PDSF Status and Overview

Eric Hjort, LBNL

STAR Collaboration Meeting

February 28, 2003

PDSF Overview

- STAR has production facilities at RCF and PDSF
  - PDSF needs to transfer data: Replicate all DSTs and some raw data from RCF to PDSF
  - PDSF computing: Embedding production, data analysis, simulations
- PDSF infrastructure people:
  - Doug Olson (STAR computing at PDSF coordinator)
  - Iwona Sakrejda (PDSF user support, STAR libraries, accounts, etc.)

    Contact by filing a PDSF support request
  - Eric Hjort (Embedding, data management)

    Contact by email or hypernews: pdsf-hn@www.star.bnl.gov, etc.
  - PDSF oversight committee: P. Jacobs (chair), D. Olson, K. Schweda, D. Hardtke, I. Sakrejda, S. Canon (PDSF project leader), J. Lauret, E. Hjort

File Replication: RCF to PDSF

(A simple example of Grid Tools)

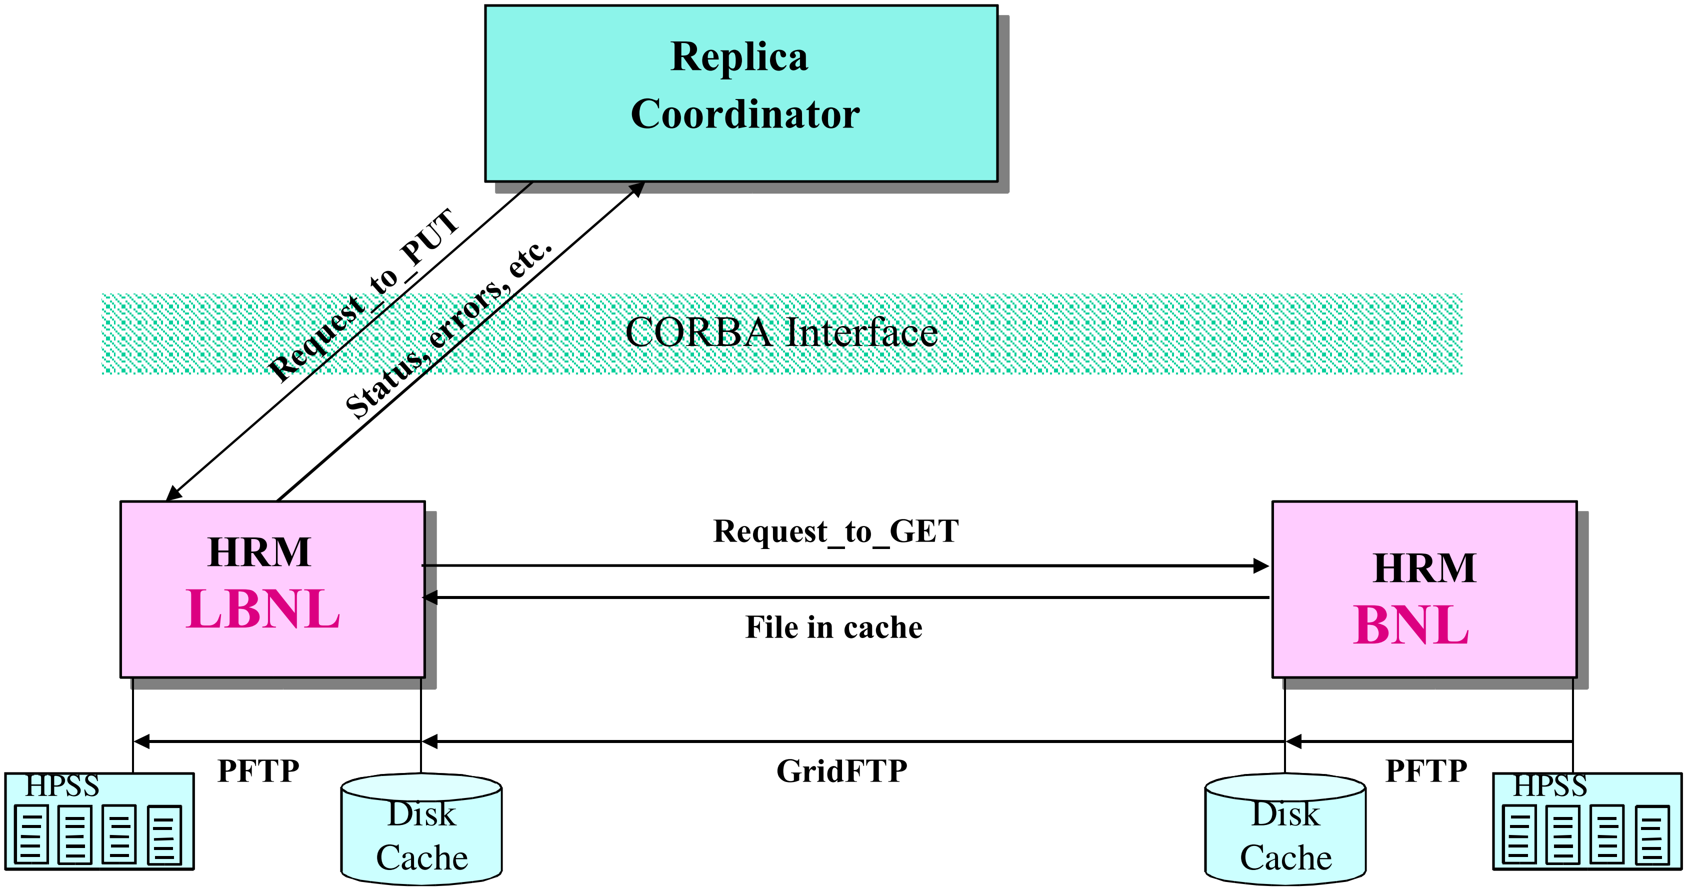

HRM = **H**ierarchical **R**esource **M**anager

Grid computing

- For STAR data transfer:
  - Authenticate with grid certificate
    - Convenient: not necessary to log in at RCF
    - Use grid-proxy-init (requires password)
  - Automatic caching of data by HRMs
    - large cache disk not necessary
    - enables continous automatic transfers
  - Result: Grid tools improve net transfer rate, reduce effort
- What does it take to do grid computing?
  - Get a DOE Science Grid certificate
  - Get it installed on STAR grid nodes at RCF and PDSF
  - Easy to do some simple, convenient things with Globus
  - Alpha testers wanted – Iwona has prepared instructions

Disk Resources at PDSF

- High performance disk (1 TB on /aztera)
  - Heavily used datasets; embedding input files
- Distributed disks (14 TB on 70 nodes)
  - For large MuDst productions.
  - Data access by the STAR job scheduler.
- Production data vaults (11 TB)
  - Embedding output; simulations data; selected MuDst's; etc.
  - Managed by production people
- PWG data vaults (5.3 TB)
  - Managed by PWG's.
  - More space available on a by-request basis.
- Scratch space (1 TB on pdsfdv15)

How to find STAR Data at PDSF

New links on the STAR PDSF help page:

Summary of STAR data at PDSF:

Embedding data on disk

MuDst's and simulation data on disk

Job scheduler instructions and data on distributed disks

These pages update automatically every 24 hours.
If you can’t find it on these pages… email: ELHjort@lbl.gov.

Job submission at PDSF

- Job scheduler in use for data on distributed disks
  - No special queue for scheduler
  - Not fully function without a PDSF file catalog
  - Uses pre-made filelists
- Job scheduler not used for data on NFS disks
  - **Data is filtered for sanity=1**
- Queues/priorities:
  - Short (1hr), medium (24hr), long (5 days)
  - Production account (user starofl) has a higher priority
  - At present production is not run on distributed disk nodes
  - Important to balance production vs. users' resources

Embedding Overview

- Embedding production done at PDSF
  - "embed" simulated particles into real data at the raw data level
  - Reconstruction yields efficiencies
  - Important test of STAR software: simulations meet real data
- 20 TB, 10 M events in 2002
- People involved:
  - Eric Hjort (infrastructure and development; spectra production)
  - Matt Lamont (strangeness-specific infrastructure, strangeness production)
  - Patricia Fachini (development; miscellaneous production)
  - Christina Markert (development; miscellaneous production)
  - Olga Barranikova (QA)
  - STAR collaborators (SOFI, calibrations, simulations, etc.)

Embedding Methods

- Year 2 AuAu TPC embedding
  - 20 GeV
  - 200 GeV (P02gd) running without problems for about 1 year
    - **At least 29 different embedded particles**
    - **Central, minbias, both fields, various pt, y ranges, etc.**
- pp embedding
  - Hijing pp -> zerobias (vertex reconstruction studies)
  - Hijing pp + embedded particle -> zerobias (Jon G.)
  - Embedded particle -> pp data (Matt/strangeness)
- RICH embedding (Boris H.): Tested and in production
- dAu embedding status
  - Working in a testing mode
  - Need to test/understand dE/dx shift
  - Many dAu and zerobias daq files are at PDSF
  - Initial production setup ready next week
- dAu FTPC (Frank S.): Working; needs details + testing

Embedding Requests

- Ask your PWG convenor to submit request to Simulations Production Request Page:
  - Organizes and documents the work
  - Specifies job parameters for reference
  - Allows for prioritization of jobs by Jamie
  - Protects against resource misuse
- If not prioritized, jobs order = submissions order, but…
  - Multiple operators mix order
  - Technical reasons mix order
  - Some requests take much longer than others

Summary/Future plans

- Data transfer
  - Grid tools serve us well
  - Data transfer needs are met in general
  - Goal for this year’s run: reduce latency to 1 week or less?
- Data management
  - PDSF data discovery webpages overhauled
  - Job scheduler in use
  - Next big step: file catalog at PDSF
- Embedding and Simulations
  - AuAu, pp, RICH embedding all working
  - dAu TPC and FTPC embedding almost in production
  - Future: new detectors, understand triggers, etc.
- Bigger picture
  - Seamless, more automated data transfer RCF <-> PDSF
  - Distributed grid computing with the Job Scheduler.