

Technical Audit for Vertica Analytics Platform at Riot Games



Background

Riot Games was established in 2006 by a couple of entrepreneurial gamers who believe player-focused game development can result in awesome games. They released their debut title, League of Legends, to critical and player acclaim in 2009. It has over 32 million plays every month.

The Estuate technical audit was undertaken to evaluate the architecture and configuration of the Vertica analytics platform and provide guidance on

technical best practices. Riot Games was out of compliance with best practices and Estuate provided advice in many areas. Non-compliant audit points were addressed in detail and the necessary recommendations were provided.

The Opportunity

- ⦿ Evaluation of the Vertica multi-node cluster by understanding their current architecture, database design, operational profile and their network architecture.
- ⦿ Perform Vertica health check for better performance by doing system audits.
- ⦿ Collect diagnostics and perform detailed analysis of audit groups.

The Estuate Solution

- ⦿ Based on the audit check best recommendations and suggestions were provided for following areas.
- ⦿ Physical and logical design.
- ⦿ Vertica configuration.
- ⦿ Maintenance and monitoring.

- ⦿ Hardware and OS configurations.
- ⦿ Data load and data retrieval mechanism.
- ⦿ Delete/Purge/Archive strategy.
- ⦿ Audit groups and calculation.
- ⦿ Provided list of all audit points and against which system it was evaluated.
- ⦿ Recommendations for each audit point was properly explained, prioritized and well documented.

Business Impact

Riot Games' issues with the current environment were addressed and necessary action items identified in detail. Below are a few of the issues, their impact and our solution.

Sl.No	Issues/Concerns	Impact	Solution
1	Sort order had high cardinality columns on projection fact tables	Having high cardinality column in projection sort order significantly impacts query performance and disk space utilization.	Ran DBD comprehensive design by providing logical schema, sample data and sample queries as input. Redesigned projections were deployed which improved the performance.
2	Dimension table projections should be replicated	If dimension tables are segmented then it might create a network join which might consume more bandwidth.	Dimension tables were replicated which created local joins. Made sure that there is no data flow between the nodes.
3	Partition granularity	It can significantly impact query performance.	Partition strategy was followed which made sure it doesn't exceed ROS limits.
4	Workload analyzer	Impacts workloads on cluster and SQL queries.	Provided recommendations which were implemented, which resulted in SQL tuning and managing workload history.
5	Workload segmentation through Resource pools	Only one custom resource pool was defined.	Multiple resource pools were defined and tuned which yielded optimal performance.
6	OS Configuration	Some of the OS parameters were not set as per Vertica recommendations. This can have negative impact on Vertica cluster.	OS configuration parameters were set like max number of file handles, max threads or processes per user, max file size etc.
7	Logical design	Data was not denormalized for the fact tables. This can create joins between the projections which can slow down the runtime queries.	Denormalized for fact tables which increased the query throughput.
8	Database statistics was not collected for around 88% projections	It causes optimizer to make poor choices resulting in reduce query performance.	Database statistics was gathered and made sure the database is up to date.
9	Disable CPU hyper threading	Hyper threading are extremely compute intensive and it cause more workloads during data retrieval process.	Hyper threading was disabled and made sure that Vertica is not overloaded for query retrieval.

Get Started Today

Feel free to contact us for more info at

sales@estuate.com

call us at

408.400.0680

visit us at:

