

Managing an Apache Hive workload

Date of Publish: 2018-07-12



Contents

Workload management.....	3
Setting up and using a resource plan.....	5
Create a resource plan.....	5
Enable a resource plan.....	6
Activate a resource plan.....	7
Query sys database for plan data.....	7
Disable a resource plan.....	8
Workload management entity data in sys.....	8

Workload management

As administrator, you manage workloads by creating resource plans to improve parallel query execution and cluster sharing of queries running on LLAP.

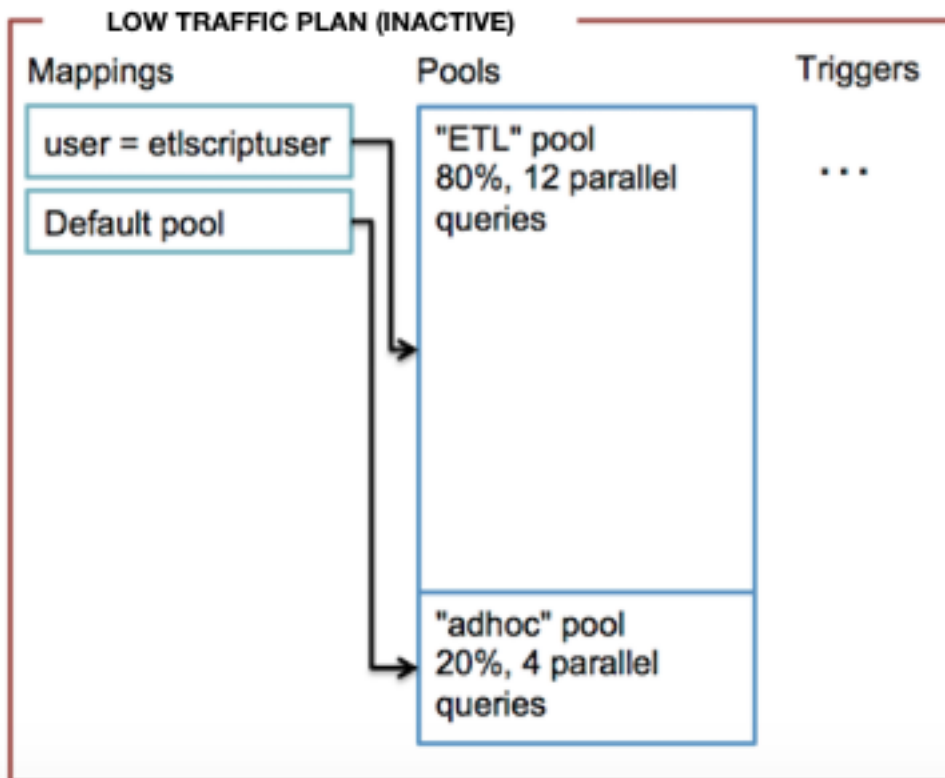
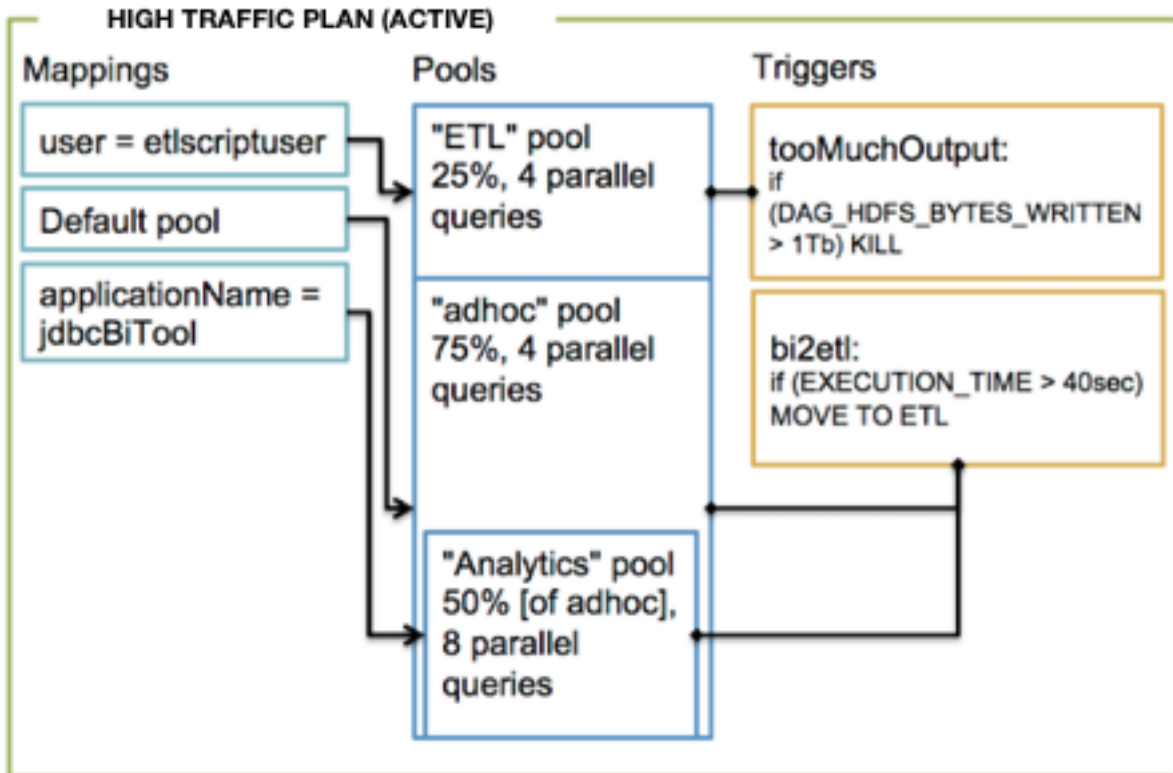
The resource plan is a self-contained, resource-sharing configuration. One resource plan is active on the cluster at a time. Typically, enabling and disabling a resource plan on a live cluster does not affect running queries. You, as administrator, or your script can apply resource plans that configure the cluster for different situations. For example, your script can apply a resource plan that configures the cluster to handle high traffic. When traffic decreases, you can switch the resource plan to support traffic for interactive data visualization, deep ad-hoc analytics, and large-scale BI reporting.

You can create a resource plan to meet a defined data processing benchmark when a workload reaches a high volume of concurrent queries. For example, consider an enterprise that has an ad-hoc analytics application exposed to around 100 analysts. Data set and query patterns dictate that a generated query executes within a few seconds. A resource plan can ensure that when up to 100 users might concurrently use the system, at least 95 percent of queries complete in fewer than 15 seconds.

A resource plan can consist of one or more query pools, mappings, and triggers:

- A query pool shares resources with LLAP cluster processes and queries within the pool, and sets the maximum concurrent queries.
- A mapping routes incoming queries to pools based on specified factors, such as user name, group, or application.
- A trigger initiates an action, such as killing queries in a pool or all queries running in a cluster, based on query metrics represented by Apache Hadoop, Tez, and Hive counters.

The following diagrams depict a simple resource plan. The first diagram shows a resource plan designed for high traffic periods and the second for low traffic periods.



Related Information

[Hive Workload Management Command Summary](#)

Setting up and using a resource plan

As administrator, you can define a resource plan using HiveQL DDL statements that typically include event-triggered resource pools for different users or applications.

About this task

After creating and enabling a resource plan, you must activate it. When you want to alter a resource plan or stop using it, you must disable it.

Related Information

[Hive Workload Management Command Summary](#)

Create a resource plan

As administrator, you can create resource plans that configure your cluster for different conditions, enabling you to improve parallel query execution and share queries on nodes running LLAP. You can allocate resources to a pool of users, groups, or applications and trigger actions in a plan.

Before you begin

You started Hive from the Beeline shell.

About this task

As administrator, you can either create a new resource plan or clone another plan and then modify it.

Procedure

1. Create a resource plan.

- Create a new resource plan (for example, one named `high_concurr_rp`):

```
CREATE RESOURCE PLAN high_concurr_rp;
```

- Clone an existing resource plan (for example, `ha_plan`) and rename it `high_concurr_rp`:

```
CREATE RESOURCE PLAN high_concurr_rp LIKE ha_plan;
```

- Create a new resource plan (for example, named `high_concurr_rp`) and specify a query concurrency limit:

```
CREATE RESOURCE PLAN high_concurr_rp WITH QUERY_PARALLELISM=15;
```

2. Optionally, add one or more query pools to the plan.

```
CREATE POOL high_concurr_rp.bi_pool  
WITH ALLOC_FRACTION=75, QUERY_PARALLELISM=5;
```

```
CREATE POOL high_concurr_rp.etl_pool  
WITH ALLOC_FRACTION=25, QUERY_PARALLELISM=10;
```

This example configuration allocates 75 percent of the LLAP cluster resources to one pool and 25 percent to the other, and allows 5 concurrent queries by users in one pool and 10 concurrent queries by users in the other.

3. Optionally, map an application to a pool.

```
CREATE APPLICATION MAPPING tableau_to_bi IN bi_pool;
```

This example mapping of the Tableau application to the `bi_pool` requires Tableau to connect to HiveServer using a JDBC connection string that includes `applicationName="Tableau"`. In a mapping of a group to a pool, group authorization depends on the HDFS group configuration on the cluster. You can configure Lightweight Directory Protocol (LDAP) and other mechanisms.

- Optionally, create a trigger. For example, create a trigger that moves a query to the `etl_pool` for execution when the execution time exceeds one minute.

```
CREATE TRIGGER high_concurr_rp.slow_query
WHEN execution_time_ms > 60000
DO MOVE TO etl_pool;
```

- If you are running the LLAP service, add a trigger to the pool or resource plan as shown in the following example. `ADD TRIGGER high_concurr_rp.slow_query ADD TO POOL bi_pool;`

You must add a trigger to a pool or resource plan when running under LLAP; otherwise, the trigger is considered global and runs only in Tez.

Related Information

[CREATE MAPPING](#)

[CREATE POOL](#)

[CREATE TRIGGER](#)

[Hive Workload Management Command Summary](#)

Enable a resource plan

You can enable, and optionally activate, a resource plan to manage a workload by running a query.

Before you begin

- You connected to HiveServer over ODBC or JDBC on the cluster you want to manage.
- You started Hive from the Beeline shell.

About this task

Part of the procedure for enabling the resource plan is to validate the plan. Enabling the plan does not apply the plan to a query workload. You cannot modify the plan while enabled. You can enable multiple plans, but you can activate only one at a time on a cluster. Enabling and activating the plan also validates the plan.

Procedure

- Validate the resource plan called `myplan`.
`ALTER RESOURCE PLAN myplan VALIDATE;`
- Enable the resource plan using one of the following methods, depending on whether or not you want to activate the plan now or later and if you already have another plan activated.

- Enable the resource plan:

```
ALTER RESOURCE PLAN myplan ENABLE;
```

- Enable and activate the resource plan:

```
ALTER RESOURCE PLAN myplan ENABLE ACTIVATE;
```

- Archive the active resource plan and replace it with the named resource plan:

```
ALTER RESOURCE PLAN myplan ENABLE ACTIVATE WITH REPLACE
```

Related Information

[Start a Hive shell locally](#)

[Hive Workload Management Command Summary](#)

Activate a resource plan

You can activate an enabled resource plan to apply the plan to a query workload to start managing the workload.

Before you begin

- You connected to HiveServer over ODBC or JDBC on the cluster you want to manage, and started Hive from the Beeline shell.
- You enabled the resource plan.

About this task

You can make only one resource plan active at a time.

Procedure

1. Activate a resource plan.
ALTER RESOURCE PLAN plan1 ACTIVATE;
This example activates plan1.
2. Deactivate one resource plan by activating another.
ALTER RESOURCE PLAN plan2 ACTIVATE;
This example activates plan2 and deactivates plan1, but plan1 is still enabled.

Related Information

[Start a Hive shell locally](#)

[Hive Workload Management Command Summary](#)

Query sys database for plan data

You query the sys database in the Hive Metastore to retrieve information about workload management entities, such as a resource plan.

Before you begin

- You established a connection to HiveServer on the cluster you want to manage, and you started Hive from the Beeline shell, for example.
- A resource plan for managing a workload on your cluster was created.

About this task

Assuming you created the high_concurr_rp resource plan as described earlier, query the sys database for information about the plan.

Procedure

1. Get views of the query pools for the resource plan, for example.

```
SELECT * FROM SYS.WM_POOLS WHERE RP_NAME = 'high_concurr_rp';
```

2. Get information about a trigger named `slow_query` in the pool named `etl_pool` pool of the plan.

```
SELECT TRIGGER_NAME.slow_query FROM SYS.WM_POOLS_TO_TRIGGERS WHERE RP_NAME = 'high_concurr_rp';
```

Disable a resource plan

You can disable a resource plan to modify it or to stop managing a workload by running a query.

Before you begin

- You connected to HiveServer over ODBC or JDBC on the cluster you want to manage.
- You started Hive from the Beeline shell.

Procedure

Disable the resource plan.

```
ALTER RESOURCE PLAN myplan DISABLE;
```

Related Information

[Start a Hive shell locally](#)

[Hive Workload Management Command Summary](#)

Workload management entity data in sys

From the Hive `sys` database, you can get information about workload management and other Hive entities.

Workload Management Views

Hive has a special database named `sys`, which is similar to a system catalog or information catalog. The `sys` database is part of the Hive Metastore. In the `sys` database you can query for views of all Hive entities, including workload management entities. The following workload management views are available:

- `SYS.WM_RESOURCEPLANS` (NAME string, STATUS string, QUERY_PARALLELISM int, DEFAULT_POOL_PATH string)
- `SYS.WM_POOLS` (RP_NAME string, PATH string, ALLOC_FRACTION double, QUERY_PARALLELISM int, SCHEDULING_POLICY string)
- `SYS.WM_MAPPINGS` (RP_NAME string, ENTITY_TYPE string, ENTITY_NAME string, POOL_PATH string, ORDERING int)
- `SYS.WM_TRIGGERS` (RP_NAME string, NAME string, TRIGGER_EXPRESSION string, ACTION_EXPRESSION string)
- `SYS.WM_POOLS_TO_TRIGGERS` (RP_NAME string, POOL_PATH string, TRIGGER_NAME string)