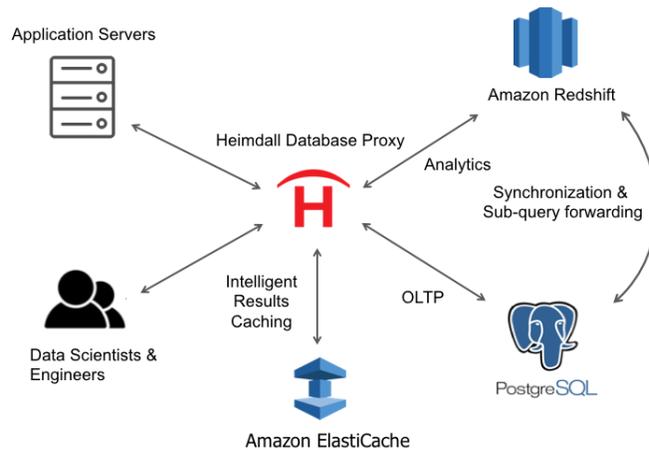


OLTP Performance at Data Warehouse Scale



Companies deploy Amazon Redshift for large-scale MPP (Massively Parallel Processing) data analytics. These systems apply parallel compute resources to answer queries quickly. However, there are still performance challenges:

- Higher latency: Distributed queries means that many nodes have to coordinate to generate an answer
- No materialized views: If a query result is used for multiple following queries, additional work is required to preserve the result or repeated calls to the same base query will be required

These challenges can be solved by modifying the application. However, the ideal solution requires no code changes. Heimdall Data is that solution providing Amazon Redshift users two benefits:

- Management of both Analytics and OLTP traffic for optimal performance
- Fast materialized views for results from Amazon Redshift

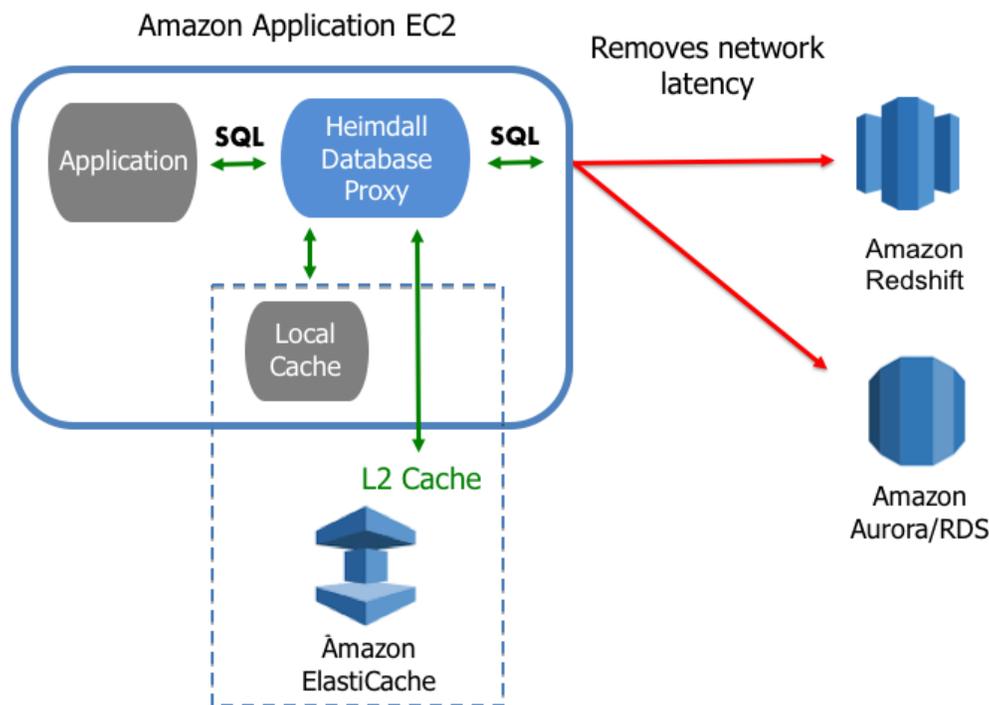
Fast materialized views are very important in analytics environments. When reports are generated, a subset of data is pulled from the back-end data store, then various operations are performed on that data. Heimdall provides the following functionality:

- Queries against a materialized view can be routed to an alternate database, typically Postgres, which acts on behalf of Amazon Redshift. Postgres answers queries offloading Amazon Redshift.
- Heimdall triggers a refresh of the view automatically. Heimdall is aware updated views from Amazon Redshift and when data was loaded that may impact the view. The net result is faster reports and a lighter load on Redshift, allowing the processing of other queries to be faster and more scalable.

OLTP Performance at Data Warehouse Scale



SQL Results Caching: The fastest query is the one that does not have to be executed. Heimdall's intelligent auto-caching and auto-invalidation works together with Amazon Redshift's query caching, but in the EC2 application tier, removing network latency. This distributed architecture allows caching to be scalable, while acting as one cache cluster. Result sets are cached in tandem from local memory and Amazon ElastiCache and are invalidated upon writes to the table. Best of all, Heimdall deployment requires zero code changes.



The Heimdall architecture was designed for ease of deployment without the need to modify the application or database. Configuration changes are updated at runtime without restarting the application. The net result is a platform that can be updated with 100% application uptime.

About Heimdall Data:

Heimdall Data is an AWS Advanced Technology partner providing SQL visibility and performance optimization as a Database Proxy. Available in the AWS Marketplace, Heimdall is transparently integrated with Amazon Redshift, Aurora, RDS, ElastiCache and CloudWatch.

For more information about how Heimdall Data can help improve your AWS environment, contact us at info@heimdalldata.com.