

Long-distance AXFR versus Database Replication

Libor Peltan

• 7. 10. 2025

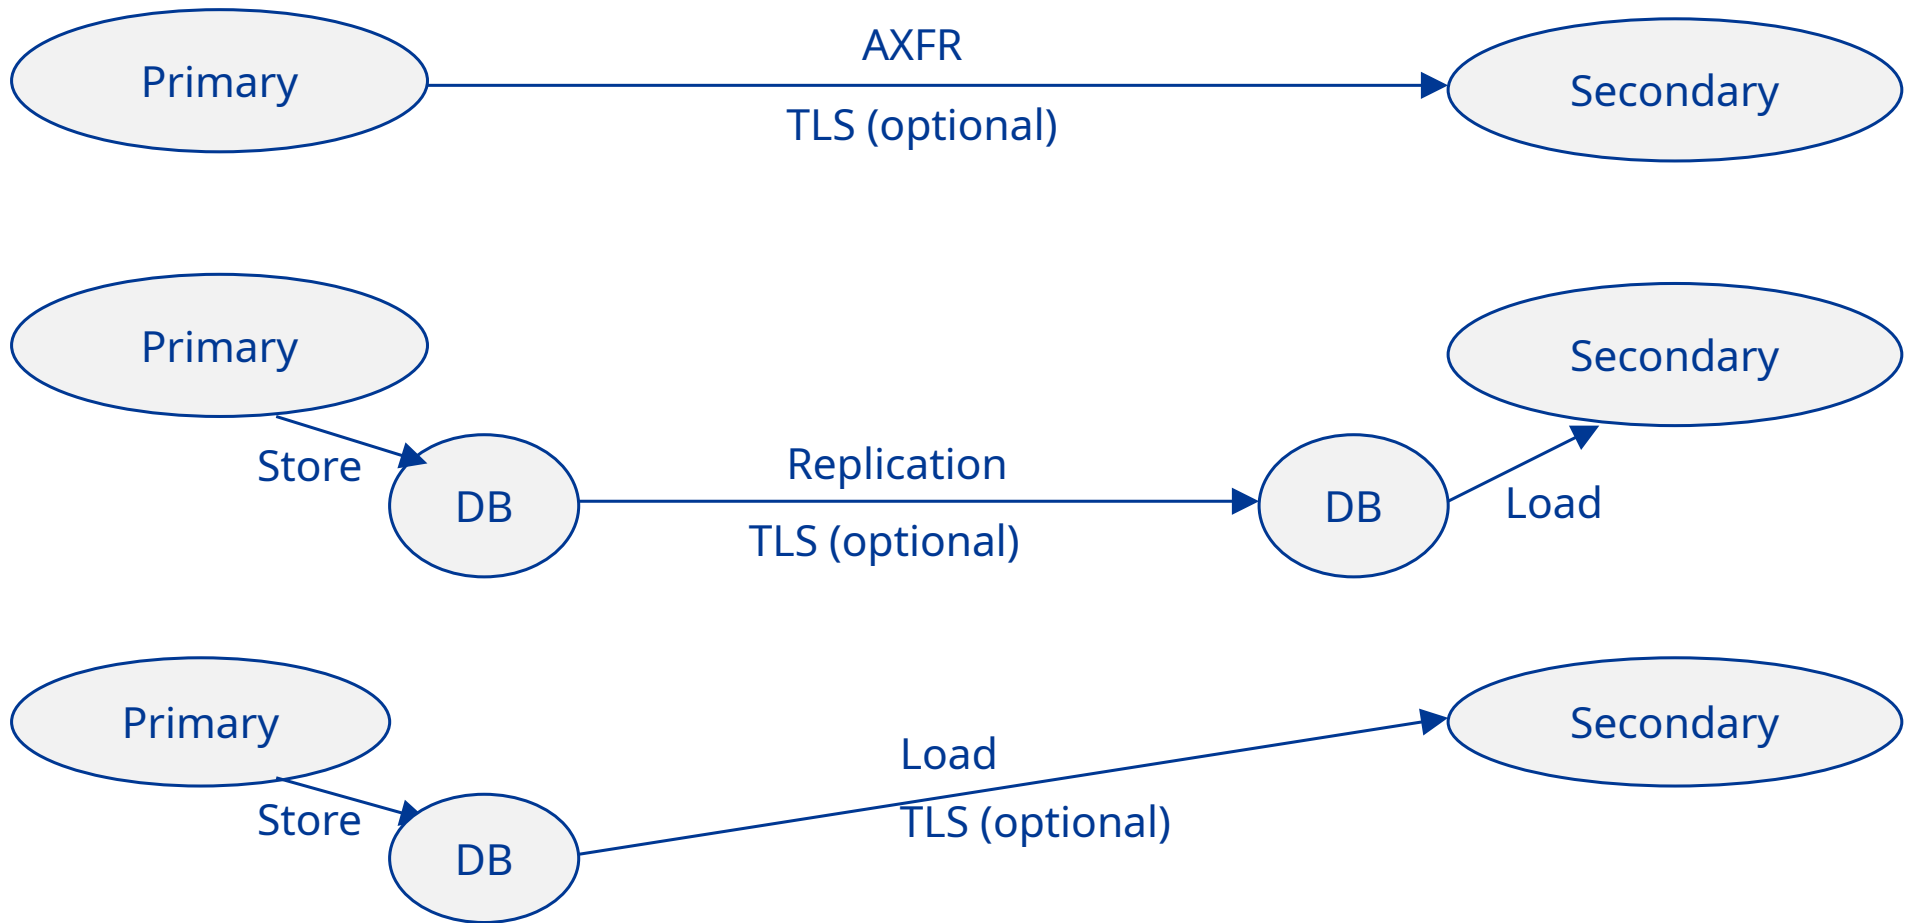
World-wide Anycast

- Transfer zones to distant locations
 - Latency + packet loss
- Huge zone
 - Strive for IXFR, prepare for AXFR
- Many zones
 - Also many conns?

DB backend for Auth DNS

- Zone contents in Database
- One DNS query → one DB request?
- Knot DNS+Redis:
 - load from DB, answer from RAM

Setups



Cases

- Full AXFR of se. zone (~8M indiv. RRs)
 - IXFR usual but not interesting
- AXFR of 100'000 tiny zones (3 RRs each)
 - Reusing TCP conns
 - TLS: not reusing TCP, but session
 - 50 worker threads
- Your setup combines both effects

Test Environment

- Distance only Simulated

```
tc qdisc add dev ens25f1np1 root netem \  
    rate 20mbit \  
    delay 125ms 50ms \  
    loss 0.1% limit 5000
```

- Both directions

- Ping 250 ± 100 ms
 - Loss 0.2%

Results Big Zone

- Big zone, full transfer:
 - AXFR: 27 min 45 s
 - With TLS: 28 min 12 s (+2%)
 - Redis replica: 32 min 12 s (+16%)
 - With TLS: 36 min 10 s (+12% ex)
 - Redis remote load: 27 min 14 s
 - With TLS: 27 min 58 s (+3%)

Results Many Zones

- 100'000 tiny zones
 - AXFR: 20 min 59 s
 - With TLS: 72 min 17 s (3,4x slow)
 - Redis replica: 0 min 30 s (42x faster!)
 - With TLS: 0 min 30 s (equal)
 - Redis remote load: 31 min 07 s (+48%)
 - With TLS: 31 min 09 s (equal)

Analysis 1/2

- Big zone: depends on data volume
 - Redis replica overhead 15-20%
 - TLS overhead small
- Many zones + Redis replica
 - Same effect => fast

`time ~ volume / speed`

Analysis 2/2

- Many zones + 50 workers + RTTs
 - 1-3 RTT per zone
 - Each worker taking zones in seq
- 50 workers only
 - Free to increase far above CPU#
 - 2x increase → ~2x speedup (AXFR)

$\text{time} \sim \# \text{zones} * \# \text{RRTs_per_zone} * \text{RTT} / \# \text{wrks}$

Overview

- Big zone AXFR: time depends on volume
 - Database overhead ~+15%
- Many zones AXFR: DB helps compacting
 - Sync magnitudes faster
 - TLS no penalty
- Other than Knot DNS + Redis -- ???

Thank you for attention