




Security Overview: Fully Automated Snowflake Optimization

May 2023




Keebo warehouse optimization security

- Access only usage metadata (76 fields total), never user data
- Run within dedicated instances, separating each customer's queries and models
- Encrypt all communications during transit and at rest (TLS1.2, Google KMS using FIPS 140-2 Level 3)
- Automated monitoring of code with every release, plus regular 3rd-party penetration testing
- All computations carried out inside DB



SOC2 certified



Keebo needs only usage metadata (1 of 2)

From SNOWFLAKE .ACCOUNT_USAGE .QUERY_HISTORY :

DATABASE_ID	BYTES_WRITTEN	QUERY_TEXT (hashed)
DATABASE_NAME	BYTES_WRITTEN_TO_RESULT	OUTBOUND_DATA_TRANSFER_CLOUD
SCHEMA_ID	BYTES_READ_FROM_RESULT	OUTBOUND_DATA_TRANSFER_REGION
SCHEMA_NAME (hashed)	ROWS_PRODUCED	OUTBOUND_DATA_TRANSFER_BYTES
QUERY_TYPE	ROWS_INSERTED	INBOUND_DATA_TRANSFER_CLOUD
SESSION_ID	ROWS_UPDATED	INBOUND_DATA_TRANSFER_REGION
USER_NAME (hashed)	ROWS_DELETED	INBOUND_DATA_TRANSFER_BYTES
ROLE_NAME (hashed)	ROWS_UNLOADED	LIST_EXTERNAL_FILES_TIME
WAREHOUSE_ID	BYTES_DELETED	RELEASE_VERSION
WAREHOUSE_NAME	PARTITIONS_SCANNED	EXTERNAL_FUNCTION_TOTAL_INVOCATIONS
WAREHOUSE_SIZE	PARTITIONS_TOTAL	EXTERNAL_FUNCTION_TOTAL_SENT_ROWS
WAREHOUSE_TYPE	BYTES_SPILLED_TO_LOCAL_STORAGE	EXTERNAL_FUNCTION_TOTAL_RECEIVED_ROWS
CLUSTER_NUMBER	BYTES_SPILLED_TO_REMOTE_STORAGE	EXTERNAL_FUNCTION_TOTAL_SENT_BYTES
QUERY_TAG	BYTES_SENT_OVER_THE_NETWORK	EXTERNAL_FUNCTION_TOTAL_RECEIVED_BYTES
EXECUTION_STATUS	COMPILATION_TIME	IS_CLIENT_GENERATED_STATEMENT
ERROR_CODE	EXECUTION_TIME	
ERROR_MESSAGE	QUEUED_PROVISIONING_TIME	
START_TIME	QUEUED_REPAIR_TIME	
END_TIME	QUEUED_OVERLOAD_TIME	
TOTAL_ELAPSED_TIME	TRANSACTION_BLOCKED_TIME	
BYTES_SCANNED	CREDITS_USED_CLOUD_SERVICES	
PERCENTAGE_SCANNED_FROM_CACHE	QUERY_LOAD_PERCENT	

Keebo needs only usage metadata (2 of 2)

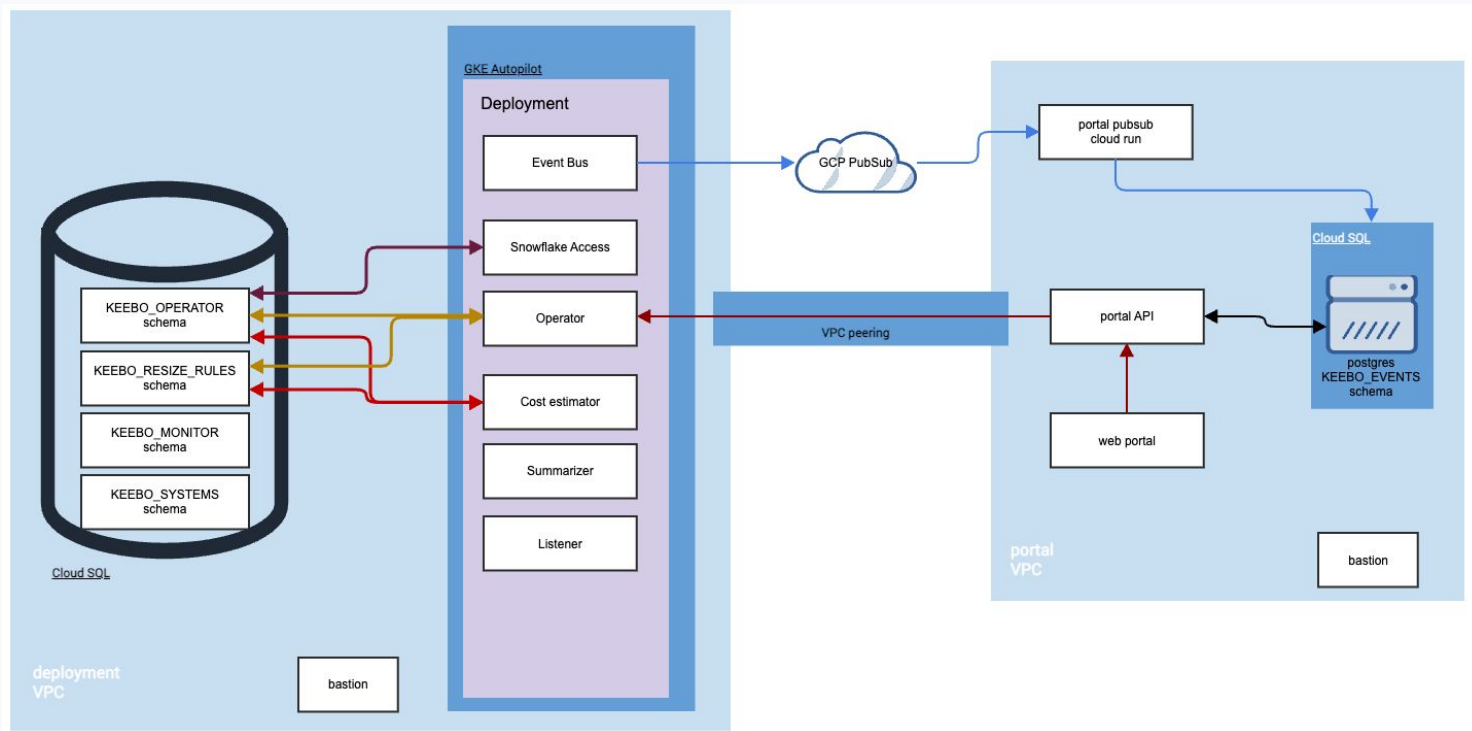
From SNOWFLAKE .ACCOUNT_USAGE .WAREHOUSE_METERING_HISTORY :

CREDITS_USED
CREDITS_USED_CLOUD_SERVICES
CREDITS_USED_COMPUTE
END_TIME
START_TIME
WAREHOUSE_ID
WAREHOUSE_NAME

From SNOWFLAKE .ACCOUNT_USAGE .WAREHOUSE_EVENTS_HISTORY :

CLUSTER_NUMBER
EVENT_NAME
EVENT_REASON
EVENT_STATE
QUERY_ID
ROLE_NAME
TIMESTAMP
USER_NAME
WAREHOUSE_ID
WAREHOUSE_NAME

API and deployment architecture



Network data flow

