


CS162 Section

Lecture 11


Project 4

- Implement a distributed key-value store that uses
 - Two-Phase Commit for atomic operations,
 - Replication for performance and fault-tolerance,
 - Encryption for security

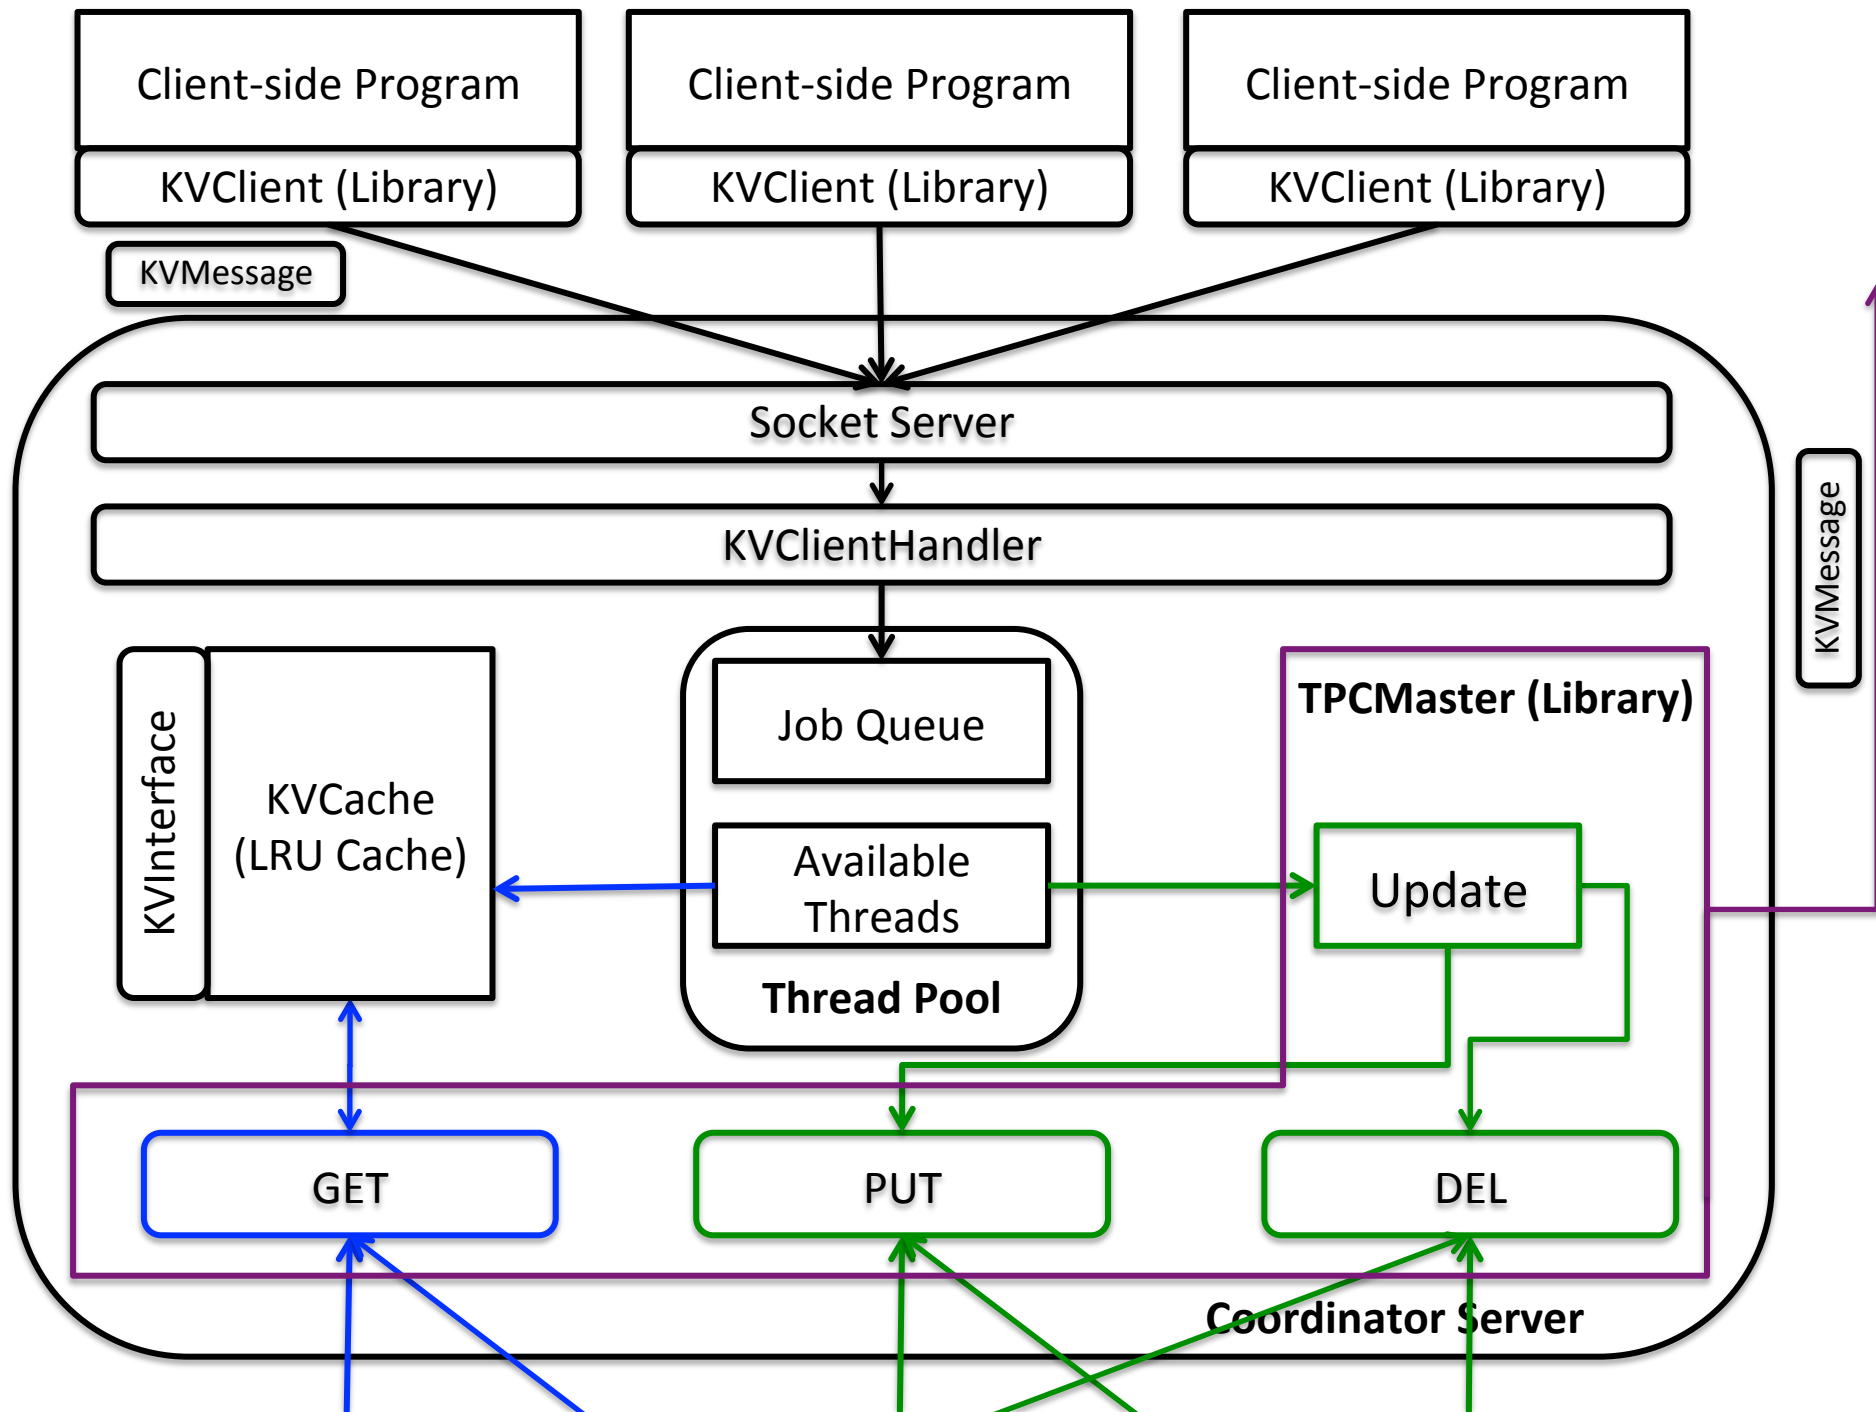
Distributed Key-Value Store

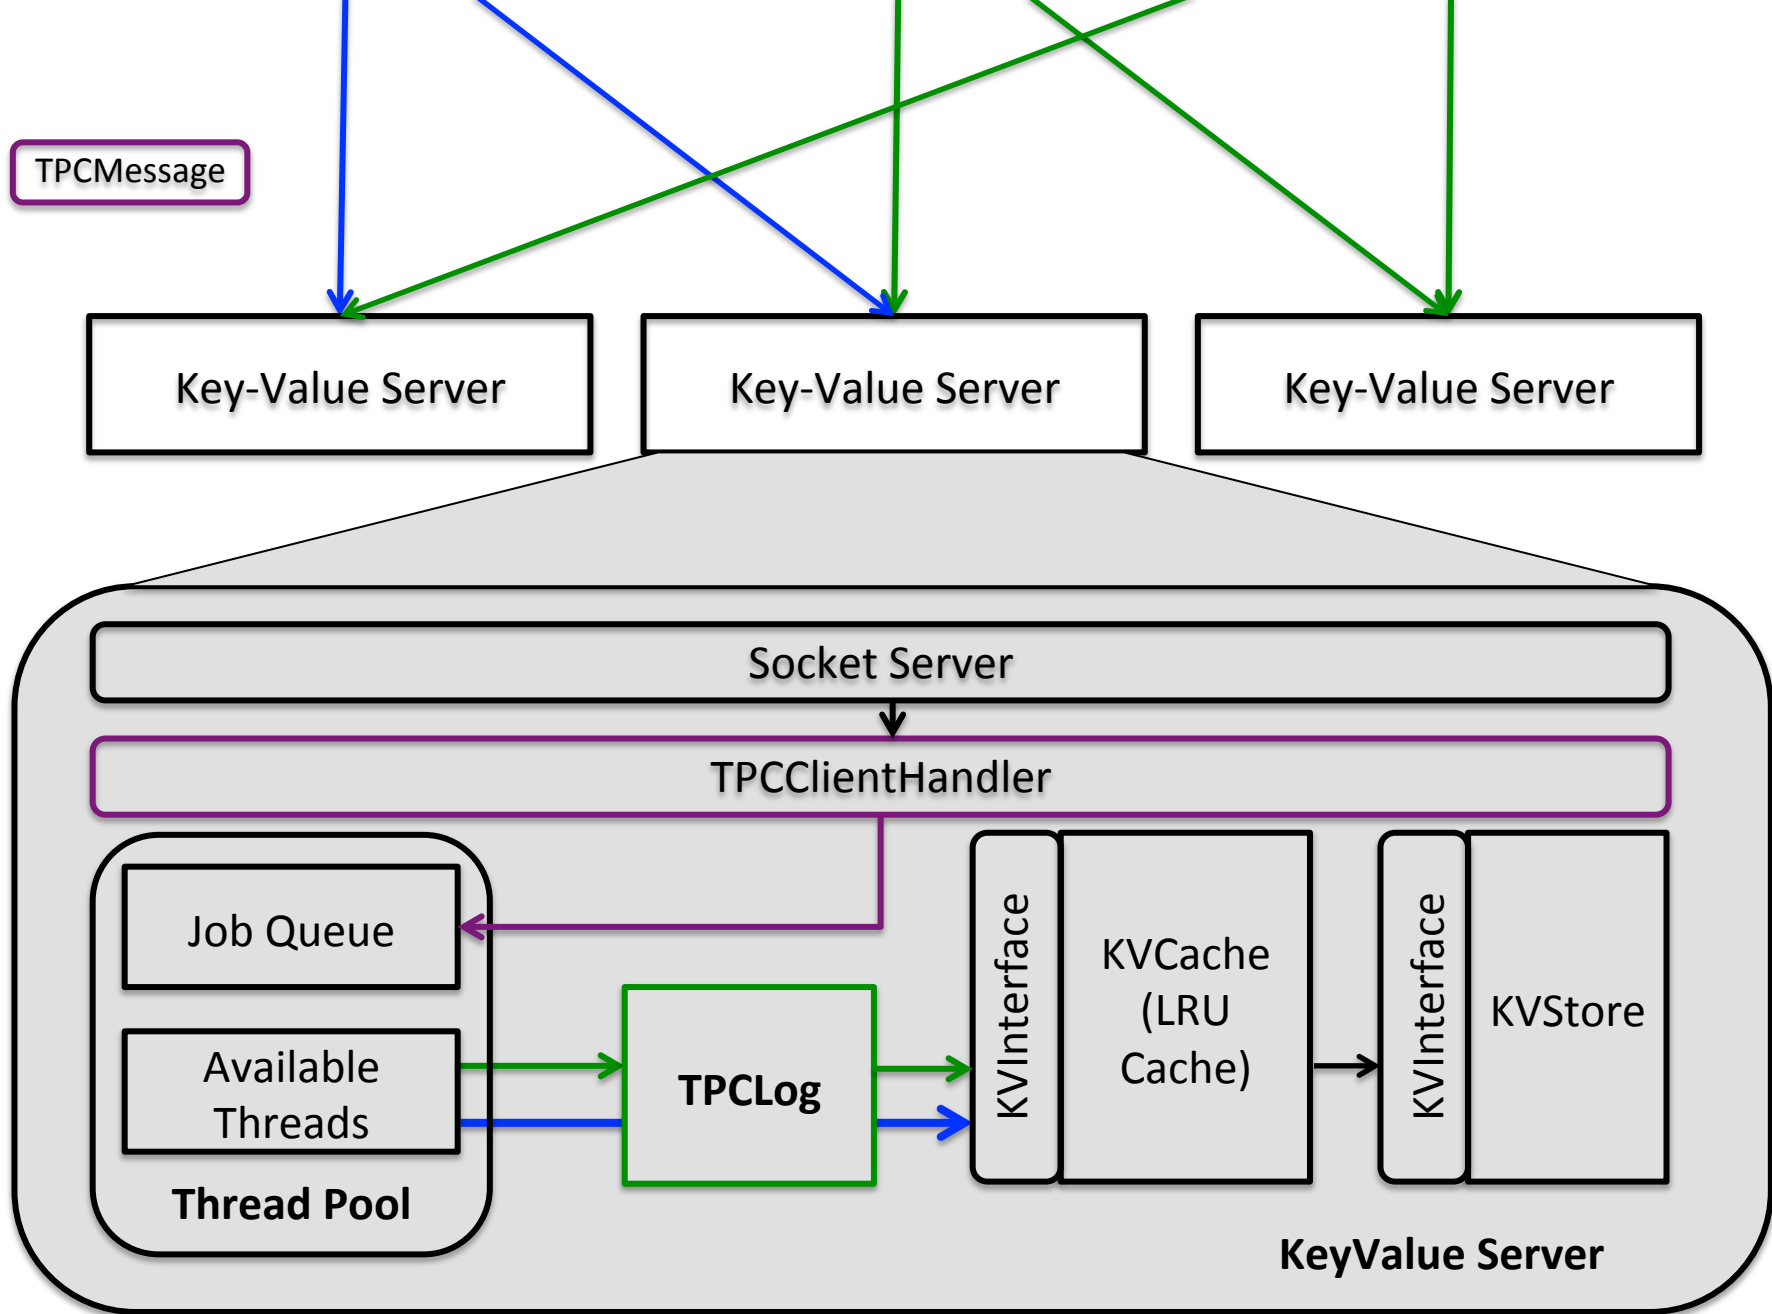


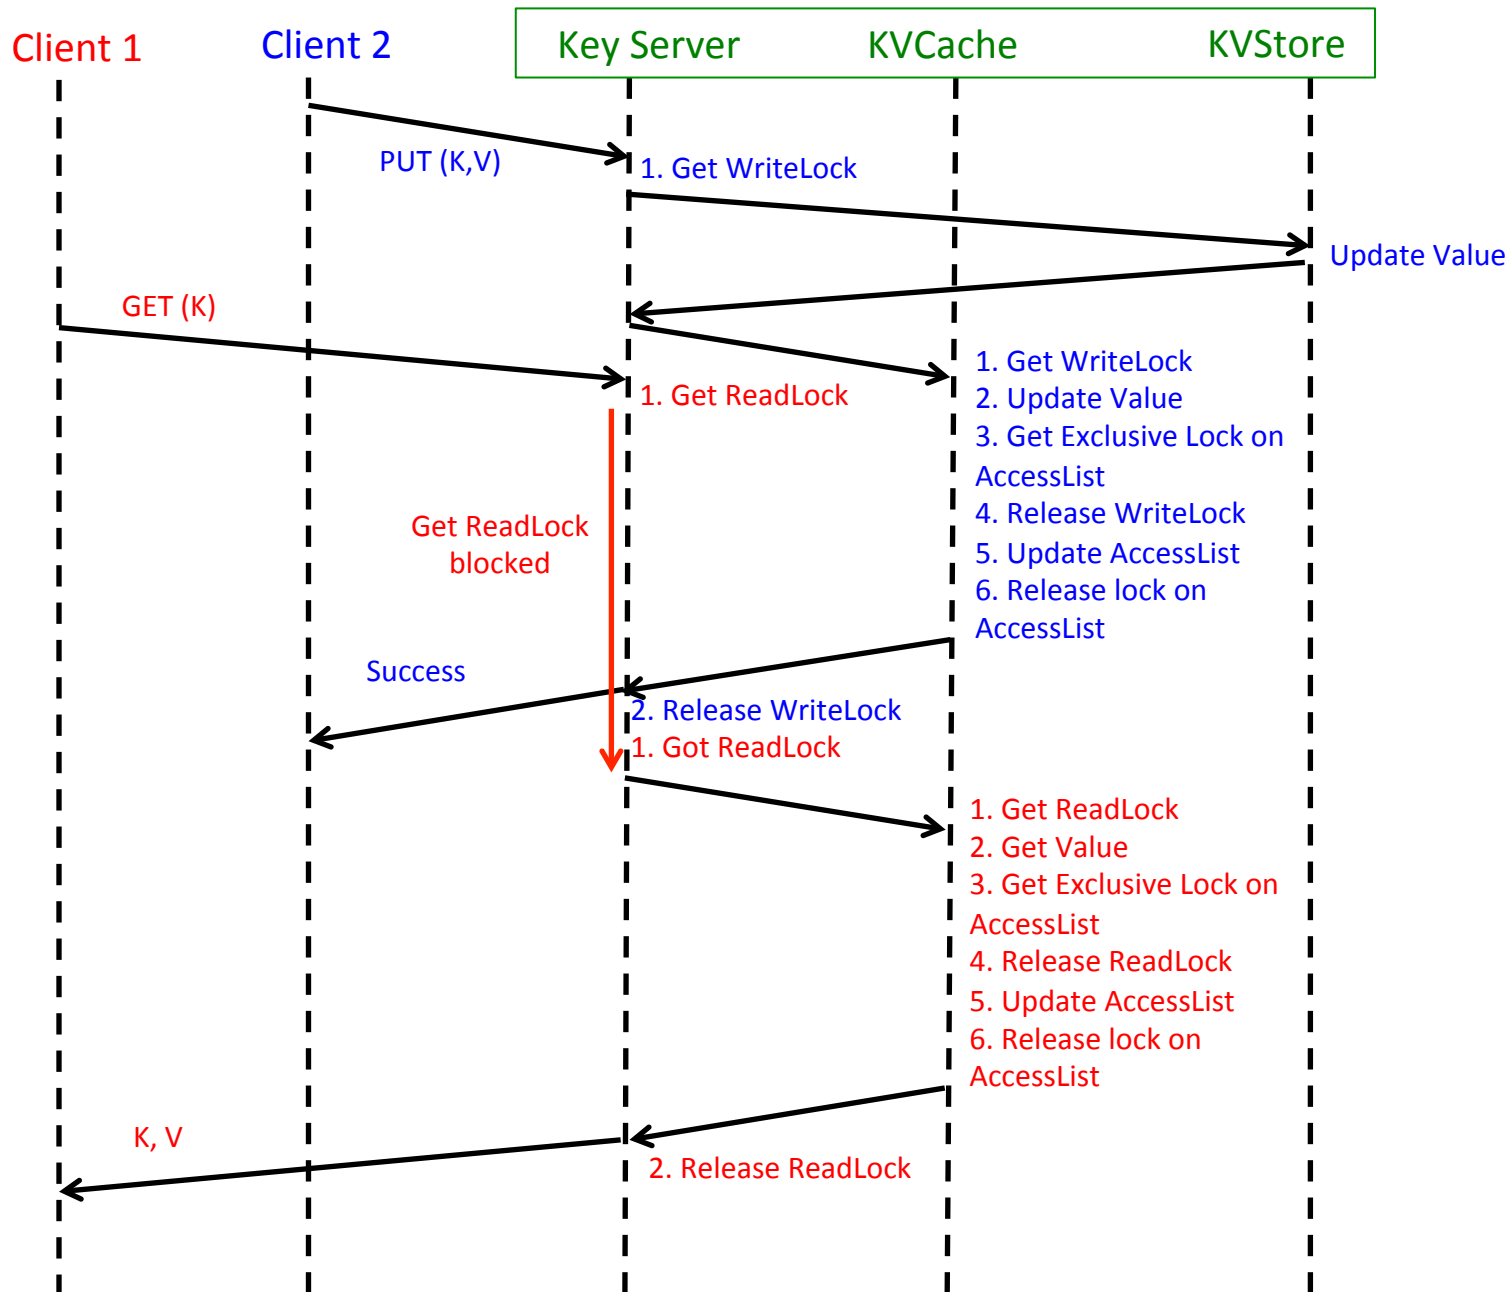
Consistent Hashing

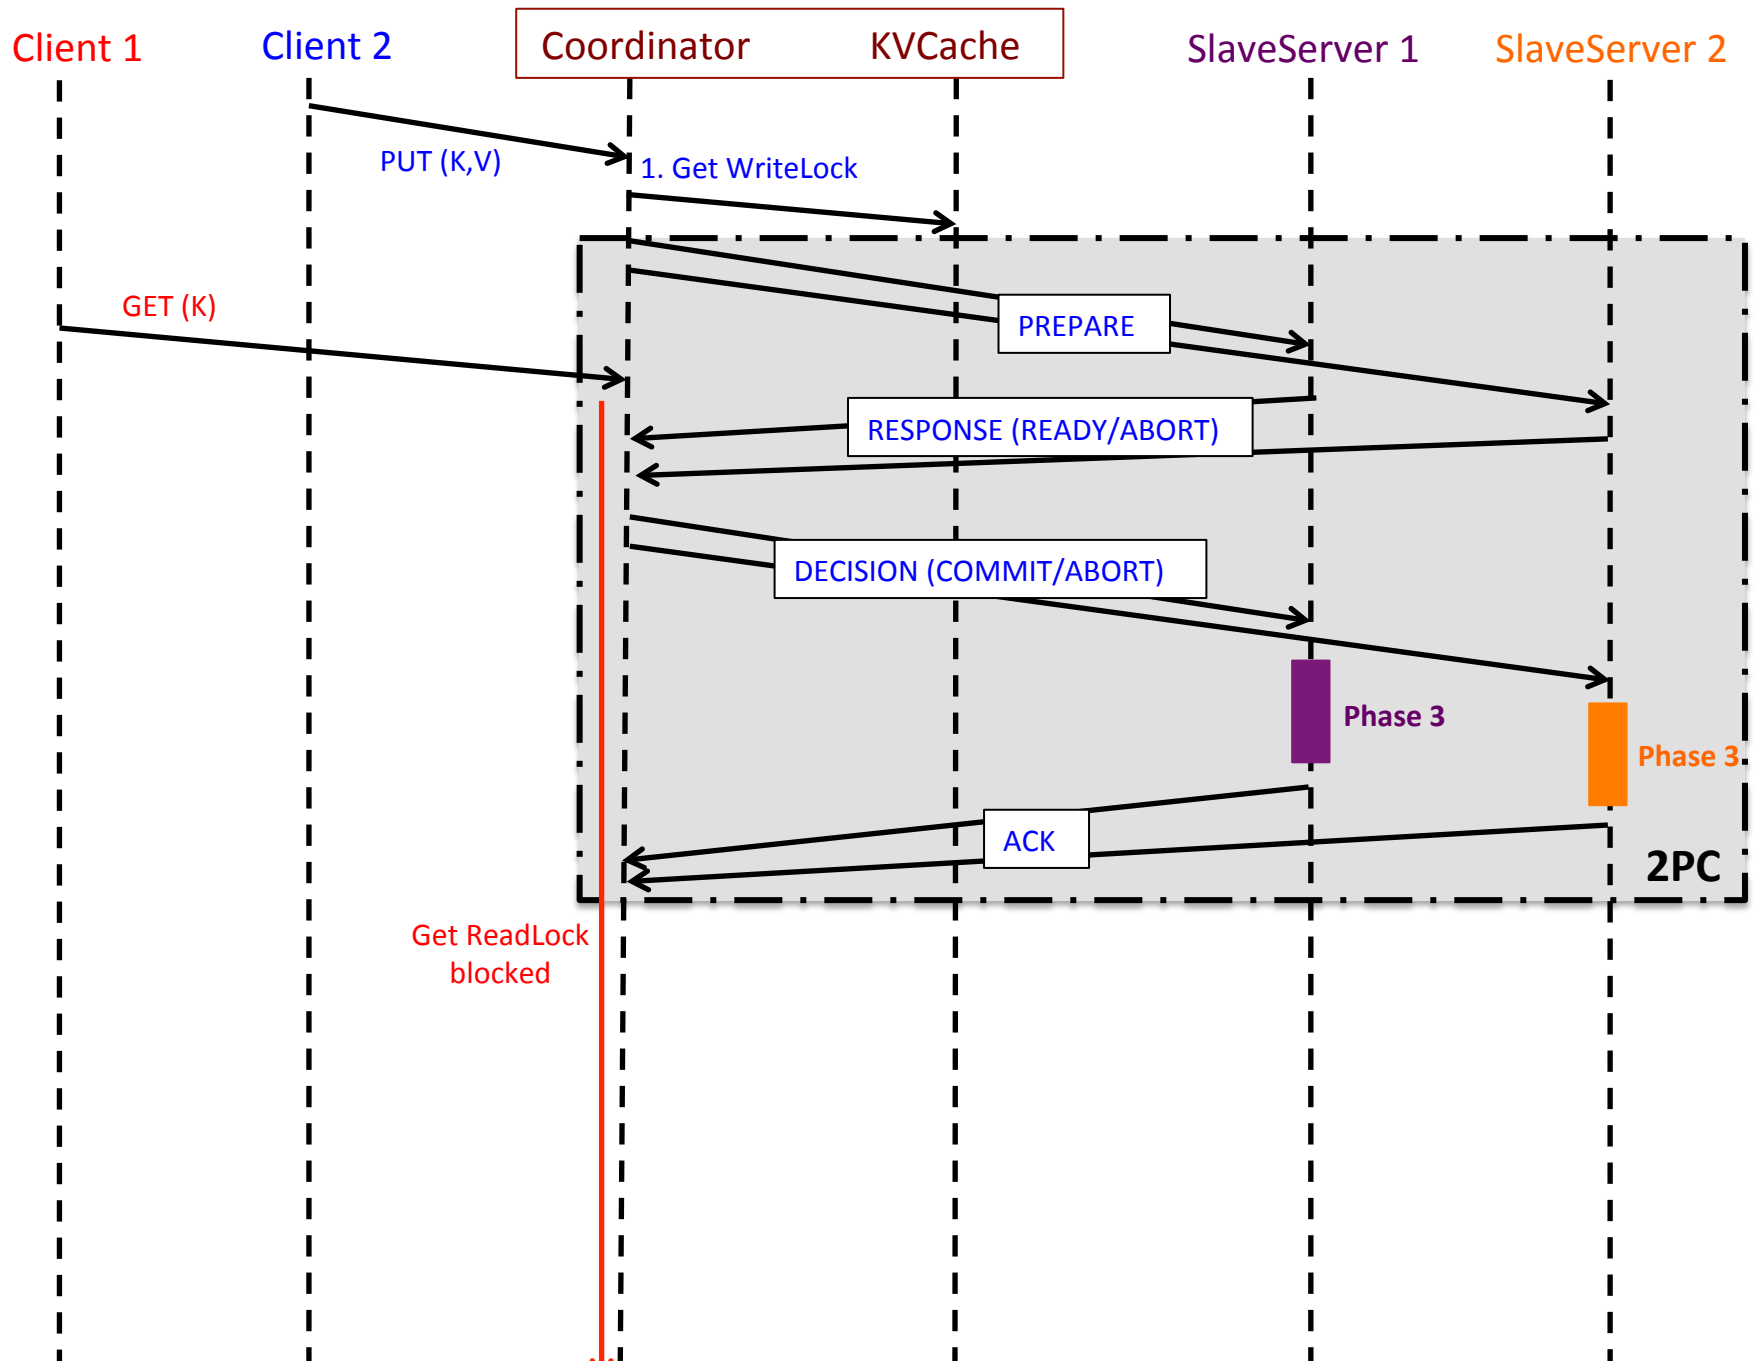


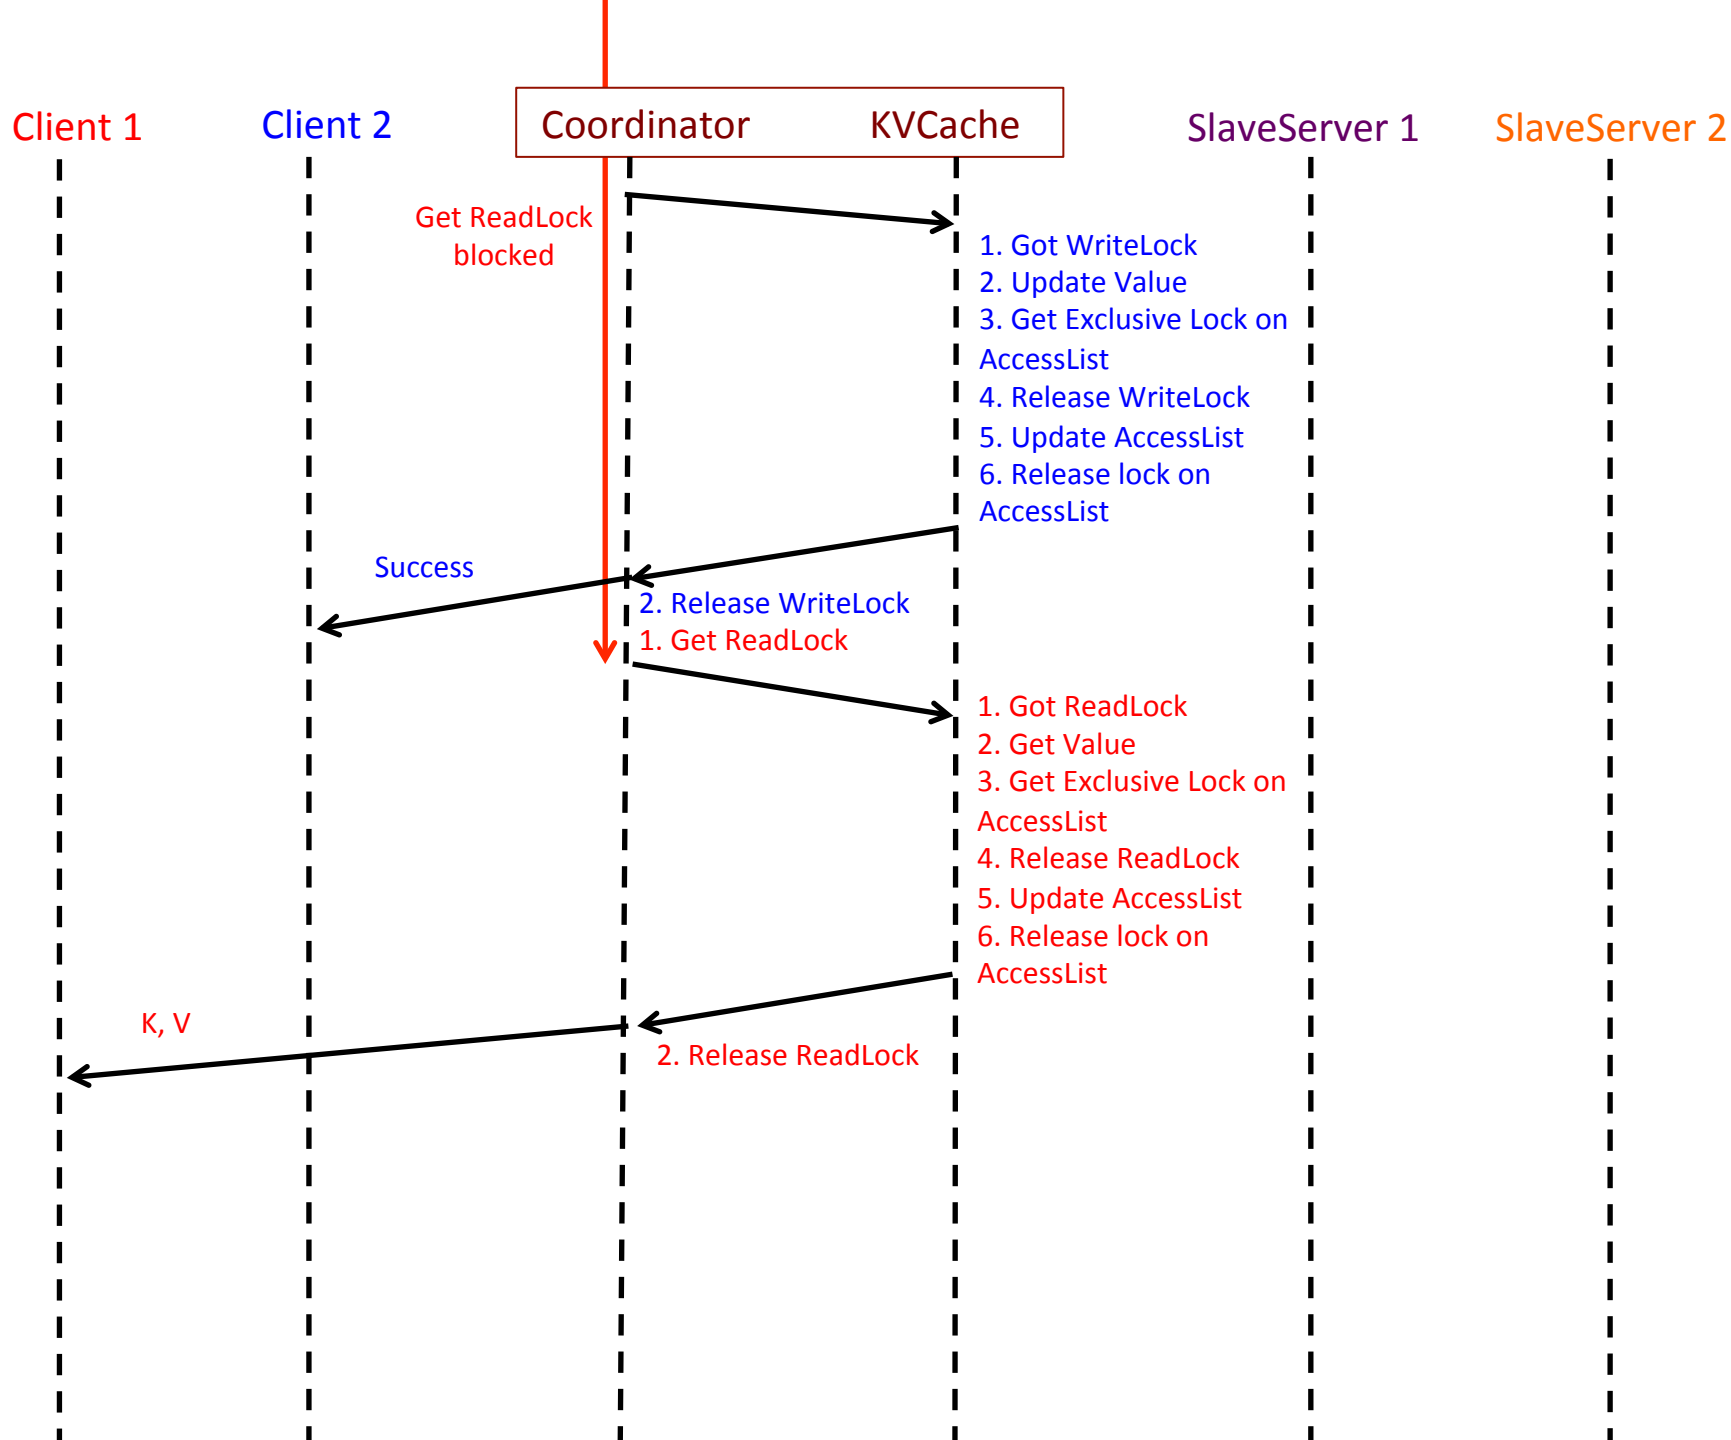
- SlaveServer 128-bit ID using `java.util.GUID`
 - Slave count known in advance
 - Nodes will come back if fails (no permanent failure)
- Hash Keys to 128-bit numbers



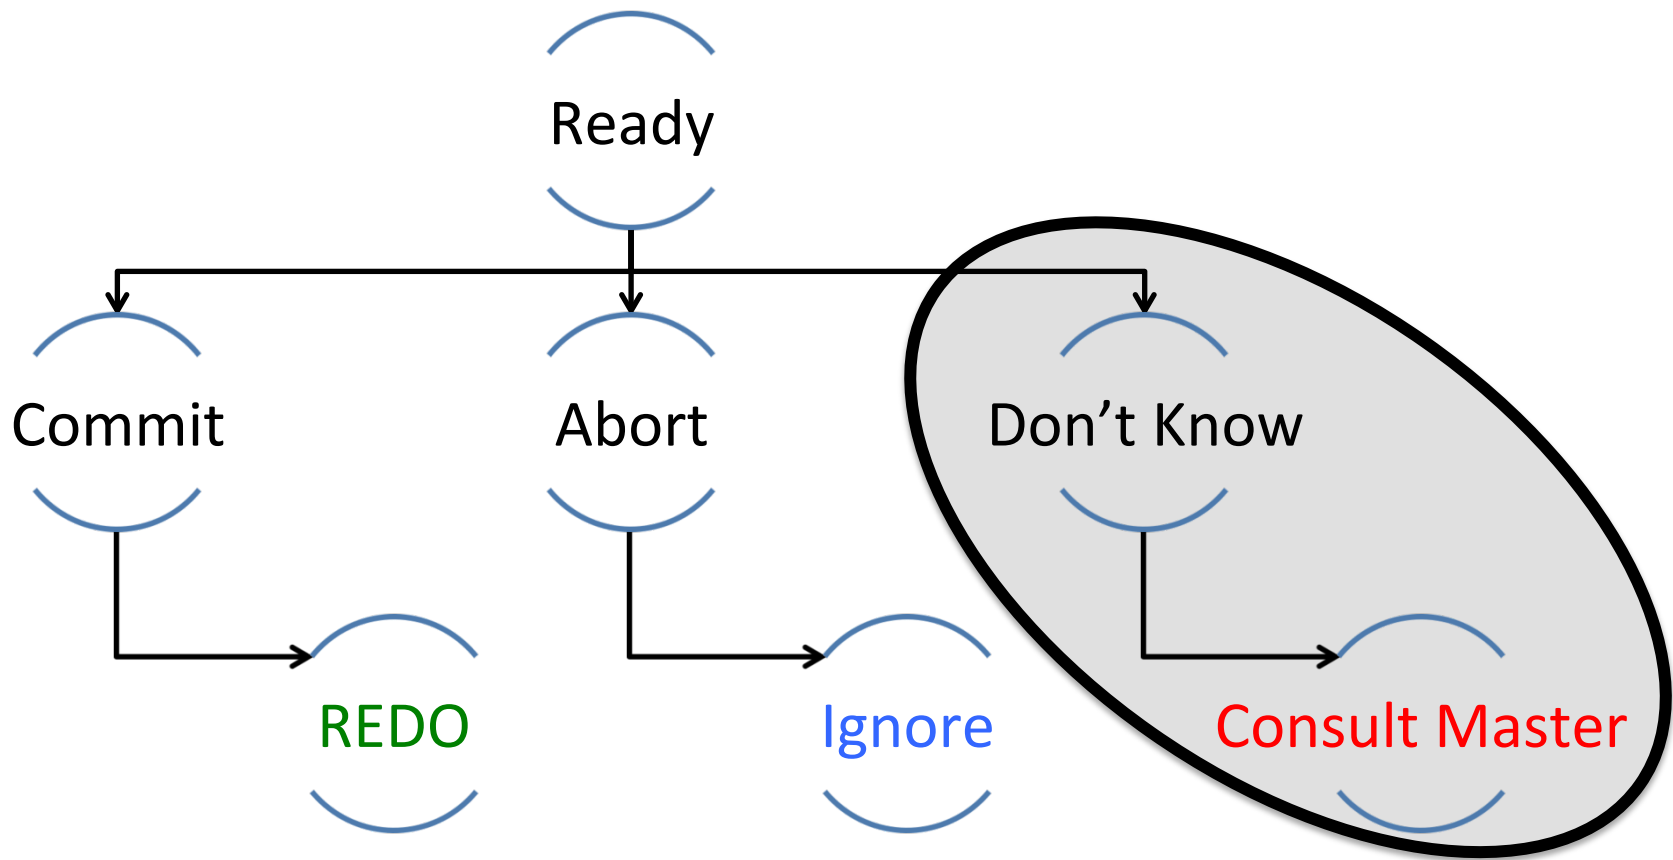








Recovery from Site Log Entries



What is MTBF?

What is MTTR?

Fault Models

- Failures are independent^{*}
So, single fault tolerance is a big win
- Hardware fails fast (blue-screen, panic, ...)
- Software fails-fast (or stops responding/hangs)
- Software often repaired by reboot:
 - Heisenbugs – Works On Retry
 - (Bohrbugs – Faults Again On Retry)
- Operations tasks: major source of outage
 - Utility operations – UPS/generator maintenance
 - Software upgrades, configuration changes

Fault Tolerance Techniques

- Fail fast modules: work or stop
- Spare modules: yield instant repair time
- Process/Server pairs: Mask HW and SW faults
- Transactions: yields ACID semantics (simple fault model)

Fault-tolerance in MapReduce

- What if
 - A task crashes
 - A node crashes
 - Worker node
 - Master node
 - A task is going slow

Key Security Properties

- Authentication
- Data integrity
- Confidentiality
- Non-repudiation

Lxvnkbgz Vhffngbvtmbhg pbma Vkrimhzktiar


- Vhffngbvtmbhg bg ikxlgvx hy twoxkltkbxi
- Maxkx bl t dxi, ihllxllbhc hy pabva teehpl
wxvhwbgz, unu pbmahnu pabva wxvhwbgz
bl bgyxtlbuex
 - Manl, dxi fnlu ux dxim **lxvkm** tgw ghm **znxlltuex**

Securing Communication with Cryptography

- Communication in presence of adversaries
- There is a key, possession of which allows decoding, but without which decoding is infeasible
 - Thus, key must be kept **secret** and not **guessable**

Using Symmetric Keys


- Same key for encryption and decryption



- Transferring keys over the network is problematic

Public Key / Asymmetric Encryption

- Sender uses receiver's **public** key
 - Advertised to everyone
- Receiver uses complementary **private** key
 - Must be kept secret



Properties of RSA

- Requires generating large, random prime numbers
 - Algorithms exist for quickly finding these (probabilistic!)
- Requires exponentiation of very large numbers
 - Again, fairly fast algorithms exist
- Overall, much slower than symmetric key crypto
 - One general strategy: use public key crypto to exchange a (short) symmetric [session key](#)
 - Use that key then with AES or such
- How difficult is recovering d , the private key?
 - Equivalent to finding prime factors of a large number
 - Many have tried - believed to be very hard (= brute force only)
 - (Though *quantum computers* can do so in polynomial time!)