Simplified Registry Replication and Its Effect on Schema Replication

EDG Heidelberg Conference 26/09/03 – 01/10/03 ptaylor@uk.ibm.com

Replication Topology

- Registry replication uses a distributed ‘gossip’ mechanism
- Schema replication uses a Master/Slave mechanism
  - Longer periods of time between replications, due to less load on Schema
- One schema per registry (with only one Master schema per VO)
- Queries aimed at ‘local’ registries and schemas

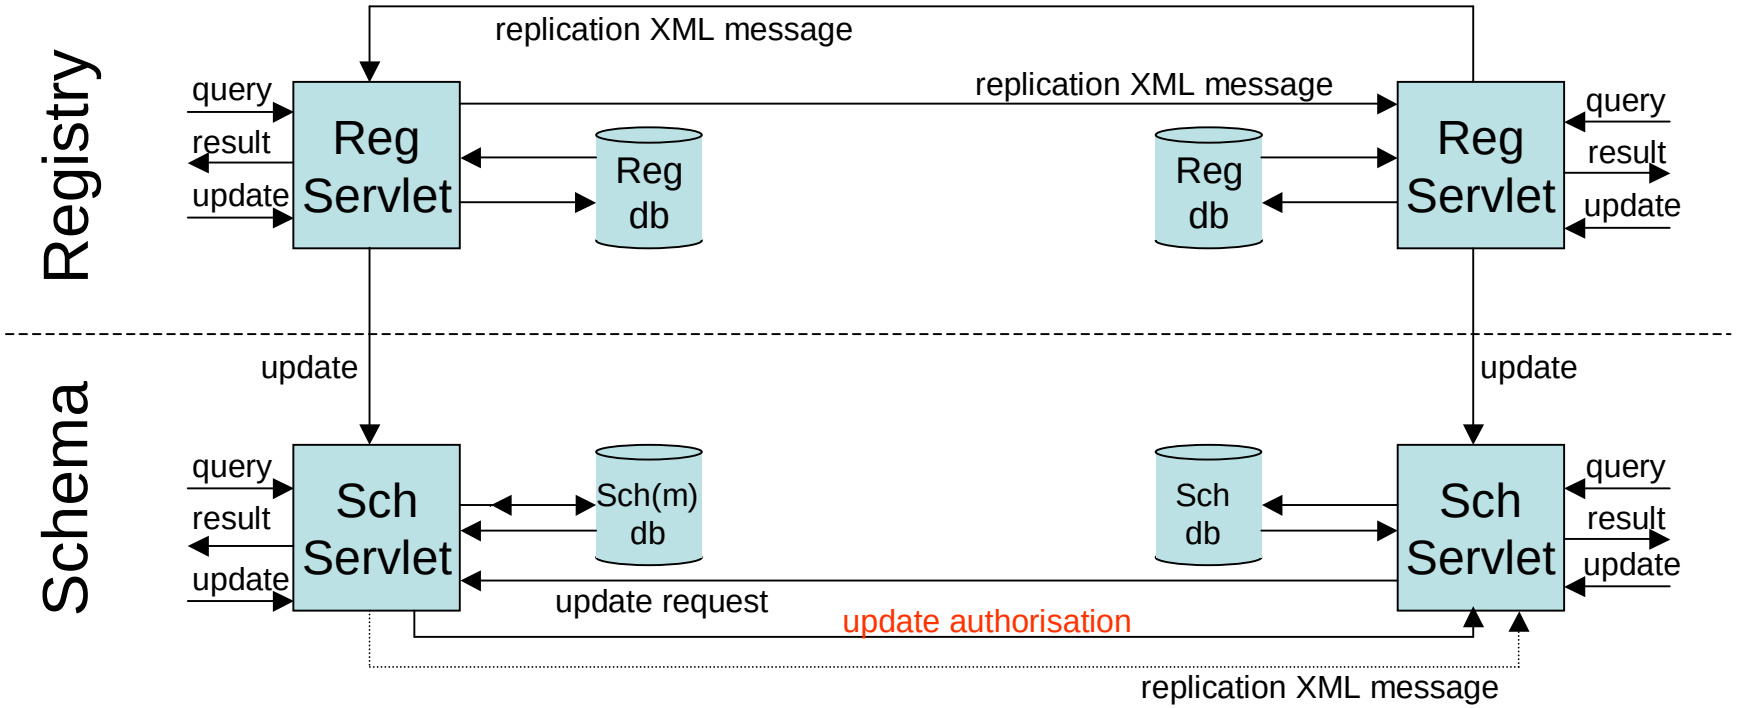

Registry Replication Components

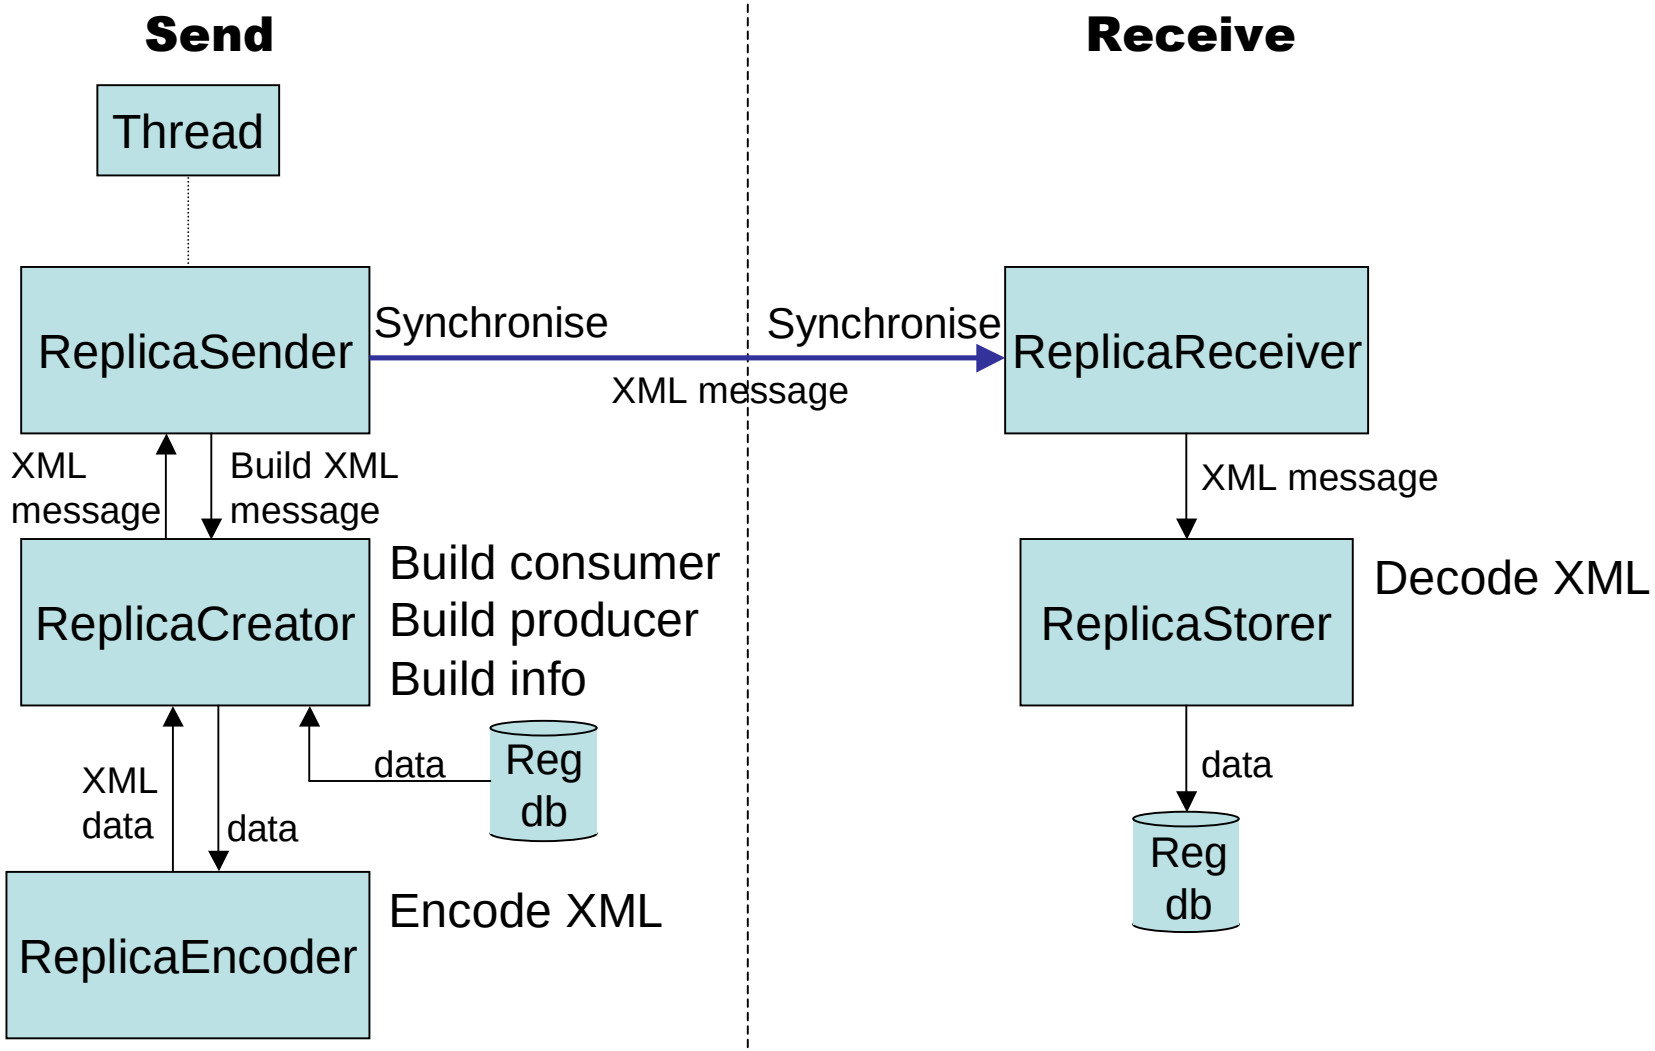

Schema Replication

- Initially, Master schema chosen from ordered list in schemaconfig.xml
  - Later, to be replaced by more elegant scheme
    - Possibly using ordering on lowest IP number
- All queries are dealt with locally
- All Schema updates are first fed to the Master for approval
  - On approval, the Master contacts remaining schemas with update
  - Write-through cache used for efficient update at source of query
- If Master dies, then next ‘live’ schema in the ordered list is new master
- For each update request, a list of ‘live’ schemas is built, using the full ordered list in schemaconfig.xml
  - First found ‘live’ schema in created list is the Master for this replication cycle
  - A restarted Master schema reasserts its Master status by being found again in the ordered list, when the ‘live’ schema list is created

Some Food for thought

- Deletion of tables
  - Timeouts
  - Propagation of deletes (related to registry entries)
- Consistent state on restart
  - If a master restarts, then it must copy the previous master's state
    - Only then can the restarted master take over the role of Master
- Merging subnets
  - Handle duplicate name clashes between schemas in the subnets to be merged
    - Could table names be appended with Producer ids for uniqueness?
      - If so, then handle Producer trying to publish to a 'changed' table name
      - Might need an internal translation table to keep track of changes
- Security
  - Validation of registries and schemas during replication
    - Do we know and/or trust who we're sending to?
  - Adding a new registry or schema to the VO
    - Is this new registry or schema known and/or trusted?

End of Presentation