



CMS Experience in Service Challenge 3

Daniele Bonacorsi for the CMS Collaboration
CHEP 2006



Introduction to the Service Challenges

The Service Challenges (SC) are designed as periods of focused activity designed to prepare the distributed computing system for the LHC experiments to be ready for the start of running

- ➔ SC1 and SC2 were mainly throughput challenges
 - SC1 concentrated on disk to disk transfers using any protocol between single sites
 - SC2 concentrated on using the Storage Resource Manager (SRM) interface and increasing the number of participating sites
- ➔ SC3 had both a throughput phase and an experiment service phase
 - The throughput phase was a test of the migration to CASTOR2 at CERN
 - Initially run in July and much more successfully rerun in January 2006.
 - The experiment specific was a more meaningful exercise for CMS and is the bulk of this presentation



SC3 Operations

Integration test of data transfer and data serving infrastructure

- ➔ Computing integration test exercising the bulk data processing portion of the CMS computing model under realistic conditions
- Test end-to-end system of both CMS-specific and LCG services
- Focused validation of data storage, transfer and serving infrastructure, plus required workload components: job submission, RB

Test complexity built in three major steps over 2005

Throughput phase: High-throughput storage+transfer system test

Service I (Sep-Oct): Concurrent transfer+grid jobs to read/write data

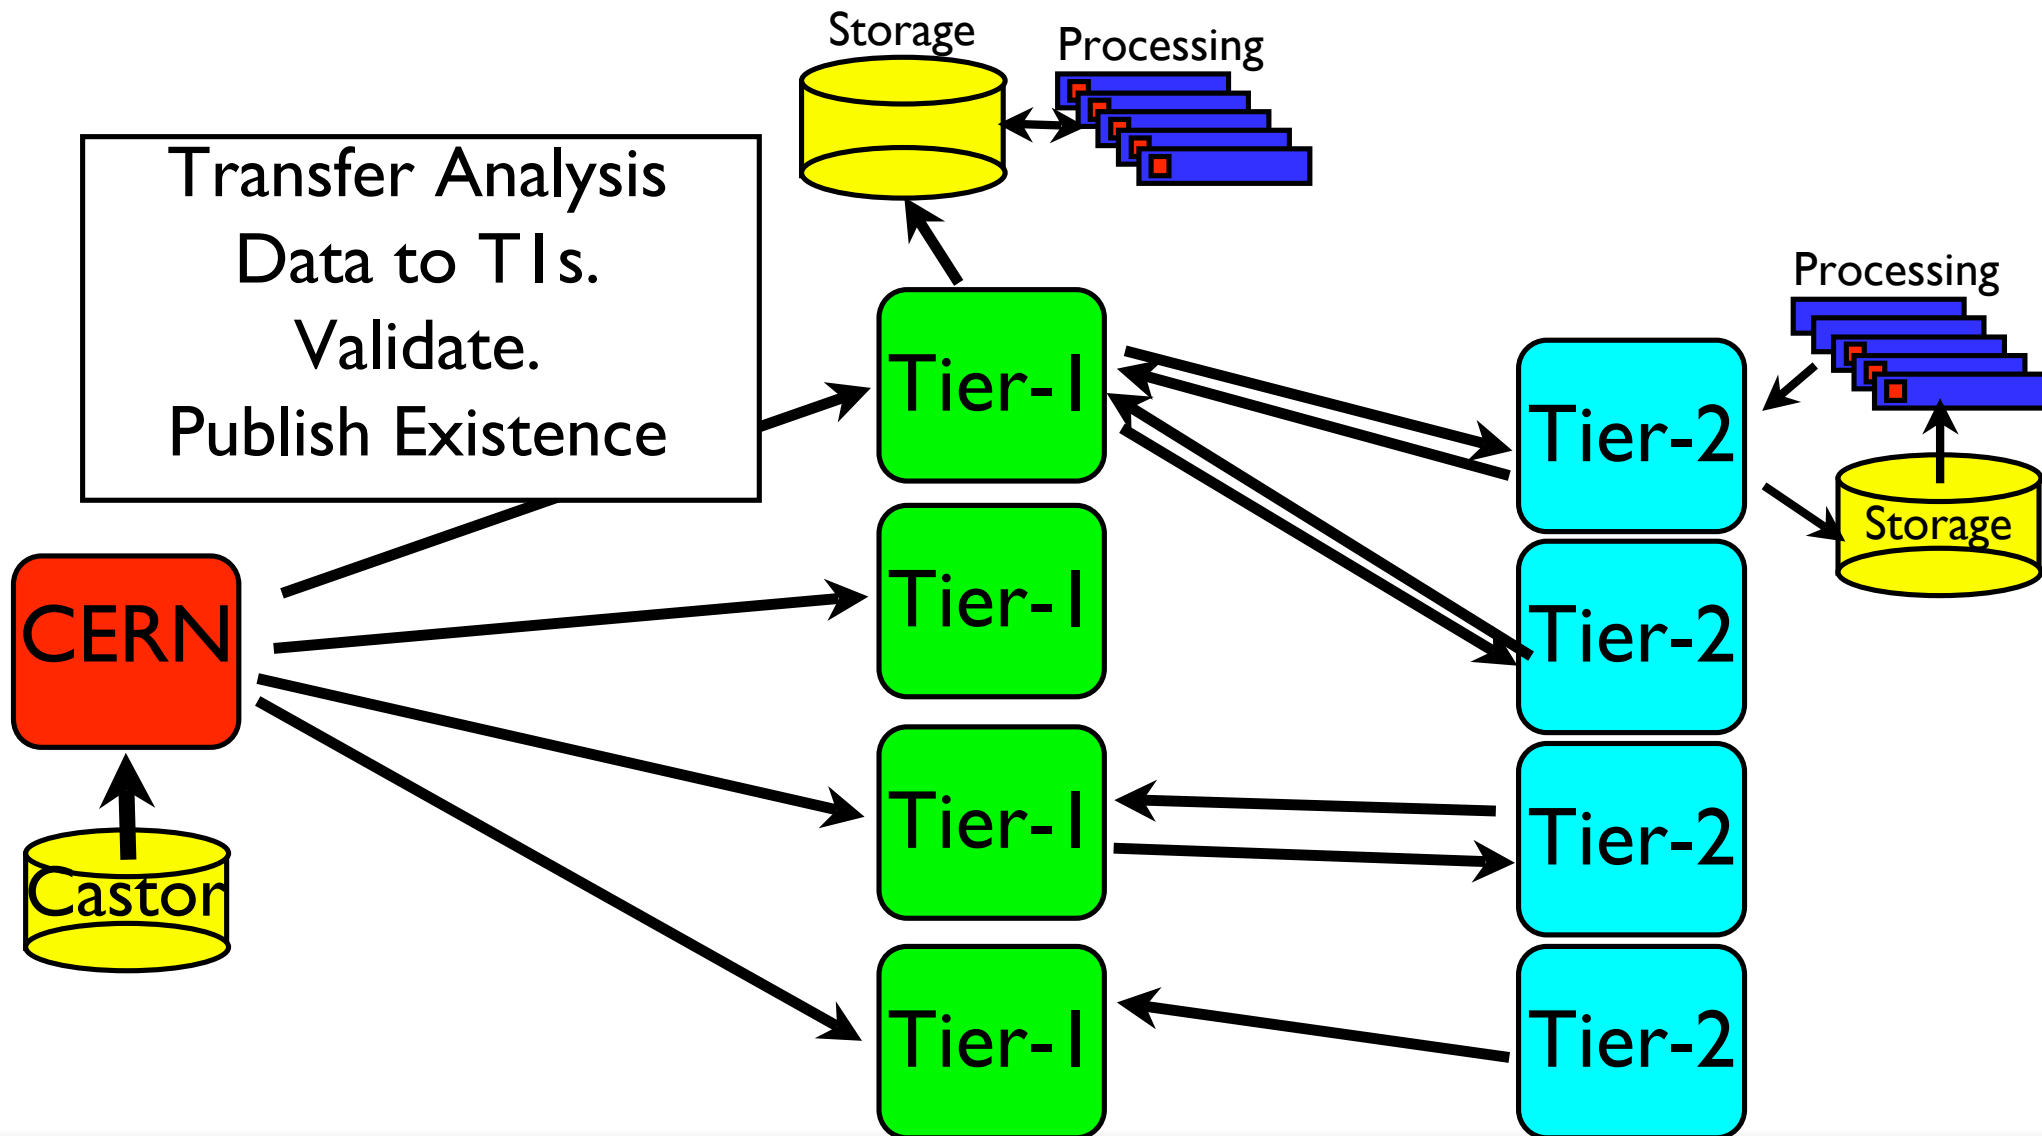
Service II (Nov-Dec): Concurrent data flow following C-TDR

Involving a significant number of Tier-1 and Tier-2 sites

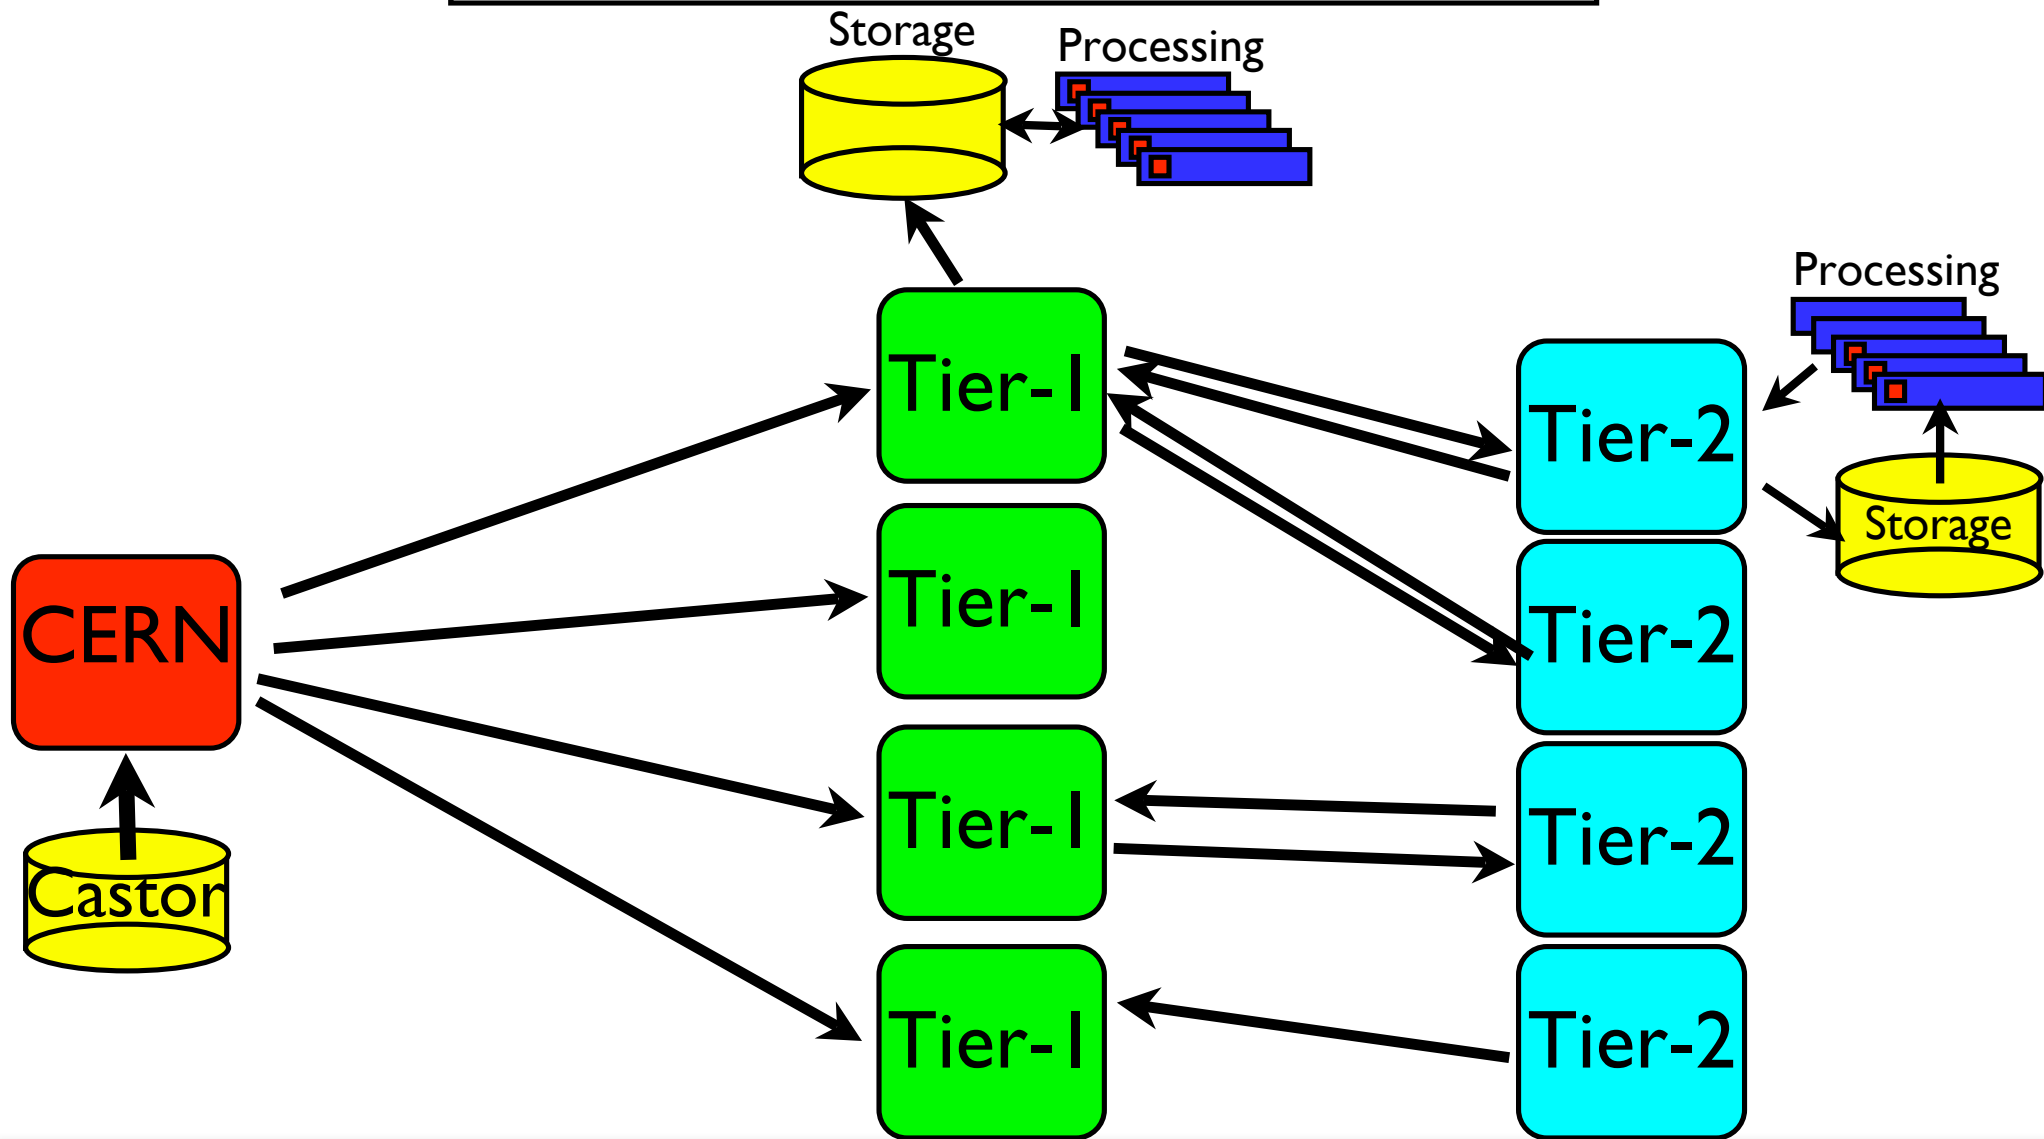
- ➔ CERN + all 7 CMS T1: ASGC, CNAF, FNAL, FZK, IN2P3, PIC, RAL
- ➔ 13 Tier-2s by main T1 affiliation: DESY (FZK); Bari, Legnaro (CNAF); CIEMAT+IFCA (PIC); NCU ASGC); Imperial (RAL); Caltech, Florida, Nebraska, Purdue, UCSD, Wisconsin (FNAL); [Plus: SINP, ITEP, GRIF,CSCS]

Original Goals

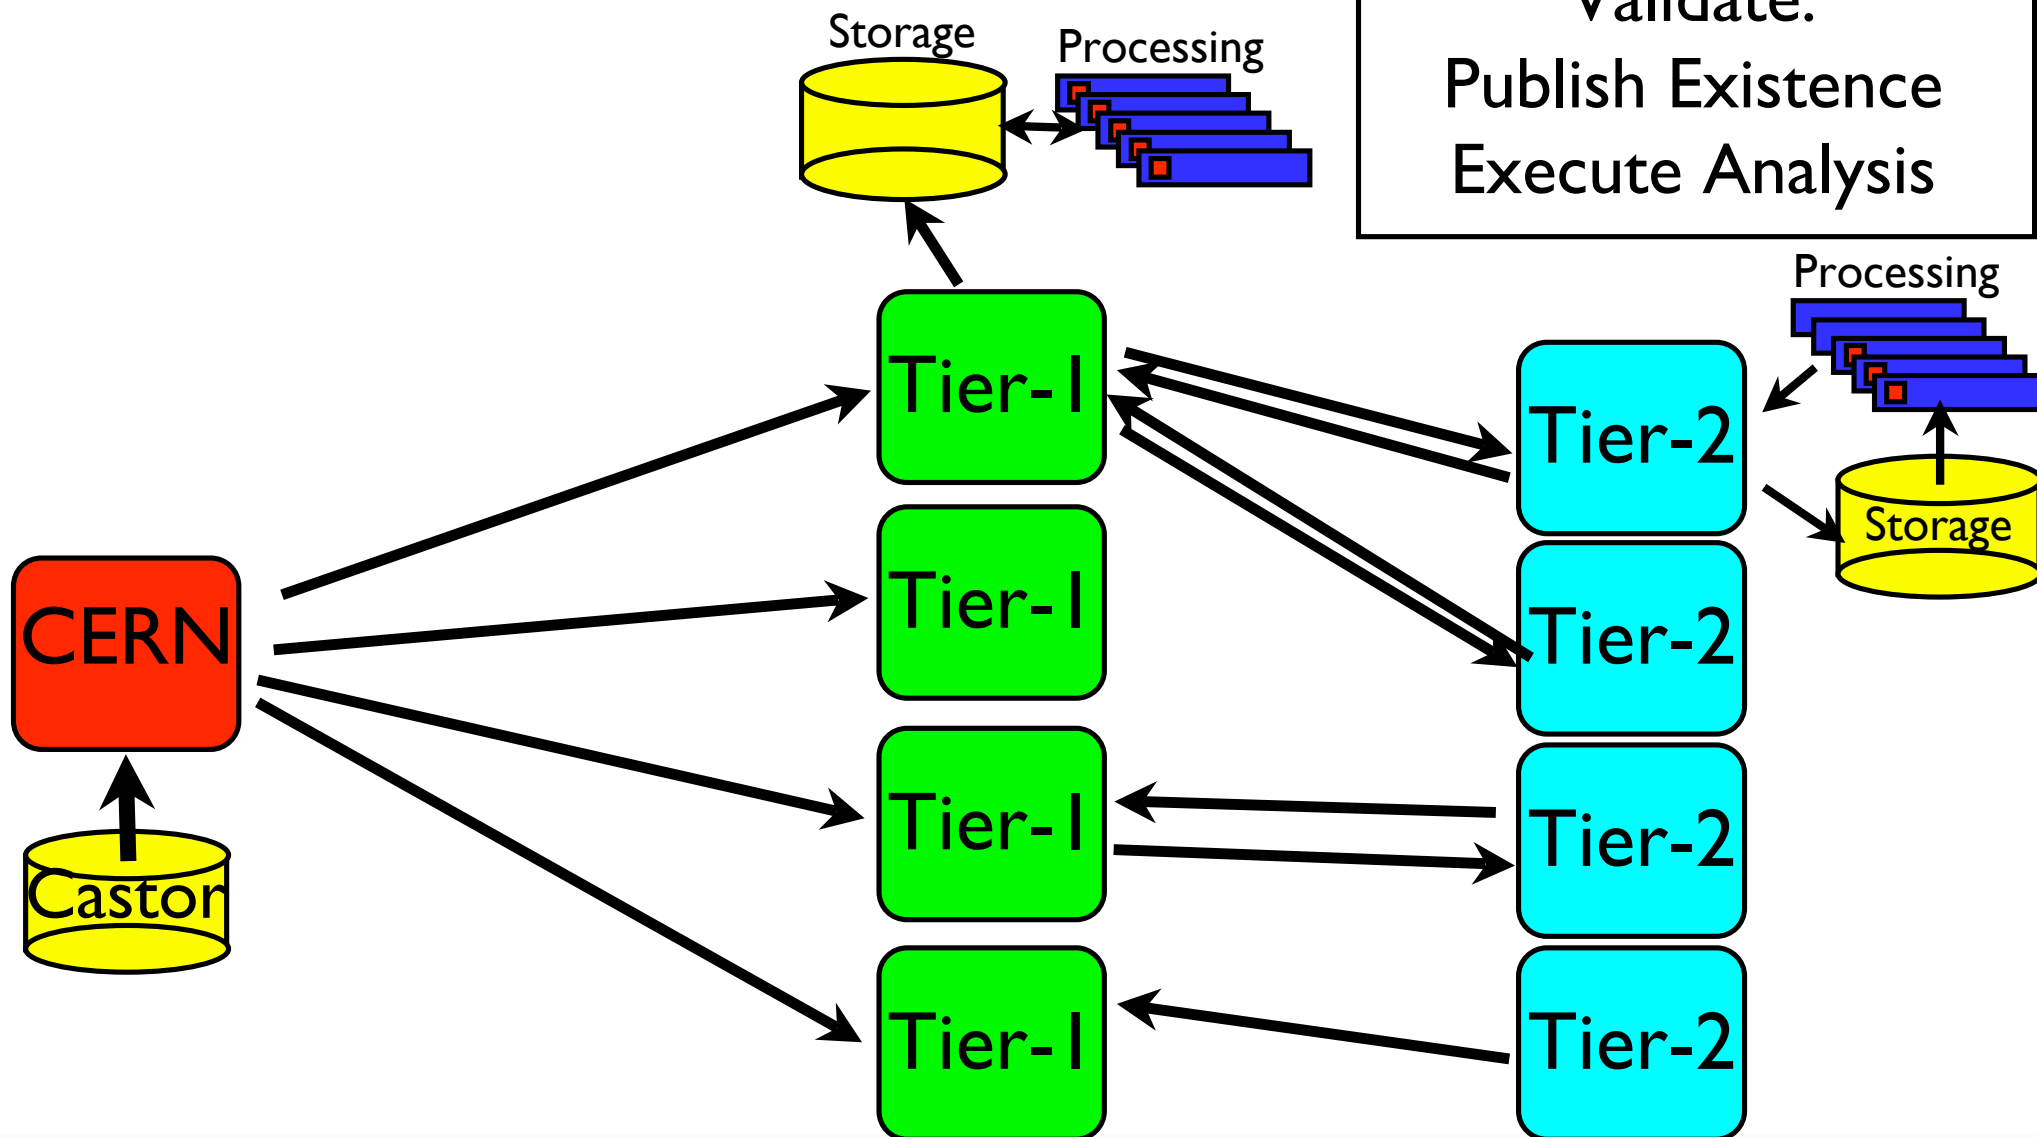
Test the very basic elements of each computing tier simultaneously

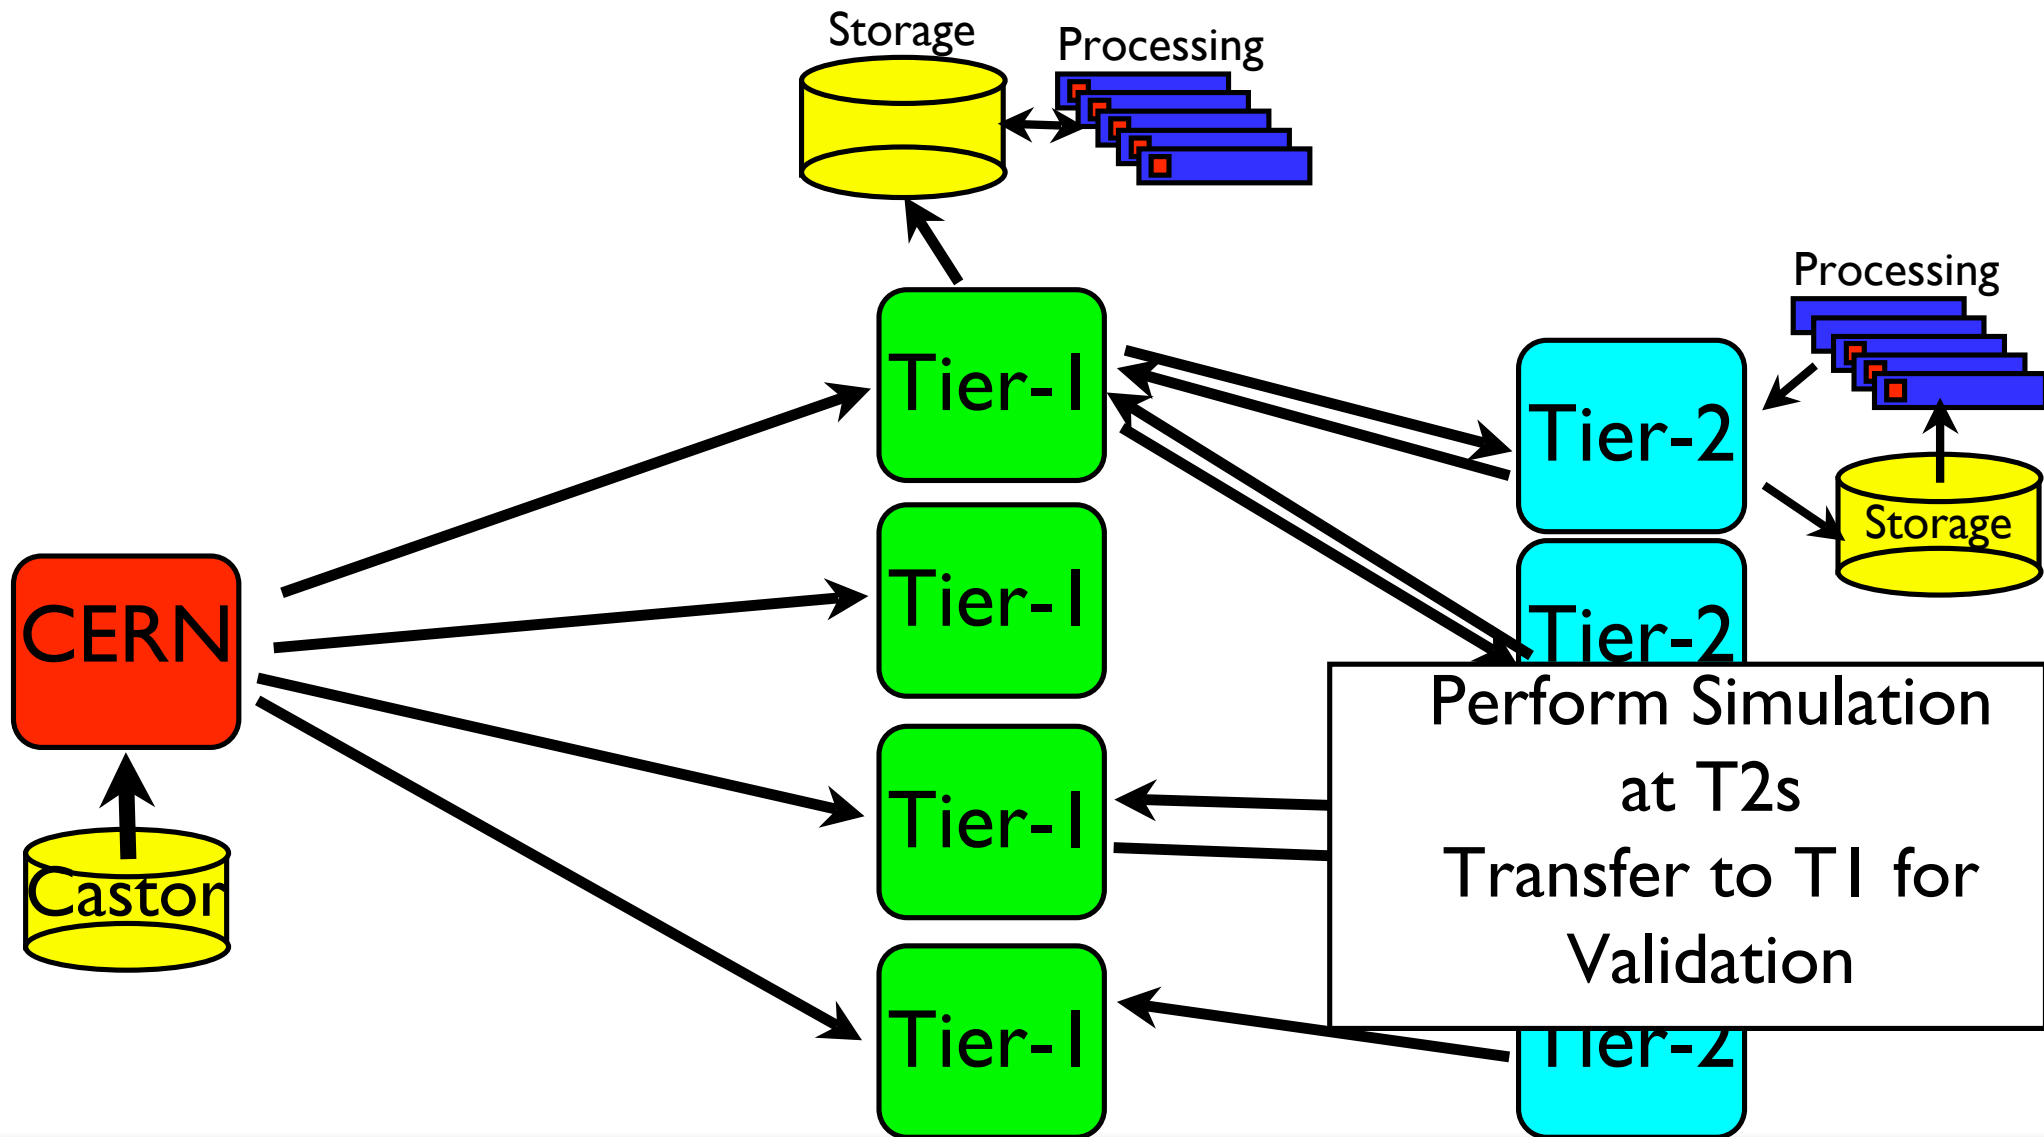


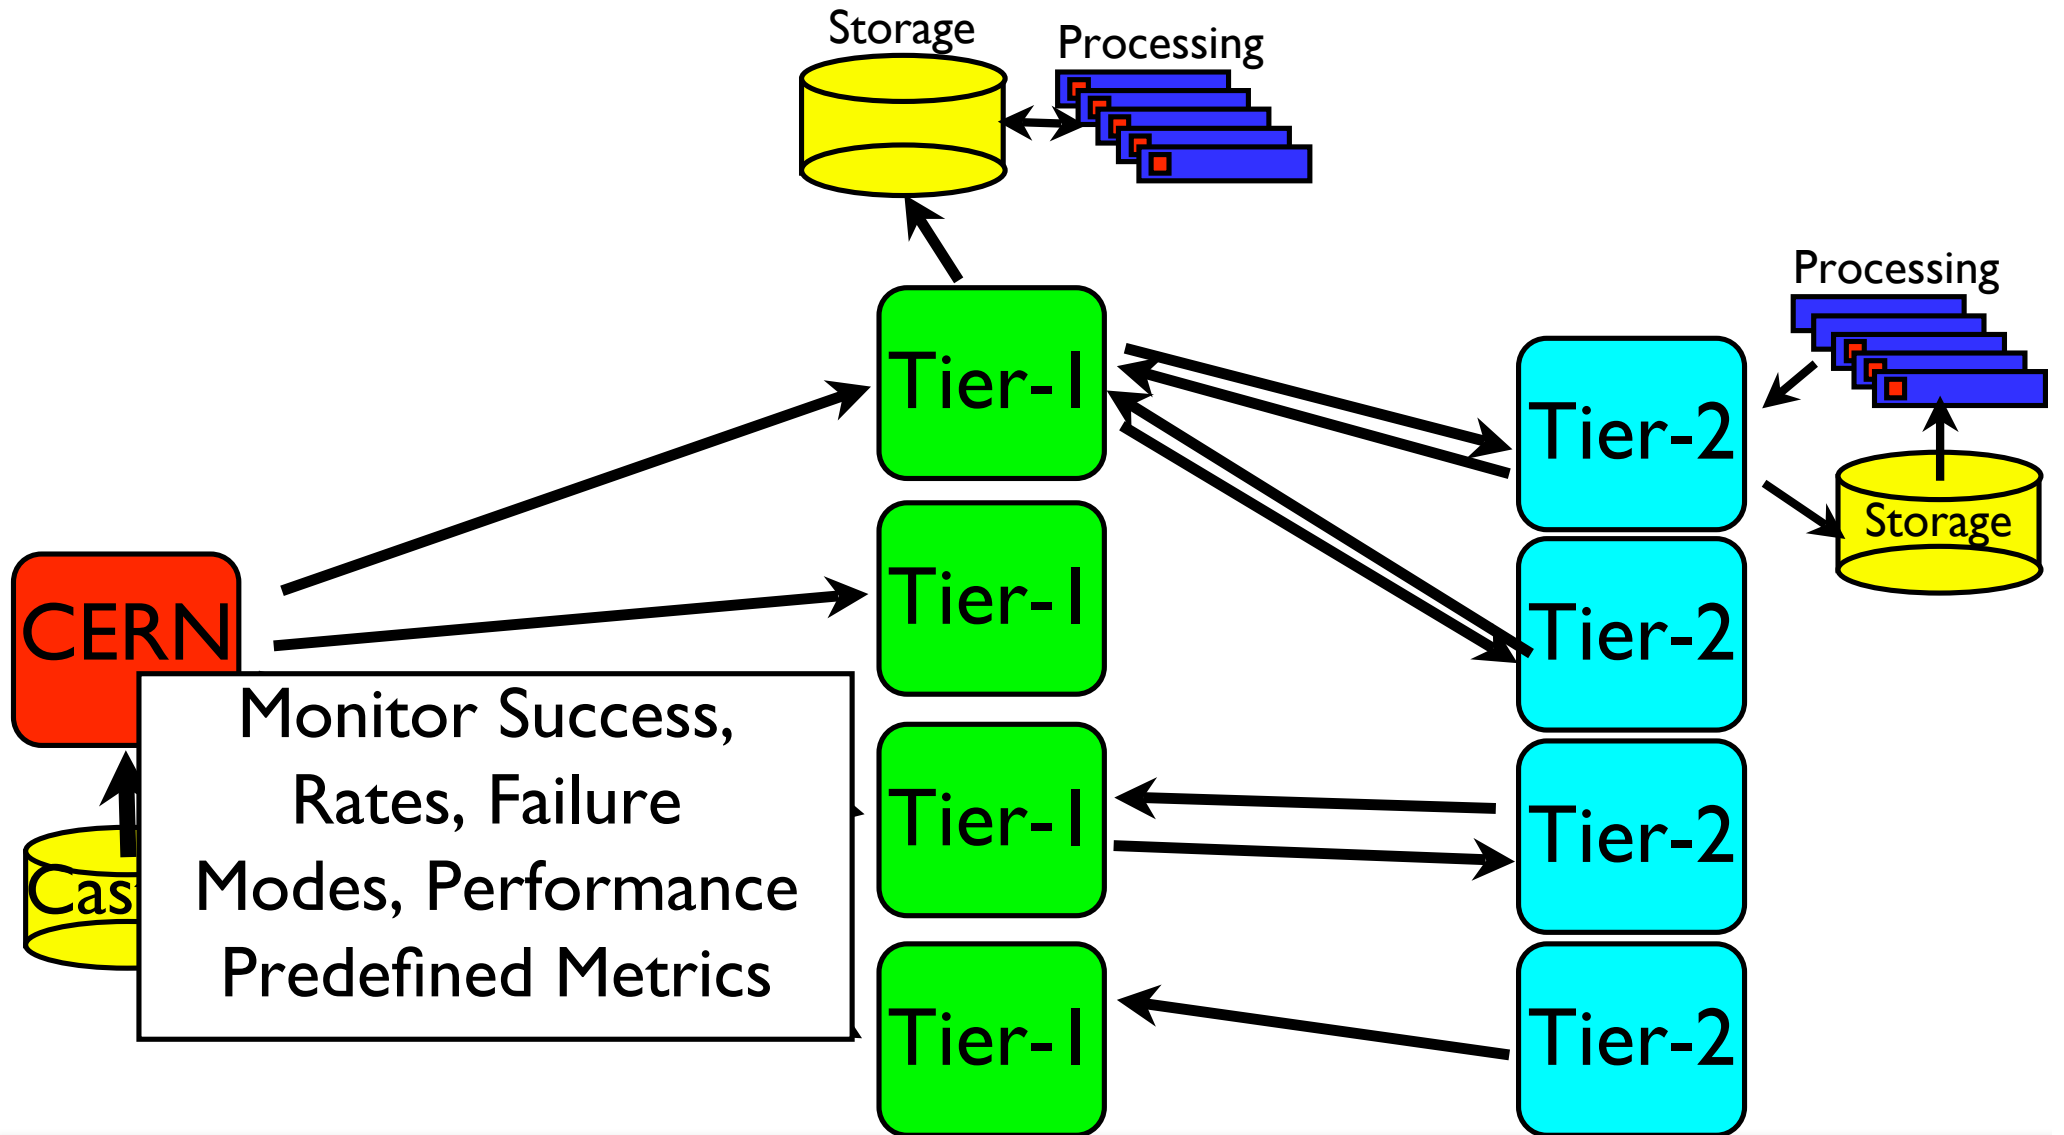
Execute Analysis Jobs on Tier-1 Resources. Select Data



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









Refined Goals

It became clear during first phase of the service challenge that not all the goals could be met

- ➔ Too many of the underlying grid and experiment services were not sufficiently well prepared to stress in a challenge environment
 - Some brand new or recently deployed services
 - A lot of effort was spent on debugging and integration
 - Useful effort, but prevented some of the more advanced challenge goals
- ➔ Not all of the sites were sufficiently shaken-out before the challenge
 - Challenge provided a lot of site validation work

In the end CMS chose to concentrate on two areas of the original goals

- ➔ Data Transfers
 - Verify CASTOR2 from CERN and improve T1-T2 transfers efficiency
- ➔ Execute Analysis Applications
 - Submit jobs through LCG and OSG interface and access the data



Progress

CMS was able to exercise some aspects of the computing model

- ➔ Separations of site and global experiment responsibilities
- ➔ Exercising elements of the support model between the Tier-1 and Tier-2 centers
- ➔ Demonstrate placement of data and drive apps. based on location

CMS was able to demonstrate data transfers

- ➔ Replicating data between the Tier-0 to Tier-1 sites and then from the Tier-1 to Tier-2 sites
- ➔ Prepare the data for analysis application access and publish the existence of the data

CMS was able to demonstrate access by analysis-like applications submitted through the grid interfaces

- ➔ Using the CMS workflow manager for analysis jobs CRAB (CMS Remote Analysis Builder), CMS was able to submit to through the LCG and OSG Grid Infrastructure and access prepared data



Site and Experiment Responsibilities

In SC3 CMS was able to begin to test the distinctions between the sites for local issues and the central experiment for global issues

Global

- ➔ Data Transfers are entirely managed through a CMS data location manager called PhEDEx.
- Steering of data from T0 to T1 to T2 was handled by the central transfer management database operated by PhEDEx operations
- ➔ Analysis jobs are submitted by a Job Robot through the CMS CRAB (CMS Remote Analysis Batch) system and monitoring information is centrally collected using MonaLisa and the ARDA Dashboard project

Local

- ➔ Sites are responsible for ensuring the site mass storage and mass storage interface is functional, the grid interfaces are responding, and the data publishing steps are succeeding
- A lot of infrastructure tools are provided to the sites, but having the whole chain hang together requires perseverance



Transfers

Substantial Effort was spent getting the transfers to perform

- ➔ During SC3 CMS Transferred ~290TB
- ➔ In the phase 2 period 140TB was transferred

Tier-1	Volume	Rate
ASGC	19.7 TB	17.8 MB/s
CNAF	7.8 TB	18.3 MB/s
FNAL	23.6 TB	17.1 MB/s
FZK	7.1 TB	8.2 MB/s
IN2P3	19.4 TB	20.5 MB/s
PIC	12.9 TB	11.4 MB/s
RAL	10.5 TB	10.8 MB/s

- ➔ CMS defined a transfer quality
 - Success rate vs attempts
- ➔ Averaged was ~50% across all sites
 - Highest was 90% and lowest was 5%


Tier-2	Volume	Rate
Bari	3.4 TB	5.6 MB/s
Caltech	1.1 TB	1.8 MB/s
DESY	6.4 TB	6.6 MB/s
UFL	1.5 TB	20.4 MB/s
Legnaro	1.9 TB	5.9 MB/s
NCU	2.9 TB	10.2 MB/s
UNL	4.2 TB	3.7 MB/s
Purdue	2.6 TB	5.8 MB/s
Spain	2.1 TB	6.5 MB/s
UCSD	5.8 TB	9.8 MB/s
UW	1.9 TB	2.2 MB/s



Transfer Results

At almost no point were the sites stressing the available network with CMS transfers.

Max Aggregate for the Tier-I centers during the second phase of the challenge reached 120MB/s



While in the end essentially all files were transferred the rate of successful transfer after one attempt was too low.

➔ Needs to be improved for the future



Accessing Data

Once the data was transferred it had to be made available to analysis applications and the existence published

- ➔ Preparing data in the current CMS data model is fragile and prone to problems
 - Frequent complaint by the sites was the difficulty of this step
 - CMS data files are tightly connected and need to be accessed to build a readable dataset
 - Being addressed in the new Event Data Model work
- ➔ The publication of the data existence was done with a series of site local web services that published catalogs of catalogs
 - Reasonably simple system that scaled well to the scale of service challenge 3
 - Well suited for publishing the existence of well defined datasets
 - Infrastructure allowed incoming grid applications to discover the local data location



Jobs Success

Analysis jobs were submitted through the LCG and OSG Grid Interfaces with CRAB

➔ Dedicated talks on the subjects from Marco Corvo and Oliver Gutsche
The job success rates and exit codes were monitored daily on both grids.

- ➔ The most common failure was related to data preparation or the CMS application
- ➔ Success of the grid infrastructure itself was better than 90% for both grids
- Still considerable room for improvement, but better than it was a year ago





Jobs Success

The infrastructure for submitting jobs was reasonable stable and the success rate was relatively high

- ➔ We were able to use the analysis jobs to formally validate the site performance only at selected sites
- Goal of CMS was to document the performance of the mass storage for analysis applications
 - 100MB/s at Tier-I centers
- Achieved (substantially over achieved in some cases) at some sites. Many sites either could not meet the performance or could not document the metric
 - a few sites had substantial data serving or data preparation problems

Rate of jobs submitted per day continues to be much lower than CMS needs to maintain an achievable operations ramp-up before the start of the experiment

- ➔ Partially to do with the speed of submission and partially to do with the available resources




Lessons Learned (1/2)

The primary experience of SC3 for CMS was that it involved a lot of debugging and shaking out. Many services were not sufficiently tested before the start of the challenge period. Sites had not verified functionality

- ➔ The goal of verifying the basic functionality and demonstrating meeting predefined metrics was quickly replaced with the goal of determining why pieces of the chain did not work as expected.

The very high failure rates for transfers, which are related both to source and destination problems, consumed a lot of effort.

- ➔ There are periods when the aggregate transfer rate exceeds 100MB/s
- ➔ Not a significant stress to the network infrastructure





Lessons Learned (2/2)

Wide area communication between sites and between the experiments and infrastructure providers remains challenging

- ➔ Maintaining priorities, communication requirements, providing feedback, and meeting the needs of each group is difficult

The challenges are hard work and involve substantial effort from key individuals

- ➔ Seemingly even more work do to the communication challenges

The central job submission for analysis validation worked well

- ➔ Collaboration with ARDA team and MonaLisa on the dashboard resulted in reasonable job monitoring with a quick development cycle
- ➔ LCG-2 interfaces and OSG interfaces performed well under job load
- ➔ Ability to submit analysis to LCG and OSG with the same tool set is a benefit of SC3 activities

Site commissioning is proceeding but there is a lot of work left




Rerun of Throughput

After a lot of debugging and shaking out during the fall of 2005 the throughput phase was rerun with FTS driving SRM in early 2006. It involved sites and LCG, but not CMS directly.

- ➔ Achieved original transfer goals of the challenge
- Good balance between sites

Weekly Report (for the last week, by default)





Future Plans

CMS will recover the elements that were descoped in SC3 and build on these in the final full year before the pilot run at LHC

- ➔ Test with simulation production at Tier-2 sites
- ➔ Introduce more varieties of analysis challenges
 - Introduction of calibration information, database access, and more complicated analysis applications
 - Introduce more users
- ➔ Test the end to end including primary reconstruction at the Tier-0 centers and re-reconstruction at the Tier-1s
- ➔ Include the scale
 - Transfers
 - Data serving
 - Jobs Submitted