STAR MySQL DB Implementation Overview

- STAR Database Strategy
- Production DB Cases
- Server Deployments
- Client Access and Server Administration

STAR Database Strategy

STAR-Objy replaced w/ light-wt MySQL -1999

Concerns with MySQL

- Data transfer bottleneck – single server socket
- Single server storage limitations

3. **Strategy: budget DB data & develop multiple server capabilities**
   - Optimize data sizes at sub-system implementation
   - Deploy task-specific servers where applicable
   - Data replication w/ MySQL mirror servers for load balancing & wide-area deployment

STAR Database Strategy Employed Online & Offline

Online: Experimental Conditions, RunLog, EvtTag

- Fine-grained indexing (~10^9/year), ~100GB/year
- Real-time system critical to small #-users during data-taking

Several task-specific servers not for general use

Offline: Geometry, Calibrations, RunInfo

- Moderate-grained indexing (~10^7/year), 10-20GB/year
- Managed updates (per production set) , 1000s of users

1 repository w/ many mirrors deployed where needed

STAR's quick ramp-up use of MySQL w/ heavy reliance on mirrors Use MySQL's native replication

Production DB Use Cases

Use Cases	Concurrent processes	DB Data Retrieval	Client Location
Reconstruction (reco)	500	~10s MB/job ~ minutes/job - 70% iin data transfer	RCF (RCRS)
Embedding	200	same as reco	PDSF
Simulation	100	~kB/job	PDSF & RCF
Code Dev.	10	same as reco	RCF, PDSF, any
Analysis	10	~kB/job	RCF, PDSF, any

Updating Offline Master Repository

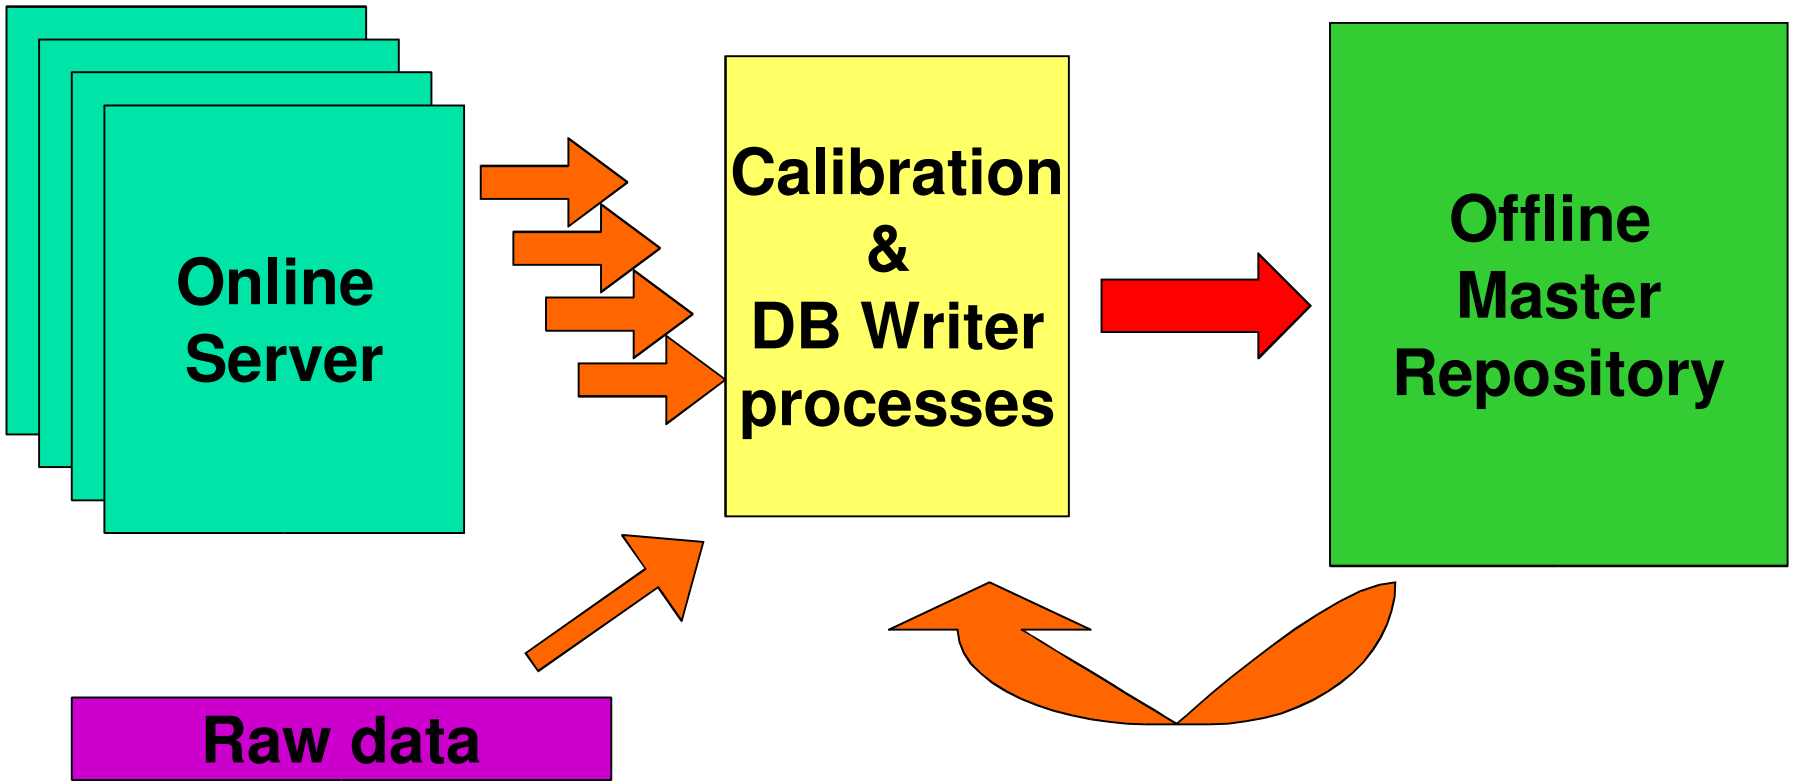

Current Status of Offline Server Deployment

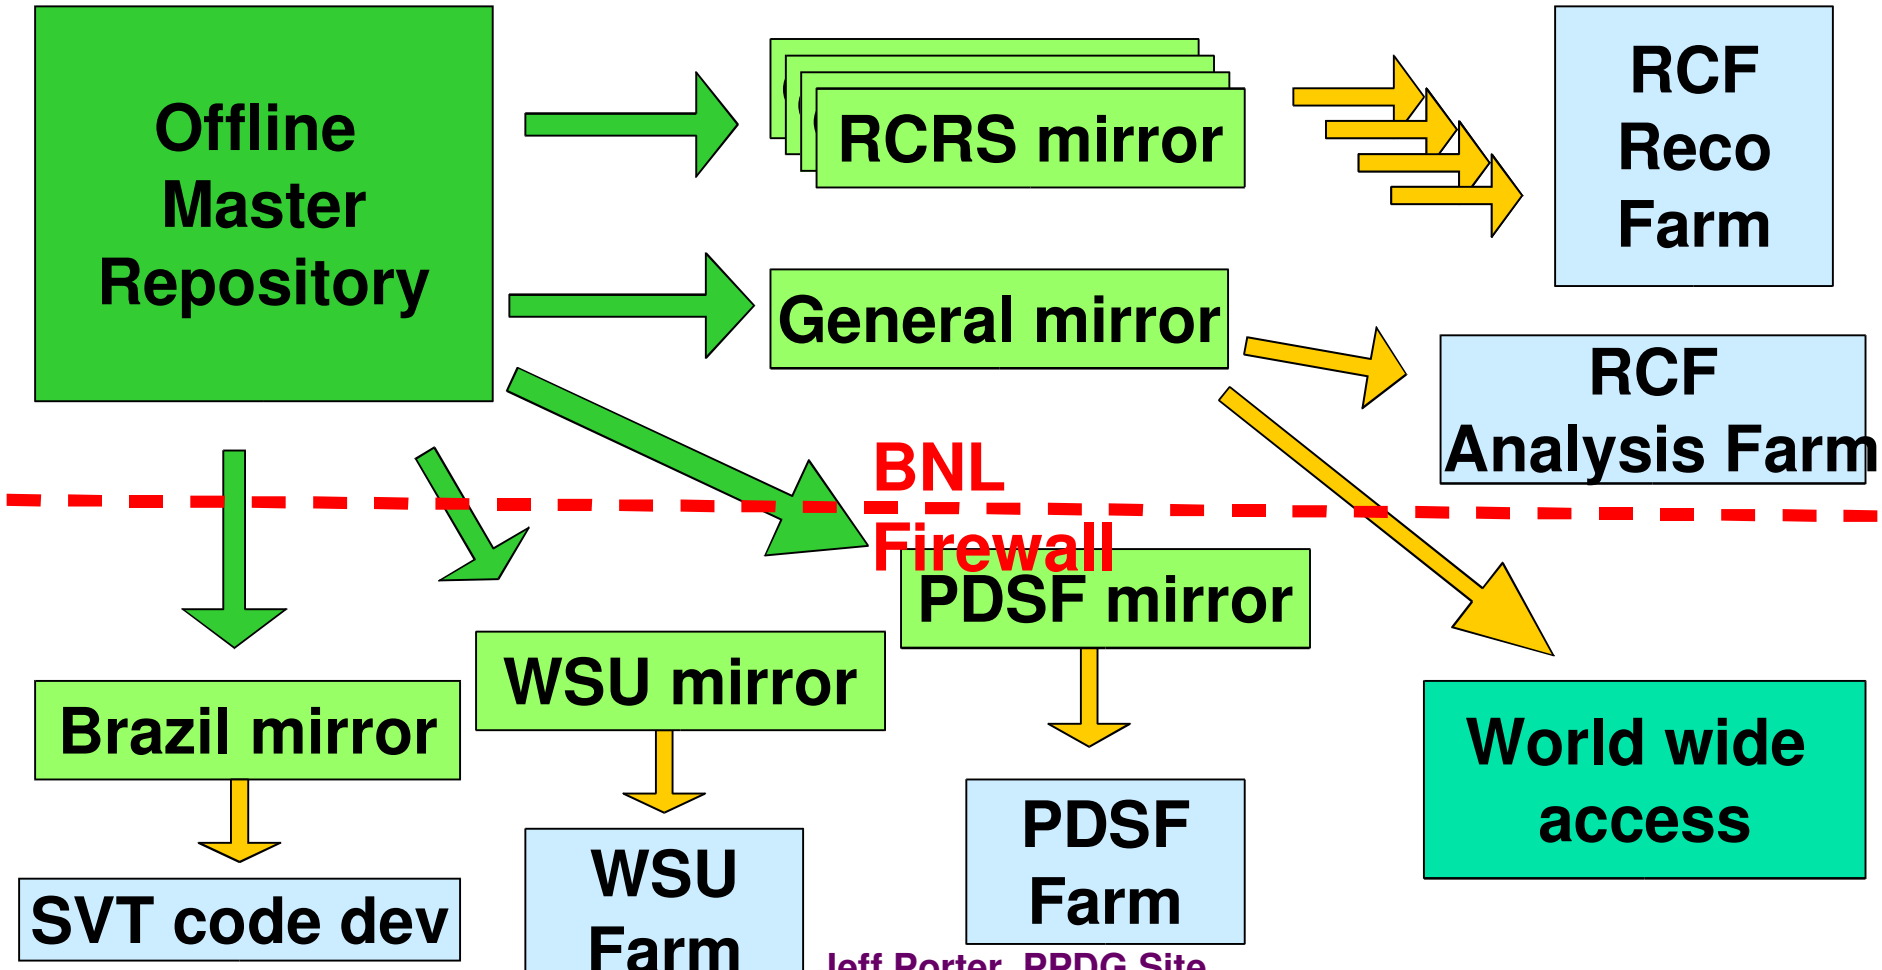

Server Selection from Client

- Servers are selected via one or more XML configuration files which include
  - A default server for all unspecified DB
  - One or more servers mapped to specified DBs
- Client codes select the XML files to use via
  - Default afs-resident dbServers.xml pointing to our general-use server
  - Environment variable $STDB_SERVERS

Mirror Server Setup & Administration

- Mirrors are set up by hand
  - tar/scp databases to mirror machine
  - Edit mirror-specific config files & run scripts to configure startup
  - Start mirror server and slave process
- Monitoring mirror
  - Run scripts to test mirror is functioning
  - Run scripts to compare global counts between mirror and master

Comments

- MySQL replication has been very robust
- Number of mirror servers will steadily increase … for next year,
  - 2-4 more at BNL
  - 2-4 more at PDSF
  - New site-specific mirrors (MPI, DUBNA,…)