

Grid Collector

Wei-Ming Zhang

Kent State University

John Wu, Alex Sim, Junmin Gu and Arie Shoshani

Lawrence Berkeley National Lab

In collaboration with

Jerome Lauret, Victor Perevoztchikov,

Valeri Faine, Jeff Porter, Sasha Vanyashin

Brookhaven National Laboratory



A View of the Analysis Process

- Users want to analyze some events of interest
- Events are stored in millions of files
- Files are distributed on many storage systems
- To perform an analysis, a user needs to
 1. Write the analysis code, run it
 2. Specify the events of interest
 3. Locate the files containing the events
 4. Prepare disk space for the files
 5. Transfer the files to the disks
 6. Recover from any errors
 7. Read the events of interest from files
 8. Remove the files



Design Goals of Grid Collector

Make analysts more productive by

- Reading only events of interest
- Automating the management of distributed files and disks



Approaches of Grid Collector

- Allow users to specify events of interest using meaningful physical quantities
 - `numberOfPrimaryTracks > 1000 AND vectorSumOfPt > 20`
 - Simplify step 2
- Automate file management tasks
 - Use File Catalogs to locate files
 - Use Storage Resource Manager to manage the disk space and file transfers
 - Remove steps 3 -- 8



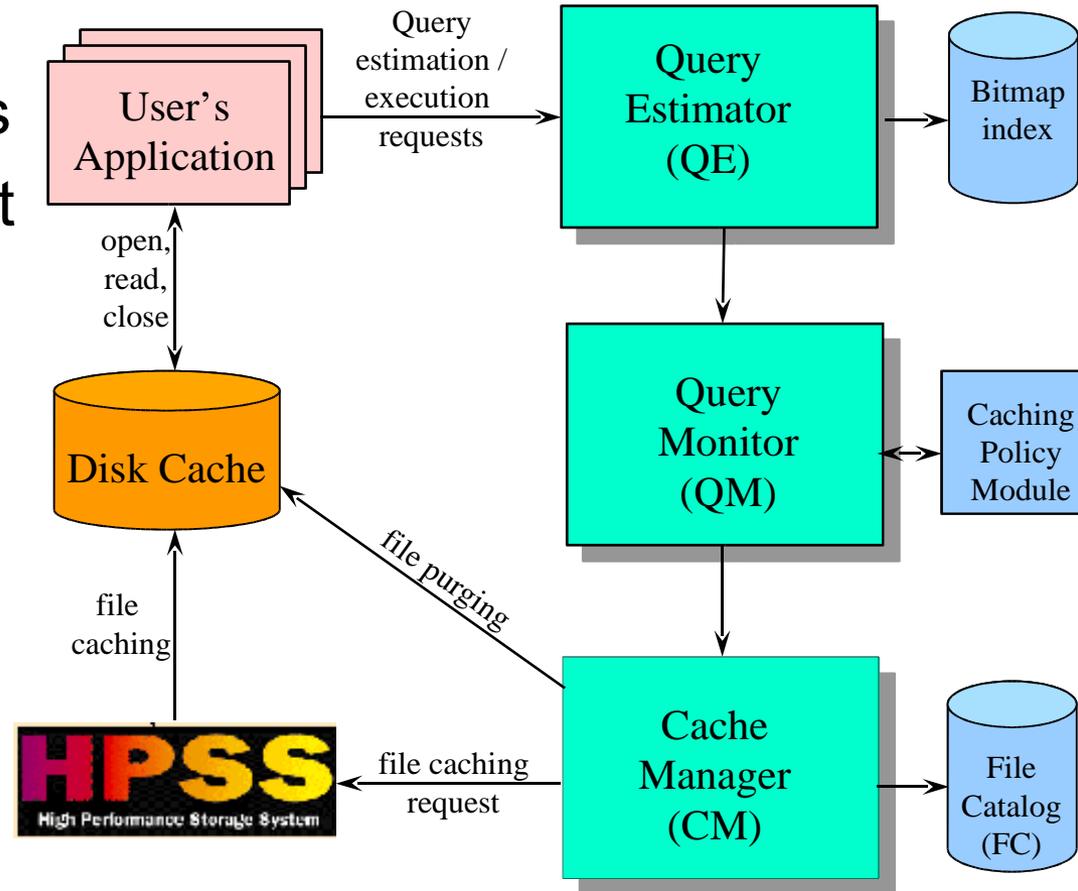
Storage Access Coordination System

Strength –

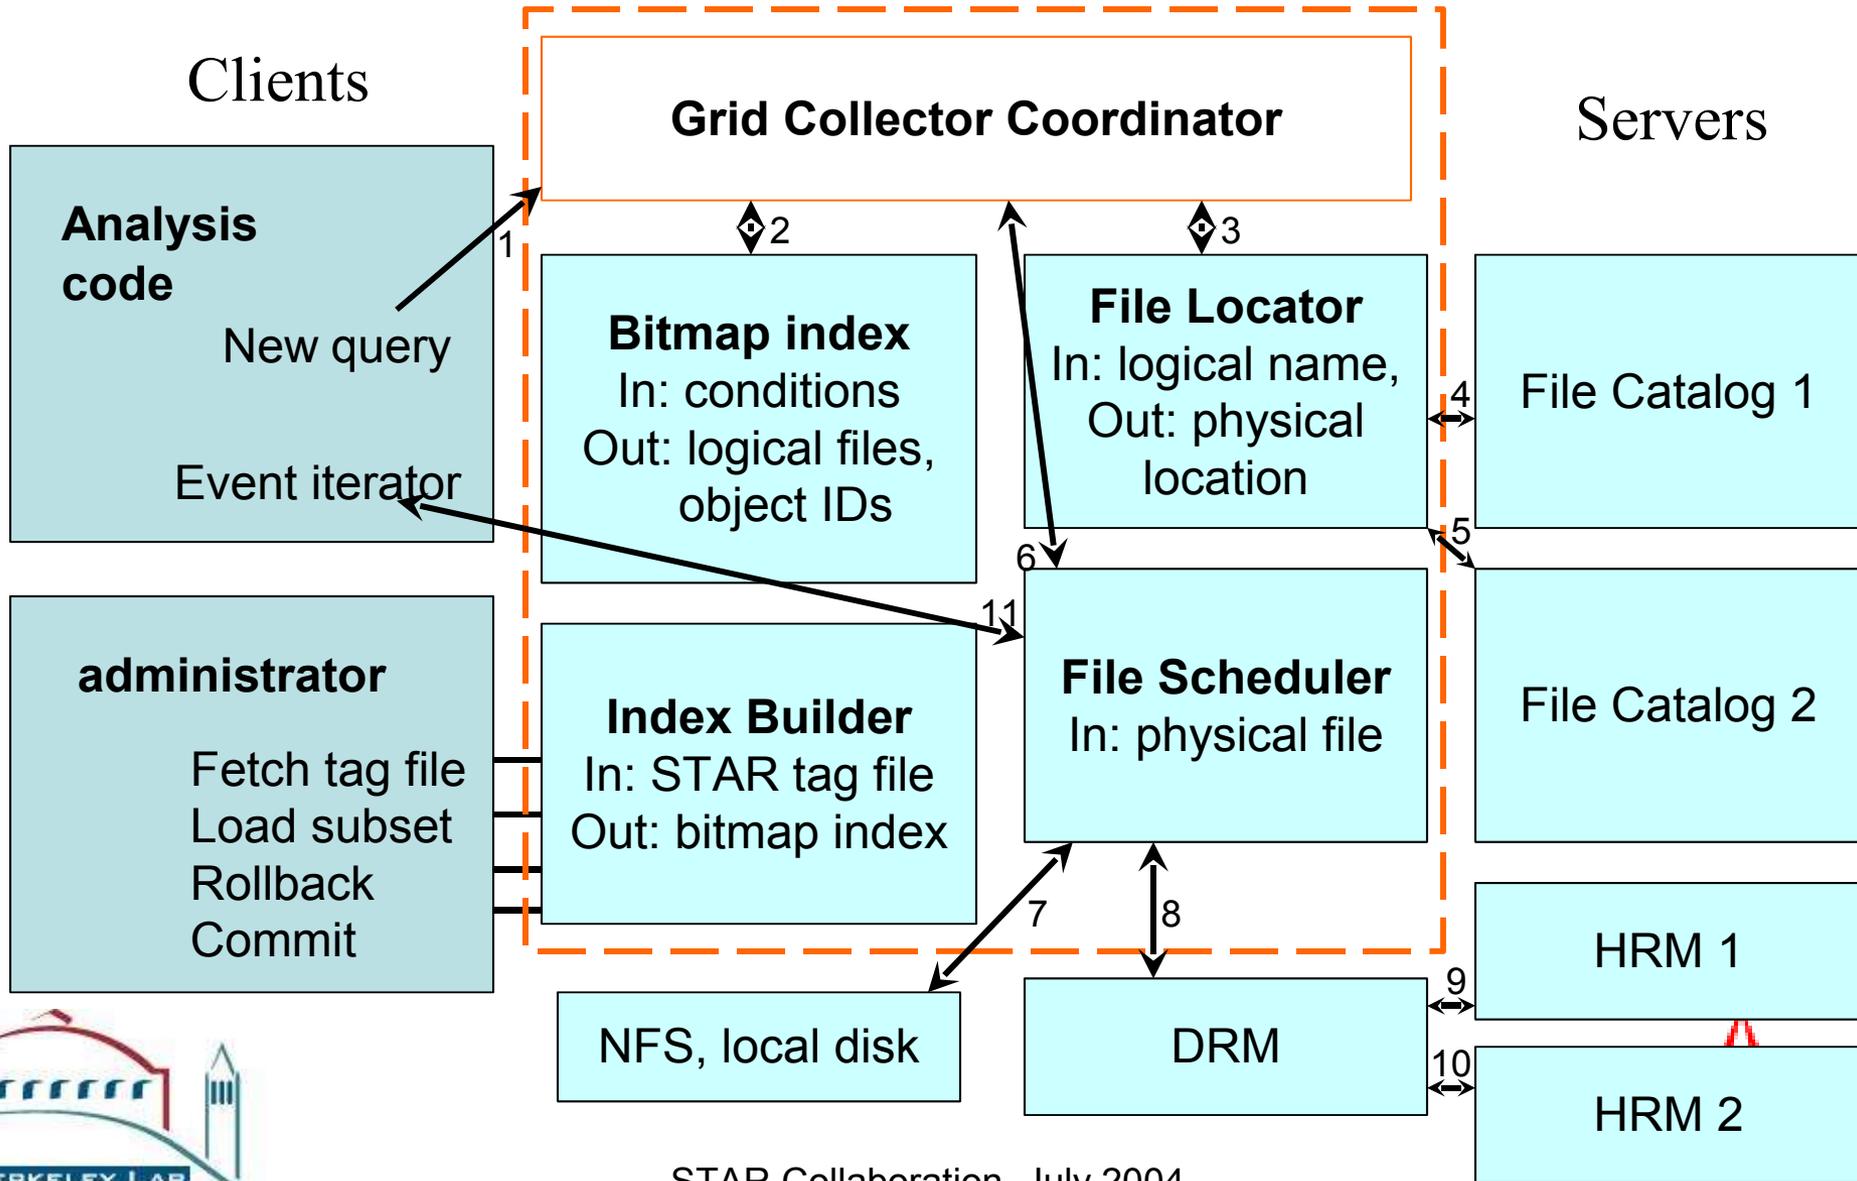
- Allow user to specify events as range conditions
- Automate file management tasks

Weakness –

- Designed for Objectivity data
- Access only one HPSS



Grid Collector: Architecture



GC vs. STAR Scheduler

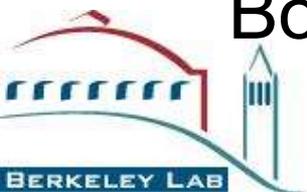
GC

- Select events with range conditions
- Read only selected events
- Automate all file and space management tasks

Scheduler

- Specify a list of files on disk
- Read all events of the files
- Use Data Carousel for HPSS files

Both can split large jobs to multiple machines



GC vs. STACS

GC

- Use multiple File Catalogs and multiple Storage Systems
- Integrate index building functions into the server
 - Improves index building speed
- Make use of distributed disk caches, clients can have their own caches

STACS

- Limit to only one File Catalog and one Storage System
- Use a separate Index Feeder to digest tag files
 - Has very low data transfer rate through CORBA
- Make use of one disk cache, clients must access the disk cache

Both select events with range conditions
Both automatically manage files and disks



This Year vs. Last Year

This Year

- Process all files, including MuDST
- Build indices fast
 - Use automated file management functions
 - Indexing 15 million events took one week
- Interact with multiple File Catalogs

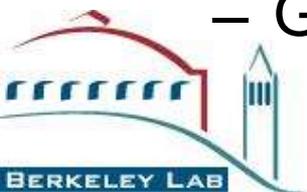
Last Year

- Process event files, but not MuDST
- Build indices slowly
 - Index feeder requires manual file transfer
 - Indexing 5 million events took 10 weeks
- Interact with only one File Catalog



What Can Grid Collector Do For You

- If you gather statistics on lots of events
 - Grid Collector allows you to work with files not already on disk
- If you search for rare events, Grid Collector allows you to
 - Specify the events with ease
 - Access only relevant files
 - Read only selected events
- If you want to try some analysis ideas outside of the main computer centers,
 - Grid Collect manages file and space for you



How To Use The Grid Collector

- Must use StIOMaker
 - StIOMaker can now handle all files including MuDST
- Replace StFile with StGridCollector
 - StIOMaker requires a StFile object
 - One currently uses “new StFile(…)” to create a StFile object
 - Grid Collector provides a new way, “StGridCollector::Create(SELECT geant, event WHERE …)”
- Iterate through events as usual



How To Use -- More Details

- External dependencies
 - Globus, ROOT, STAR Software
 - Storage Resource Manager (DRM, HRM)
 - ORBACUS
- Servers
 - Main Grid Collector Coordinator
 - DRM/HRM
 - File Catalogs
- Client library
 - User need to load this in the macros



How To Select Events

- `SELECT [MuDst|event|...] WHERE NV0>100 AND ...`
- The `WHERE` clause consist of range conditions joined with logical operators `AND`, `OR`, `NOT`.
- All tags and a few File Catalog key words can be used in the `WHERE` clause
- Variables with multiple values can be addressed with index, e.g., `scaAnalysisMatrix[7]`



Status Of Grid Collector

- One version in production mode at BNL
- An updated version in final testing stages
- Brave early adopters still needed

- Contact information
- Wei-Ming Zhang zhang@hpaq.kent.edu
- Jerome Lauret lauret@bnl.gov
- John Wu John.Wu@nersc.gov

