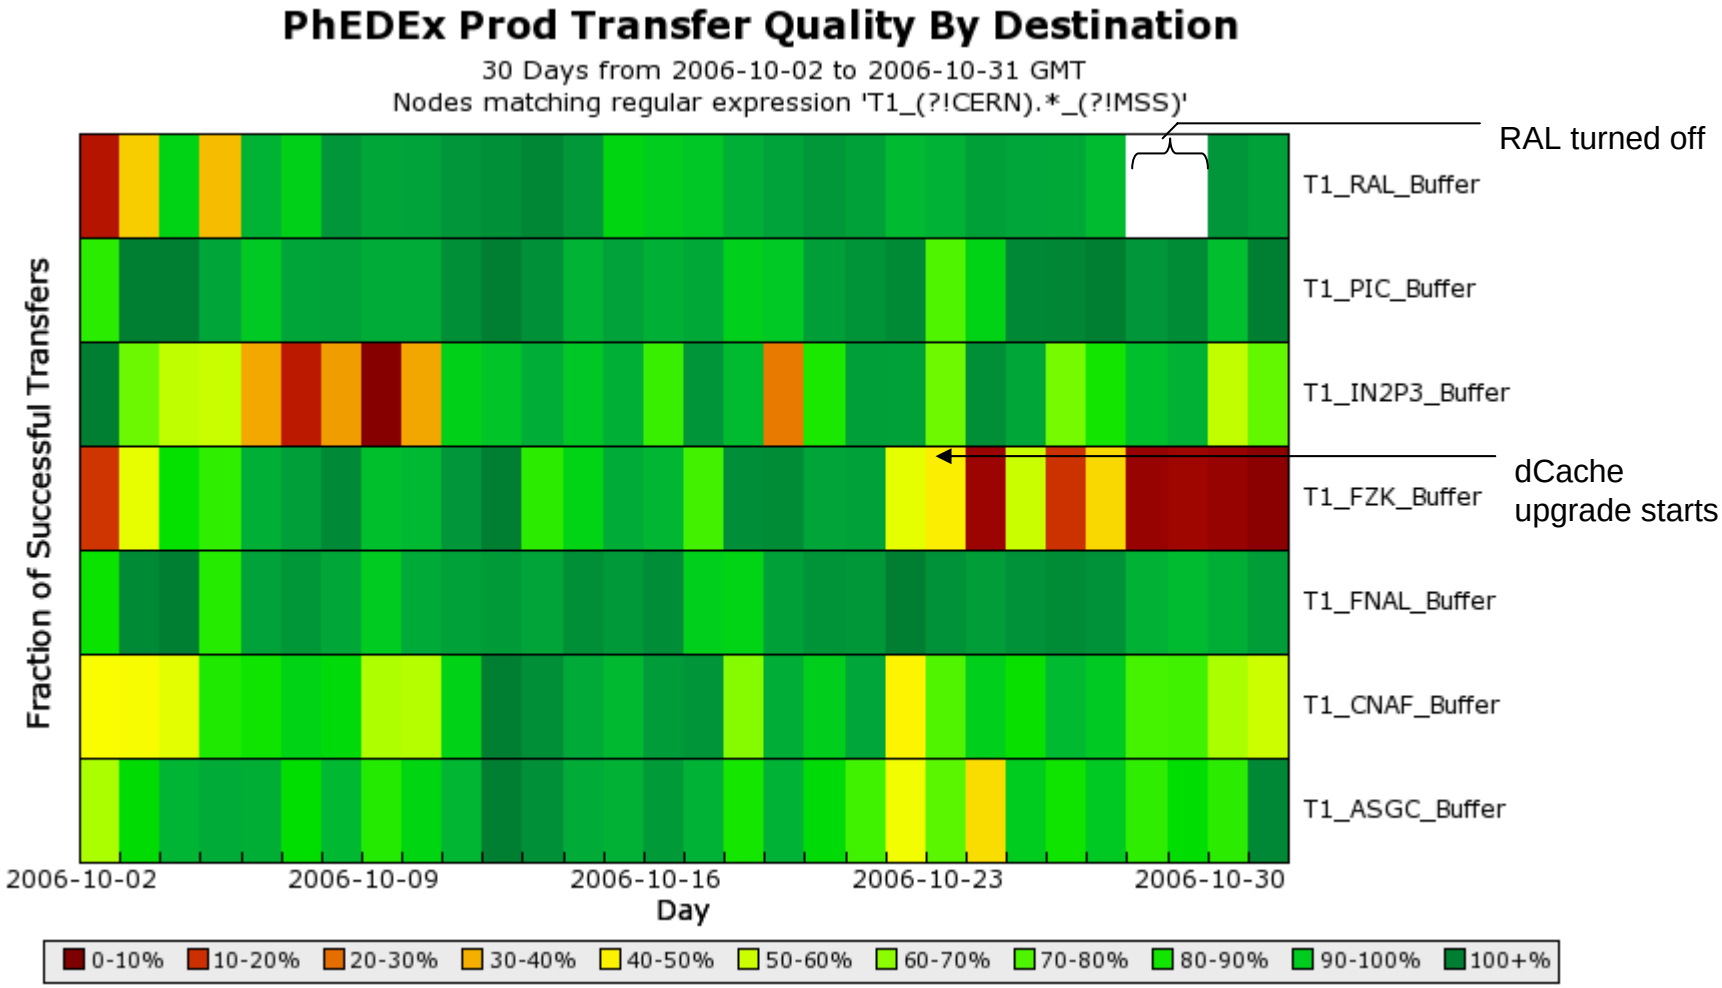

Data Storage and Data Access at CERN

- Castor-2 at CERN
  - Very stable operation throughout CSA06 – only two minor glitches
    - Configuration adjustment after Castor upgrade
    - Restart of the Castor stager (T0Export buffer) triggered an error leading the Tier-0 ingestion process and the data export to Tier-1 centers to hang/timeout
    - Both incidents were resolved within a few hours and Sites quickly cleared the backlog
  - Performed well, not limiting the transfer rates to Trier-1's
  - No noticeable interference with MTCC transfers
    - Exports from different Castor pools, but
    - Using the same queues / scheduling elements in Castor
  - No noticeable interference with traffic from other VOs
    - … but we haven't hit the full target rates for 2008 while other Experiments are running
      - Will start running "multi-VO" transfers w/ ALICE, ATLAS & CMS at nominal 2008 rates this week
  - Excellent support from CERN-IT/FIO
    - Castor Team provided new Castor client lib within a day resolving an interoperability issue between Castor and DPM
    - Good communication at all times (in particular w/ Olof, Jan and Miguel)

Questions and Concerns

- Support
  - CMS has been served well in SC4 and CSA06 through special arrangements – Sustainable in the long run?
- Deployment of Castor2 outside CERN
  - Castor2 services at CERN excellent throughout CSA06
  - While RAL is doing well INFN had/has(?) a lot of problems
- Interface to Large Scale DM Tools
  - Administration tools
  - Service and System Monitoring is crucial
    - System level monitoring is good and very detailed
    - User level Monitoring not available to CMS
      - User applications
      - Tier-0 (Merge and Export Buffers)
    - Disk retention and Garbage Collection policies
    - File Replication & Pinning
    - Latencies (Open & Staging)
    - Active Users

Questions and Concerns

- Performance and Scalability
  - Castor2 at CERN performed well throughout CSA06
    - But CMS has limited activities to the essentials
      - No User activities
      - Added throttling to PhEDEx agents
  - Request processing in a mixed application environment
    - Prioritization
    - ...
- Wrong data delivered on Reads
  - No unique request ID used to ensure “right” client/server connection

Questions and Concerns

- Interoperability
  - CMS hit by two different versions of RFIO for DPM and Castor
  - SRM/GridFTP is (besides RFIO for local data access) the primary interface
- Client libraries
  - Forward/backward Compatibility
  - Validation by experiments?
- Castor Dependency on IT Infrastructure Components
  - AFS
  - DNS
  - other?

Conclusions

- Significant improvements since the review in June
- A number of issues we know of (!)
  - Request processing
  - User level monitoring
  - RFIO compatibility
- Scalability
  - Need multi-VO exercises w/ all experiments running the entire workflows – including the Tier-0
    - Ideally a continuous activity rather than short Challenges
- Support
  - No complaints today, but is there enough effort allocated to it for the significant ramp in 2007, and beyond?