

System Design Interview

- main reqs of system
- reduce scope of problem
- get interviewer buy-in
- consistency vs availability
- tradeoffs

- different levels
- hashing algorithms
- eviction policy (LRU, etc)
- cache hit / miss
- session persistence

- Relational (ACID, normalization, indexing)*
- NoSQL (KV store, Blob, Timeseries, Graph)
- Horizontal scaling vs sharding
- Write Master + Read Replicas
- Multi-tenancy vs Single-tenancy
- MapReduce (distributed querying)
- ETL pipeline (extract, transform, load)

- types / roles, demographics
- current quantity & future-proofing
- peaks (times of day, days of year)
- frequency, avg duration
- web vs mobile
- accessibility / needs

- Redundancy, Five 9s
- Fault tolerance, circuit breakers
- Leader election (consensus algo)
- Container orchestration
- Side cars / service mesh
- Observability (Logs, Metrics, Tracing, Alerts)

- regions, CDNs
- distance, RTT
- network latency
- DNS lookups

- TCP vs UDP
- Queues (FIFO, in-order)
- Pull vs Push
- Polling vs Streaming
- Publish-Subscribe
- Idempotency

- Reverse proxy (request sorting)
- Forward proxy (hacking, VPN)
- Load Balancers (Round robin, Least Outstanding Requests, path-based; level 4 vs level 7)

- SOAP, REST, gRPC, GraphQL*
- (Tier-based) rate limiting
- Attacks: DoS, DDoS, XSS, MITM*
- Authentication vs authorization, JWT
- Symmetric vs Asymmetric Encryption
- TLS/HTTPS, SSL certificates
- CRUD, HTTP methods, pagination
- Static vs Dynamic Configuration

- *ACRONYMS**
- ACID: atomicity, consistency, isolation, durability
 - SOAP: simple object access protocol
 - REST: representational state transfer
 - DoS: denial of service; DDoS: distributed ...
 - XSS: cross-site scripting (form injection)
 - MITM: man in the middle (interception)

- CPU, RAM, Storage
- Vertical scaling
- HDD vs SSD
- Requests/sec, QPS
- Parallelization / threads
- VMs, partitions

