

# Rapid Replication of Multi-Petabyte File Systems

Justin Sybrandt  
Jason Hick



(NSF award number 1157075)



# NERSC



- Stores PB of scientific data.
- Needed to replicate whole file systems.

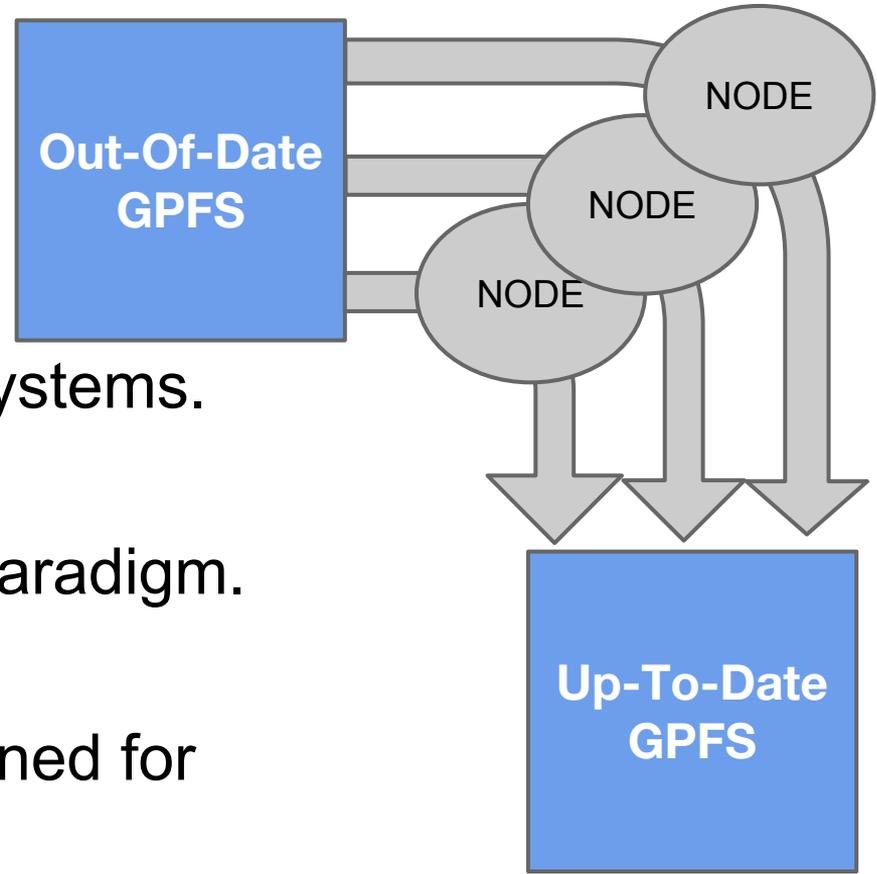
# The Problem

- Need to freshen a stale copy.
  - File system backups.
  - Disaster recovery.
  - Moving locations.



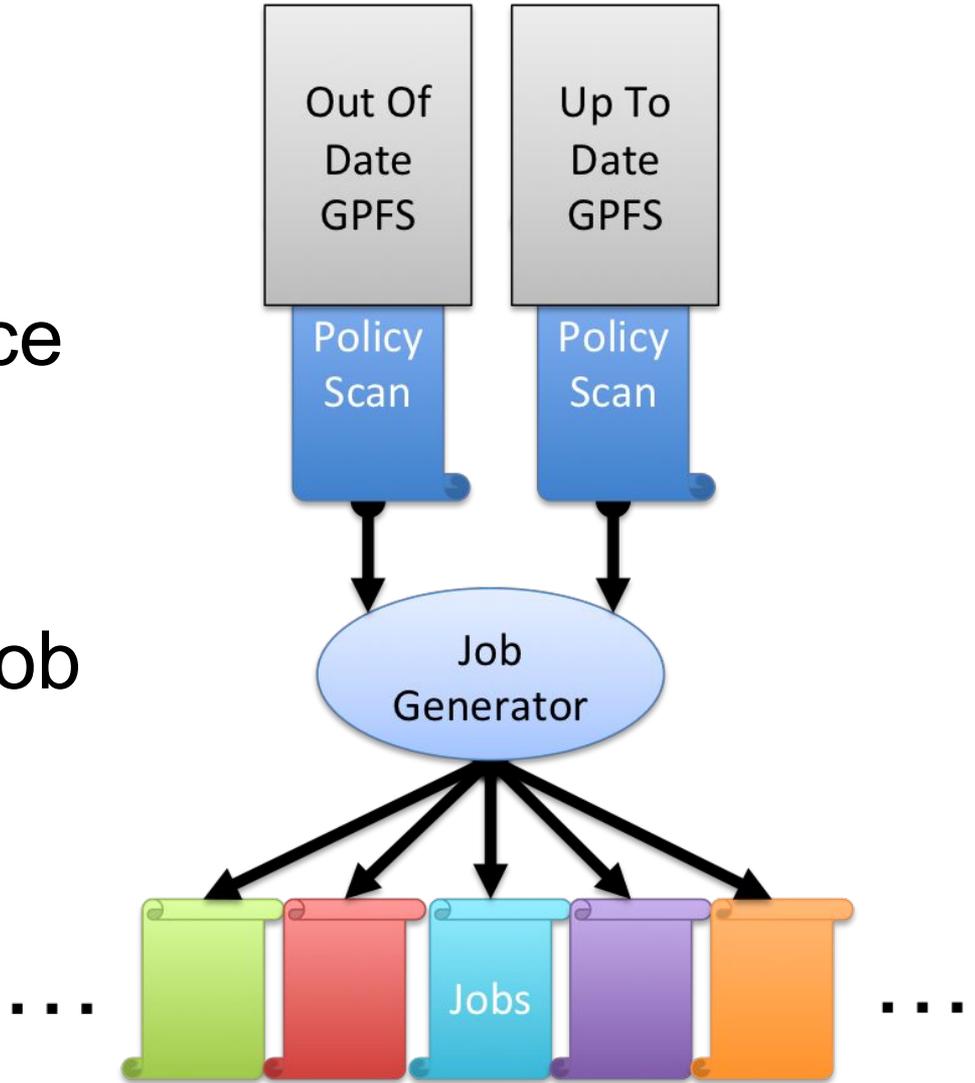
# Distsync

- Quickly determines the changes between two file systems.
- Follows the Master-Slave Paradigm.
- Similar to Shift, but streamlined for large synchronizations.



# Job Generation

- Policy scans produce lists of all files.
- Generator creates job files in linear time.



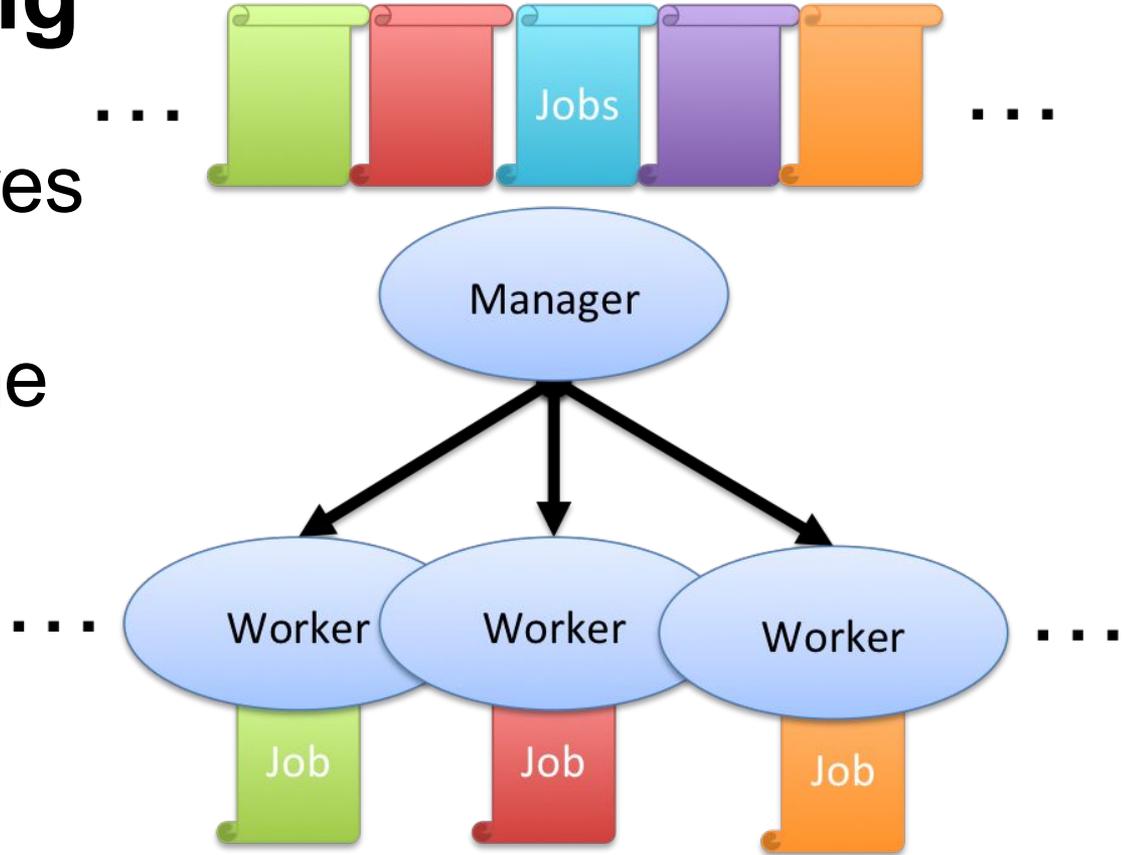
# Job File

- Contains a list of file paths.
- Limited in size.  
(when possible)
- Type implies action.



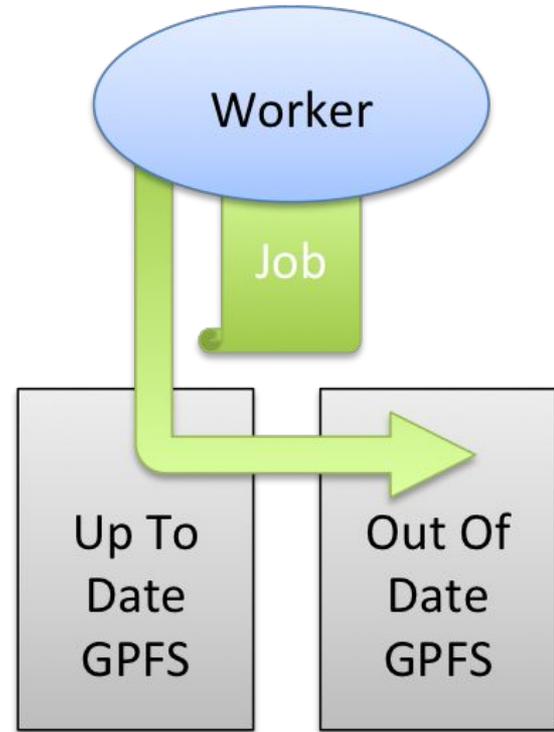
# Job Scheduling

- Manager ensures that jobs are completed in the right order.

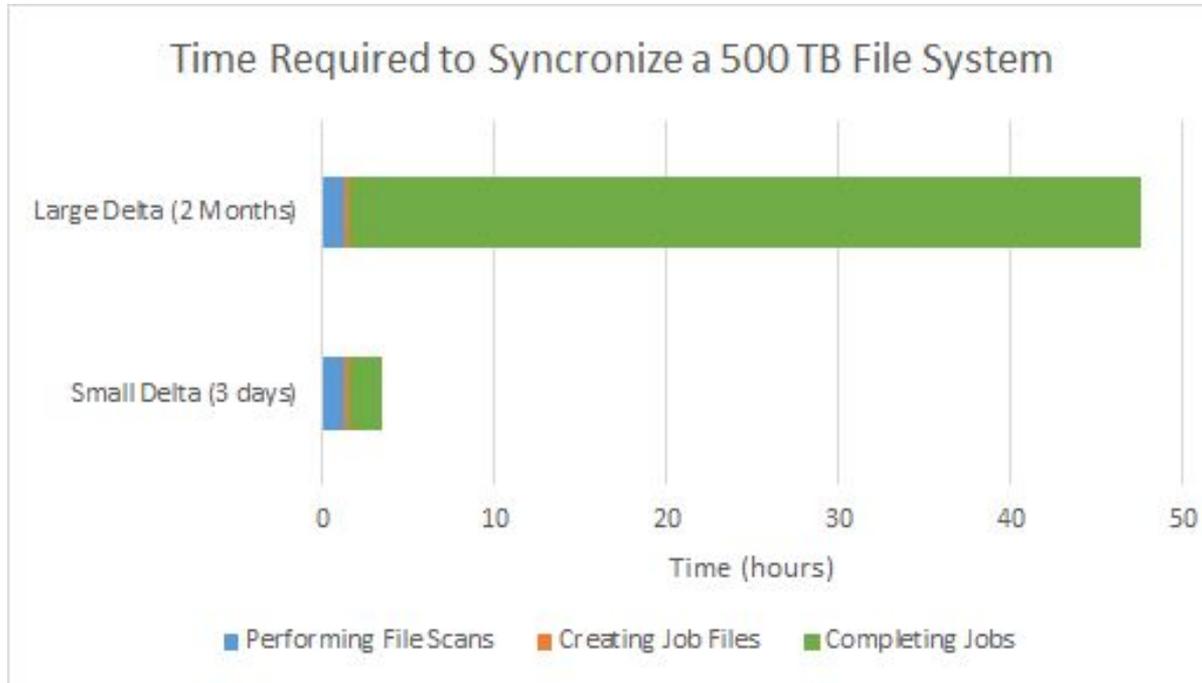


# Job Completion

- Workers start processing jobs in parallel.
- Utilizes system commands.



# Micro Benchmark



# Conclusions

## DistSync

- Processes file system scans.

- Creates job files.

- Maximises file system bandwidth.

Frequent syncs lead to faster syncs.

# Thank You!

Questions?