


① Requirements

- show all video favorable to the user (past watched videos) (demographic)
- region specific
- genre specific
- language
- editorial content / promotional videos
- bookmarked video / watch later
- continue watching
- most recent videos
- feedback loop for recommended content
- customize

② Scale (IPL: 28M users)



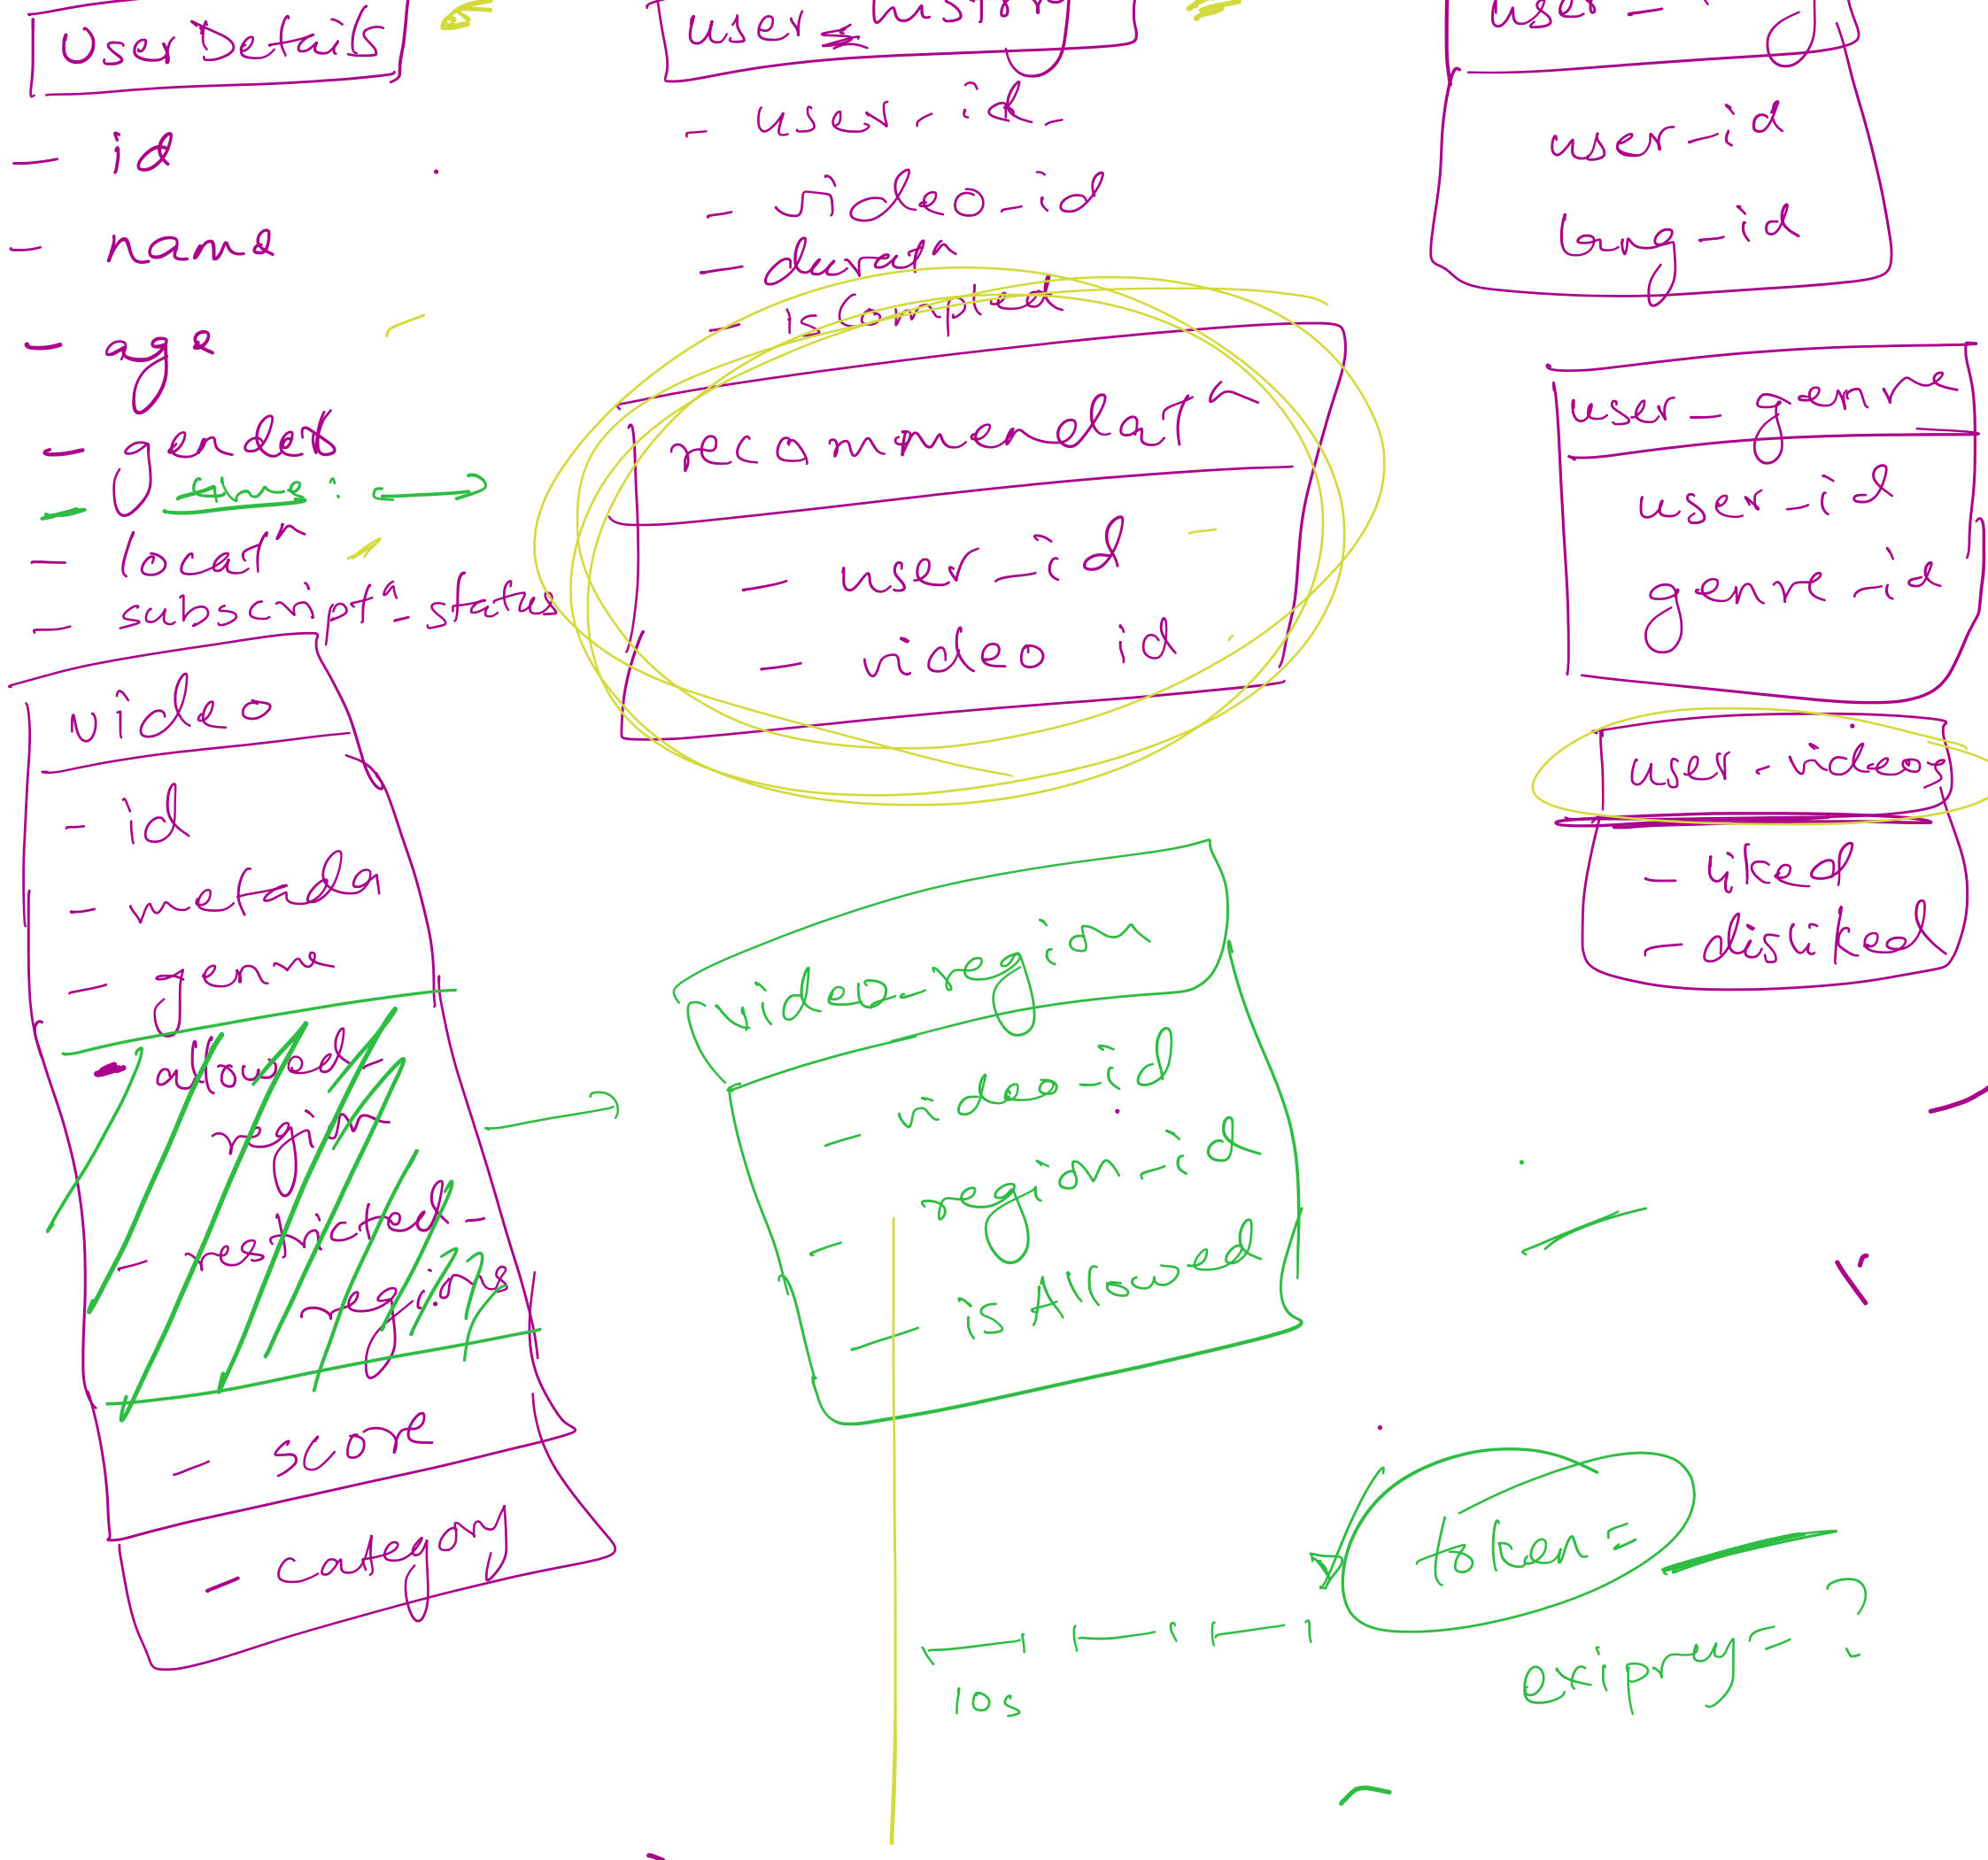
reads = $\frac{1}{n} \times 1,000,000$

2.5M - 5 mins
 $\frac{2.5 \times 10^6}{5 \times 60}$ 10K tps

writes = 10% of reads
 = 1K writes/sec

storage: $1 \text{ KB} \times 1000 \times 3600 \times 24 \times 365$
 = $3 \times 10^8 \text{ KB}$
 = 300 GB \rightarrow 1TB (user)

③ Schema (100 users)



④ Trade-Offs

- P: ✓
- C: .
- A: ✓
- latency: super-fast



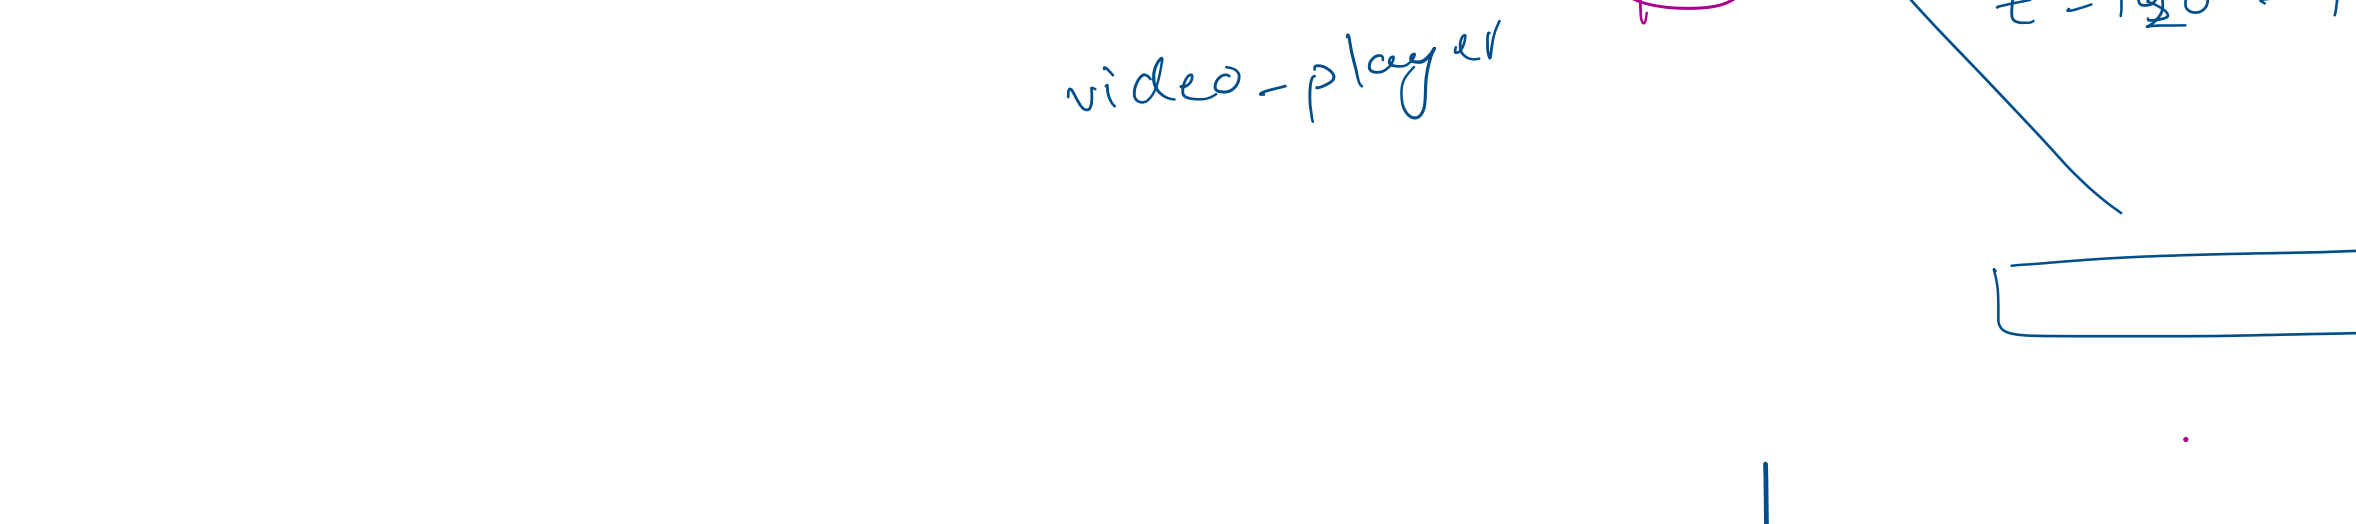

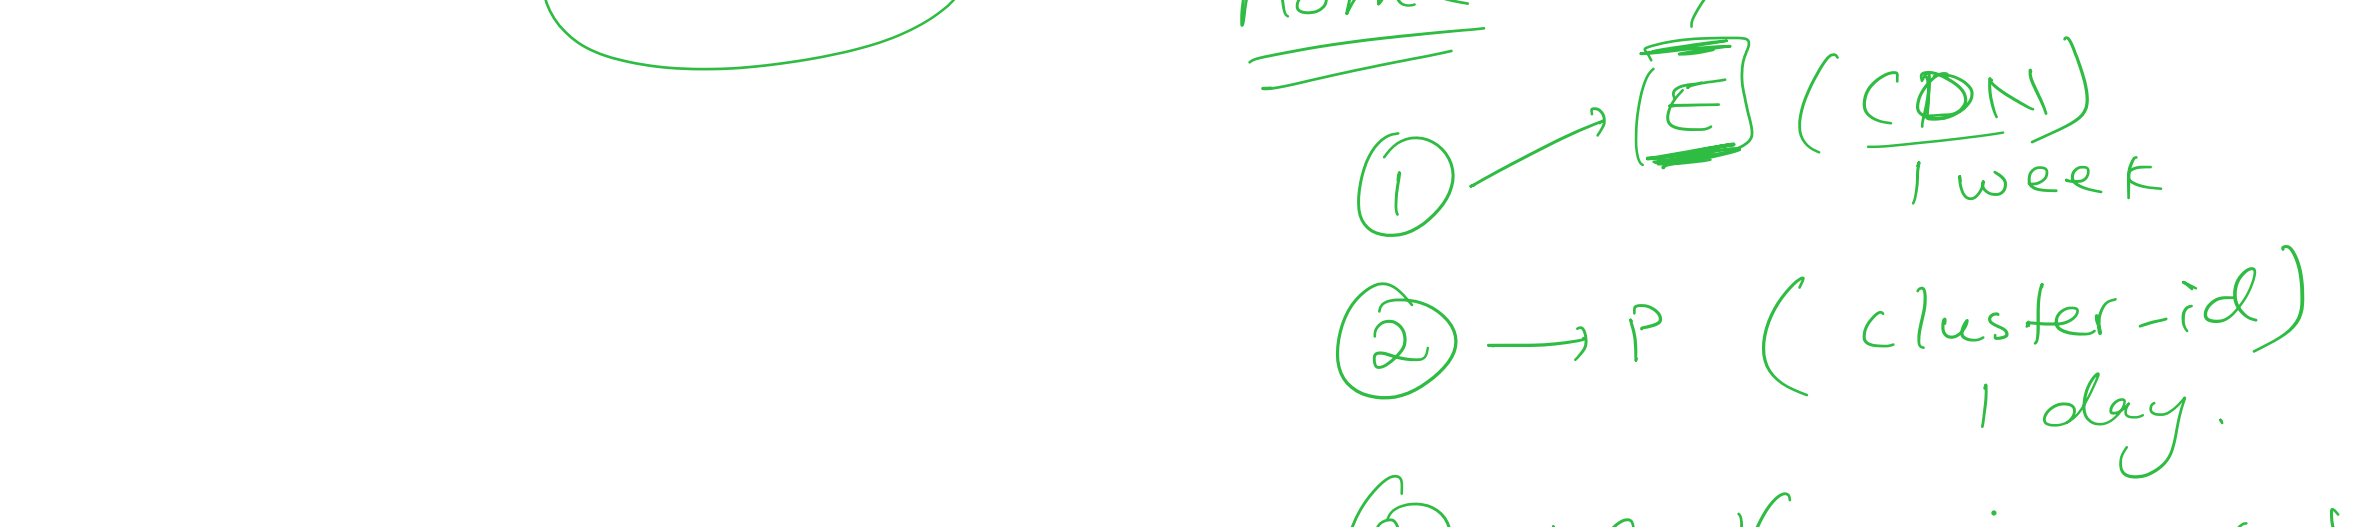
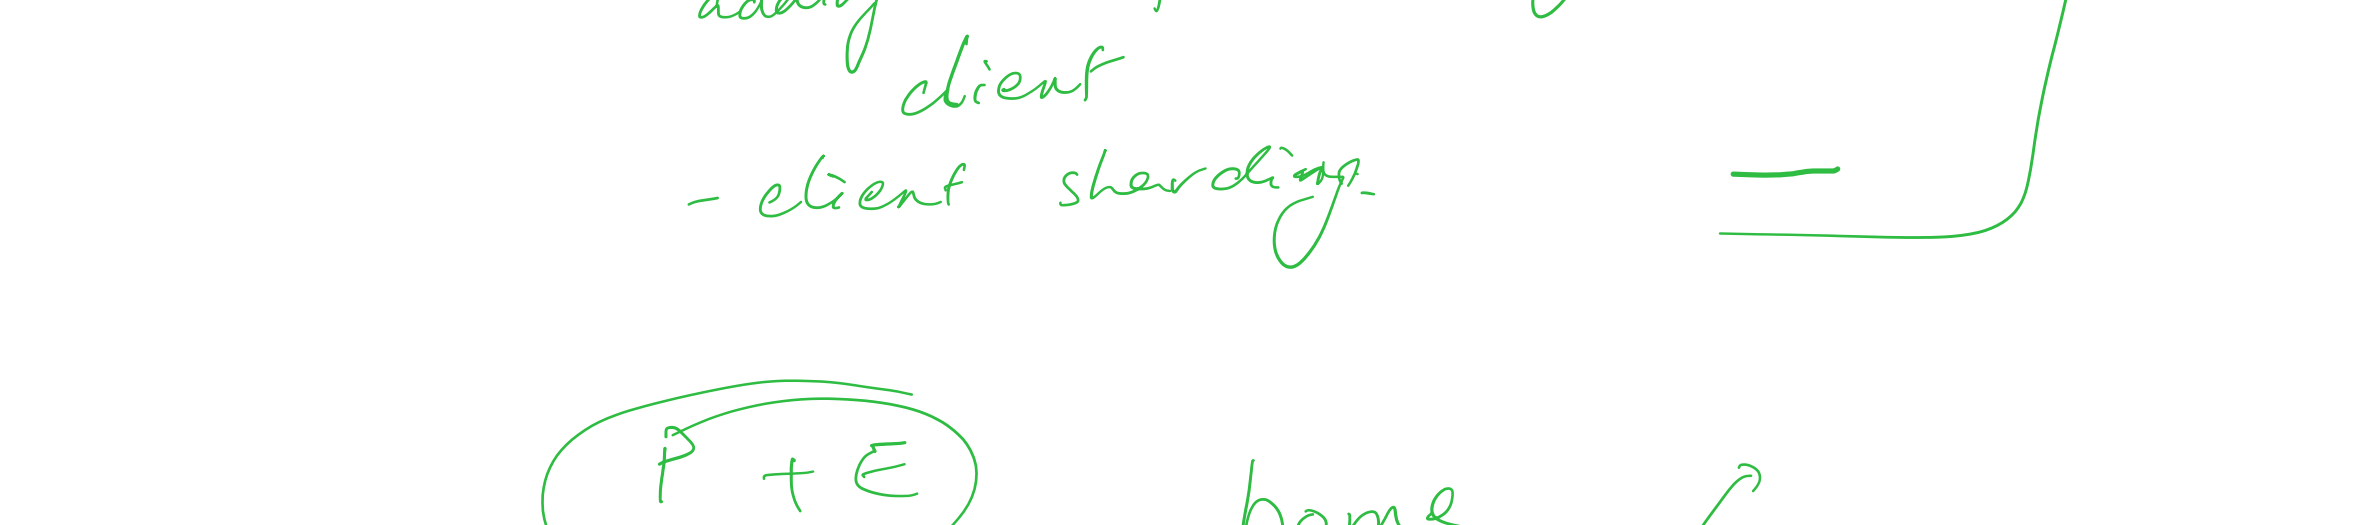

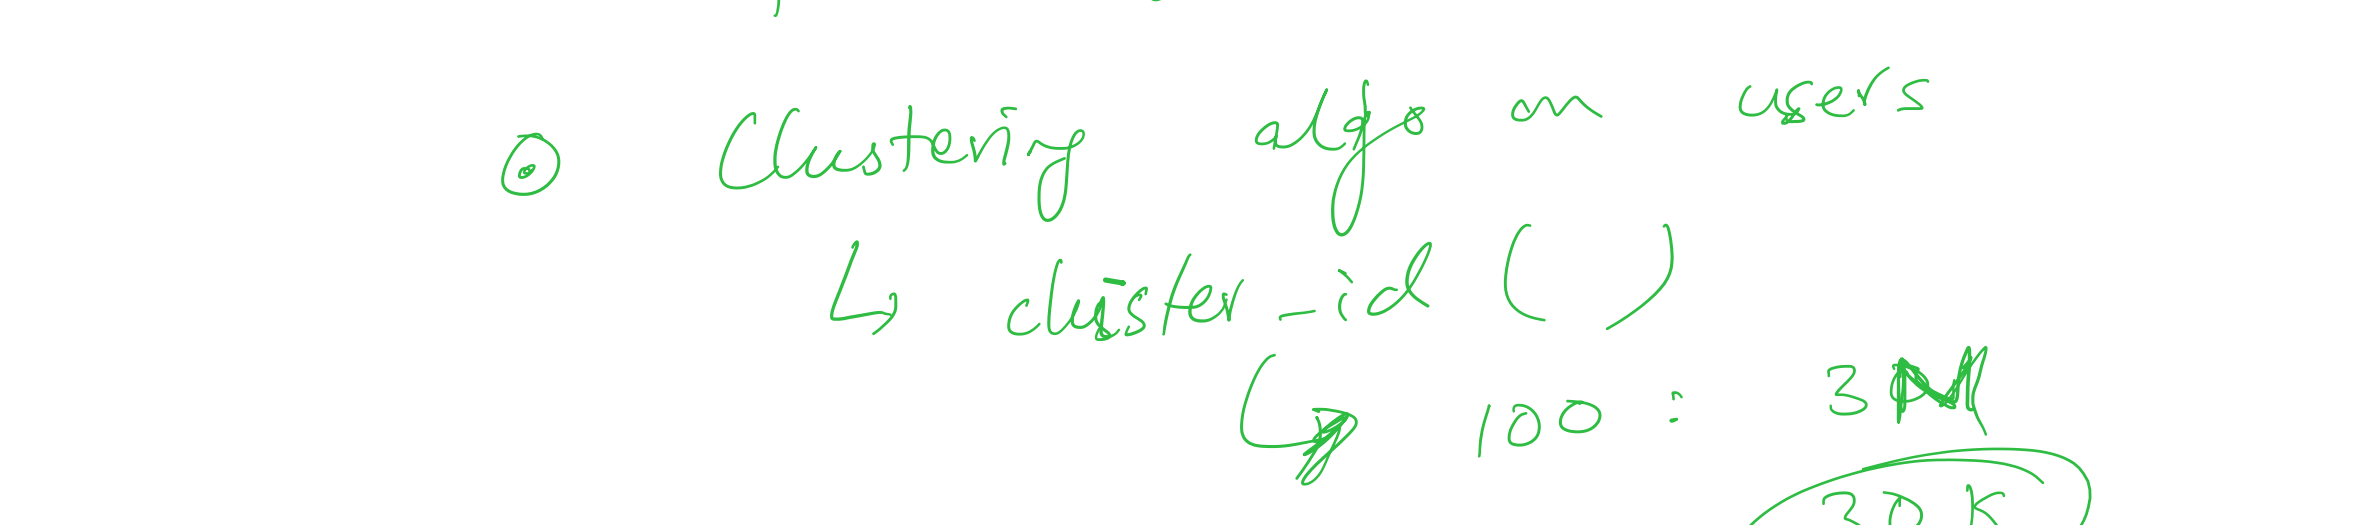
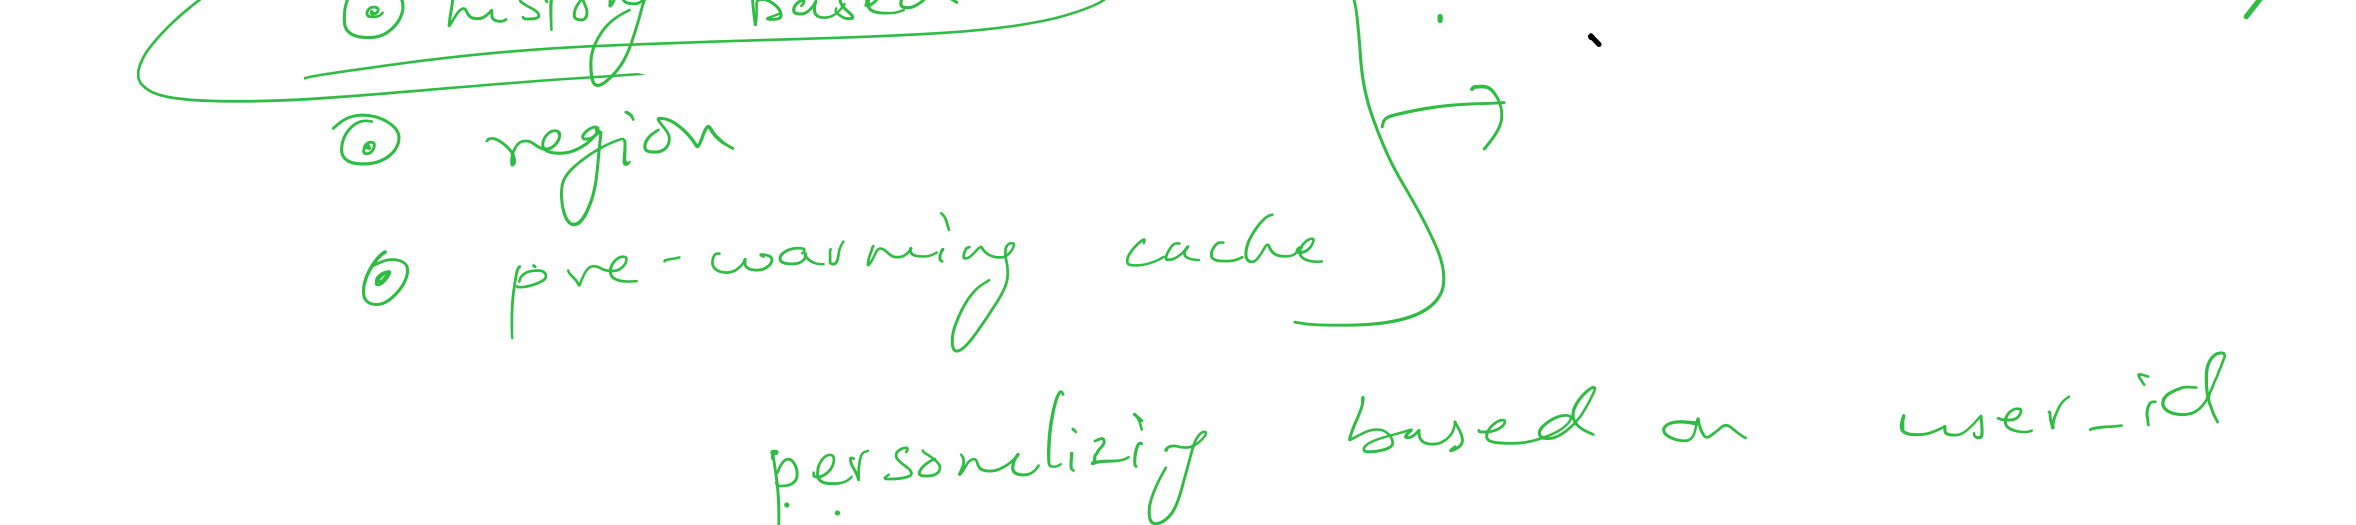
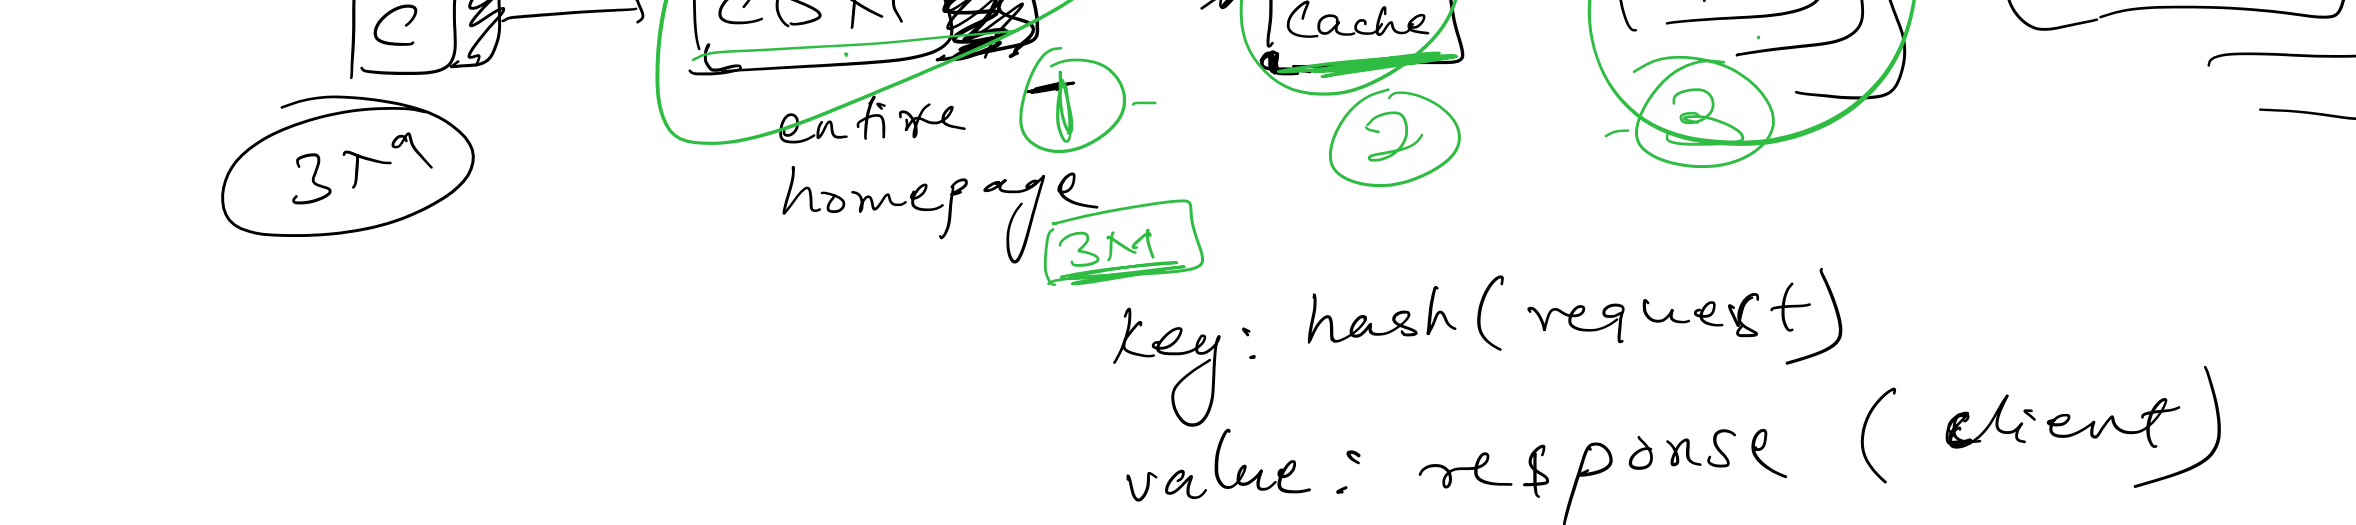




⑤ Bottlenecks

① Throughput

we set up m-s replicatn.

① Client Cache:

- Editorial: 1 week
 - Personalised: 1 day
- expiry: 5 mins



② history based

- ① region
- ② pre-warming cache

personalizing based on user_id

① Clustering algo on users

↳ cluster-id ()
 ↳ 100 : 3M
 30K
 100 servers

② 10sec → user % 10 = 5 → LBS

adding artificial delay at the client - client sharding

P + E

- ① → E (CDN) 1 week
- ② → P (cluster-id) 1 day
- ③ → CW (No-cache) (LB + CDN) scale app-server.

(hotstar.com) DNS resolution

