

R.A.I.D.F.S

Randomized Aggregation Independent Distributed File System

P2P Distributed File System with an API for Map-Reduce Integration

**Sven Reber, Jérémy Gotteland, David Froelicher,
Alban Marguet, Pascal Cudré, Valérian Pittet**

Context

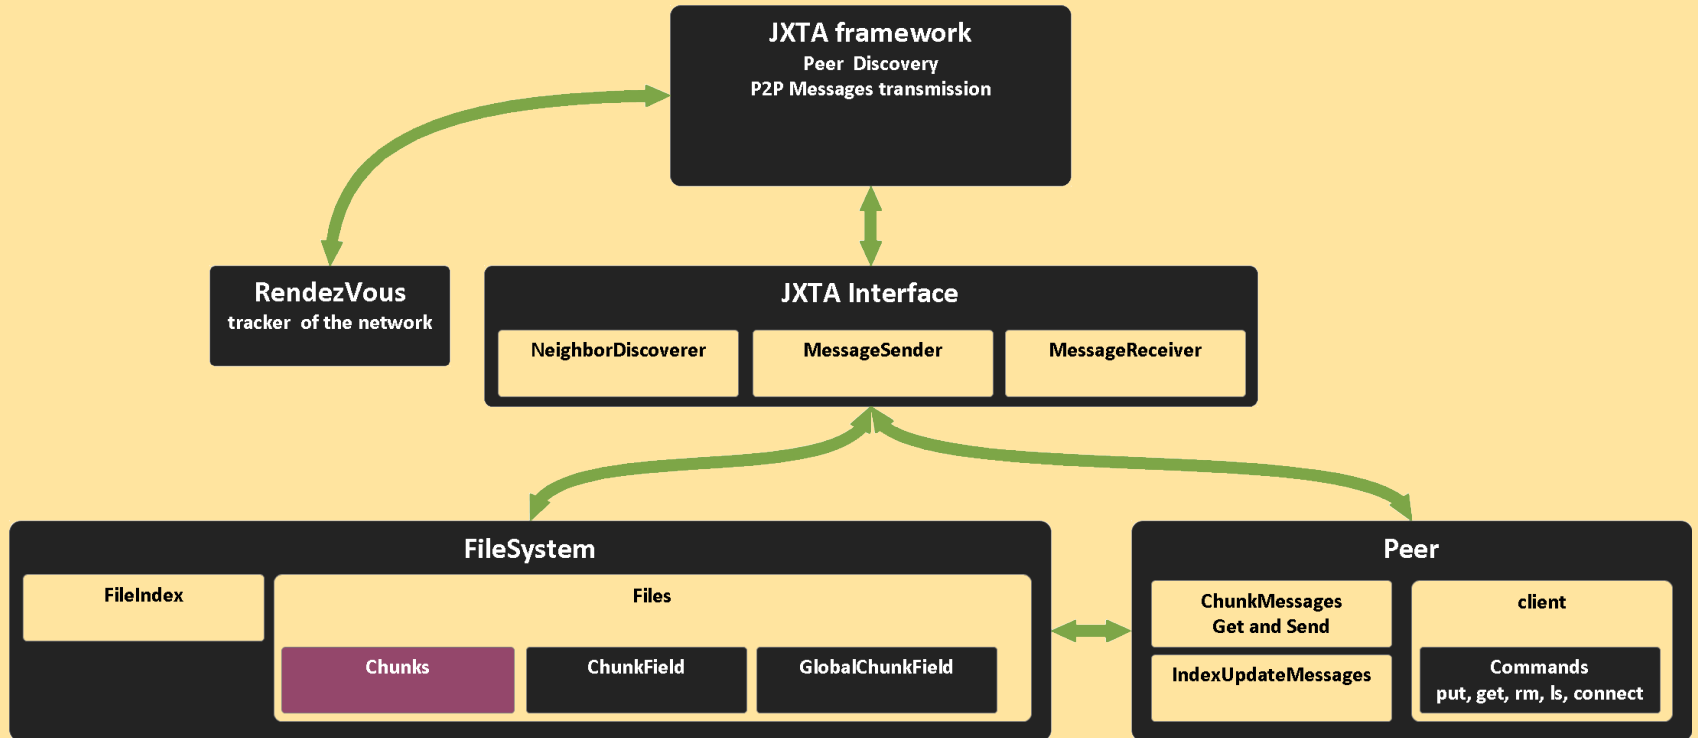
- clusters hard to configure and expensive to maintain
- everyone has a computer
- lots of unused storage and computational resources on end-user machine
- network connexions are improving

Goals



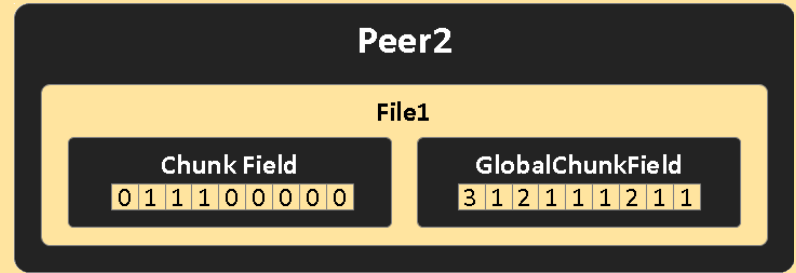

Peer to peer DFS that is

- designed to support Map-Reduce operations
 - chunking by line blocks
 - text files
- resilient
- easy to configure (dynamic configuration)
 - simply connect to the network and run your jobs

Architecture




DFS - Stabilization




GlobalChunkField ≤ 3 (arbitrary) is an unstable state

DFS - Stabilization




Look at its neighbors chunkfields

DFS - Stabilization




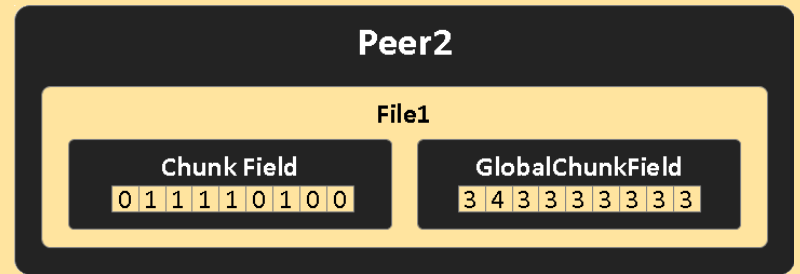

Randomly gets one of the insufficiently replicated chunk

DFS - Stabilization





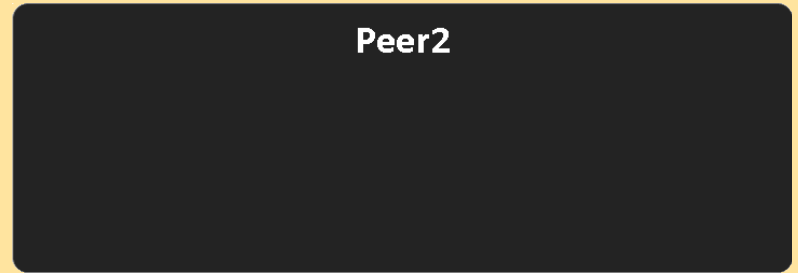

Do not download chunk if it finds enough replicas

DFS - Stabilization



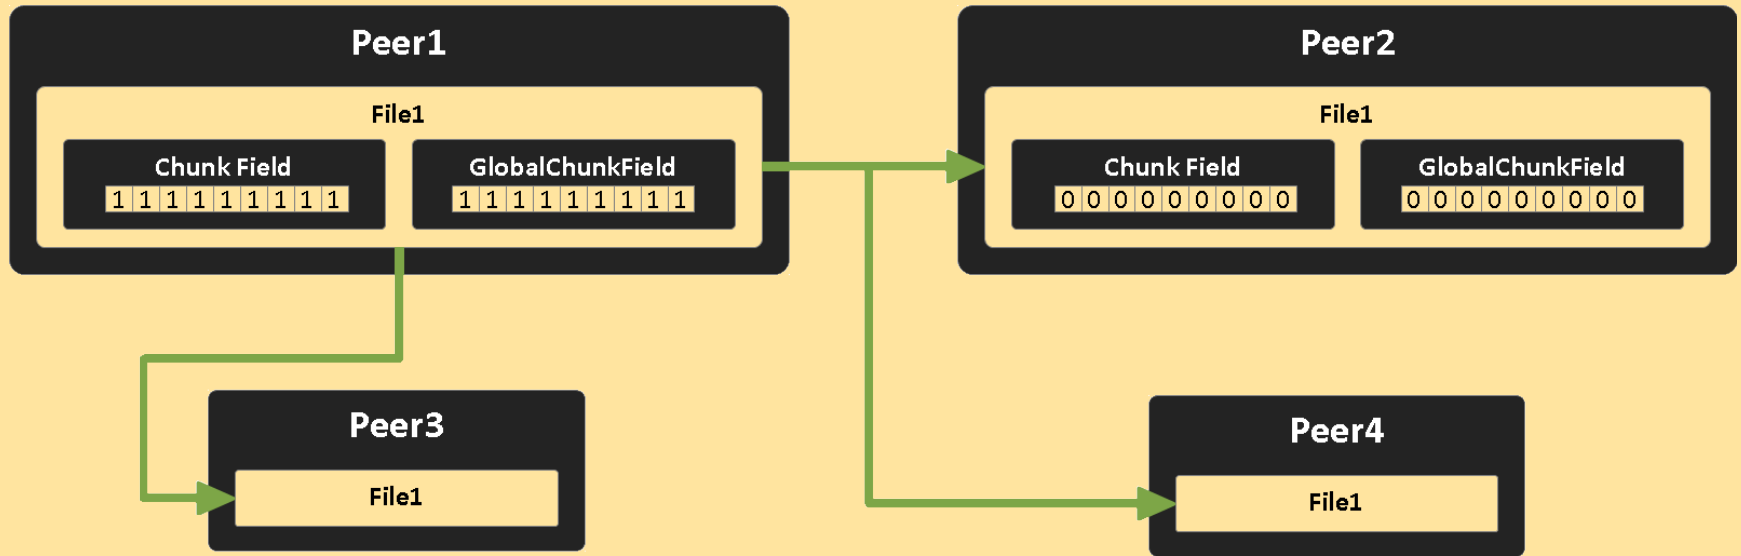
File is “stable” when there is enough replicas

DFS - put



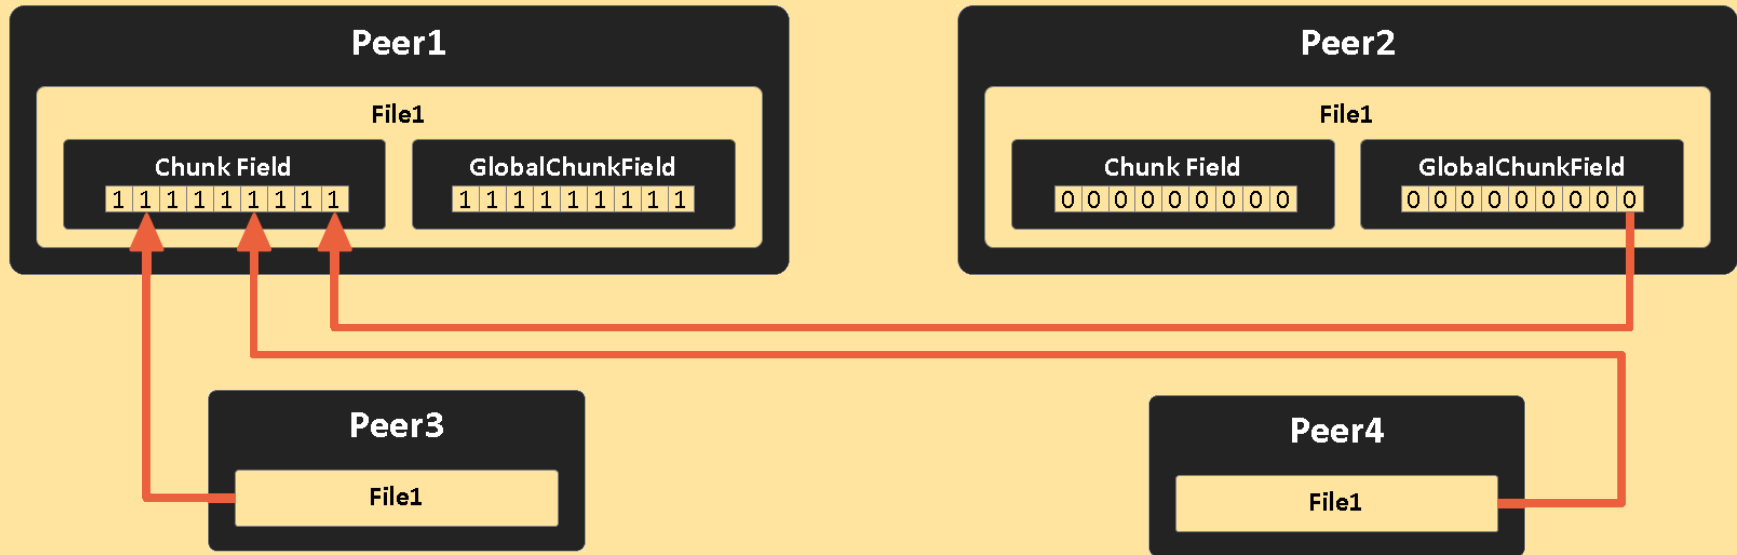
New file : “put” command

DFS - put



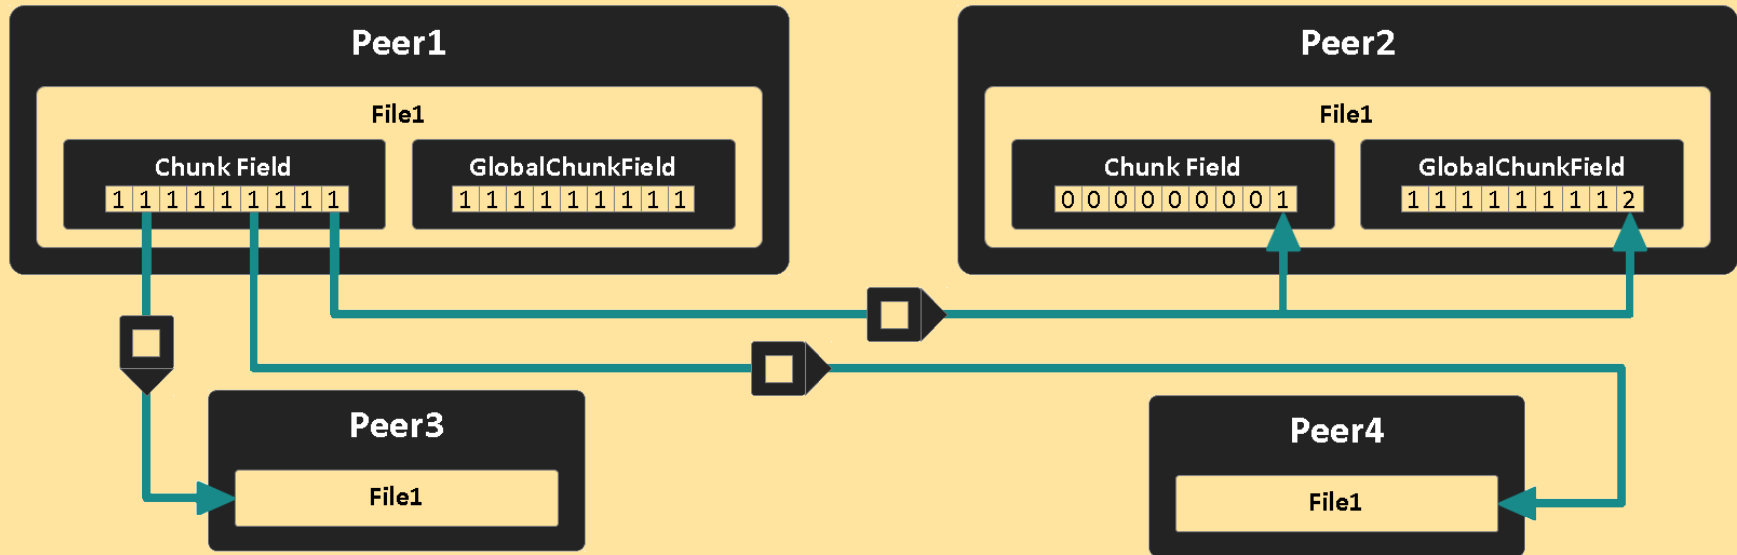
publish an index update, then neighbors discover every 20s

DFS - put



neighbors try to stabilize file (same process as before)

DFS - put



neighbors get missing chunks randomly to complete their GCF

DFS - other commands

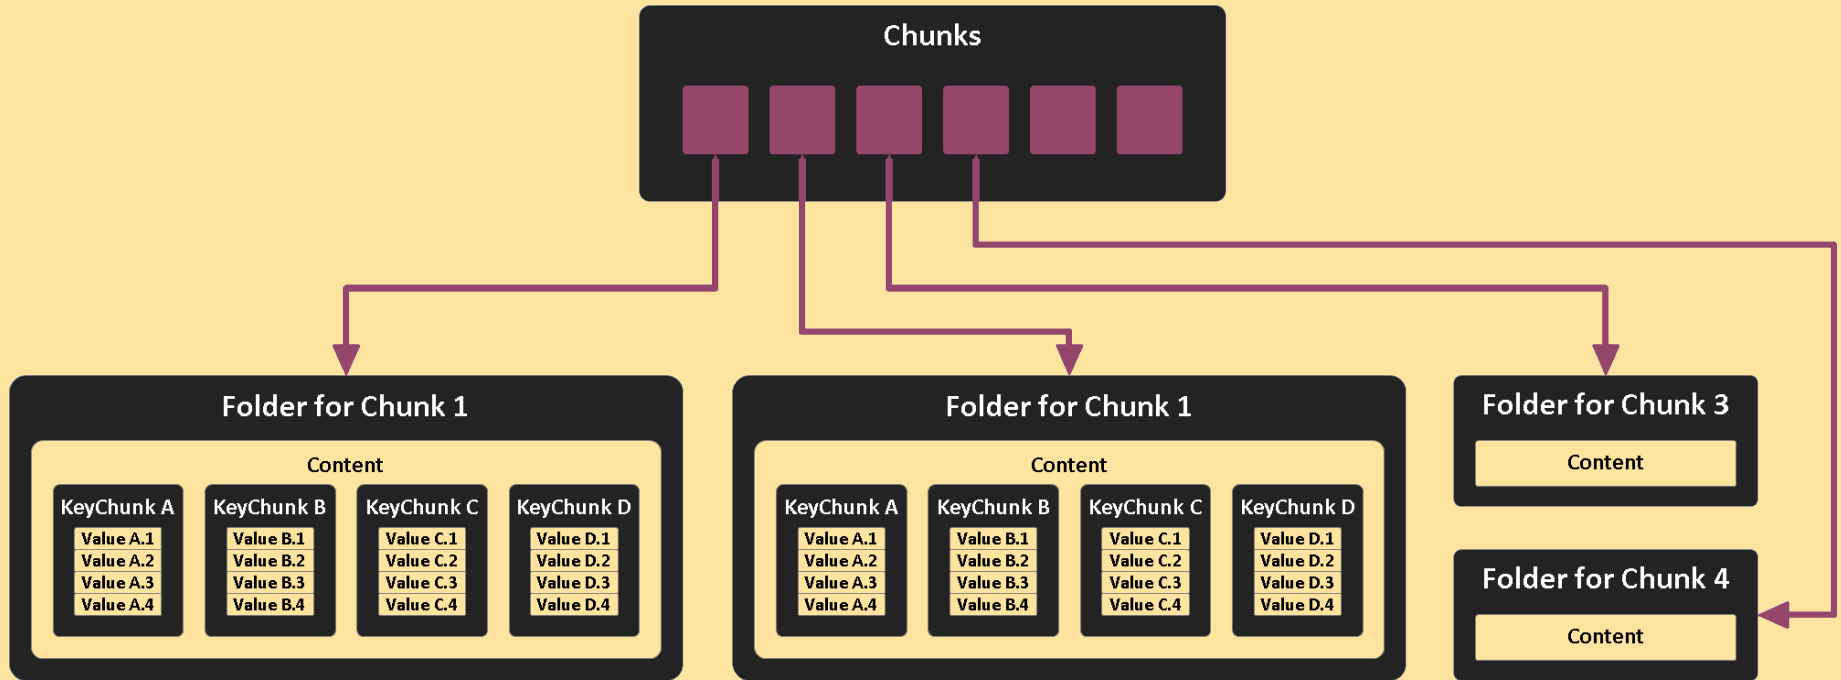
commands available

- ls
- put
- get
- rm

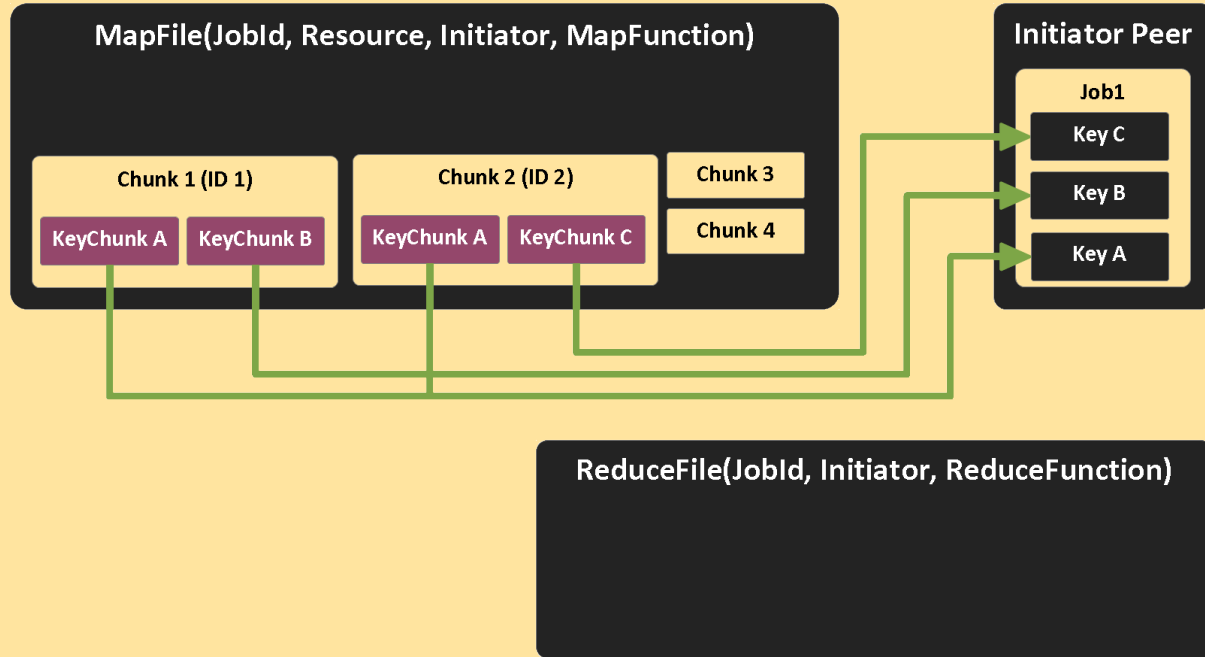
Map operation

- Some peer starts a Job
- MapFiles (jobid, Resource, Initiator, MapFunction)
 - Each chunk mapped to its result files (can be created in advance) -> One folder for each mapped chunk
 - One key chunk for each key discovered in the original chunk

MapFile

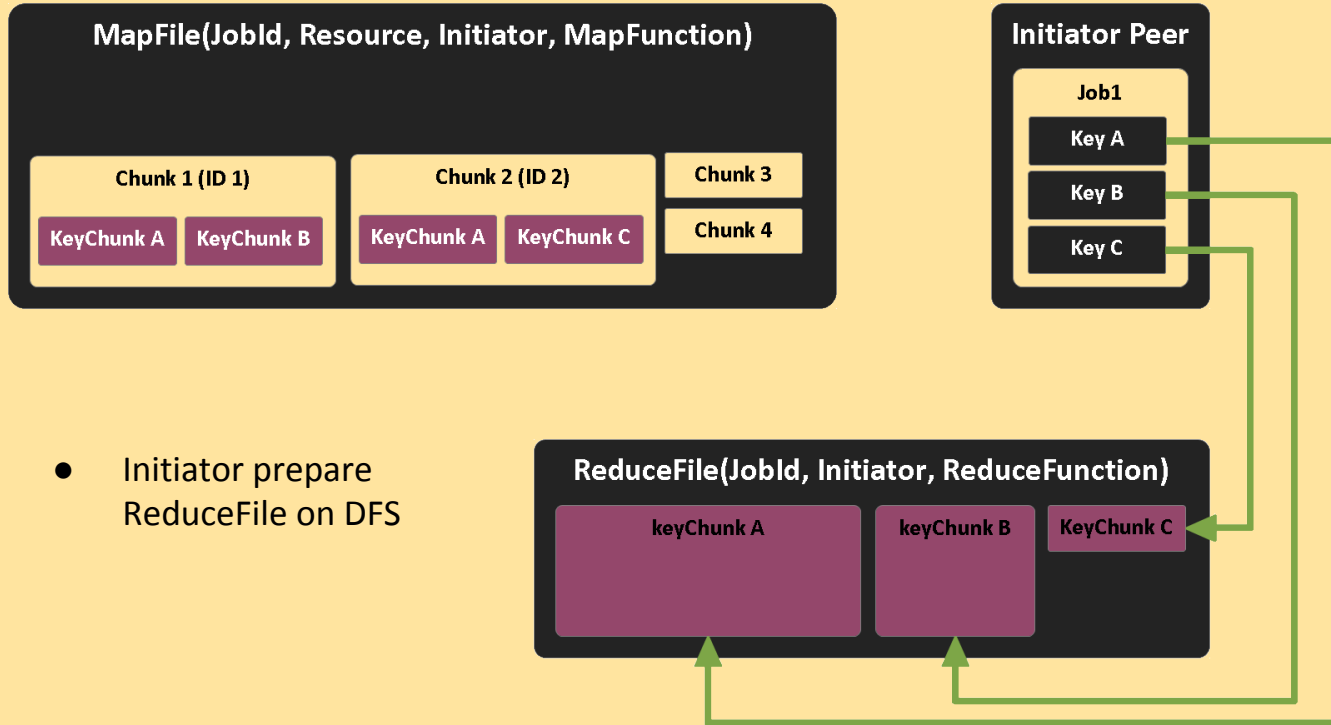


Reduce operation

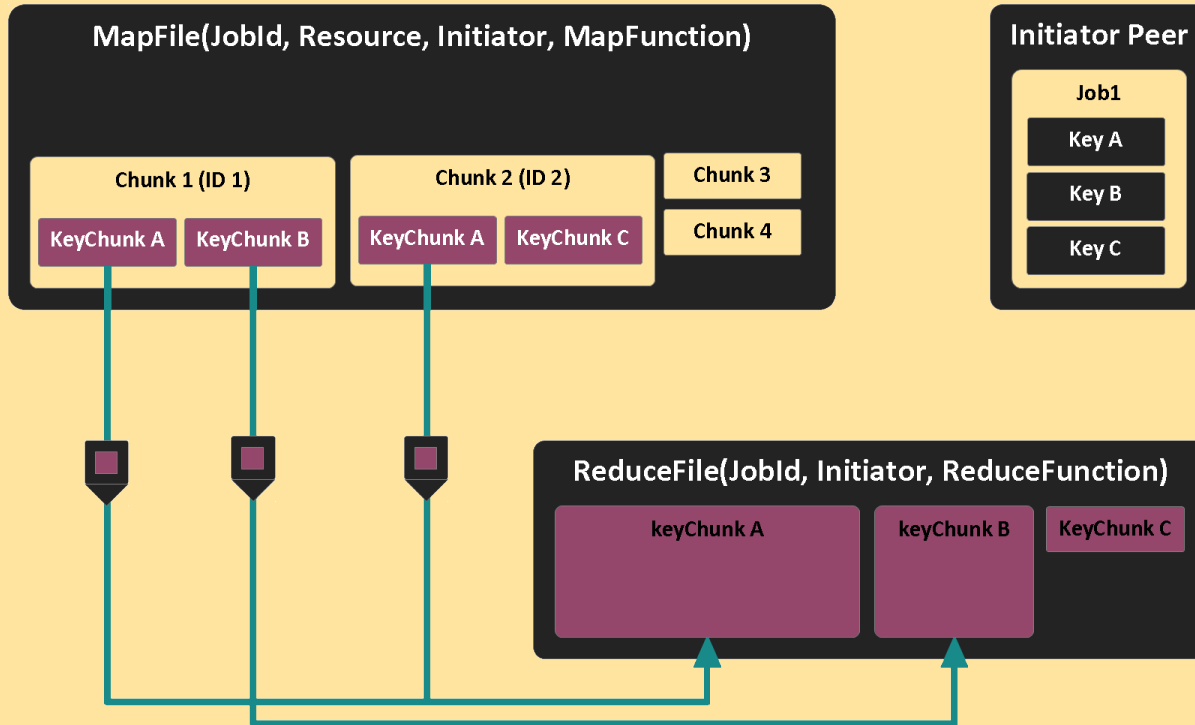


- Keys discovered during map
- Keys sent to initiator

ReduceFile

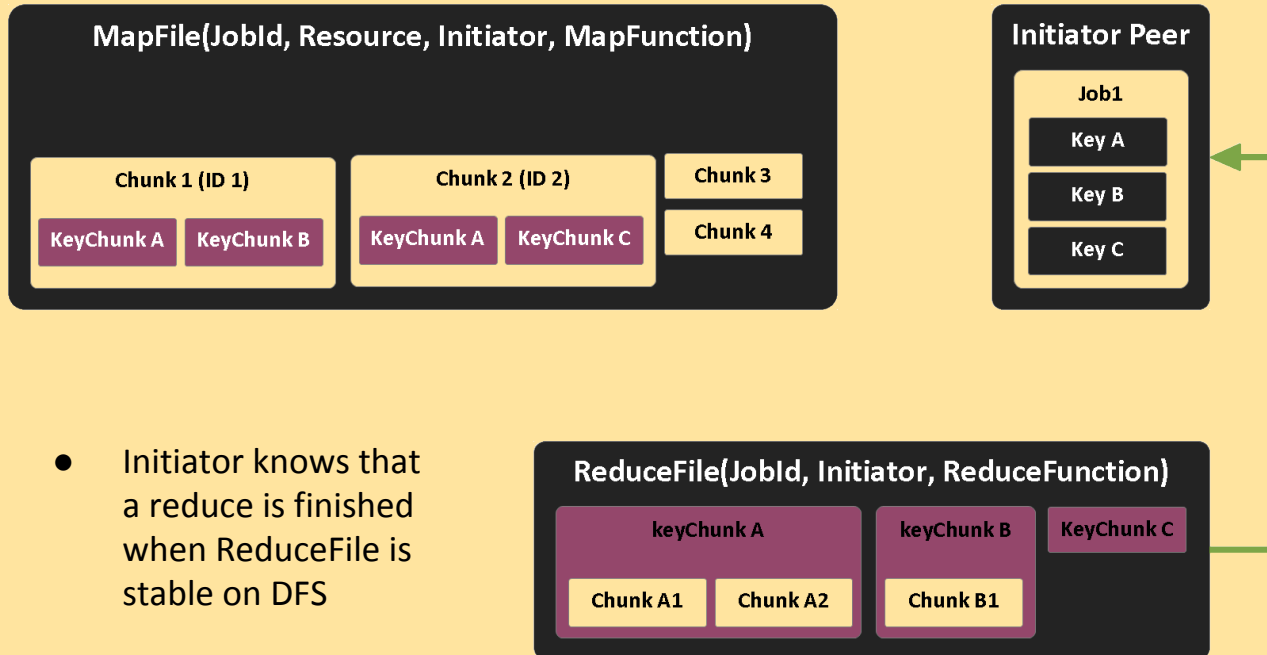


ReduceFile



- Peer that wants to create a ReduceFile chunk download the needed keyChunks

ReduceFile



- Initiator knows that a reduce is finished when ReduceFile is stable on DFS

What's Next

- Large Scale & Stress Tests of DFS
- Implement the Map and Reduce files
- Include multi-master management (results from the MRp2p paper)