

Auditing Queries

Use the audit logs to find out which users have run or tried to run queries and the tables that those queries have accessed or tried to access.

Each analytical query against a cluster is logged. An analytical query is a query against a published project, that is issued from a Thrift endpoint, or that is issued from an XMLA endpoint. Queries from the Design Center for data previews and queries for setting up or testing connections are not audited.

How Log Files Are Managed

The log file, `audit.log`, is located in the path `/opt/atScale/log/engine/` on the AtScale server. It is rotated daily. Old log files are compressed and given the name `audit.yyyy-mm-dd.log.gz`. Ninety days of old log files are retained.

Content Of Log Entries

Each entry in the audit log contains the following key/value pairs.



Note: Due to the removal of environments in AtScale 7.4.0, the Environment ID field is no longer written to the query audit logs.

Key	Value
queryID	A string that identifies the query.
allowed	A Boolean value that indicates whether the user had permission to execute the query. Values: true false
isCanary	A Boolean value that indicates whether the query was a canary query. Values: true false
service or user	The name of the service or the user ID that executed or attempted to execute the query.
ip	The IP address of the client that executed or attempted to execute the query. This name/value pair does not appear for queries that are issued by services.

org_id	The name of the organization under which the query was executed or attempted to be executed.
project_id	The name of the project against which the query was executed or attempted to be executed. This name/value pair might be missing for some queries.
tables_read	The names of the tables that the query accessed or attempted to access. If the query is against a query dataset, the value is the text of the query.

Example Log Entries

These examples are formatted for better readability, with each key/value pair on a separate line. Actual log entries separate key/value pairs with spaces.

Figure 1. Raw query

```
2016-07-29T21:55:28.373Z atscale-query-audit: queryId=e06d6077-a422-4e1e-83f7-ccdb9b9fb9ab allowed=true
isCanary=true user=user_ID ip=/192.168.5.115 orgId=default projectId=1f8ef67a-b237-4ed9-7958-b17ff09e0755
tables_read=database_a.factinternetsales
```

Figure 2. Canary query

```
2016-07-29T21:42:19.949Z atscale-query-audit: queryId=52b5ac09-6d3c-4499-b6ef-a6abca677ff0 allowed=true
isCanary=true user=user_ID ip=/192.168.5.115 orgId=default projectId=1f8ef67a-b237-4ed9-7958-b17ff09e0755
tables_read=database_a.dimgender,database_a.dimcustomer,database_a.factinternetsales
```

Figure 3. Use of aggregates

Use of a system aggregate

```
2016-08-01T03:27:26.874Z atscale-query-audit: queryId=6b4fbe97-9d06-46a1-a64c-5d803a88b100 allowed=true
isCanary=false user=user_ID ip=/192.168.99.1 orgId=default projectId=demo
tables_read=as_adventure.as_agg_37b34995_none
```

Use of a user-defined aggregate

```
2016-07-29T21:42:21.201Z atscale-query-audit: queryId=52b5ac09-6d3c-4499-b6ef-a6abca677ff0 allowed=true
isCanary=false user=user_ID ip=/192.168.5.115 orgId=default projectId=1f8ef67a-b237-4ed9-7958-b17ff09e0755
tables_read=as_adventure.as_agg_2c479178_uda_hdp2sec
```

Figure 4. Drill-through

```
2016-08-01T03:28:02.380Z atscale-query-audit: queryId=e7870e98-3608-4789-88a5-9091c97e5cb1 allowed=true
isCanary=false user=user_ID ip=/192.168.99.1 orgId=default projectId=demo
tables_read=as_adventure.dim_geo_state,as_adventure.dim_geo_city,as_adventure.dimdate,
as_adventure.as_agg_ff188f43_clr,as_adventure.dim_geo_postalcode,as_adventure.dimcustomer 2016-08-
01T03:28:17.433Z atscale-query-audit: queryId=181561db-f53f-44dd-ac1b-7b63ba4c69c3 allowed=true isCanary=false
user=ouser_ID ip=/192.168.99.1 orgId=default projectId=demo
tables_read=as_adventure.as_agg_06dfc994_clr_sz_stl,as_adventure.dimdate,as_adventure.dimproduct,
as_adventure.dimcustomer
```

Figure 5. Query dataset (two fact datasets and one dimension)

```
2016-07-30T22:42:28.043Z atscale-query-audit: queryId=a20b5eac-23cb-4392-bb7b-1b6642c57045 allowed=true
isCanary=false user=user_ID ip=/192.168.99.1 orgId=default projectId=975b5a31-acef-40a9-4466-7e3fbd32beb9
tables_read="select \* from as_adventure.sales_log",as_adventure.factinternetsales,as_adventure.customer_file
```

Figure 6. Use of query datasets

Query dataset with delegated authorization and impersonation

```
2016-07-31T22:57:01.726Z atscale-query-audit: queryId=72128767-8e67-4f5b-90e9-7c60bdd75472 allowed=true
isCanary=true user=user_ID ip=/192.168.5.5 orgId=default projectId=02f8c203-27c2-4449-7e68-bc04a5cb35d8
tables_read="select \* from as_adventure.factinternetsales"
```

Canary query, then use of an aggregate table with delegated authorization and impersonation

```
2016-07-29T21:52:31.470Z atscale-query-audit: queryId=506c35e1-85b5-4507-9f82-fd15d22bf8cd allowed=true
isCanary=true user=user_ID ip=/192.168.5.115 orgId=default projectId=1f8ef67a-b237-4ed9-7958-b17ff09e0755
tables_read=database_a.factinternetsales 2016-07-29T21:52:32.411Z atscale-query-audit: queryId=506c35e1-85b5-
4507-9f82-fd15d22bf8cd allowed=true isCanary=false user=user_ID ip=/192.168.5.115 orgId=default
projectId=1f8ef67a-b237-4ed9-7958-b17ff09e0755 tables_read=as_adventure.as_agg_06ddb2d1_none
```

Figure 7. Query from the AggregationService

```
2016-08-01T03:34:03.450Z atscale-query-audit: queryId=a7a20dd7-1527-493c-974c-29e7fc3c738a allowed=true
isCanary=false service=AggregationService orgId=default projectId=demo tables_read=as_adventure.dimproduct
```

Figure 8. Query from the StatsService

```
2016-08-01T03:33:59.801Z atscale-query-audit: queryId=e562b649-d479-47ce-be55-913cda1974ae allowed=true
isCanary=false service=StatsService orgId=default projectId=demo tables_read=as_adventure.dimproduct
```