# Changing the Face of Database Cloud Services with Personalized Service Level Agreements

Jennifer Ortiz, Victor Teixeira de Almeida, Magdalena Balazinska

University of Washington, Computer Science and Engineering
PETROBRAS S.A., Rio de Janerio, RJ, Brazil

**CIDR 2015**

Astronomy

Biology

Geoscience

Oceanography

TPC-H

# Many Data Management & Analytics Systems Available







# Many Systems are Available as Cloud Services


**EMR**
Managed Hadoop Framework

**RDS**
MySQL, Postgres, Oracle, SQL Server, and Amazon Aurora

**Redshift**
Managed Petabyte-Scale Data Warehouse Service


Microsoft®
SQL Azure™

# Cloud Services Today

## Amazon EMR

**Which Hadoop Version?**

**Pig or Hive?**

**How many instances of the service?**

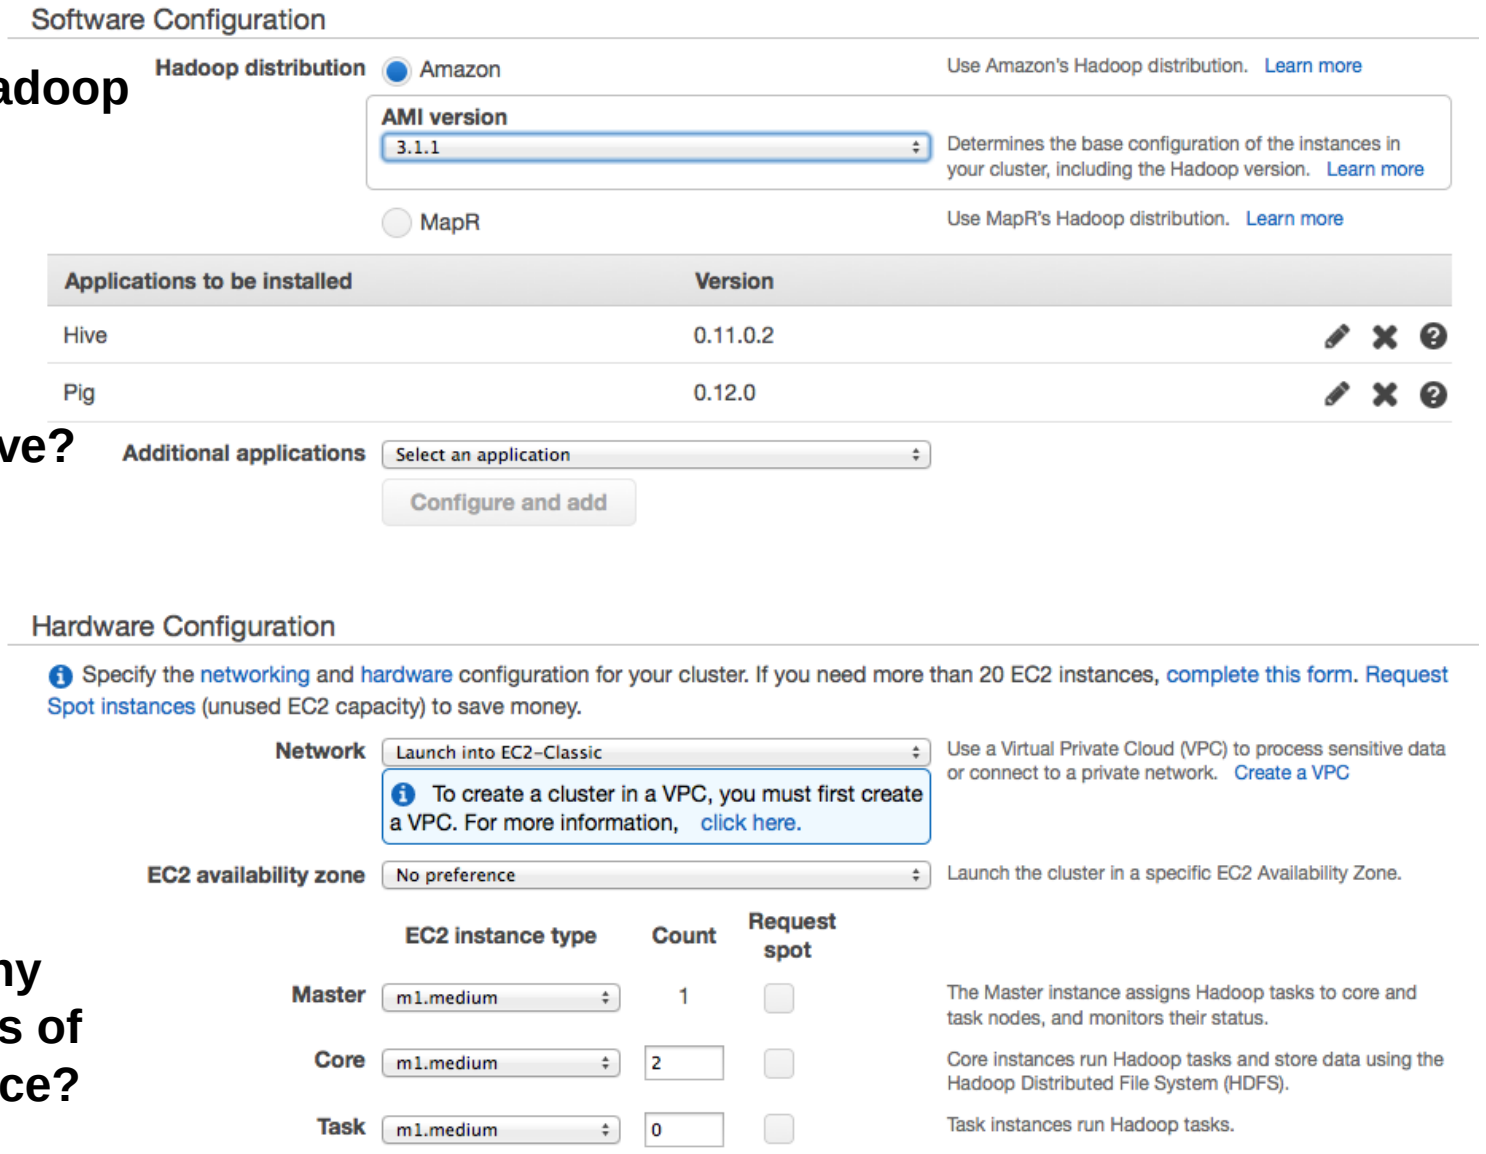

# Cloud Services Today

## Amazon EMR

**Which Hadoop Version?**

**Pig or Hive?**

**How many instances of the service?**

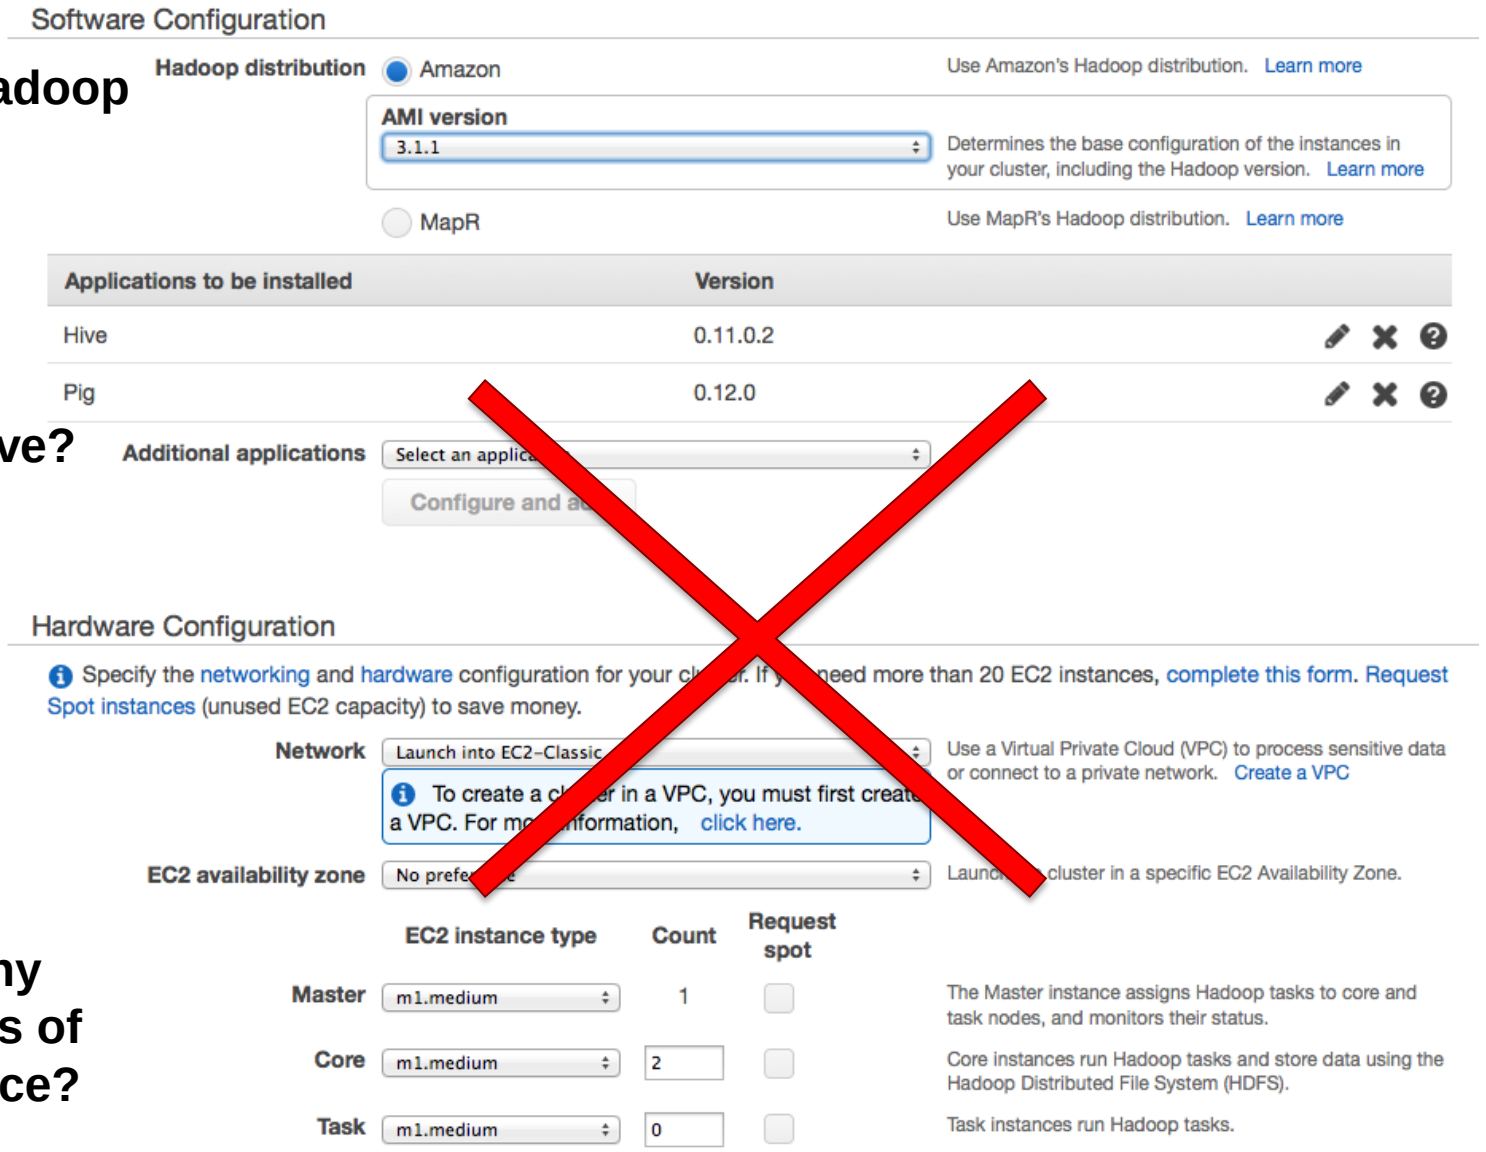

# Cloud Services Today

## BigQuery

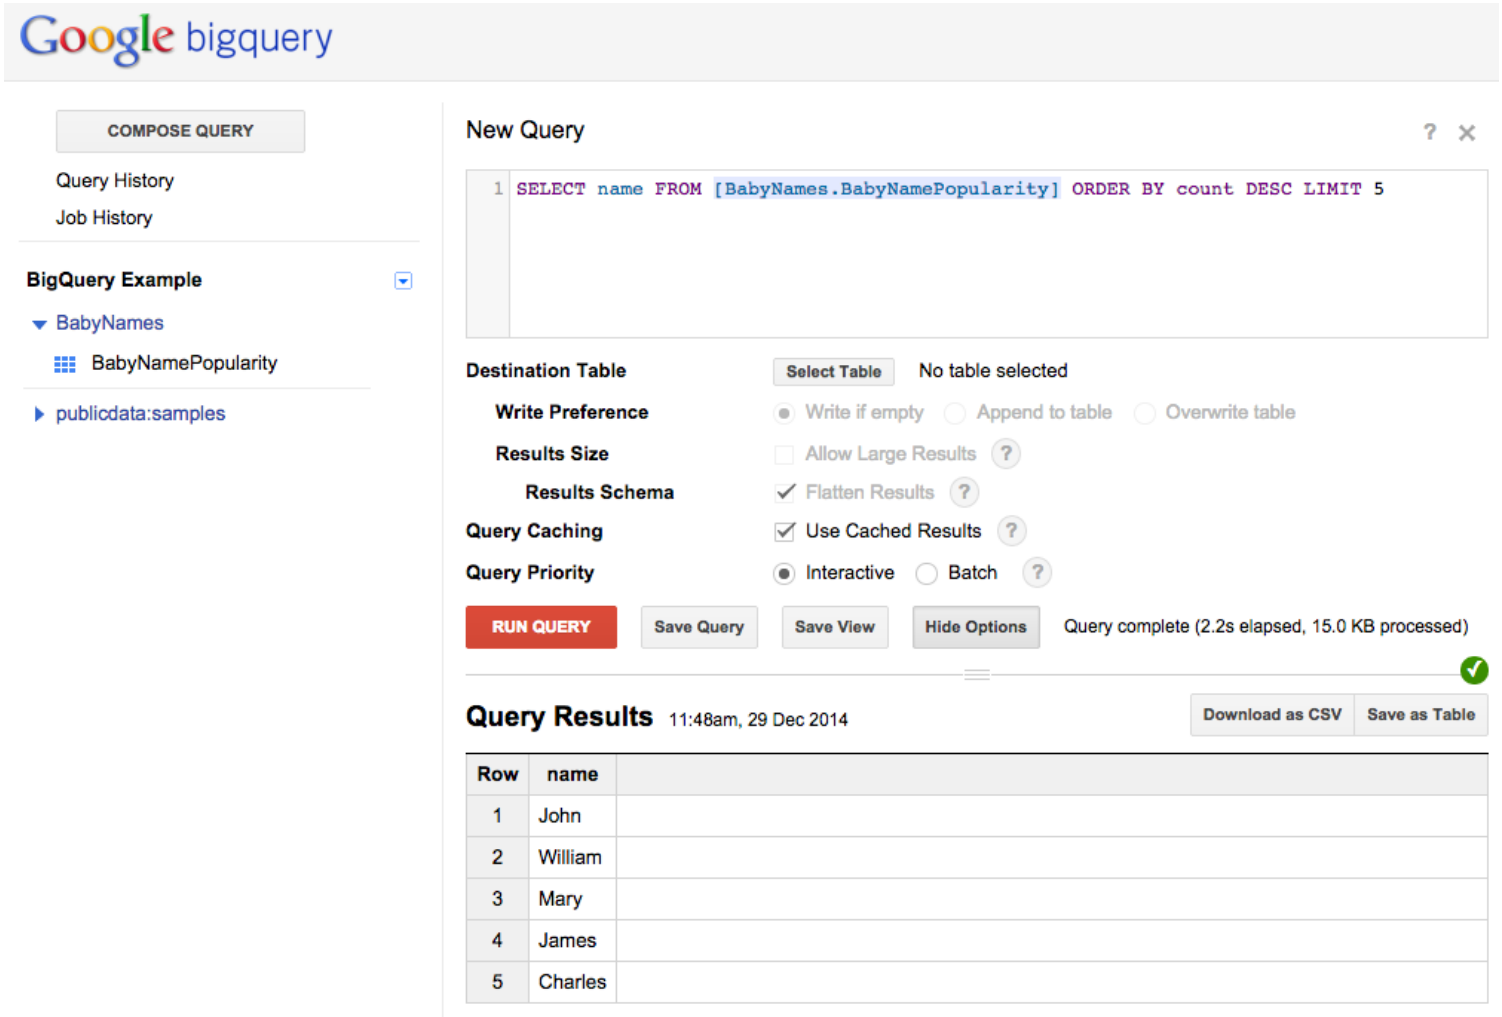


How long will my query take?

# Cloud Services Can Do Better!

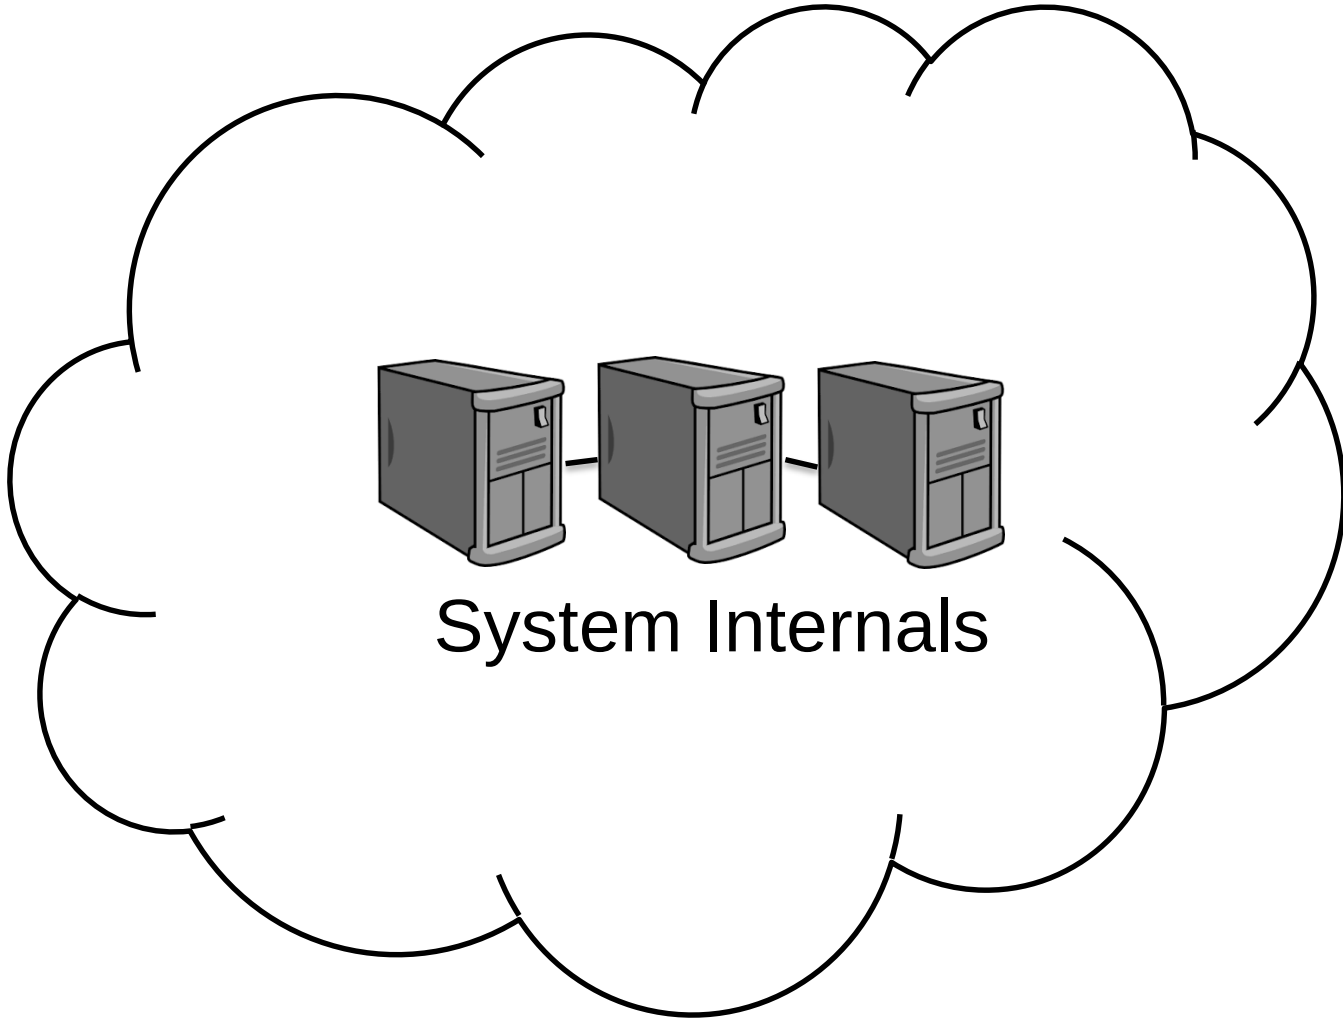


# Cloud Services Can Do Better!

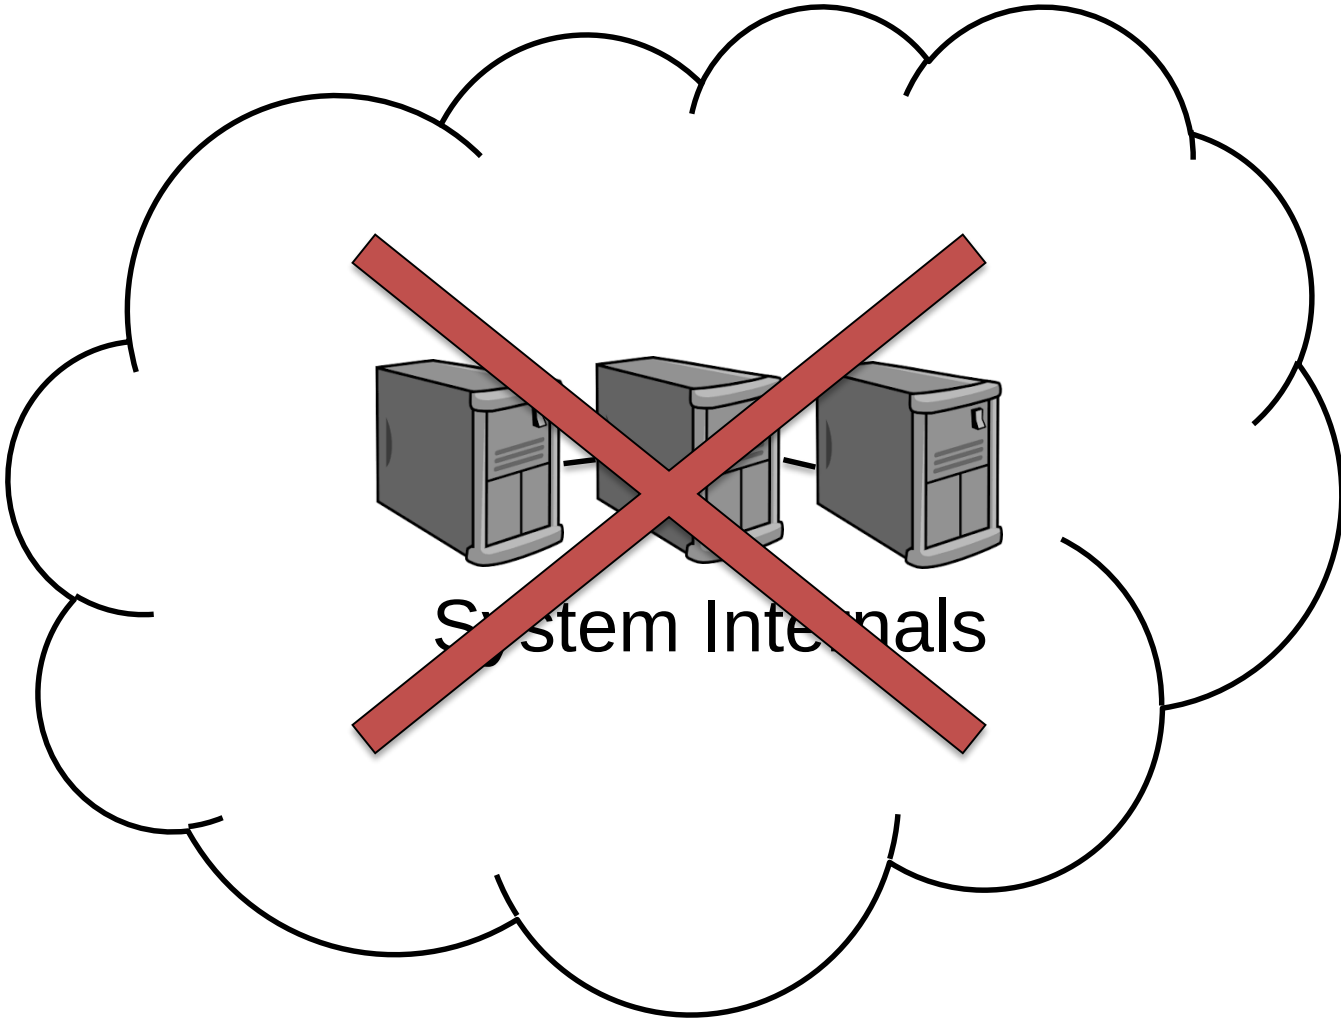


# Cloud Services Can Do Better!

Query:

- Query Capabilities
- Time
- Money


# A new proposal

Time to Re-think the interface…

- Hide details of cluster deployment and resources
- Show users monetary costs and performance estimates on their data
- Let users pick the desired trade-off between options shown

**Personalized Service Level Agreements**

# A PSLA Example

**Tier 1: $0.10/hour**

Fixed, hourly **price**

***Within 20 seconds***:

SELECT <up to 10 attributes>
FROM <Fact | Dimension>
WHERE <up to 100% of data>

Expected **performance**

***Within 1 minute***:

SELECT <up to 5 attributes>
FROM <JOIN Fact + 4 Dimensions>
WHERE <up to 10% of data>

Templates capture **capabilities**

***Within 10 minutes***:

SELECT <up to 10 attributes>
FROM <JOIN Fact + 8 Dimensions>
WHERE <up to 100% of data>

# A PSLA Example

Fixed, hourly **price**

**Tier 1: $0.10/hour**

***Within 20 seconds***:
SELECT <up to 10 attributes>
FROM <Fact | Dimension>
WHERE <up to 100% of data>

Expected **performance**

***Within 1 minute***:
SELECT <up to 5 attributes>
FROM <JOIN Fact + 4 Dimensions>
WHERE <up to 10% of data>

Templates capture **capabilities**

***Within 10 minutes***:
SELECT <up to 10 attributes>
FROM <JOIN Fact + 8 Dimensions>
WHERE <up to 100% of data>

**Tier 2: $0.50/hour**

***Within 1 second***:
SELECT <up to 10 attributes>
FROM <Fact | Dimension>
WHERE <up to 100% of data>

Different **tiers** of service

# Goals

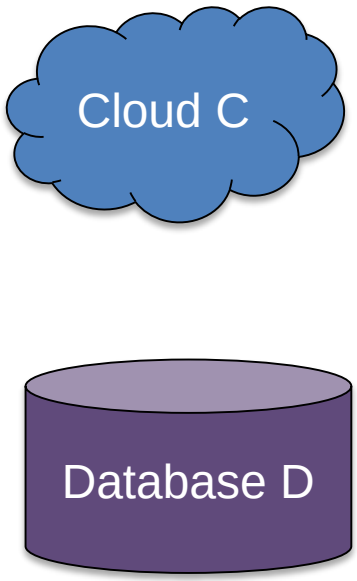

# Goals

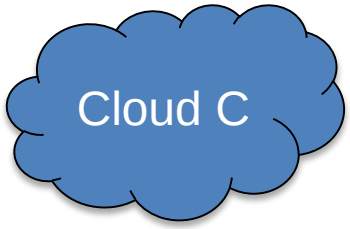


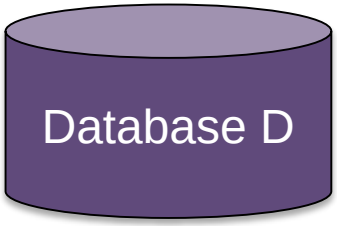



Money,
Time,
Capabilities

# Goals

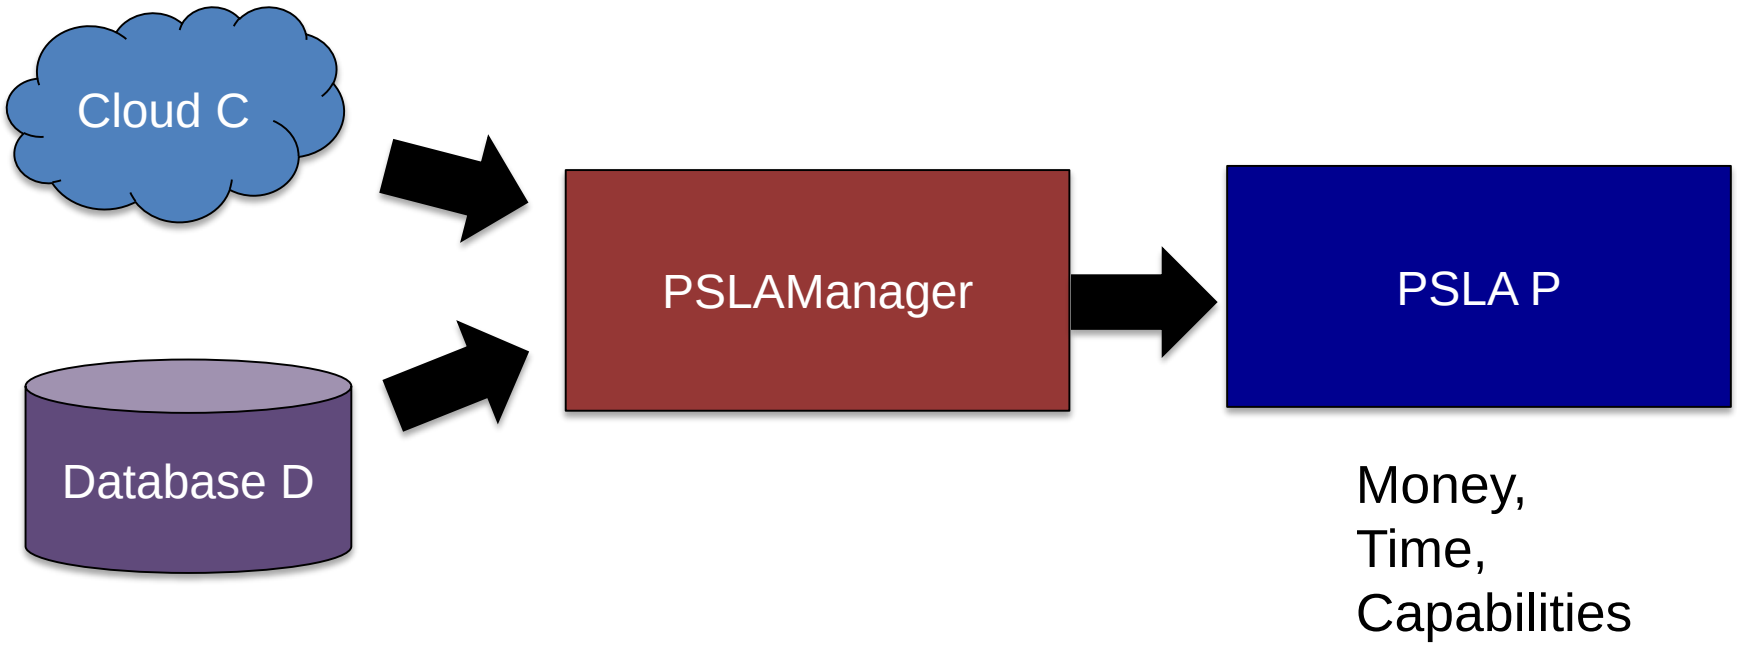
Cloud C
Database D
PSLAManager
PSLA P
Money,
Time,
Capabilities

# Goals

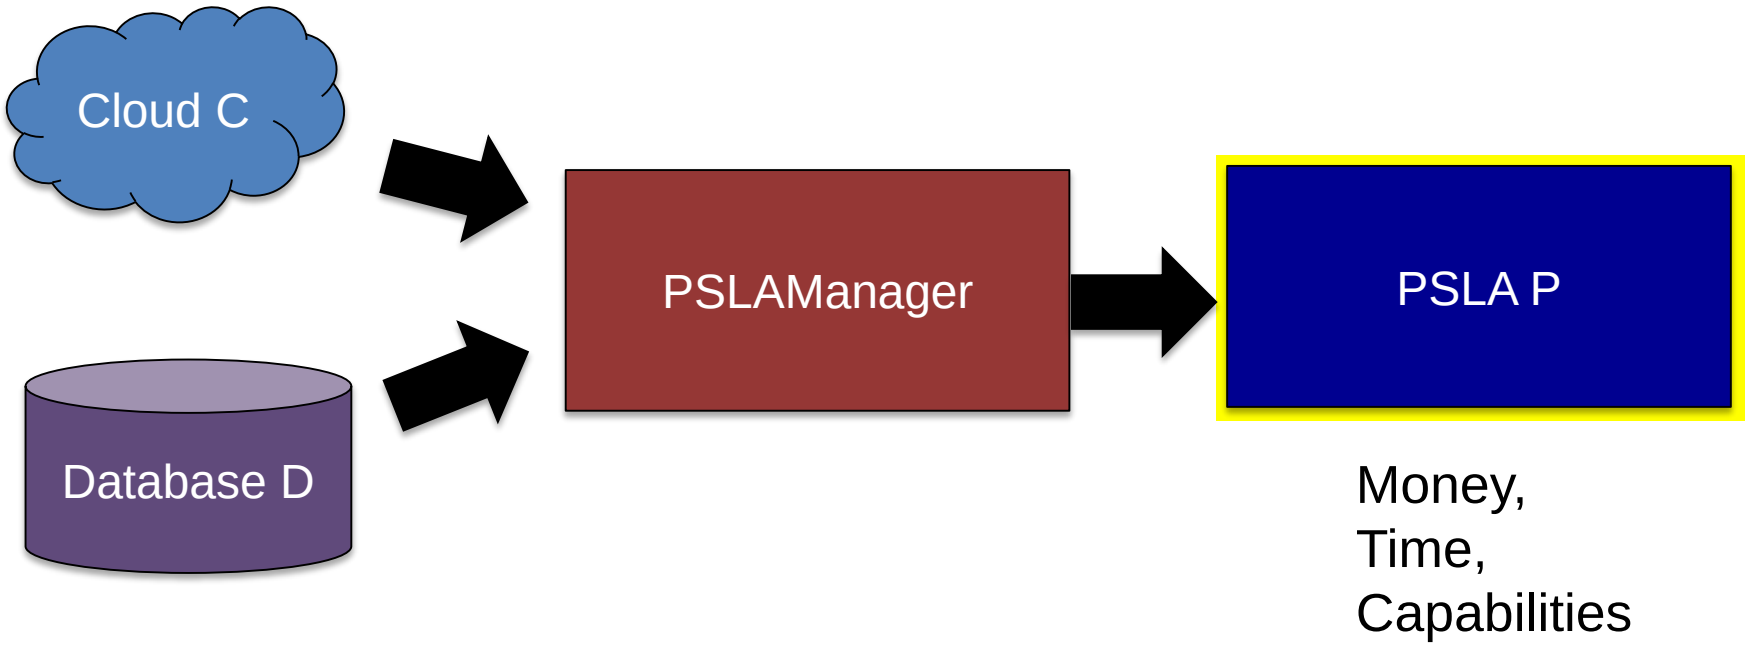
Cloud C
Database D
PSLAManager
PSLA P
Money,
Time,
Capabilities

# Example of a Real PSLA

# TPC-H Star Schema Benchmark

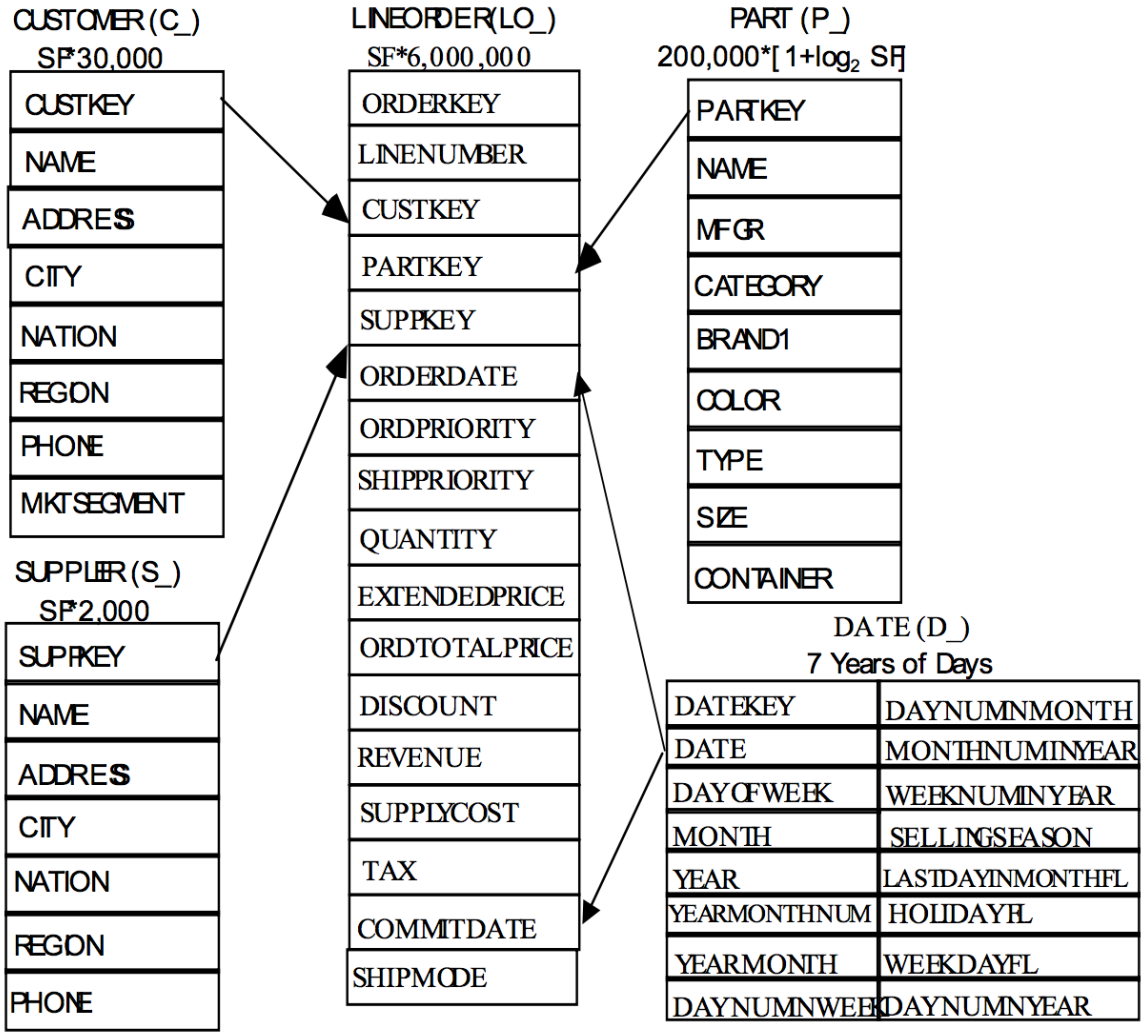


- Based on TPC-H
- 10GB


Myria is a data management service in the cloud that we built at UW.

It has a parallel, shared-nothing back-end query execution engine called MyriaX

# PSLA for Myria

Tier #1

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (1 ATTR.) FROM LINEITEM WHERE 0.1%<br>SELECT (9 ATTR.) FROM PART<br>SELECT (17 ATTR.) FROM DATE<br>SELECT (9 ATTR.) FROM CUSTOMER | 10 |
| SELECT (8 ATTR.) FROM (2 TABLES)<br>SELECT (3 ATTR.) FROM (3 TABLES)<br>SELECT (2 ATTR.) FROM (4 TABLES)<br>SELECT (59 ATTR.) FROM (5 TABLES) WHERE 10%<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 1% | 60 |
| SELECT (43 ATTR.) FROM (4 TABLES)<br>SELECT (42 ATTR.) FROM (5 TABLES)<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 10% | 300 |
| SELECT (60 ATTR.) FROM (5 TABLES) | 600 |

Purchase @ $0.16/hour

Tier #2

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (13 ATTR.) FROM LINEITEM<br>SELECT (11 ATTR.) FROM (2 TABLES)<br>SELECT (10 ATTR.) FROM (3 TABLES)<br>SELECT (8 ATTR.) FROM (4 TABLES)<br>SELECT (2 ATTR.) FROM (5 TABLES)<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 10% | 60 |
| SELECT (60 ATTR.) FROM (5 TABLES) | 300 |

Purchase @ $0.24/hour

Tier #3

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (17 ATTR.) FROM (3 TABLES)<br>SELECT (14 ATTR.) FROM (4 TABLES)<br>SELECT (8 ATTR.) FROM (5 TABLES) | 60 |

Purchase @ $0.32/hour

Tier #4

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (3 ATTR.) FROM LINEITEM<br>SELECT (2 ATTR.) FROM (2 TABLES)<br>SELECT (25 ATTR.) FROM (2 TABLES) WHERE 10%<br>SELECT (24 ATTR.) FROM (4 TABLES) WHERE 10%<br>SELECT (4 ATTR.) FROM (5 TABLES) WHERE 10%<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 1% | 10 |
| SELECT (35 ATTR.) FROM (3 TABLES)<br>SELECT (32 ATTR.) FROM (4 TABLES)<br>SELECT (31 ATTR.) FROM (5 TABLES) | 60 |

Purchase @ $0.64/hour

# PSLA for Myria

Tier #1

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (1 ATTR.) FROM LINEITEM WHERE 0.1%<br>SELECT (9 ATTR.) FROM PART<br>SELECT (17 ATTR.) FROM DATE<br>SELECT (9 ATTR.) FROM CUSTOMER | 10 |
| SELECT (8 ATTR.) FROM (2 TABLES)<br>SELECT (3 ATTR.) FROM (3 TABLES)<br>SELECT (2 ATTR.) FROM (4 TABLES)<br>SELECT (59 ATTR.) FROM (5 TABLES) WHERE 10%<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 1% | 60 |
| SELECT (43 ATTR.) FROM (4 TABLES)<br>SELECT (42 ATTR.) FROM (5 TABLES)<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 10% | 300 |
| SELECT (60 ATTR.) FROM (5 TABLES) | 600 |

Purchase @ $0.16/hour

Tier #2

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (13 ATTR.) FROM LINEITEM<br>SELECT (11 ATTR.) FROM (2 TABLES)<br>SELECT (10 ATTR.) FROM (3 TABLES)<br>SELECT (8 ATTR.) FROM (4 TABLES)<br>SELECT (2 ATTR.) FROM (5 TABLES)<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 10% | 60 |
| SELECT (60 ATTR.) FROM (5 TABLES) | 300 |

Purchase @ $0.24/hour

Tier #3

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (17 ATTR.) FROM (3 TABLES)<br>SELECT (14 ATTR.) FROM (4 TABLES)<br>SELECT (8 ATTR.) FROM (5 TABLES) | 60 |

Purchase @ $0.32/hour

Tier #4

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (3 ATTR.) FROM LINEITEM<br>SELECT (2 ATTR.) FROM (2 TABLES)<br>SELECT (25 ATTR.) FROM (2 TABLES) WHERE 10%<br>SELECT (24 ATTR.) FROM (4 TABLES) WHERE 10%<br>SELECT (4 ATTR.) FROM (5 TABLES) WHERE 10%<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 1% | 10 |
| SELECT (35 ATTR.) FROM (3 TABLES)<br>SELECT (32 ATTR.) FROM (4 TABLES)<br>SELECT (31 ATTR.) FROM (5 TABLES) | 60 |

Purchase @ $0.64/hour

# PSLA for Myria

Tier #1

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (1 ATTR.) FROM LINEITEM WHERE 0.1%<br>SELECT (9 ATTR.) FROM PART<br>SELECT (17 ATTR.) FROM DATE<br>SELECT (9 ATTR.) FROM CUSTOMER | 10 |
| SELECT (8 ATTR.) FROM (2 TABLES)<br>SELECT (3 ATTR.) FROM (3 TABLES)<br>SELECT (2 ATTR.) FROM (4 TABLES)<br>SELECT (59 ATTR.) FROM (5 TABLES) WHERE 10%<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 1% | 60 |
| SELECT (43 ATTR.) FROM (4 TABLES)<br>SELECT (42 ATTR.) FROM (5 TABLES)<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 10% | 300 |
| SELECT (60 ATTR.) FROM (5 TABLES) | 600 |

Purchase @ $0.16/hour

Tier #2

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (13 ATTR.) FROM LINEITEM<br>SELECT (11 ATTR.) FROM (2 TABLES)<br>SELECT (10 ATTR.) FROM (3 TABLES)<br>SELECT (8 ATTR.) FROM (4 TABLES)<br>SELECT (2 ATTR.) FROM (5 TABLES)<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 10% | 60 |
| SELECT (60 ATTR.) FROM (5 TABLES) | 300 |

Purchase @ $0.24/hour

Tier #3

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (17 ATTR.) FROM (3 TABLES)<br>SELECT (14 ATTR.) FROM (4 TABLES)<br>SELECT (8 ATTR.) FROM (5 TABLES) | 60 |

Purchase @ $0.32/hour

Tier #4

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (3 ATTR.) FROM LINEITEM<br>SELECT (2 ATTR.) FROM (2 TABLES)<br>SELECT (25 ATTR.) FROM (2 TABLES) WHERE 10%<br>SELECT (24 ATTR.) FROM (4 TABLES) WHERE 10%<br>SELECT (4 ATTR.) FROM (5 TABLES) WHERE 10%<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 1% | 10 |
| SELECT (35 ATTR.) FROM (3 TABLES)<br>SELECT (32 ATTR.) FROM (4 TABLES)<br>SELECT (31 ATTR.) FROM (5 TABLES) | 60 |

Purchase @ $0.64/hour

# PSLA for Myria

Tier #1

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (1 ATTR.) FROM LINEITEM WHERE 0.1%<br>SELECT (9 ATTR.) FROM PART<br>SELECT (17 ATTR.) FROM DATE<br>SELECT (9 ATTR.) FROM CUSTOMER | 10 |
| SELECT (8 ATTR.) FROM (2 TABLES)<br>SELECT (3 ATTR.) FROM (3 TABLES)<br>SELECT (2 ATTR.) FROM (4 TABLES)<br>SELECT (59 ATTR.) FROM (5 TABLES) WHERE 10%<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 1% | 60 |
| SELECT (43 ATTR.) FROM (4 TABLES)<br>SELECT (42 ATTR.) FROM (5 TABLES)<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 10% | 300 |
| SELECT (60 ATTR.) FROM (5 TABLES) | 600 |

Purchase @ $0.16/hour

Tier #2

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (13 ATTR.) FROM LINEITEM<br>SELECT (11 ATTR.) FROM (2 TABLES)<br>SELECT (10 ATTR.) FROM (3 TABLES)<br>SELECT (8 ATTR.) FROM (4 TABLES)<br>SELECT (2 ATTR.) FROM (5 TABLES)<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 10% | 60 |
| SELECT (60 ATTR.) FROM (5 TABLES) | 300 |

Purchase @ $0.24/hour

Tier #3

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (17 ATTR.) FROM (3 TABLES)<br>SELECT (14 ATTR.) FROM (4 TABLES)<br>SELECT (8 ATTR.) FROM (5 TABLES) | 60 |

Purchase @ $0.32/hour

Tier #4

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (3 ATTR.) FROM LINEITEM<br>SELECT (2 ATTR.) FROM (2 TABLES)<br>SELECT (25 ATTR.) FROM (2 TABLES) WHERE 10%<br>SELECT (24 ATTR.) FROM (4 TABLES) WHERE 10%<br>SELECT (4 ATTR.) FROM (5 TABLES) WHERE 10%<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 1% | 10 |
| SELECT (35 ATTR.) FROM (3 TABLES)<br>SELECT (32 ATTR.) FROM (4 TABLES)<br>SELECT (31 ATTR.) FROM (5 TABLES) | 60 |

Purchase @ $0.64/hour

# PSLA for Myria

Tier #1

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (1 ATTR.) FROM LINEITEM WHERE 0.1%<br>SELECT (9 ATTR.) FROM PART<br>SELECT (17 ATTR.) FROM DATE<br>SELECT (9 ATTR.) FROM CUSTOMER | 10 |
| SELECT (8 ATTR.) FROM (2 TABLES)<br>SELECT (3 ATTR.) FROM (3 TABLES)<br>SELECT (2 ATTR.) FROM (4 TABLES)<br>SELECT (59 ATTR.) FROM (5 TABLES) WHERE 10%<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 1% | 60 |
| SELECT (43 ATTR.) FROM (4 TABLES)<br>SELECT (42 ATTR.) FROM (5 TABLES)<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 10% | 300 |
| SELECT (60 ATTR.) FROM (5 TABLES) | 600 |

Purchase @ $0.16/hour

Tier #2

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (13 ATTR.) FROM LINEITEM<br>SELECT (11 ATTR.) FROM (2 TABLES)<br>SELECT (10 ATTR.) FROM (3 TABLES)<br>SELECT (8 ATTR.) FROM (4 TABLES)<br>SELECT (2 ATTR.) FROM (5 TABLES)<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 10% | 60 |
| SELECT (60 ATTR.) FROM (5 TABLES) | 300 |

Purchase @ $0.24/hour

Tier #3

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (17 ATTR.) FROM (3 TABLES)<br>SELECT (14 ATTR.) FROM (4 TABLES)<br>SELECT (8 ATTR.) FROM (5 TABLES) | 60 |

Purchase @ $0.32/hour

Tier #4

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (3 ATTR.) FROM LINEITEM<br>SELECT (2 ATTR.) FROM (2 TABLES)<br>SELECT (25 ATTR.) FROM (2 TABLES) WHERE 10%<br>SELECT (24 ATTR.) FROM (4 TABLES) WHERE 10%<br>SELECT (4 ATTR.) FROM (5 TABLES) WHERE 10%<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 1% | 10 |
| SELECT (35 ATTR.) FROM (3 TABLES)<br>SELECT (32 ATTR.) FROM (4 TABLES)<br>SELECT (31 ATTR.) FROM (5 TABLES) | 60 |

Purchase @ $0.64/hour

# PSLAManager

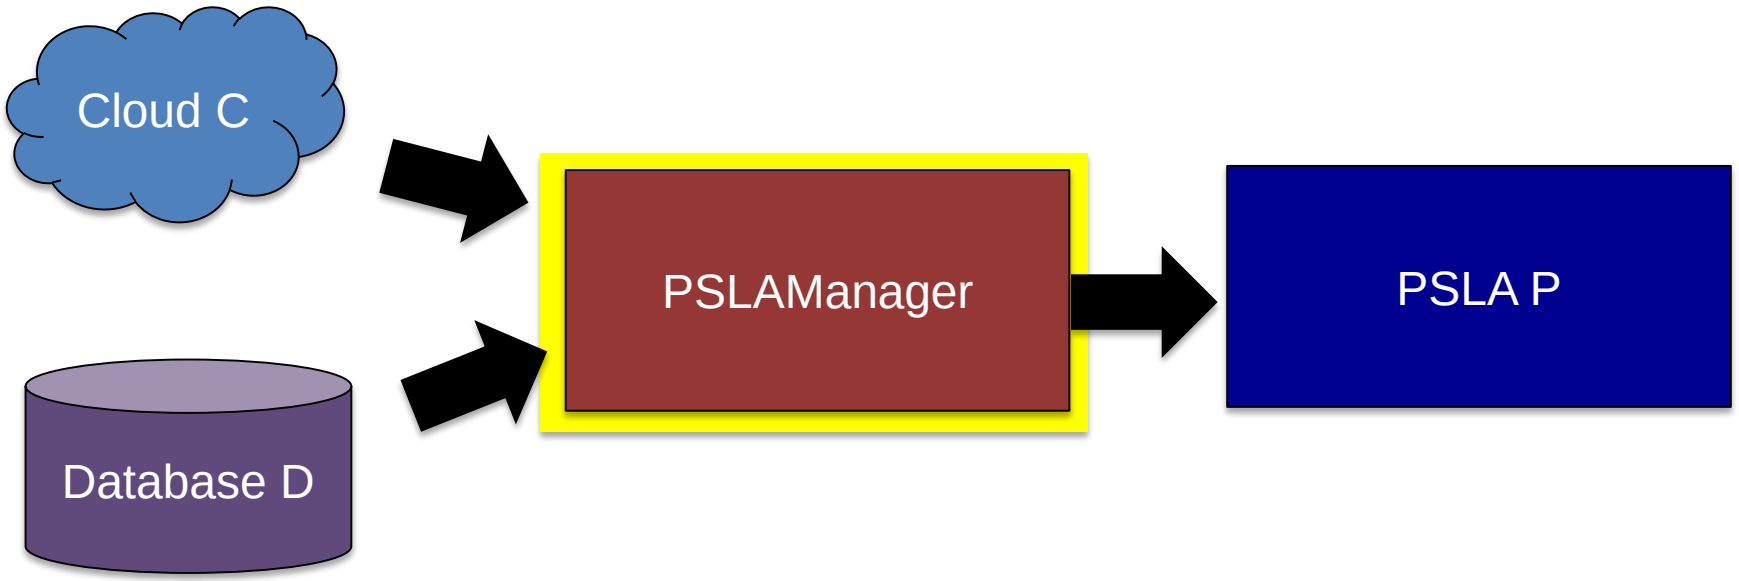

# PSLAManager Workflow

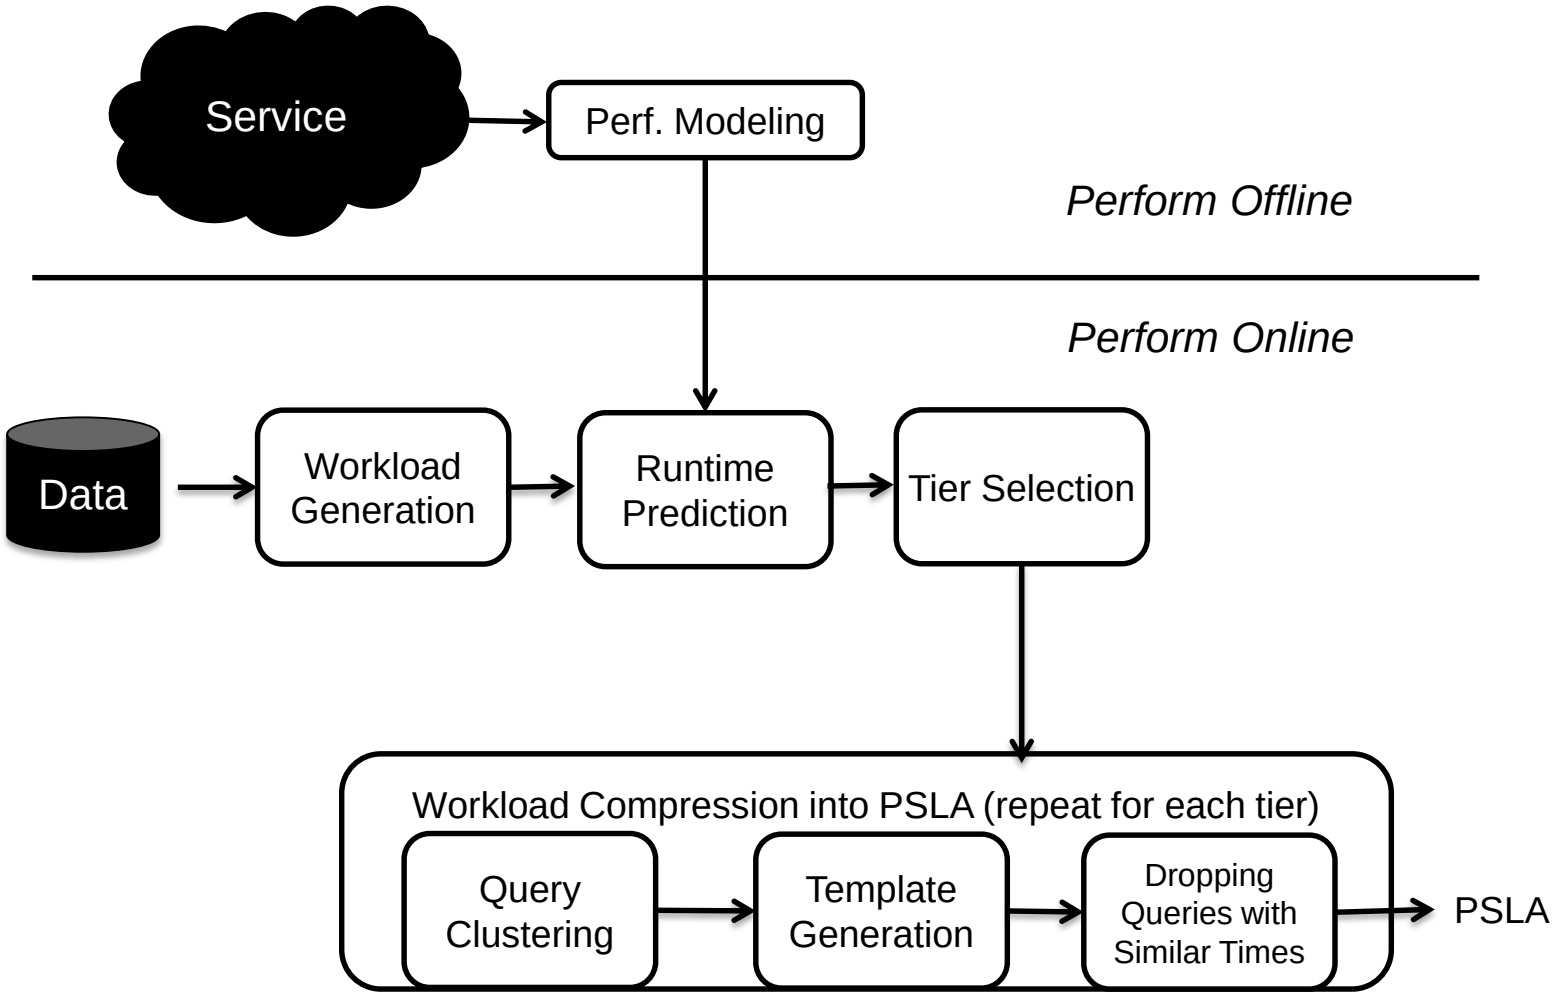

# PSLAManager Workflow

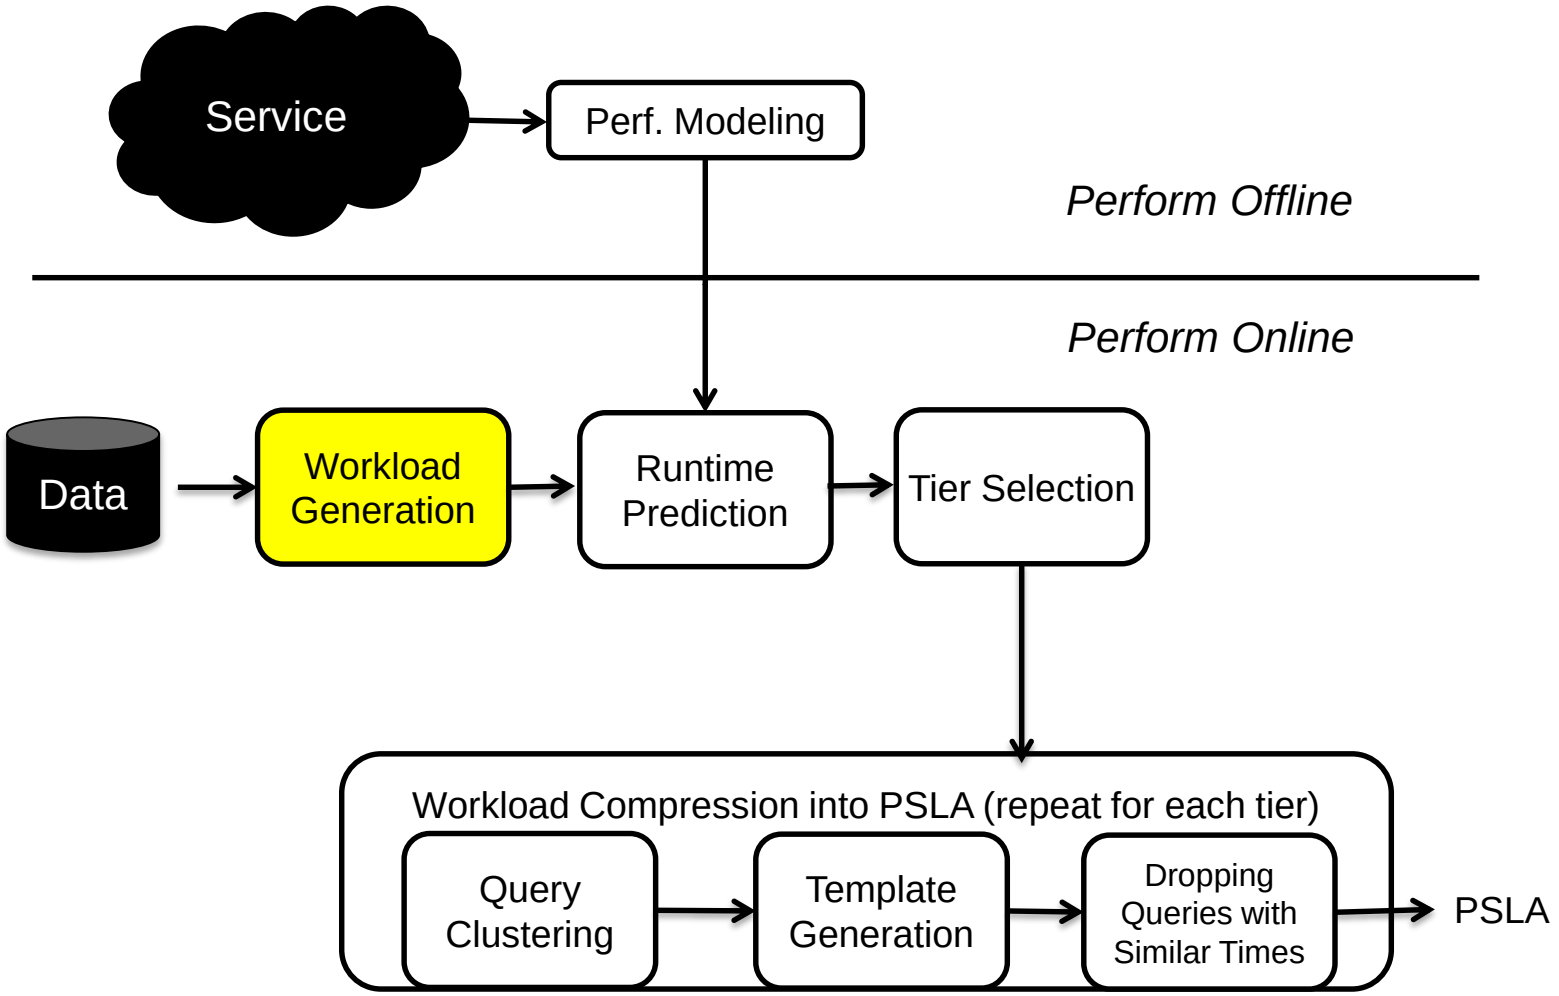

# PSLAManager Workflow

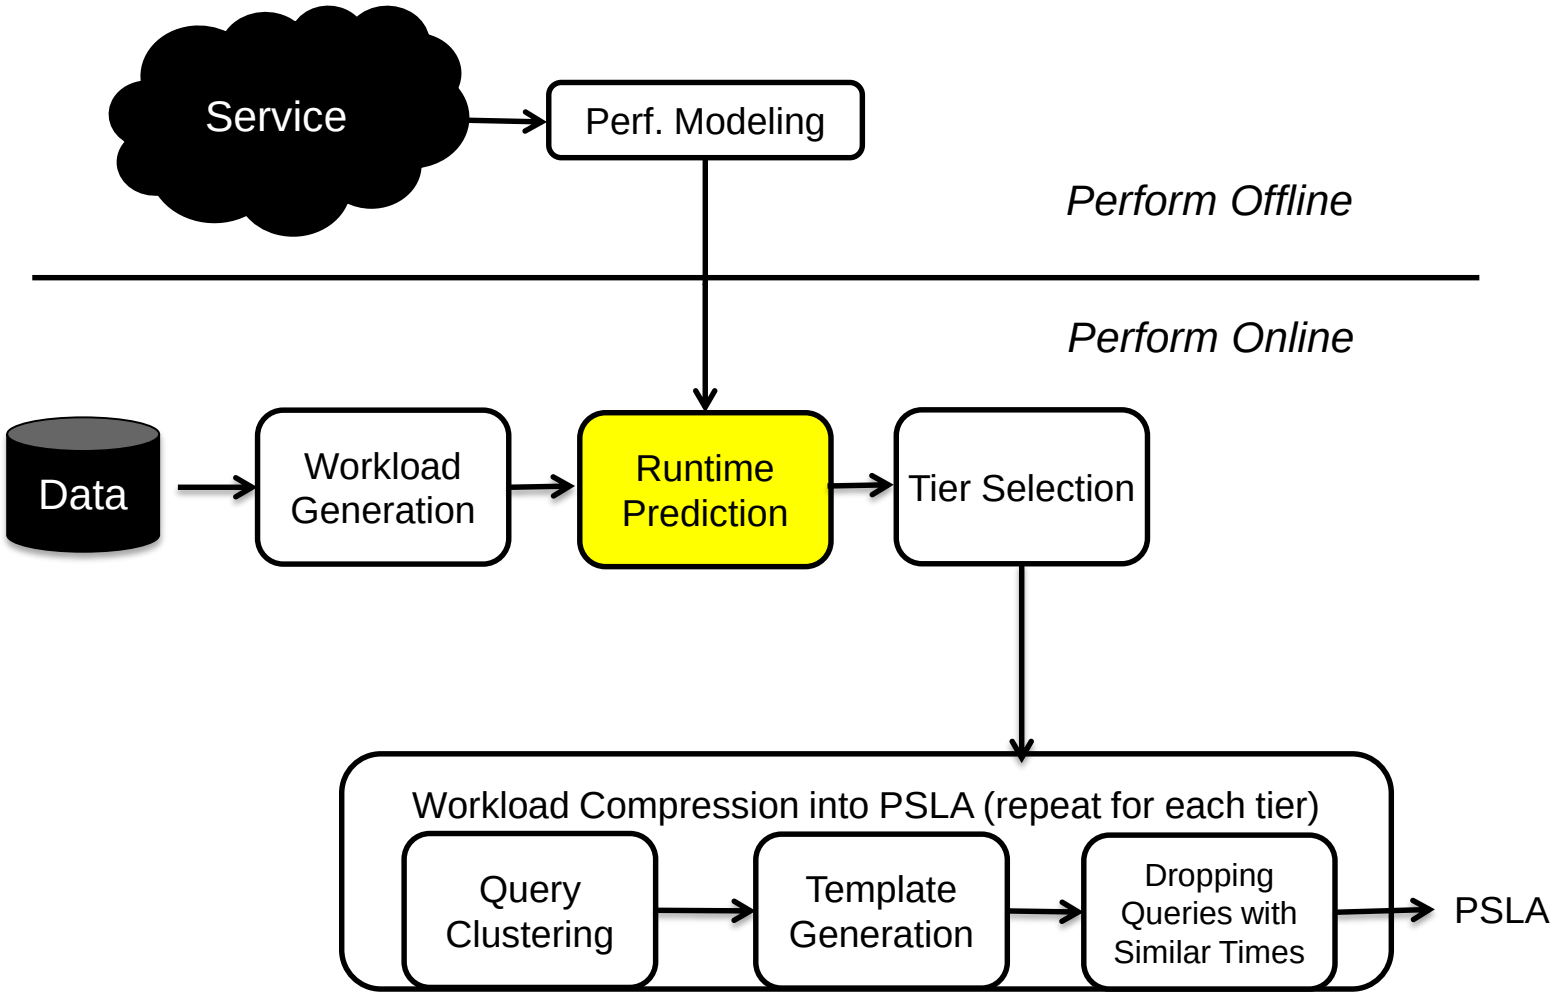

# PSLAManager Workflow

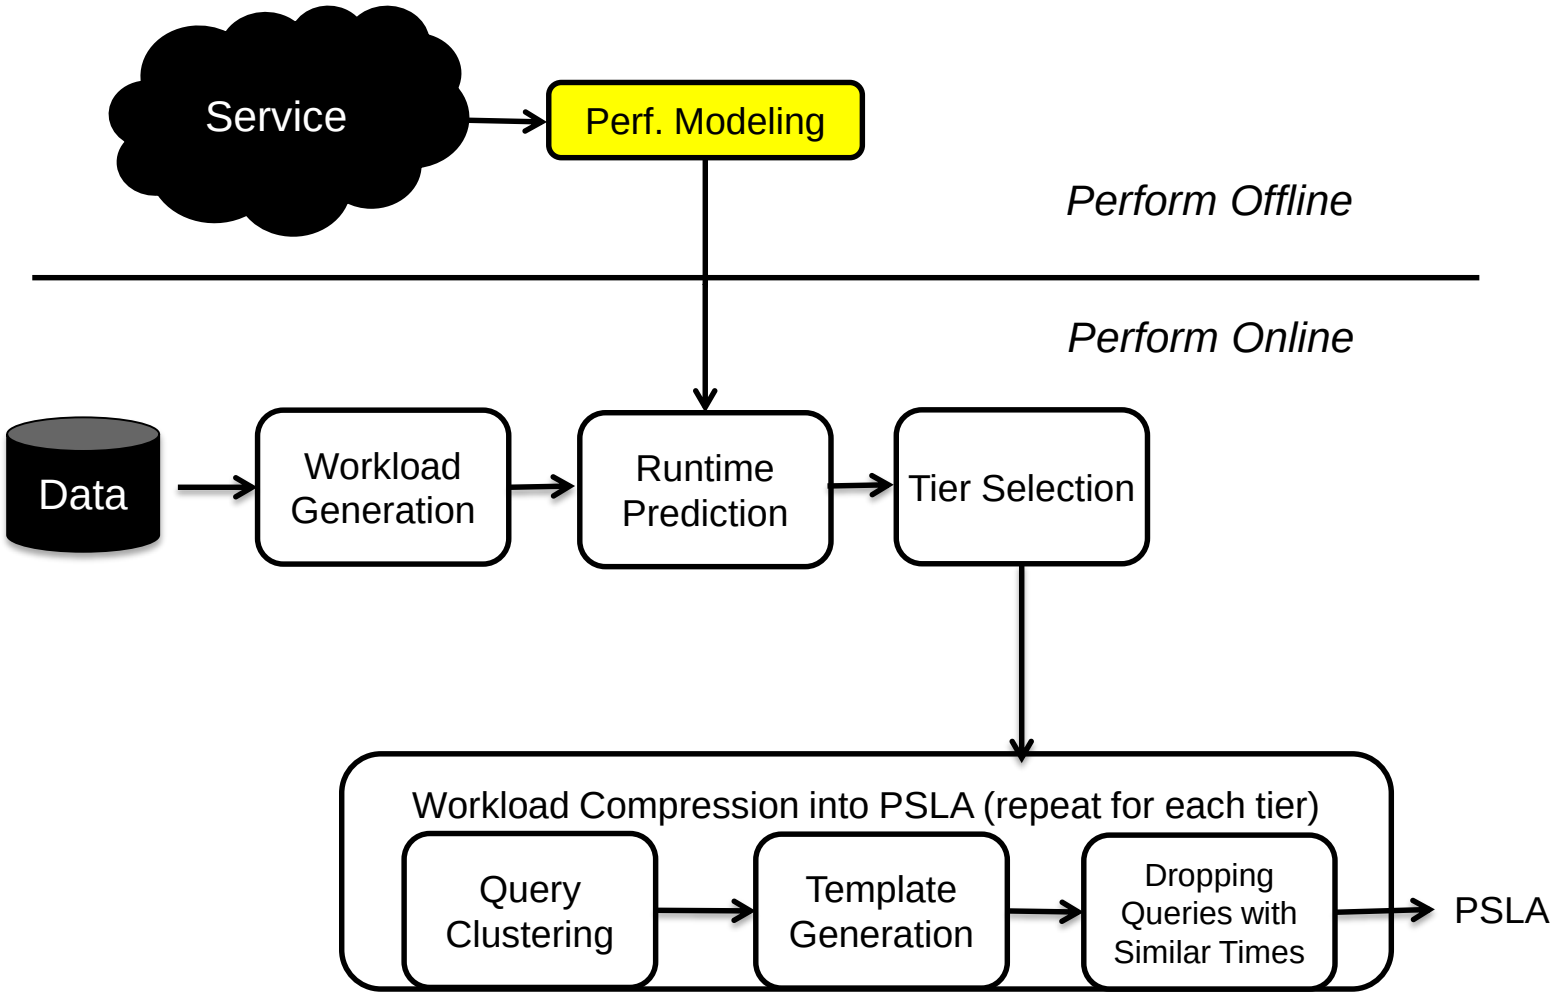

# PSLAManager Workflow

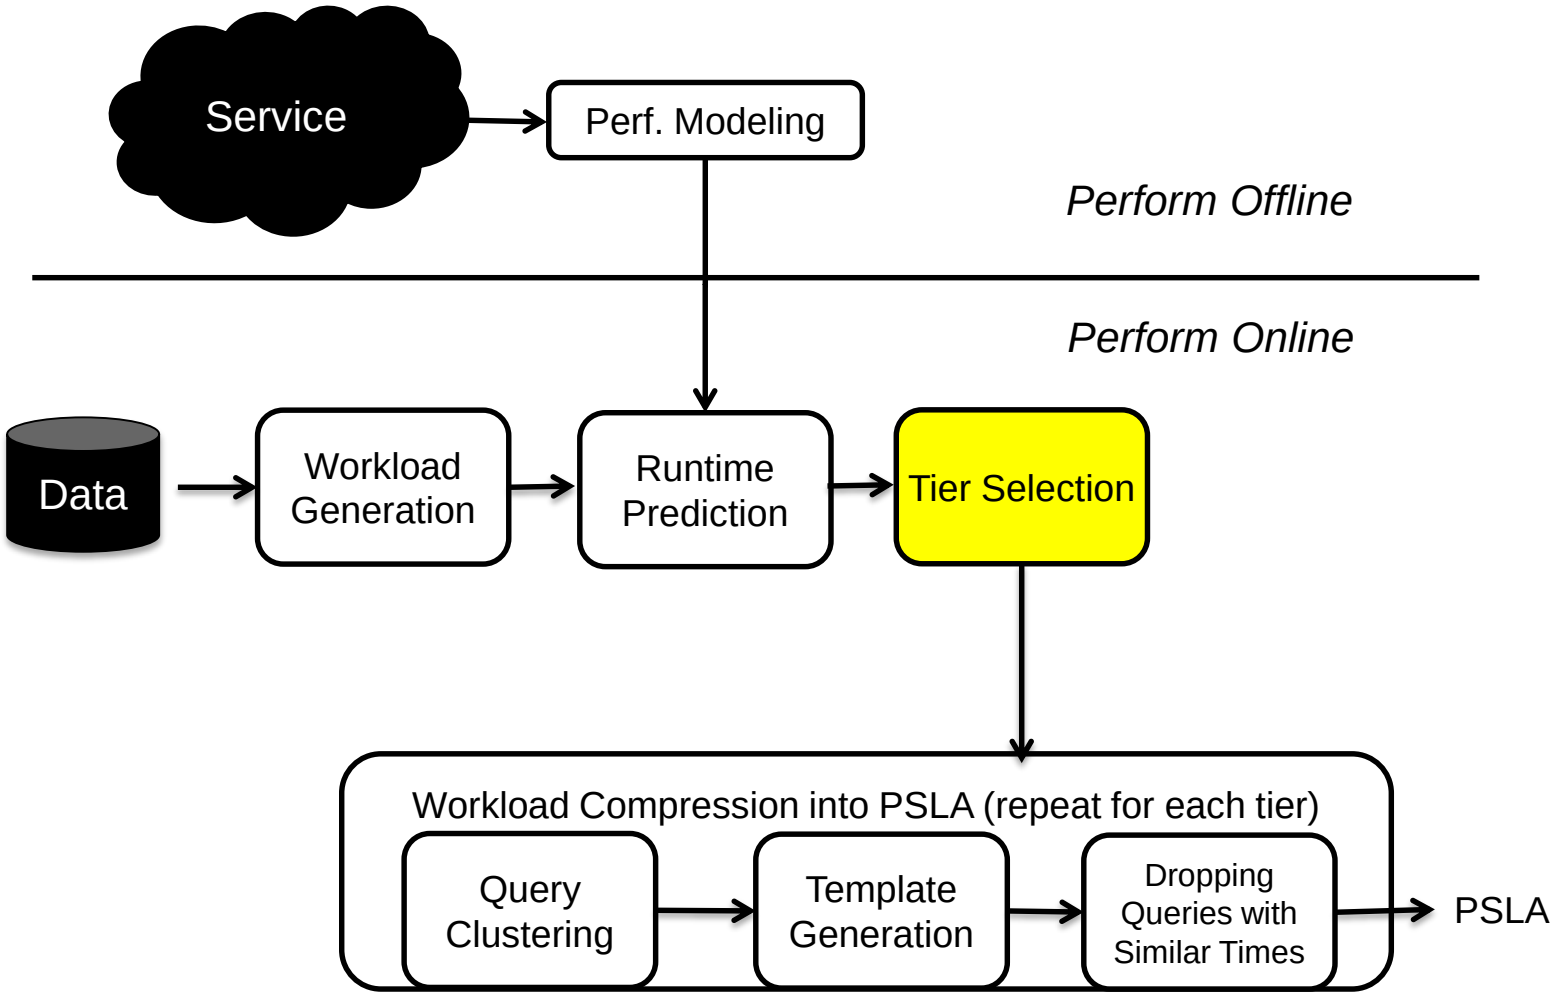

# PSLAManager Workflow

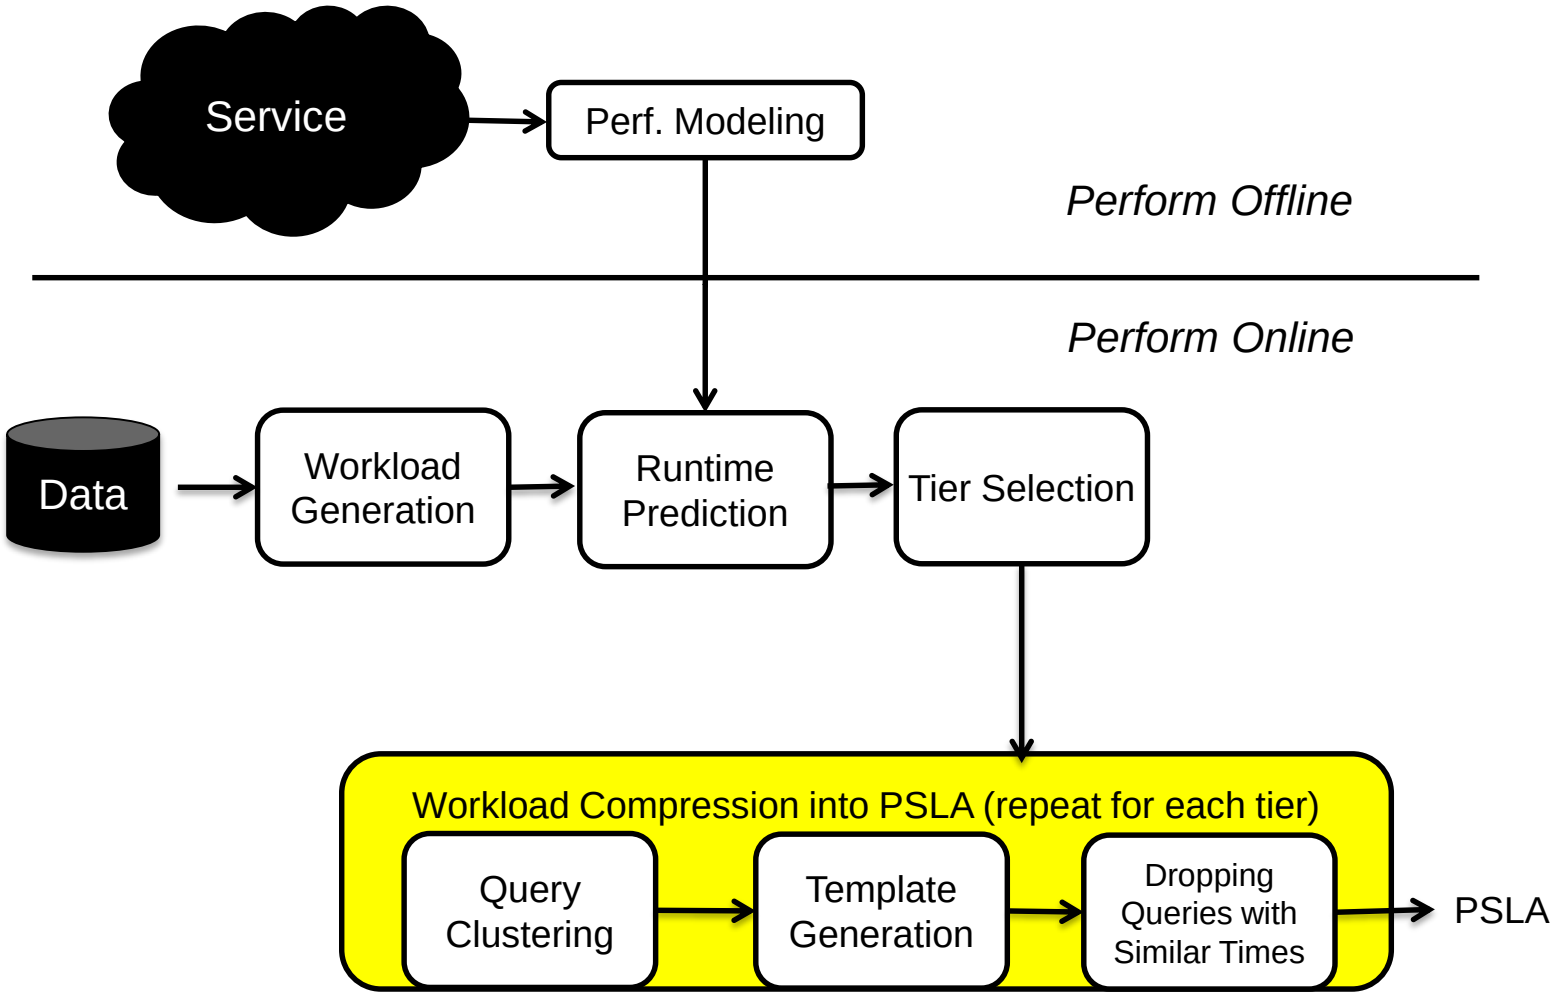

# PSLAManager Workflow

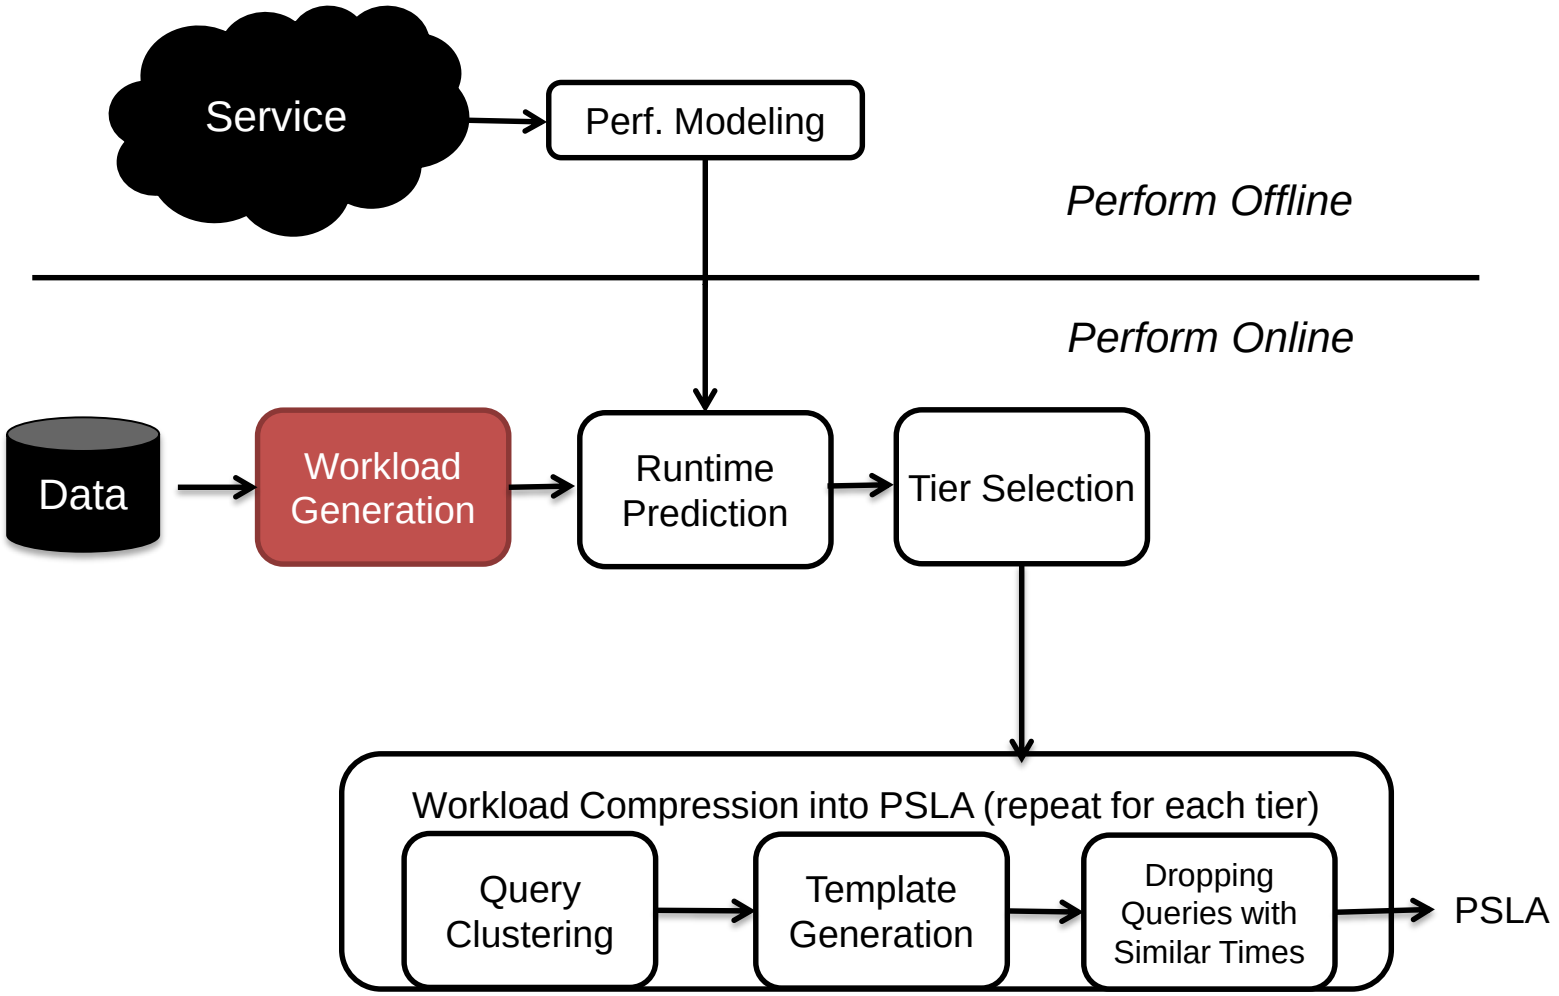

# Query Workload Generation

- Which queries to generate?
  - Joins drive performance
    - Think about possible combinations of joins

**Consider: All possible 2-way joins**

Tables in Order by Size: Lineitem, Part, Customer, Supplier, Date

(Lineitem ⋈ Part) ≥ (Customer ⋈ Date), (Lineitem ⋈ Supplier), etc.

  - Only consider most expensive queries
  - Build toward more complex queries, include selections and projections

# PSLAManager Workflow

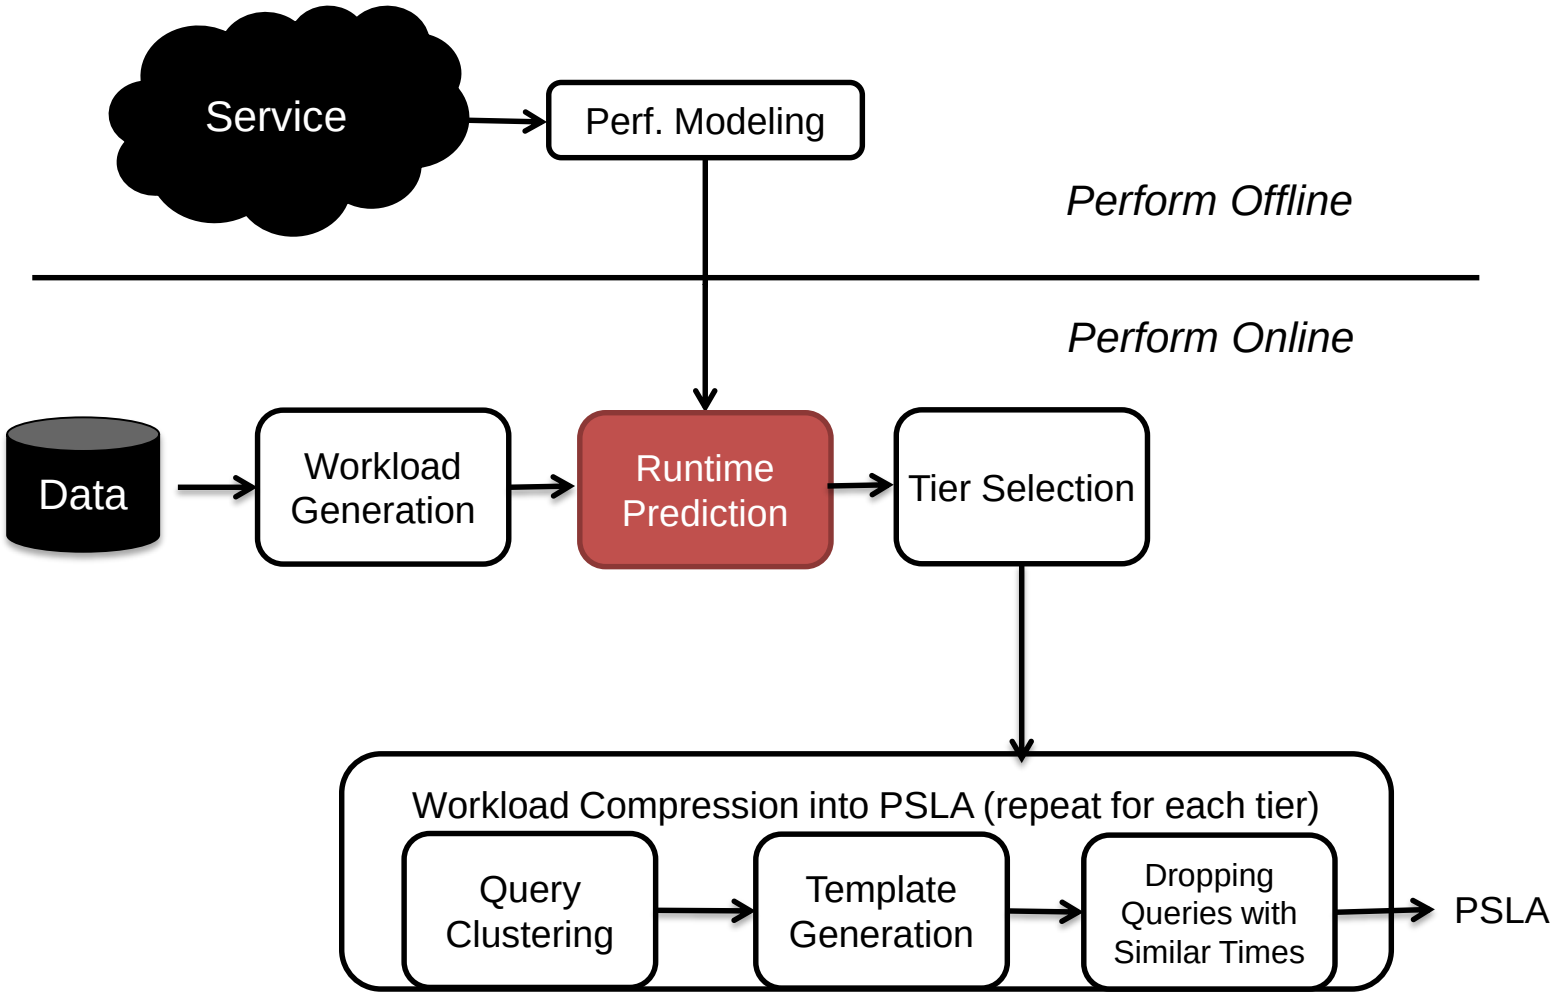

# PSLAManager Workflow

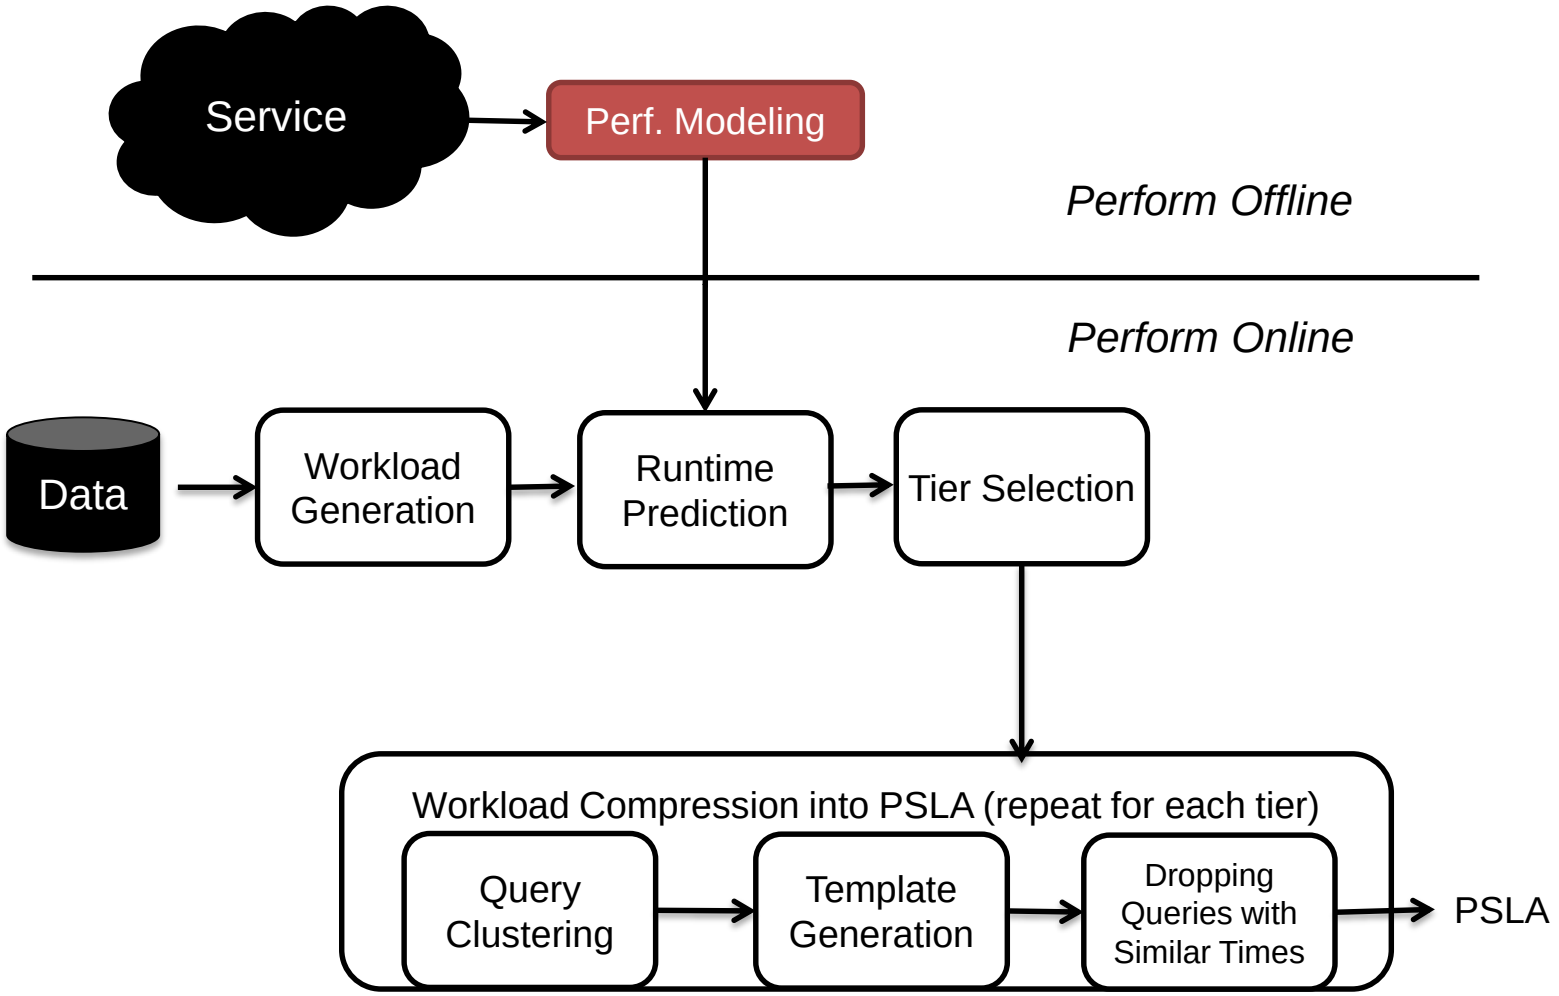

# PSLAManager Workflow

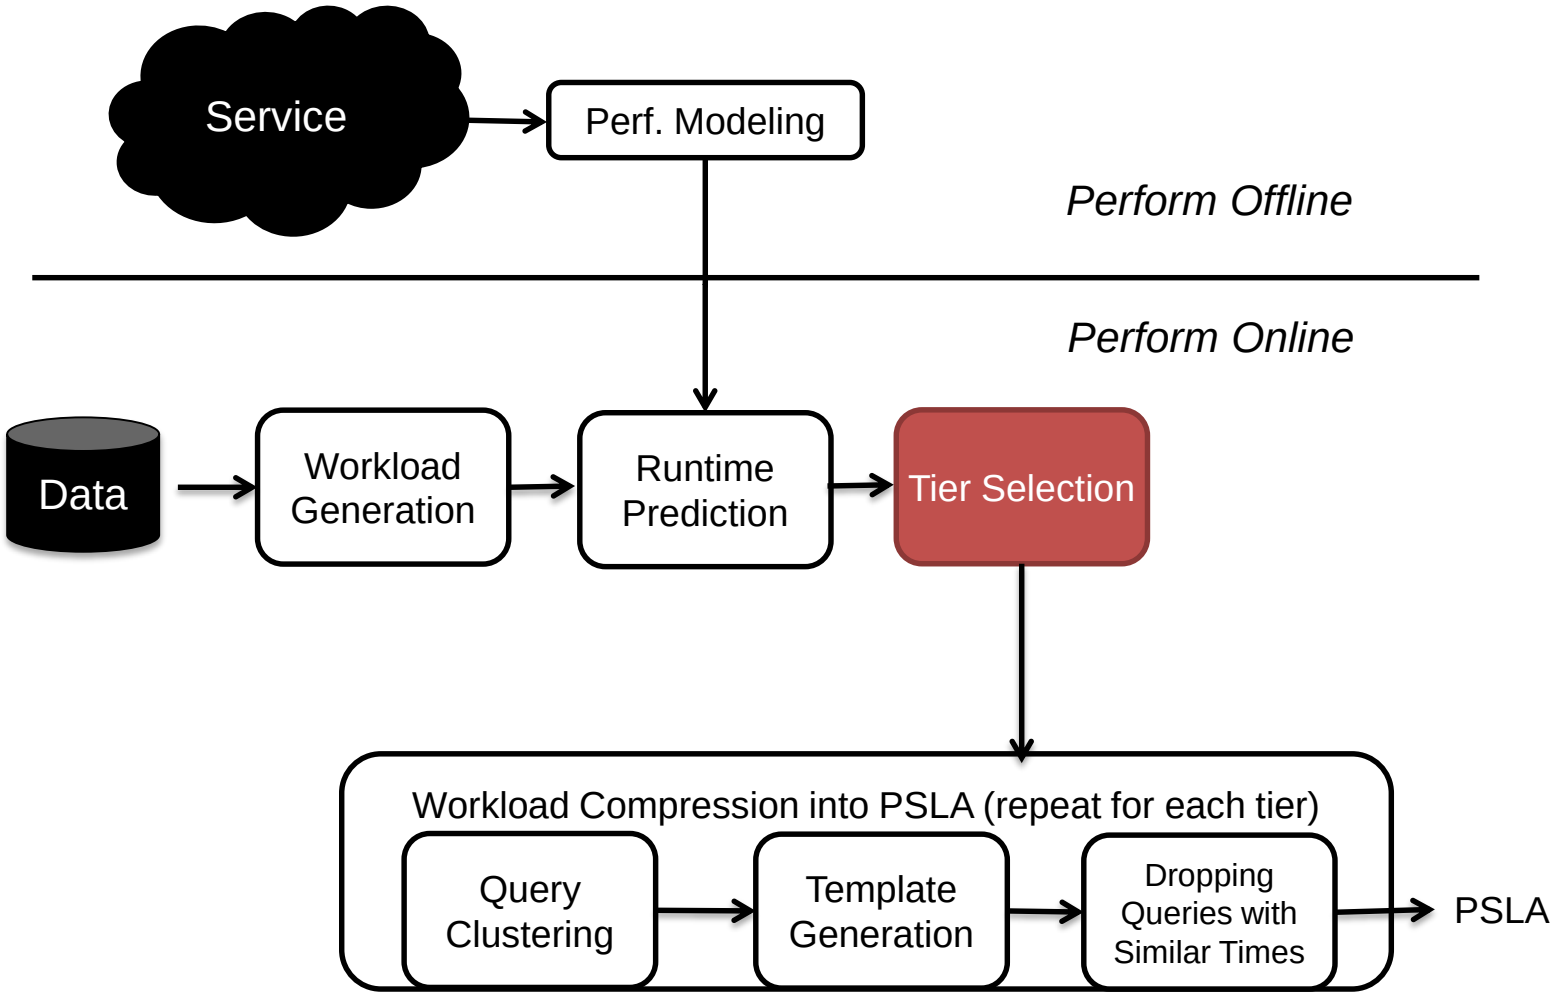

# Tier Selection

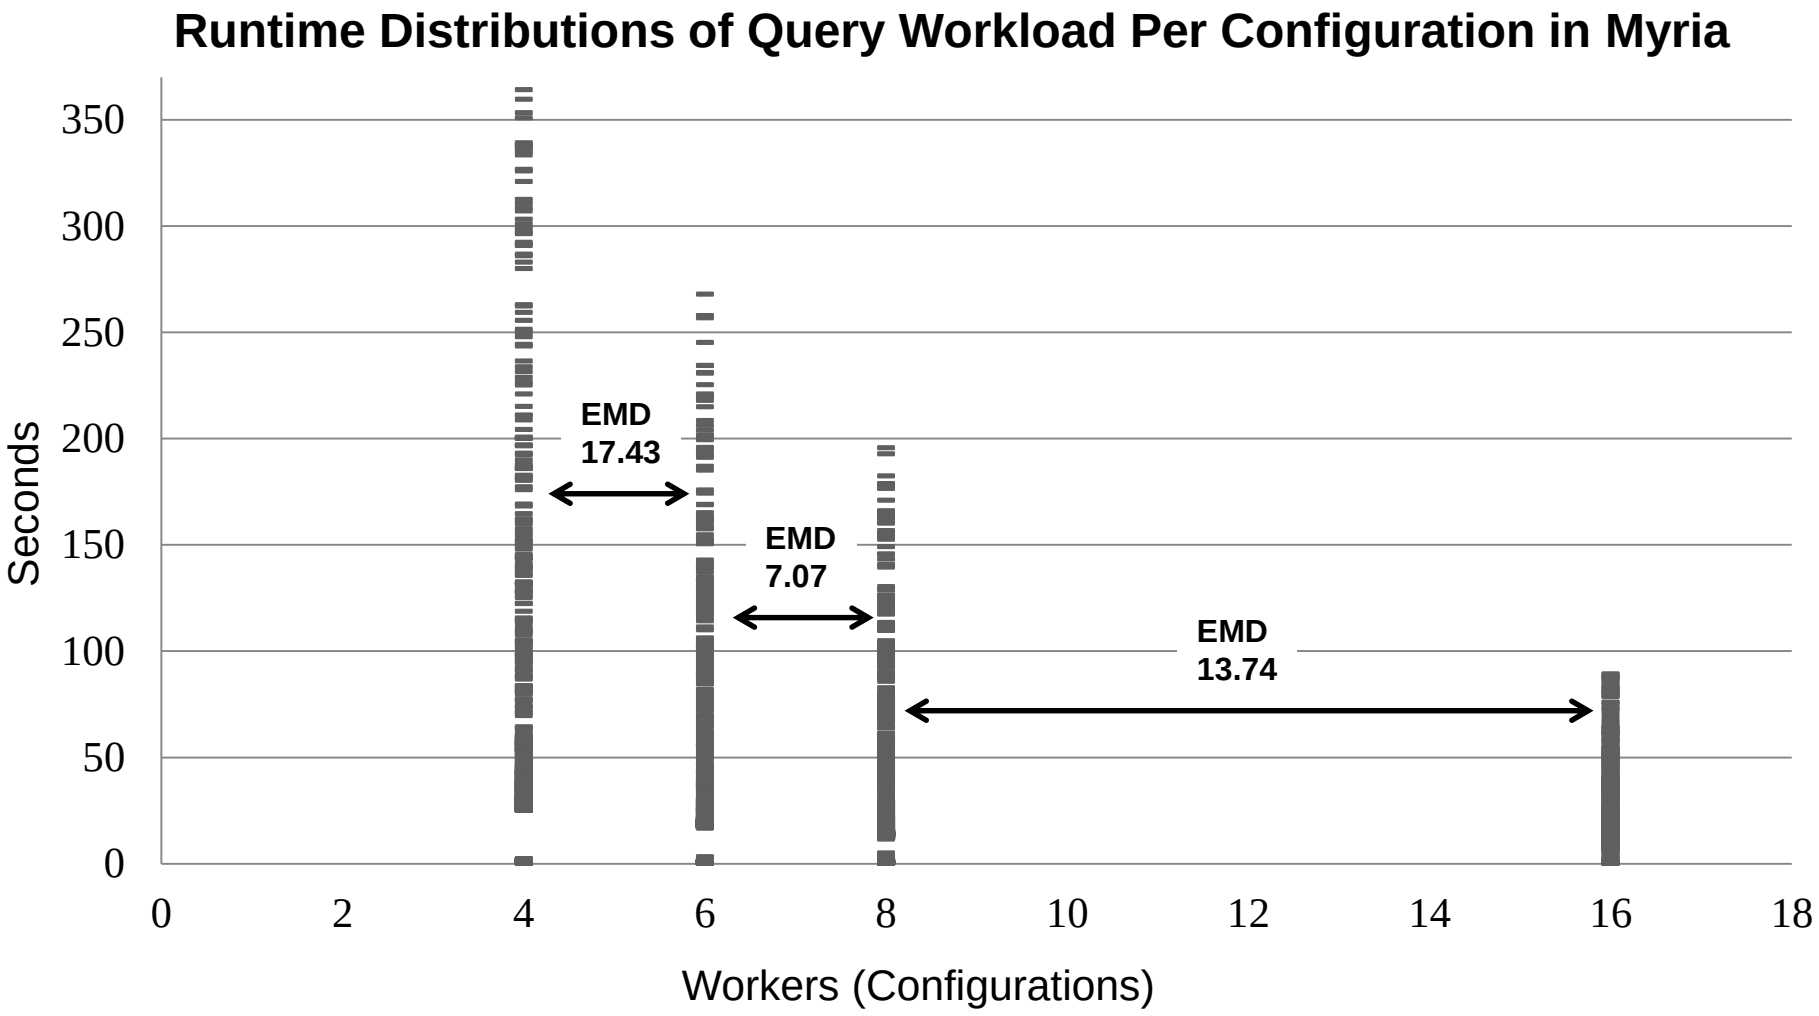

# Tier Selection

Tier #1

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (1 ATTR.) FROM LINEITEM WHERE 0.1%<br>SELECT (9 ATTR.) FROM PART<br>SELECT (17 ATTR.) FROM DATE<br>SELECT (9 ATTR.) FROM CUSTOMER | 10 |
| SELECT (8 ATTR.) FROM (2 TABLES)<br>SELECT (3 ATTR.) FROM (3 TABLES)<br>SELECT (2 ATTR.) FROM (4 TABLES)<br>SELECT (59 ATTR.) FROM (5 TABLES) WHERE 10%<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 1% | 60 |
| SELECT (43 ATTR.) FROM (4 TABLES)<br>SELECT (42 ATTR.) FROM (5 TABLES)<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 10% | 300 |
| SELECT (60 ATTR.) FROM (5 TABLES) | 600 |

Purchase @ $0.16/hour

Tier #2

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (13 ATTR.) FROM LINEITEM<br>SELECT (11 ATTR.) FROM (2 TABLES)<br>SELECT (10 ATTR.) FROM (3 TABLES)<br>SELECT (8 ATTR.) FROM (4 TABLES)<br>SELECT (2 ATTR.) FROM (5 TABLES)<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 10% | 60 |
| SELECT (60 ATTR.) FROM (5 TABLES) | 300 |

Purchase @ $0.24/hour

Tier #3

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (17 ATTR.) FRO TABLES)<br>SELECT (14 ATTR.) FRO TABLES)<br>SELECT (8 ATTR.) FR M (5 T LES) | 60 |

Purchase @ $0.32/hour

Tier #4

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (3 ATTR.) FROM LINEITEM<br>SELECT (2 ATTR.) FROM (2 TABLES)<br>SELECT (25 ATTR.) FROM (2 TABLES) WHERE 10%<br>SELECT (24 ATTR.) FROM (4 TABLES) WHERE 10%<br>SELECT (4 ATTR.) FROM (5 TABLES) WHERE 10%<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 1% | 10 |
| SELECT (35 ATTR.) FROM (3 TABLES)<br>SELECT (32 ATTR.) FROM (4 TABLES)<br>SELECT (31 ATTR.) FROM (5 TABLES) | 60 |

Purchase @ $0.64/hour

# PSLAManager Workflow

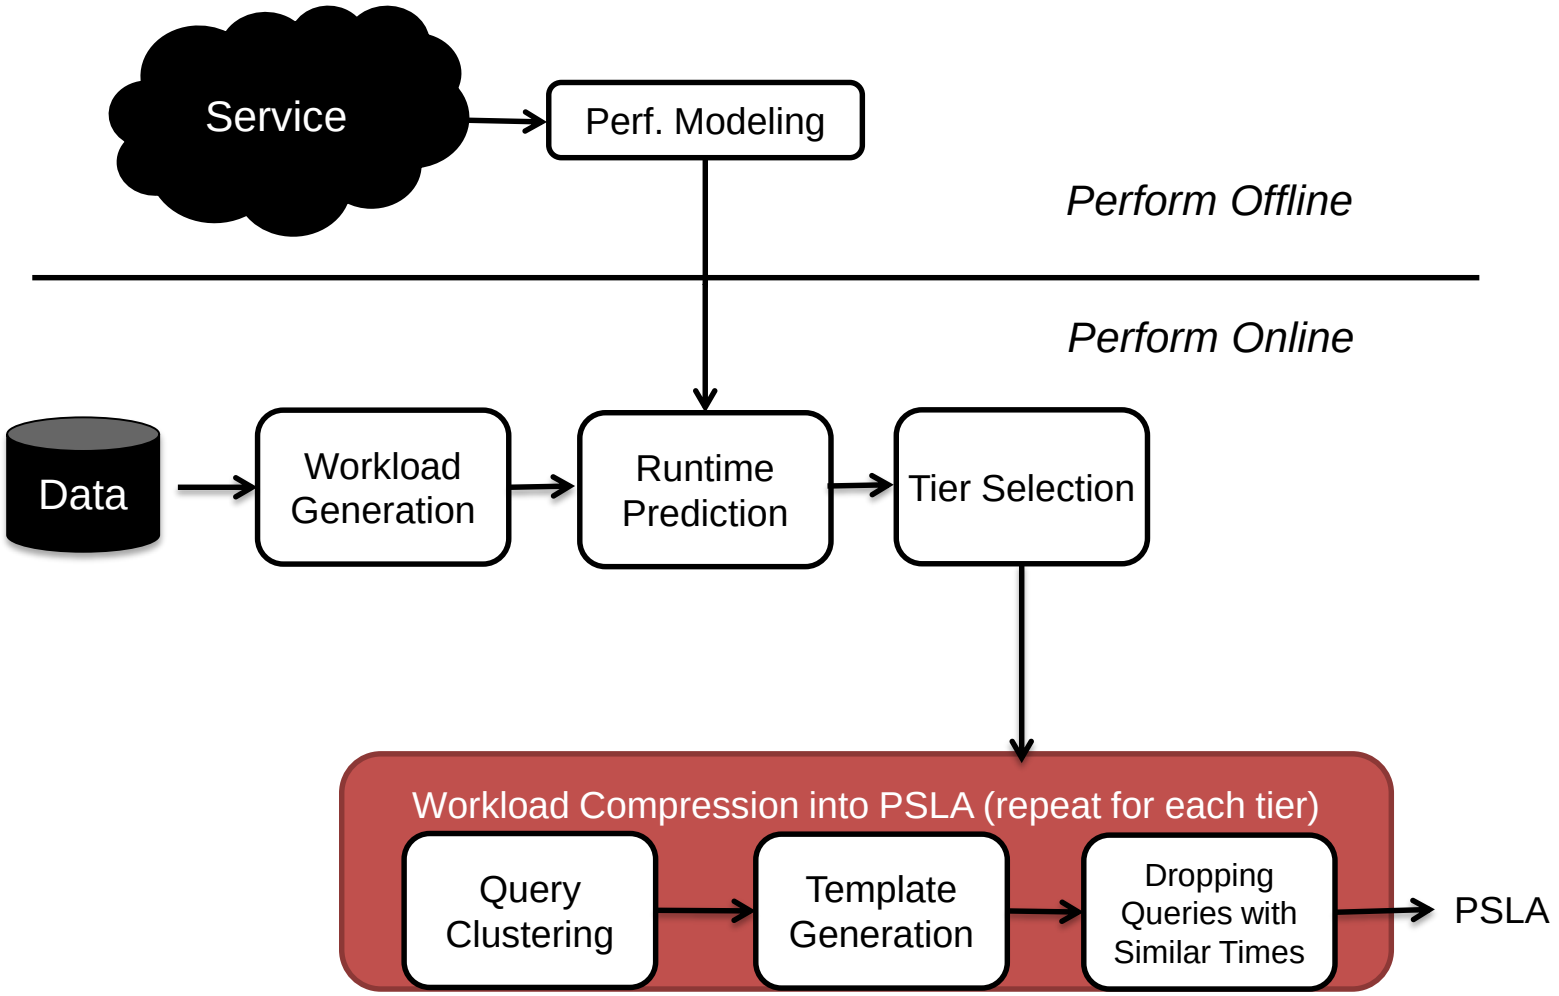

# Workload Compression

## STEP 1: Query Clustering

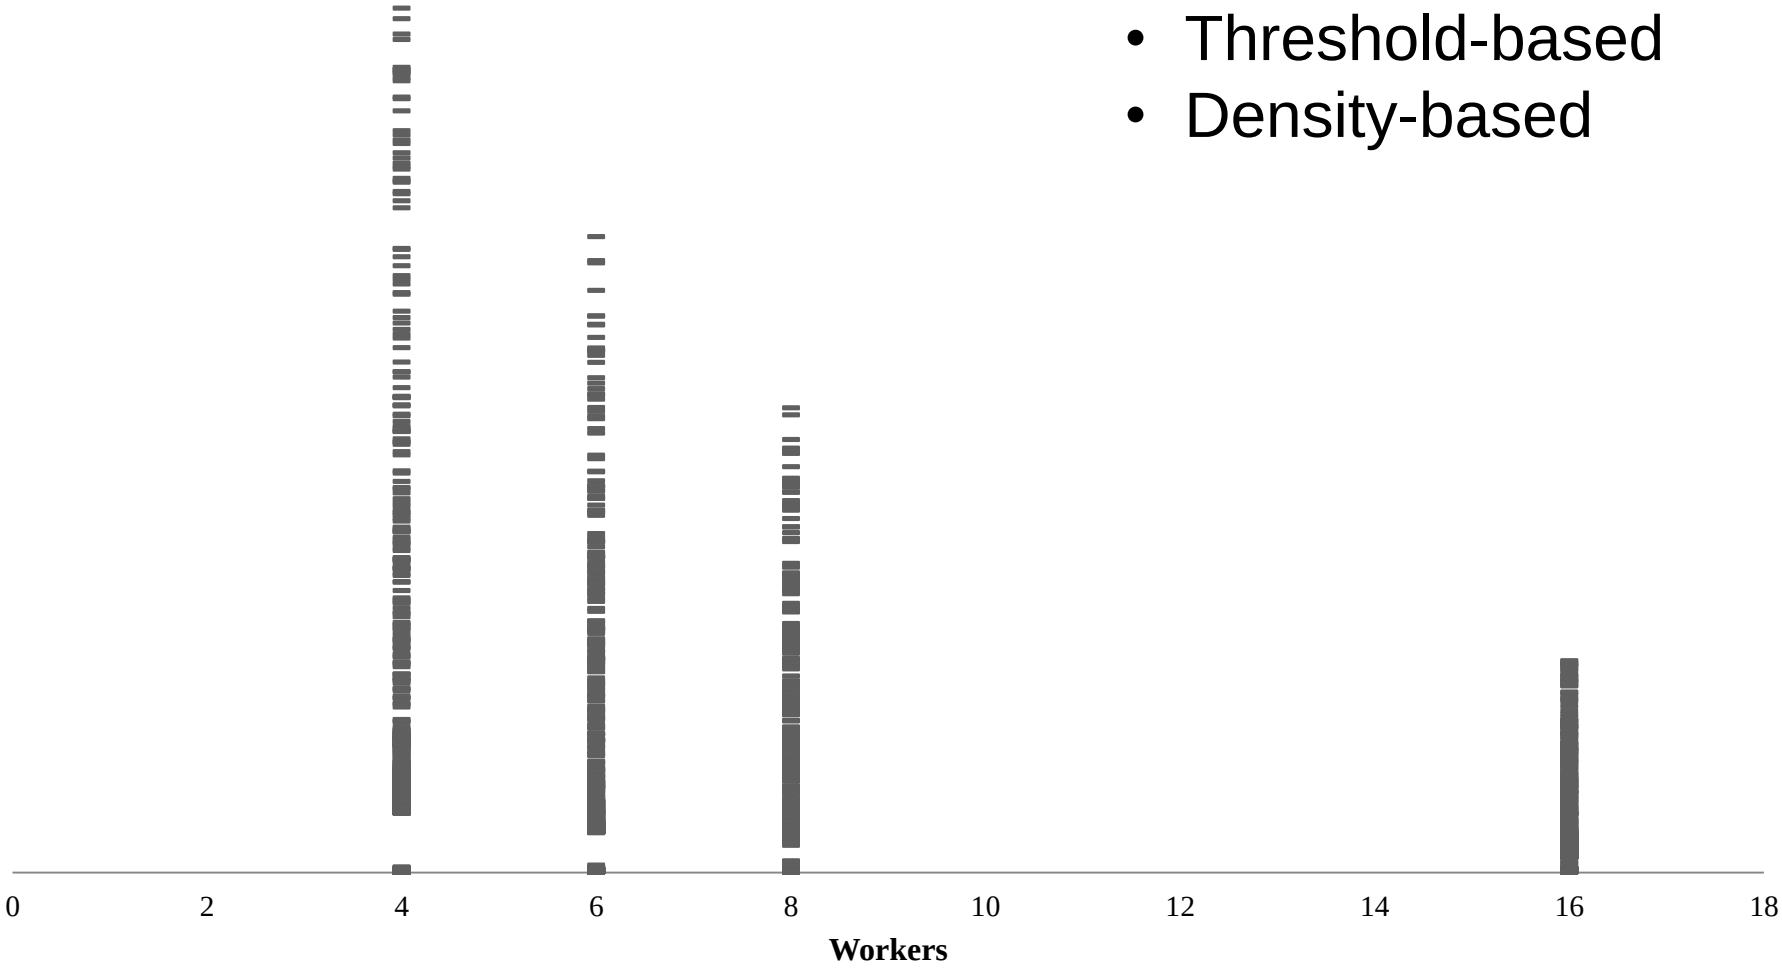


- Threshold-based
- Density-based

# Workload Compression

## STEP 1: Query Clustering

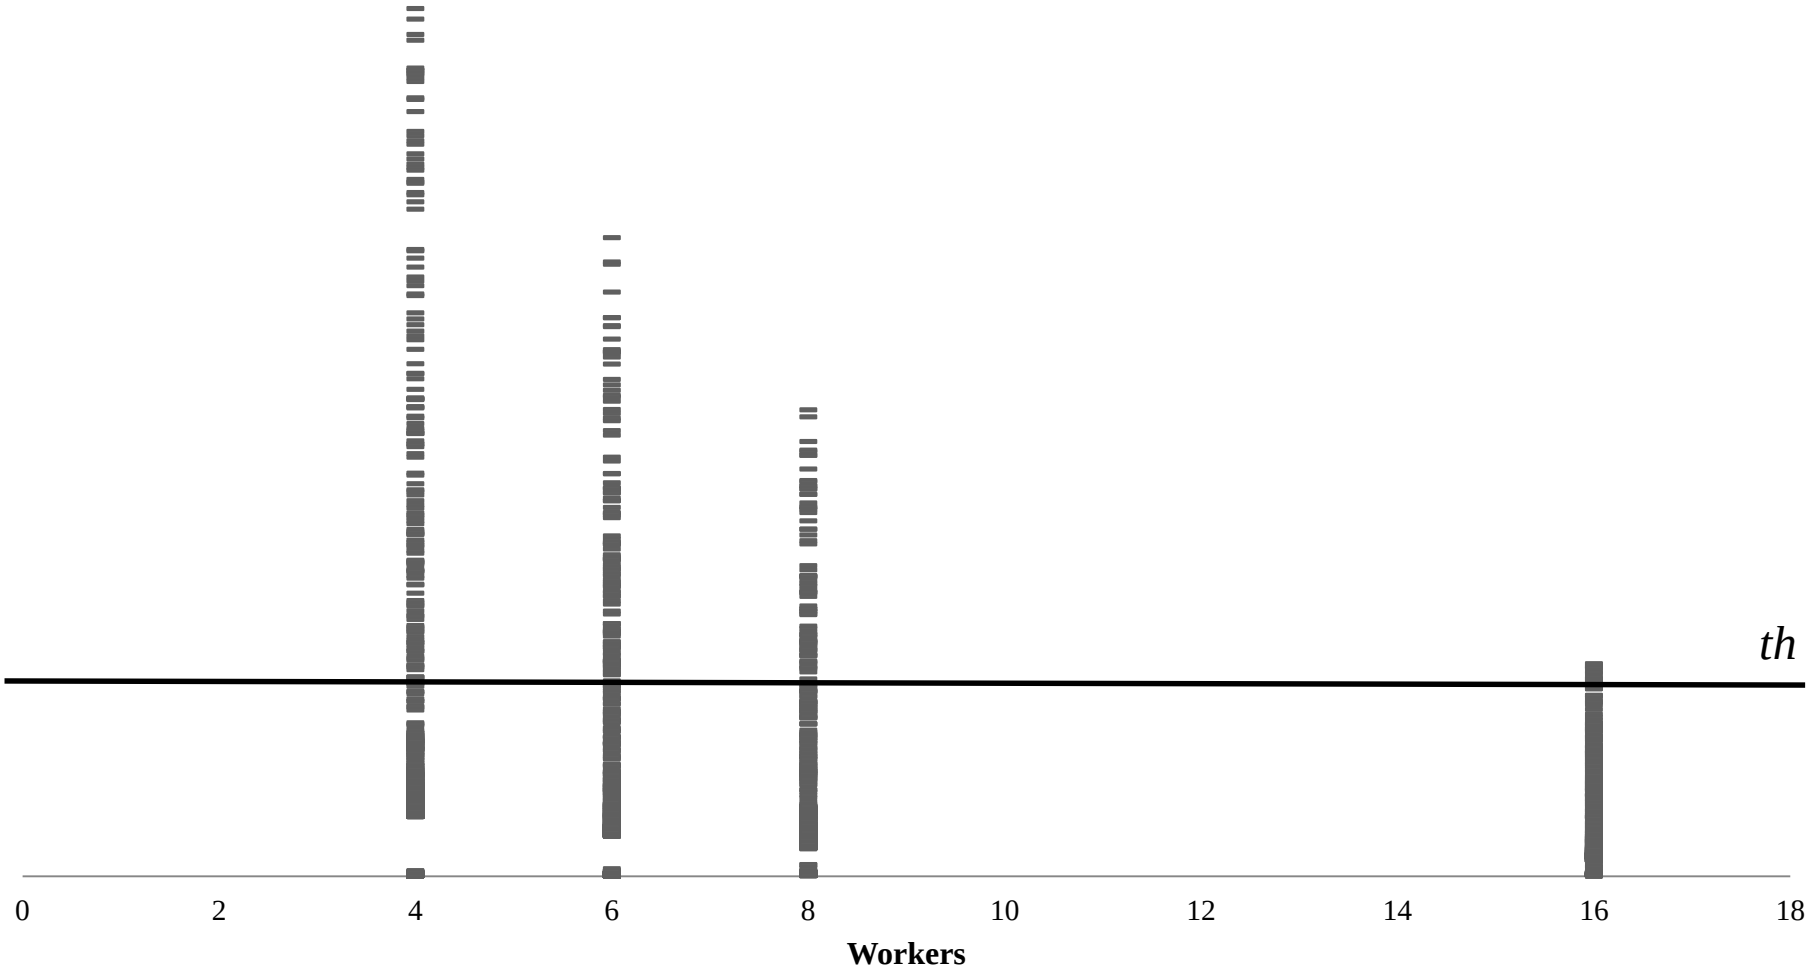

# Workload Compression

## STEP 1: Query Clustering

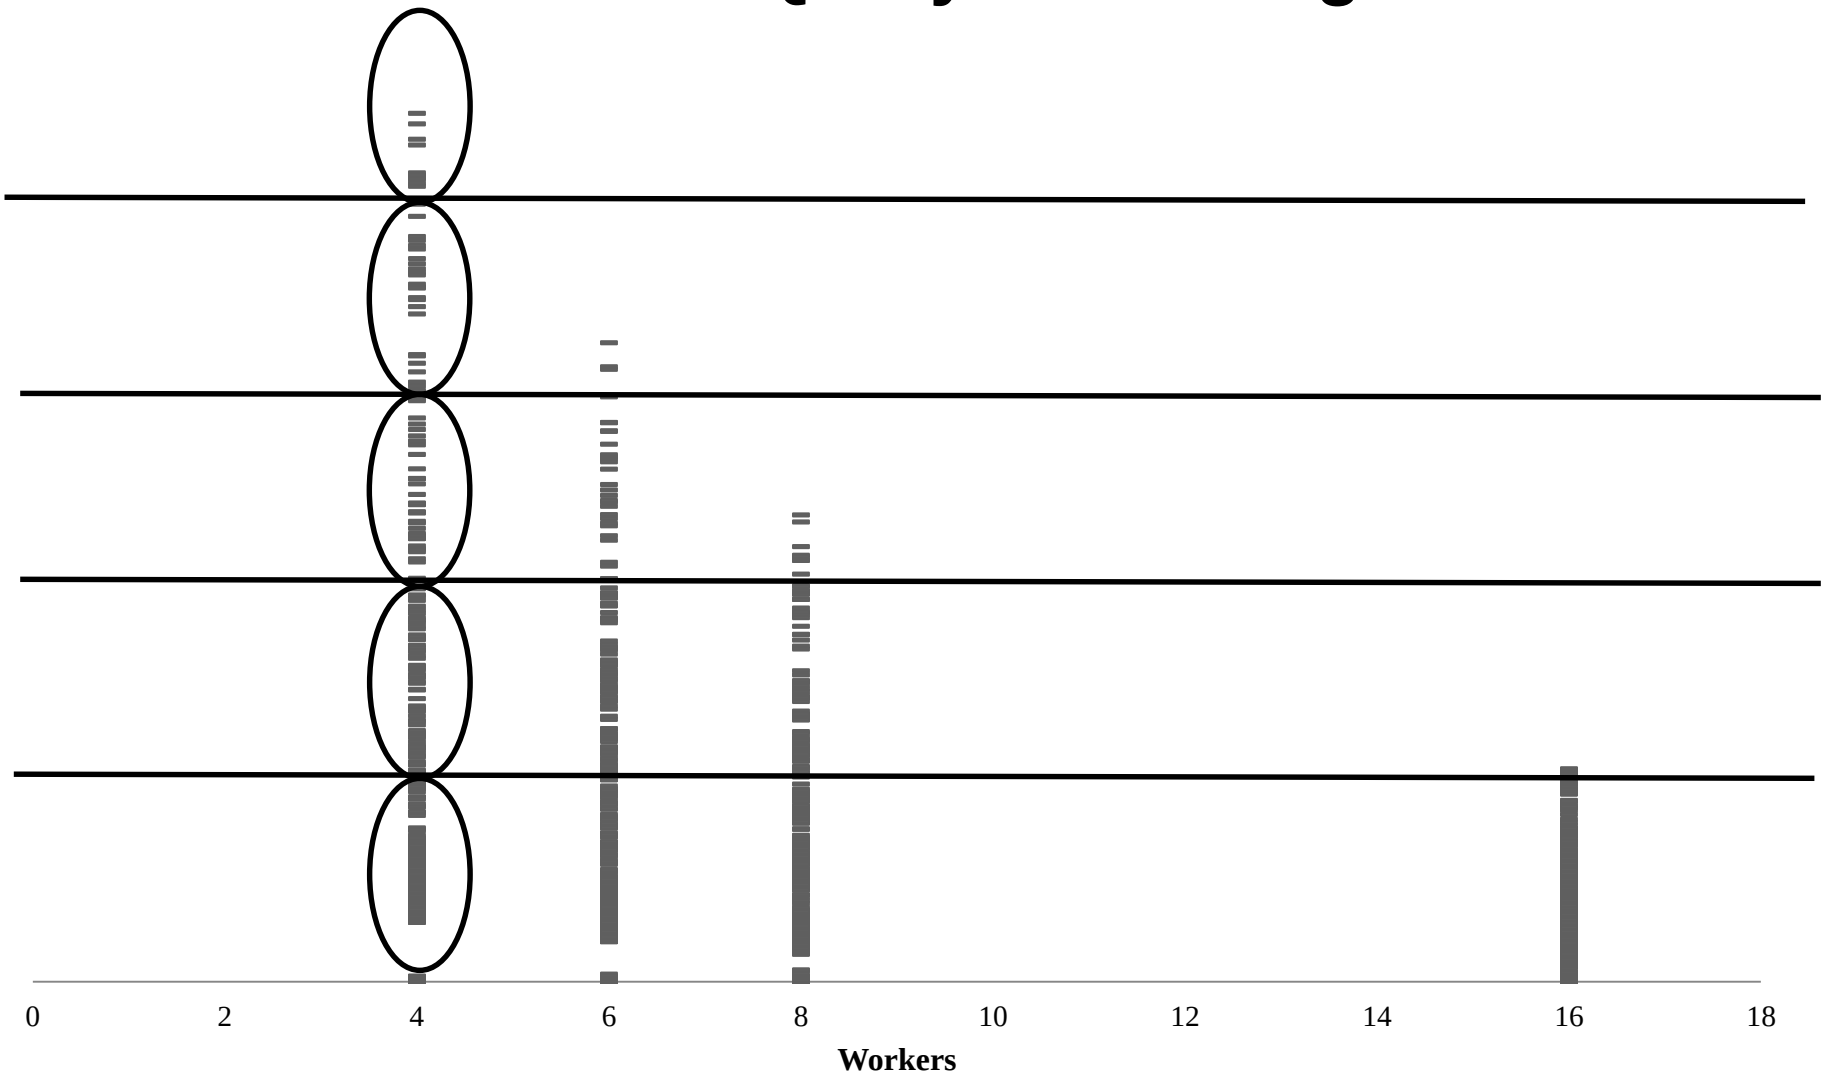

# Workload Compression

## STEP 1: Query Clustering

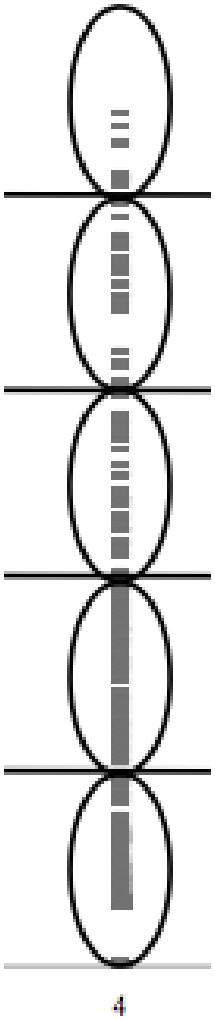

**Tier 1: $0.XX/hour**

***Within Cluster Max Threshold***:
Query Templates in Cluster…

***Within Cluster Max Threshold***:
Query Templates in Cluster…

***Still too many queries!***

# Workload Compression

## STEP 2: Template Generation

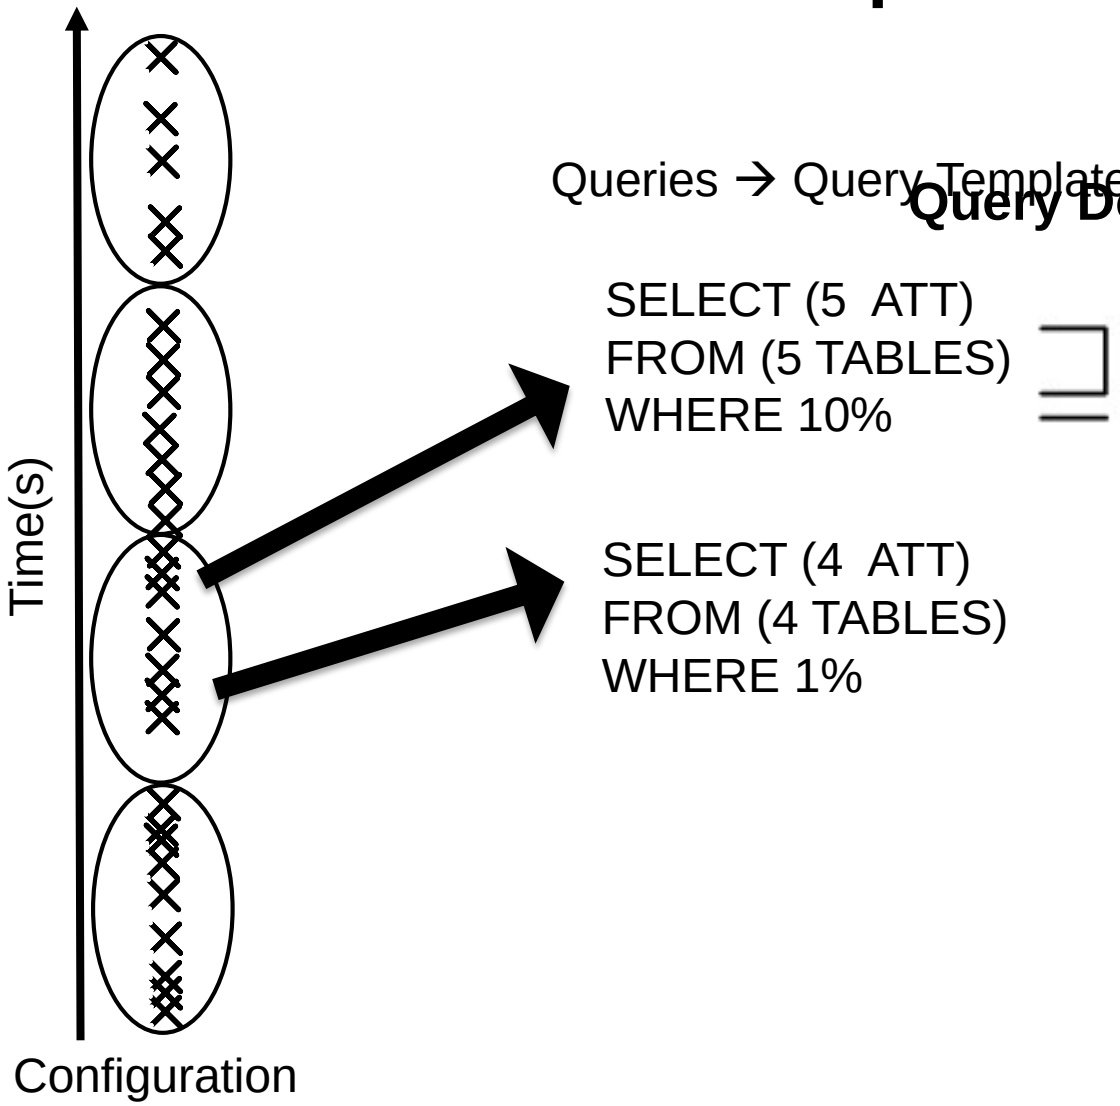

# Workload Compression

## STEP 2: Template Generation

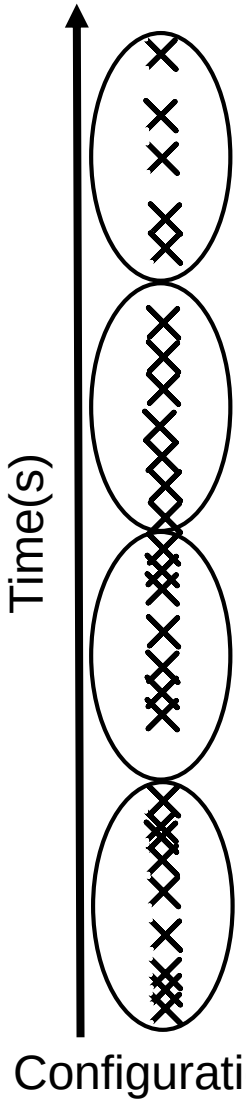


**Query Dominance**

SELECT (5 ATT)
FROM (5 TABLES)
WHERE 10%

⊒

SELECT (4 ATT)
FROM (4 TABLES)
WHERE 1%

Attributes Projected

Tables

Selectivity

# Workload Compression

## STEP 2: Template Generation

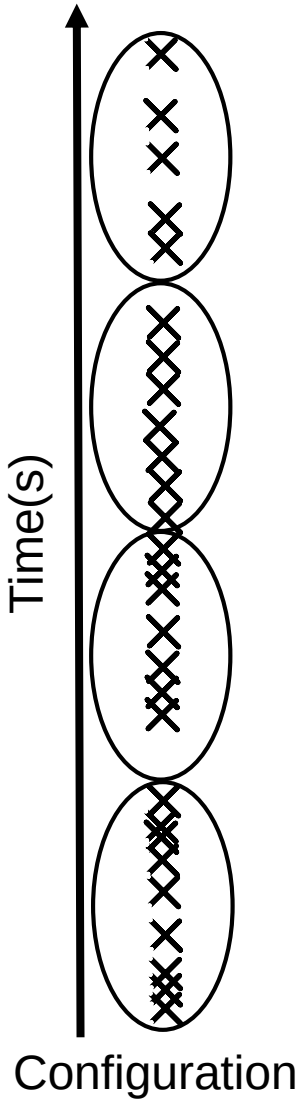


**Query Dominance**

Attributes Projected

Tables

SELECT (5 ATT)
FROM (5 TABLES)
WHERE 10%

$\sqsupseteq$

SELECT (4 ATT)
FROM (4 TABLES)
WHERE 1%

Selectivity

Given: $q_i, q_j$

$$q_i = (T_i, A_i, e_j) \qquad q_j = (T_j, A_j, e_j)$$

$$q_i \sqsupseteq q_j \iff T_i \sqsupseteq T_j \wedge A_i \sqsupseteq A_j \wedge e_i \supseteq e_j$$

# Workload Compression

## STEP 2: Template Generation

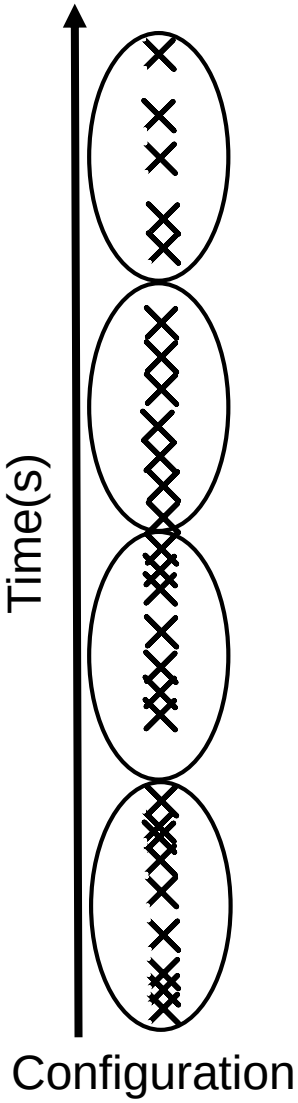

# Workload Compression

## STEP 2: Template Generation

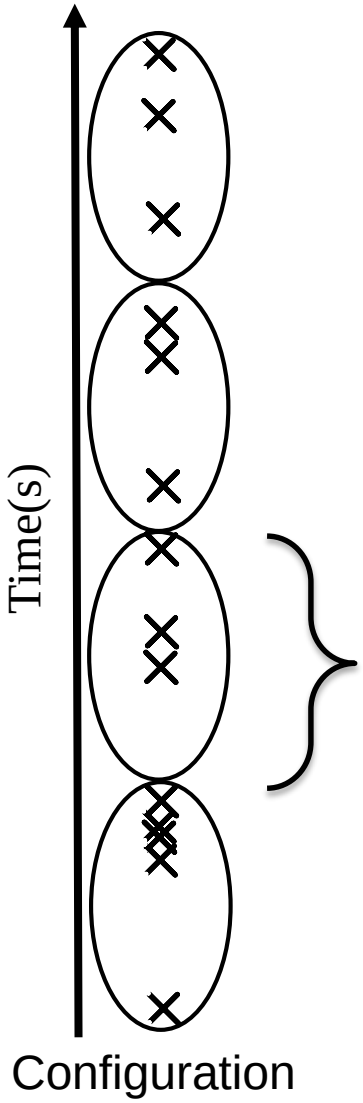


***Root Query Template***:

We call a query template a root query template if no other query template in the same cluster dominates it.

# Workload Compression

## STEP 3: Dropping Queries with Similar Times

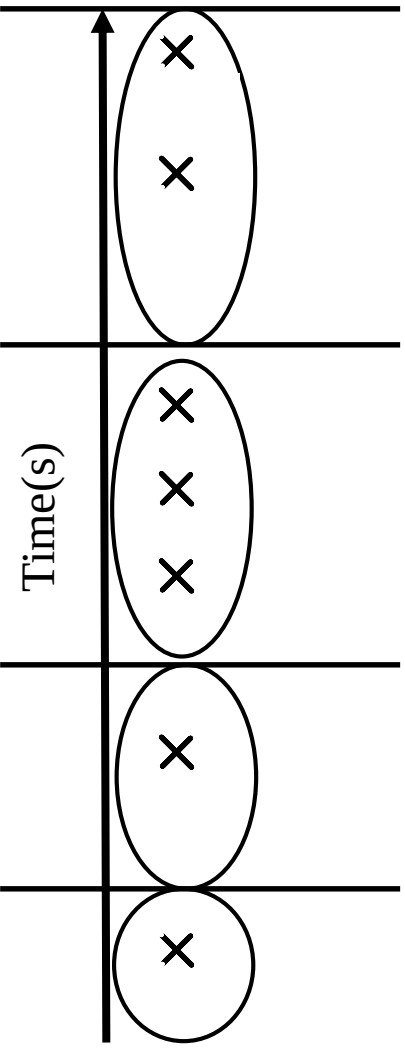


Configuration 1

# Workload Compression

## STEP 3: Dropping Queries with Similar Times

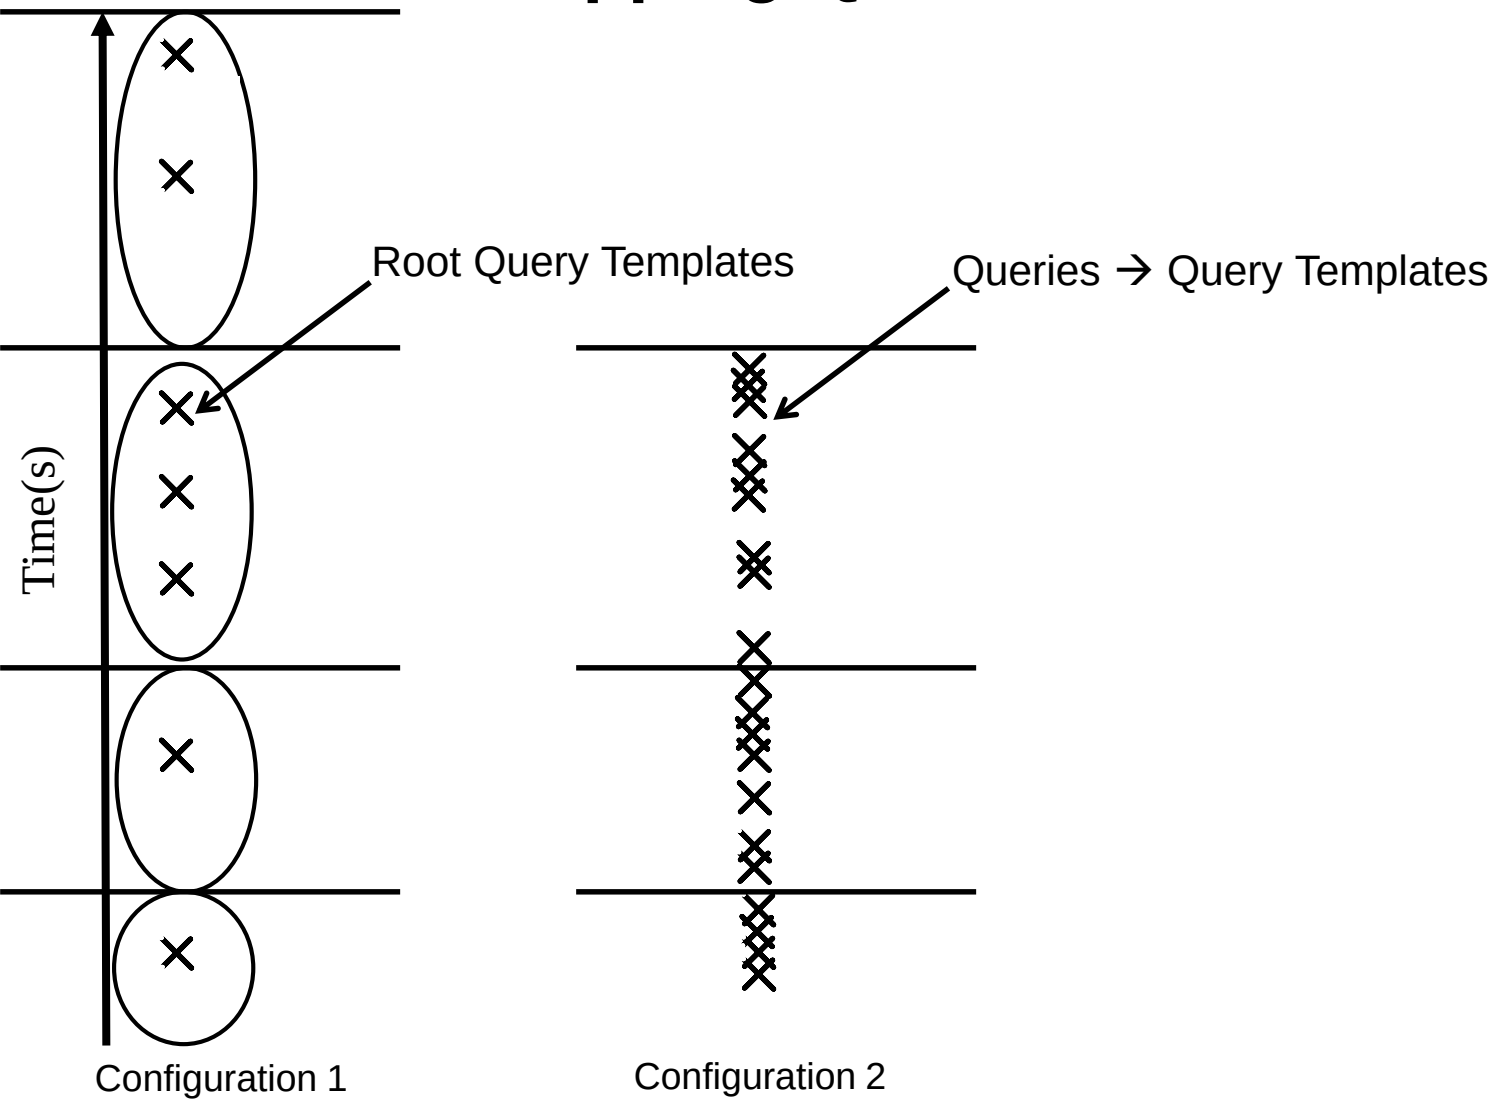

# Workload Compression

## STEP 3: Dropping Queries with Similar Times

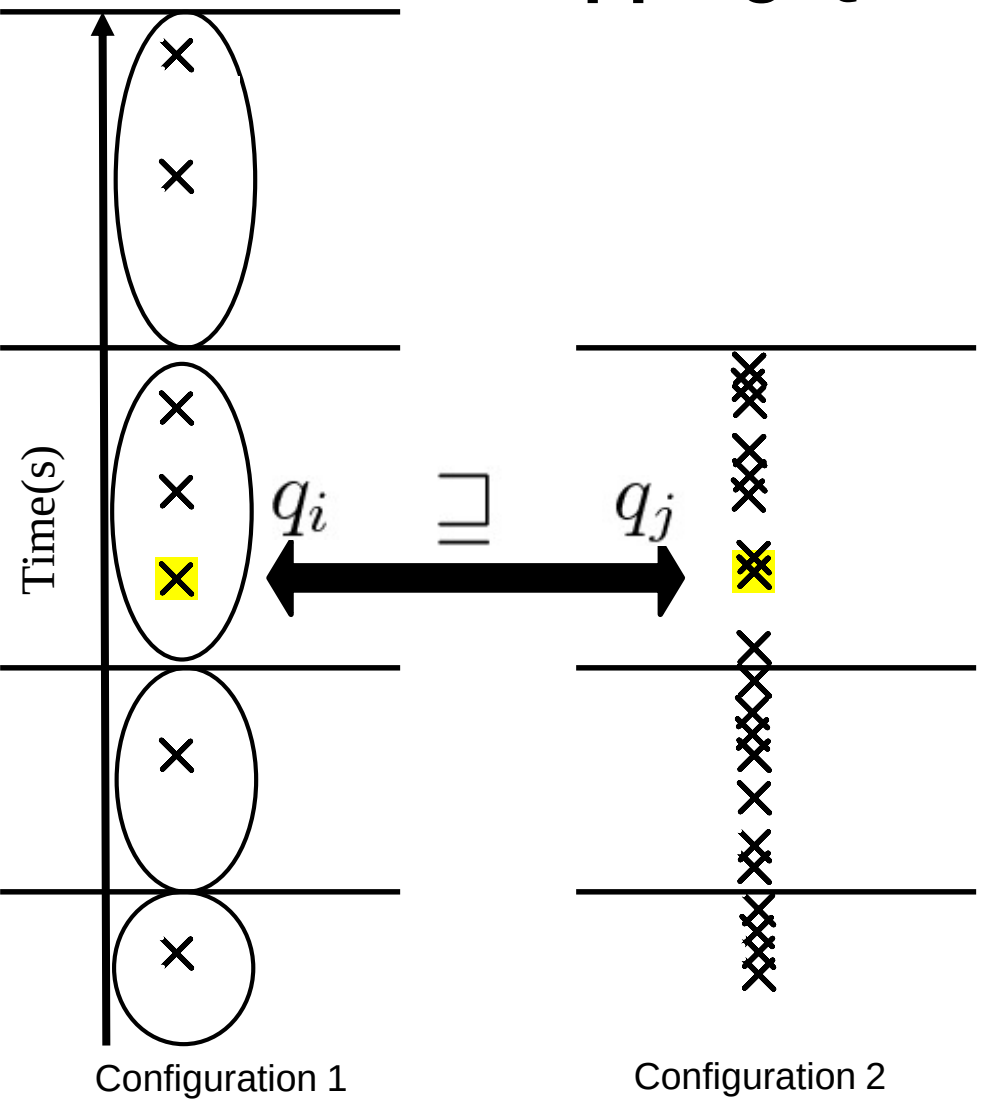

# Workload Compression

## STEP 3: Dropping Queries with Similar Times

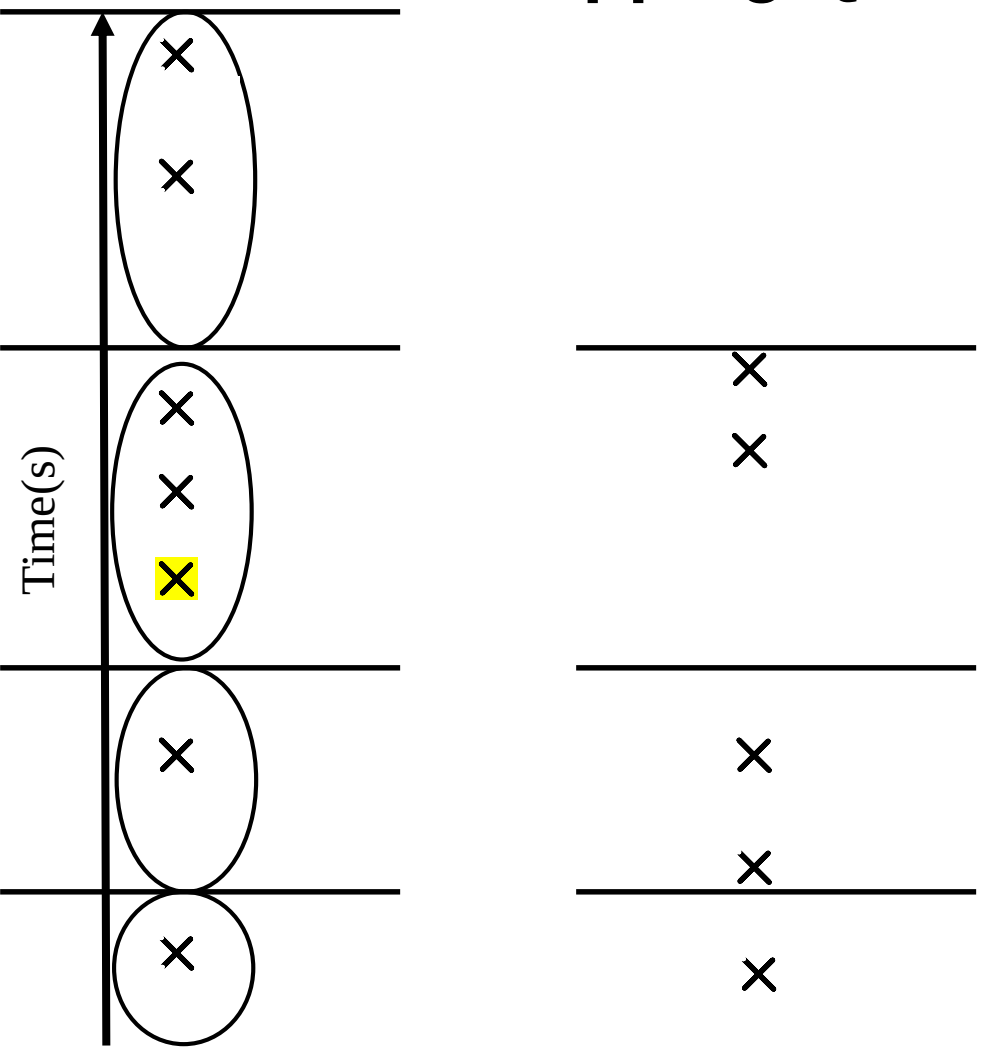


Configuration 1

Configuration 2

# Workload Compression

## STEP 3: Dropping Queries with Similar Times

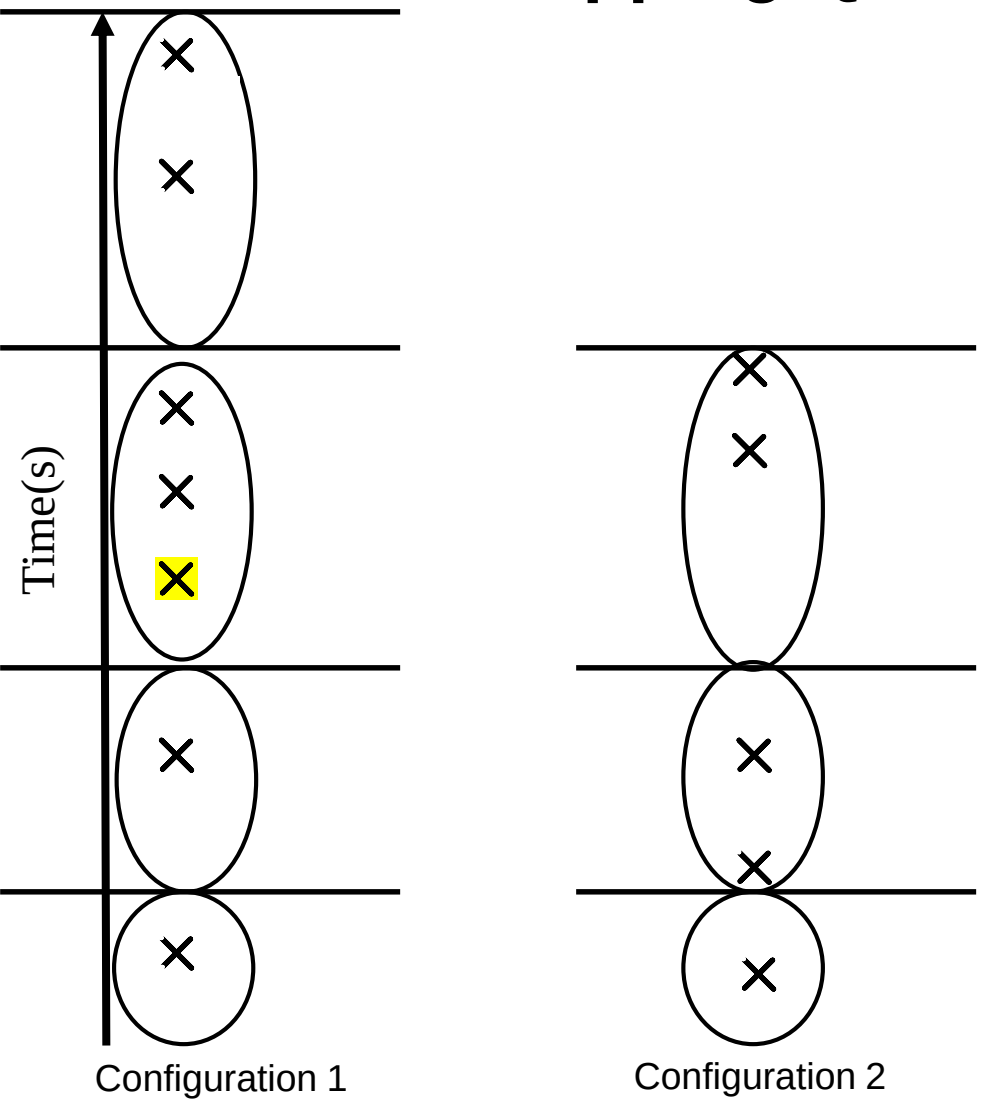

# PSLA Quality Assessment

# PSLA Quality Metrics

- PSLA Query Capabilities
- PSLA Complexity
- PSLA Performance Error Metric

$$RMSE(\{q_1, \ldots, q_k\}, th) = \sqrt{\frac{1}{k} \sum_{i=1}^{k} \left( \frac{q_i - th}{th} \right)^2}$$

# Quality Metric Trade-offs

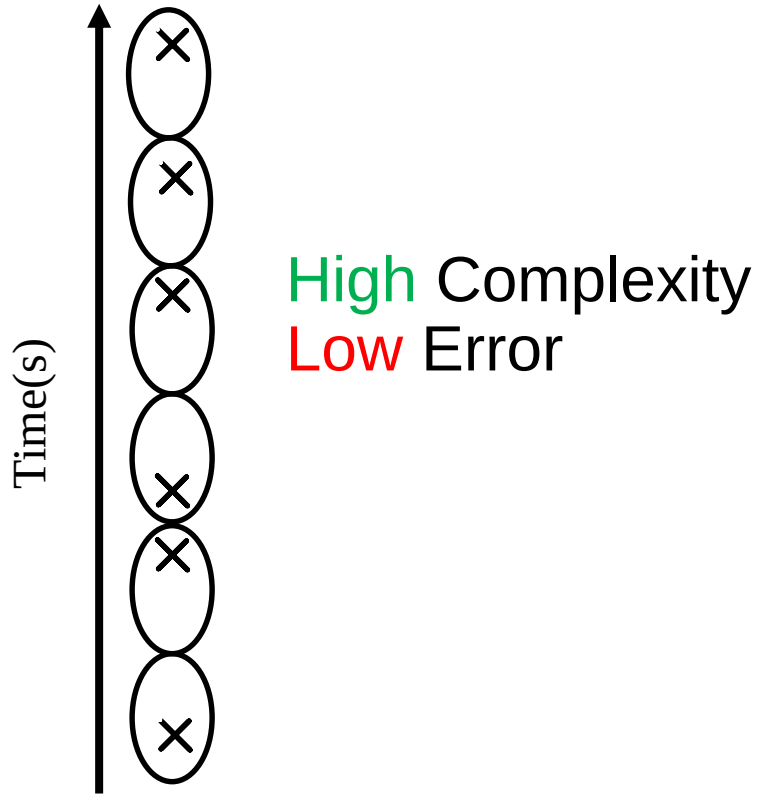


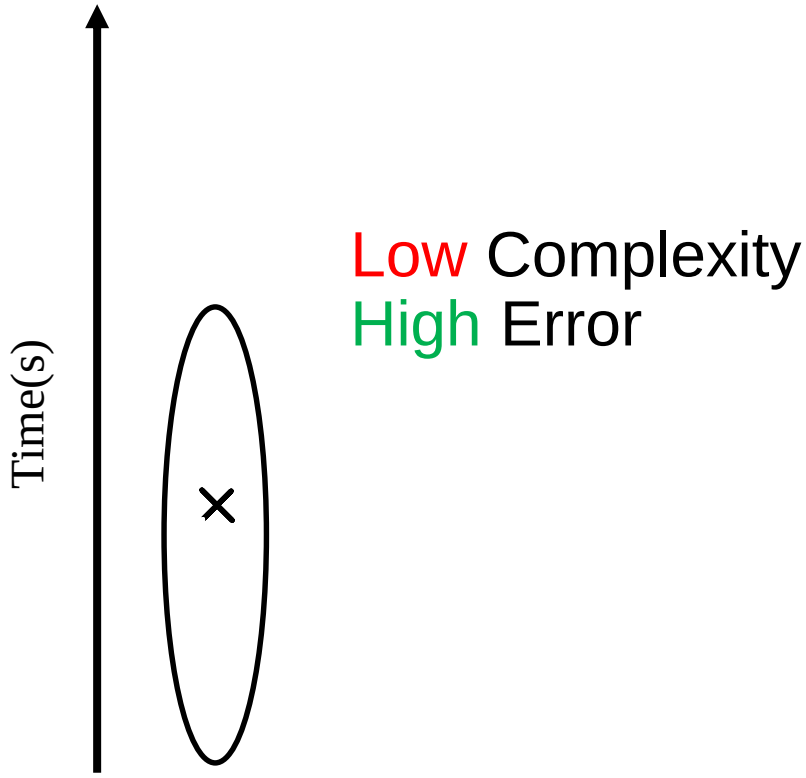


Found that a log interval is best without tuning

# PSLA Evaluation on Predicted Runtimes

# PSLAManager Workflow

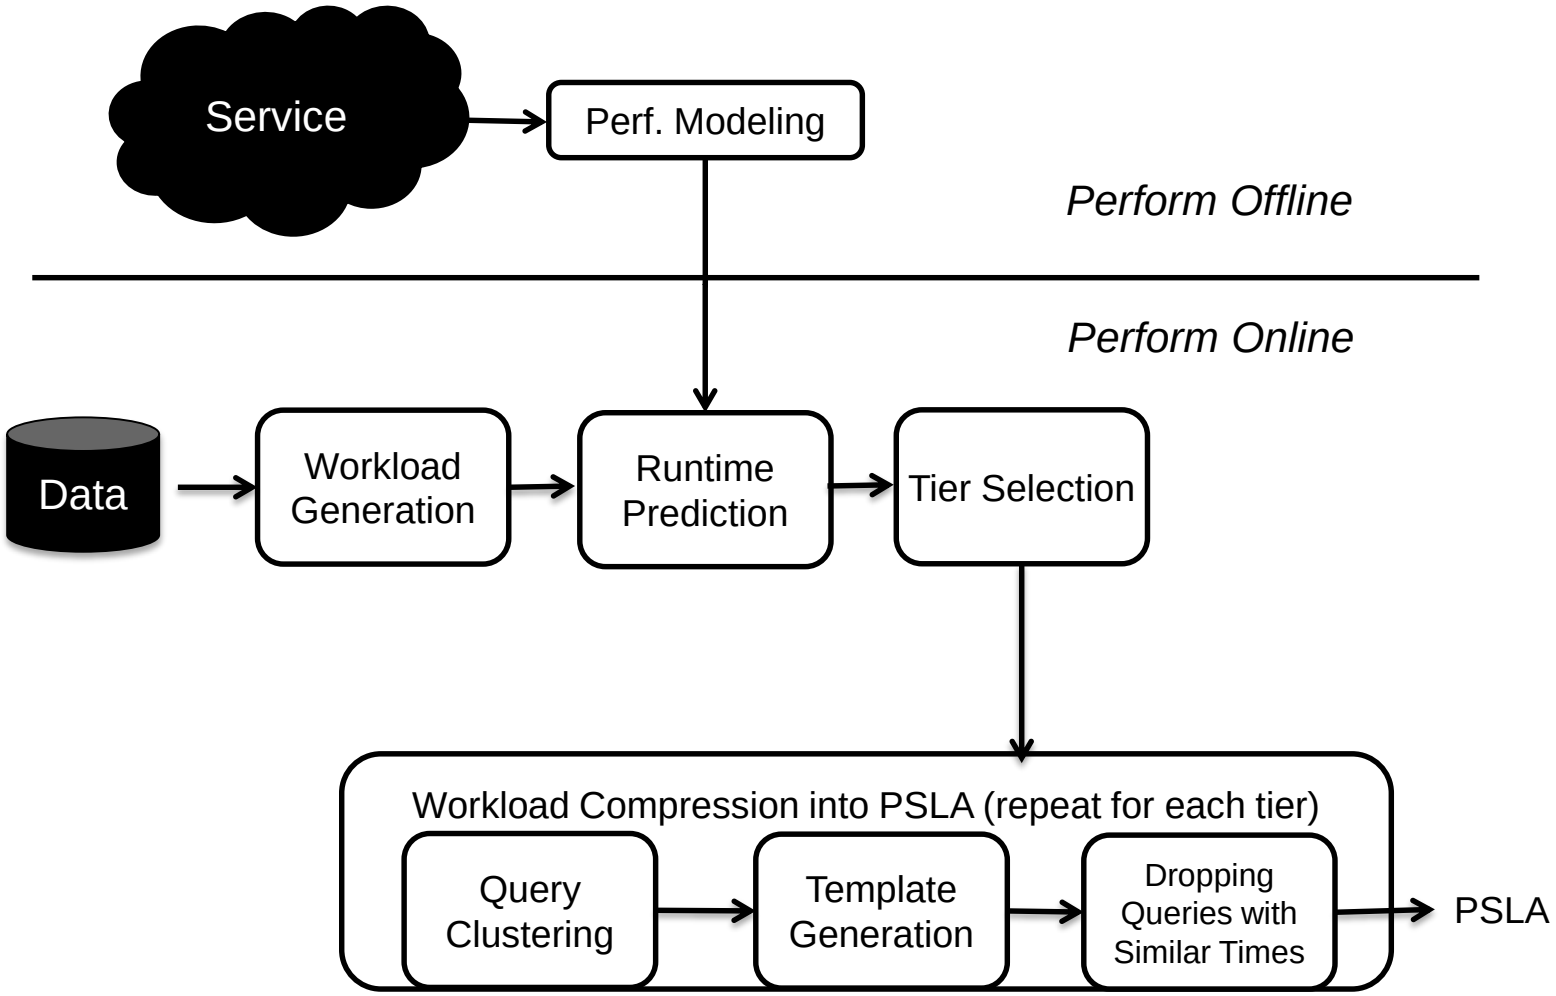

# PSLAManager Workflow

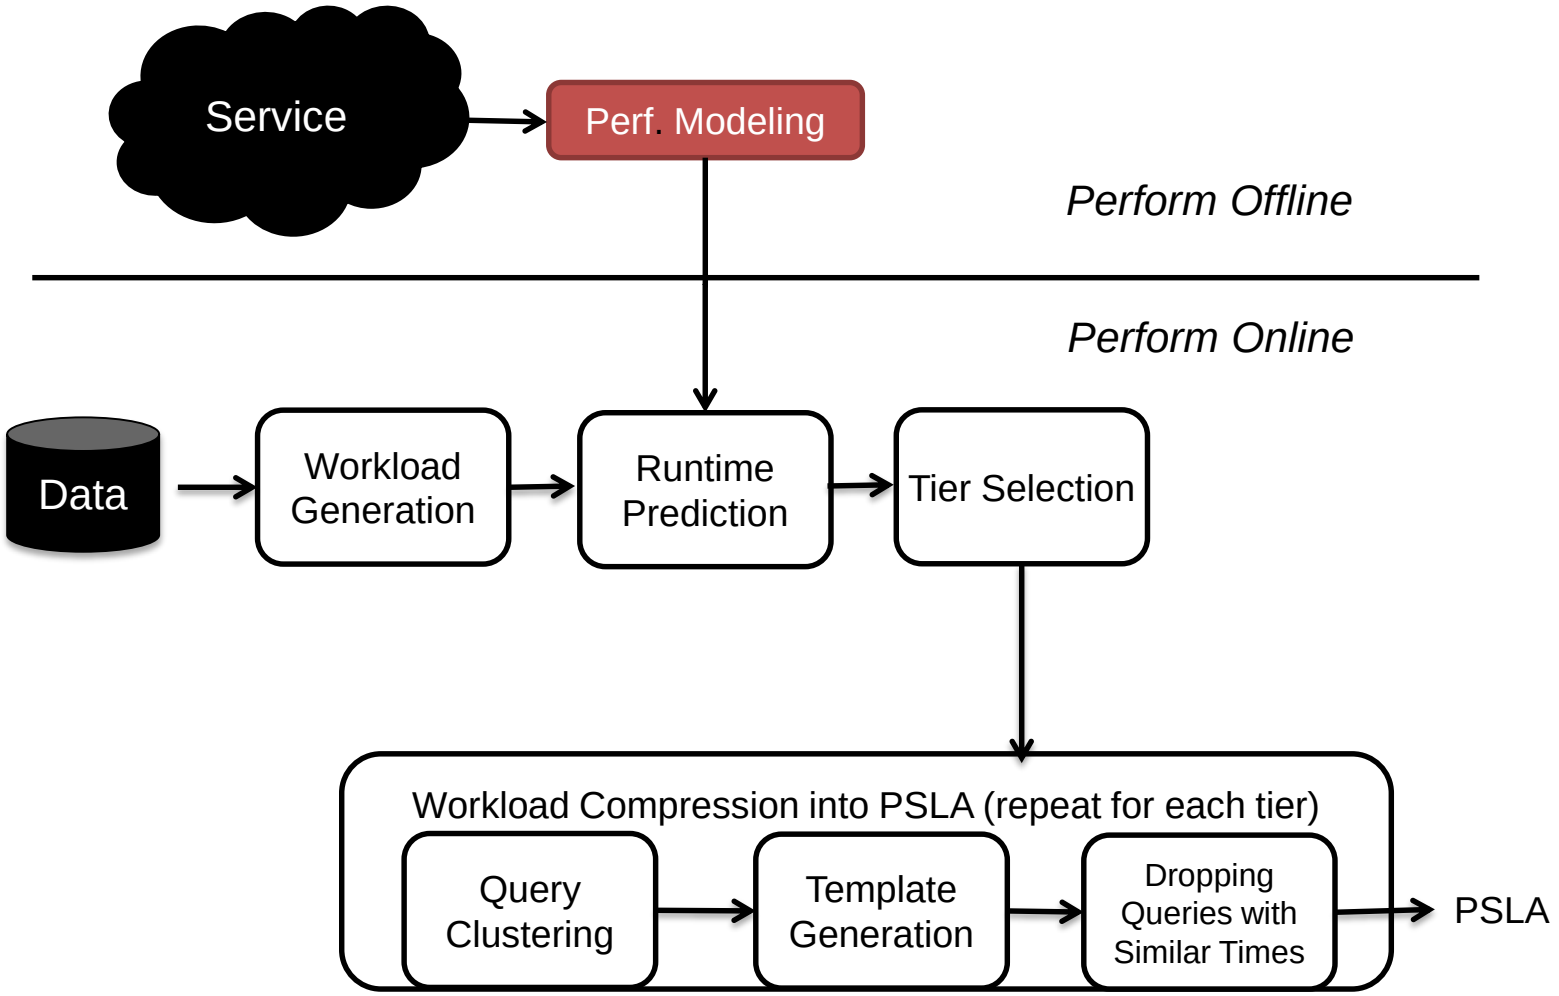

# Performance Model (Offline)

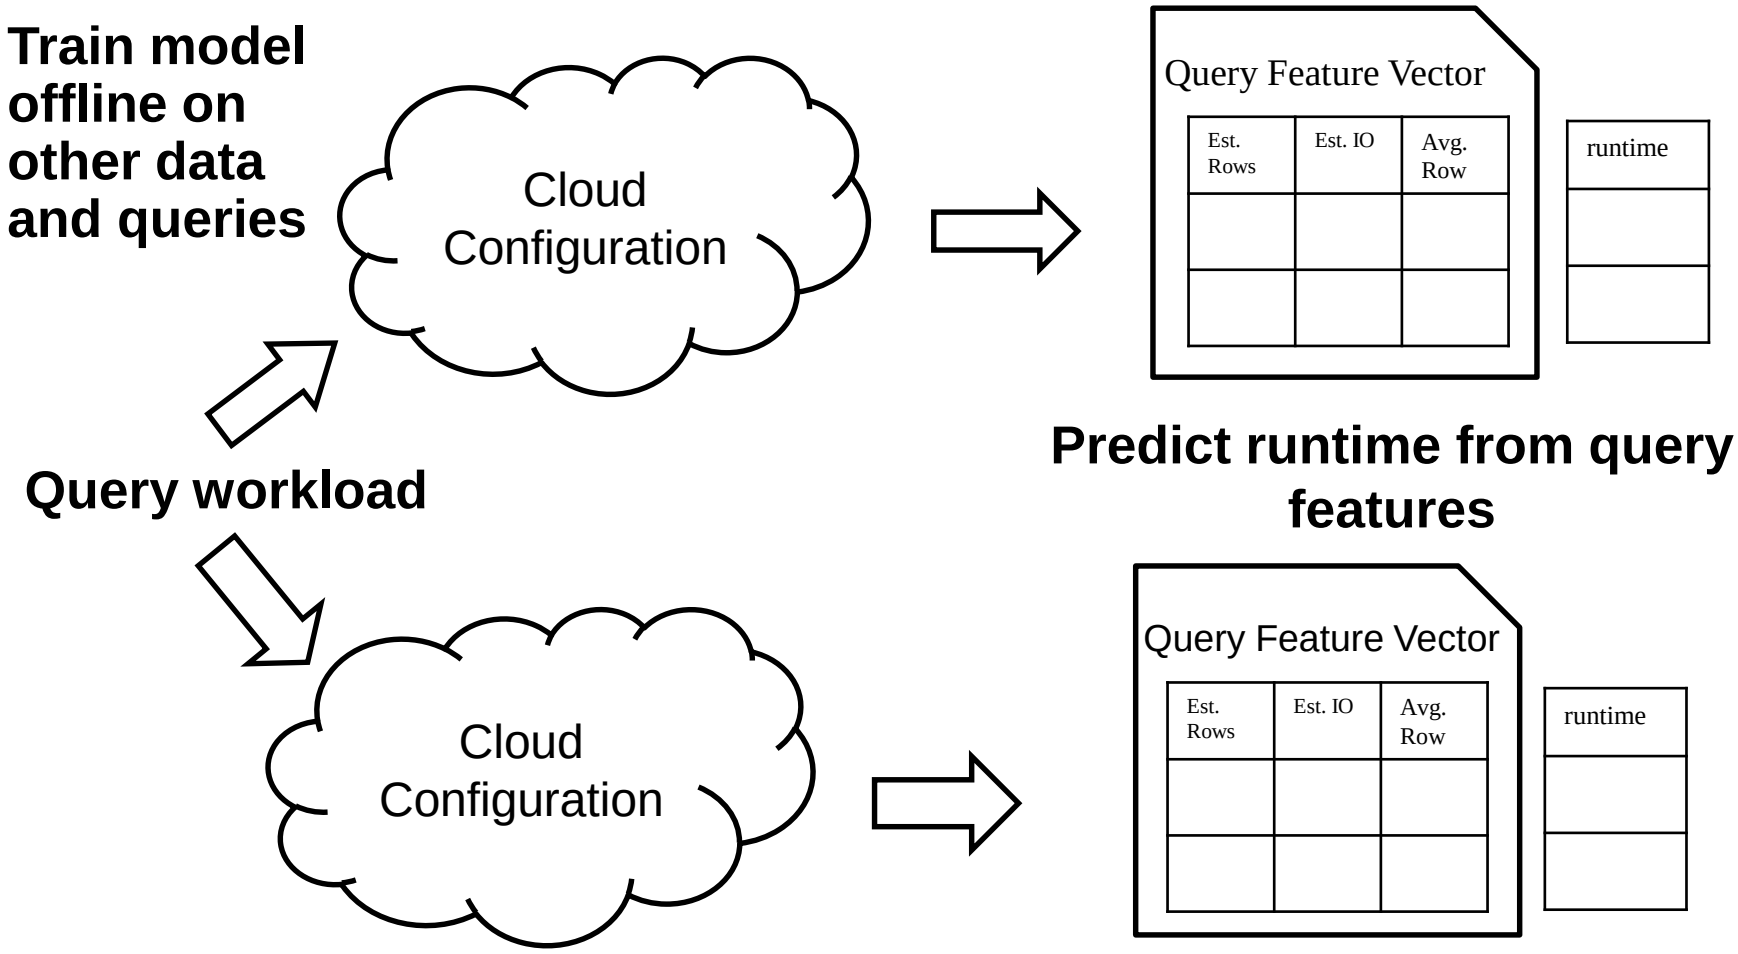


Based on Predicting Multiple Metrics for Queries: Better Decisions enabled by Machine Learning [Ganapathi et. al. 2009]

# Training Dataset

- Synthetic Dataset
  - 10GB
  - 6 Tables
  - 61 Attributes

# PSLAManager Workflow

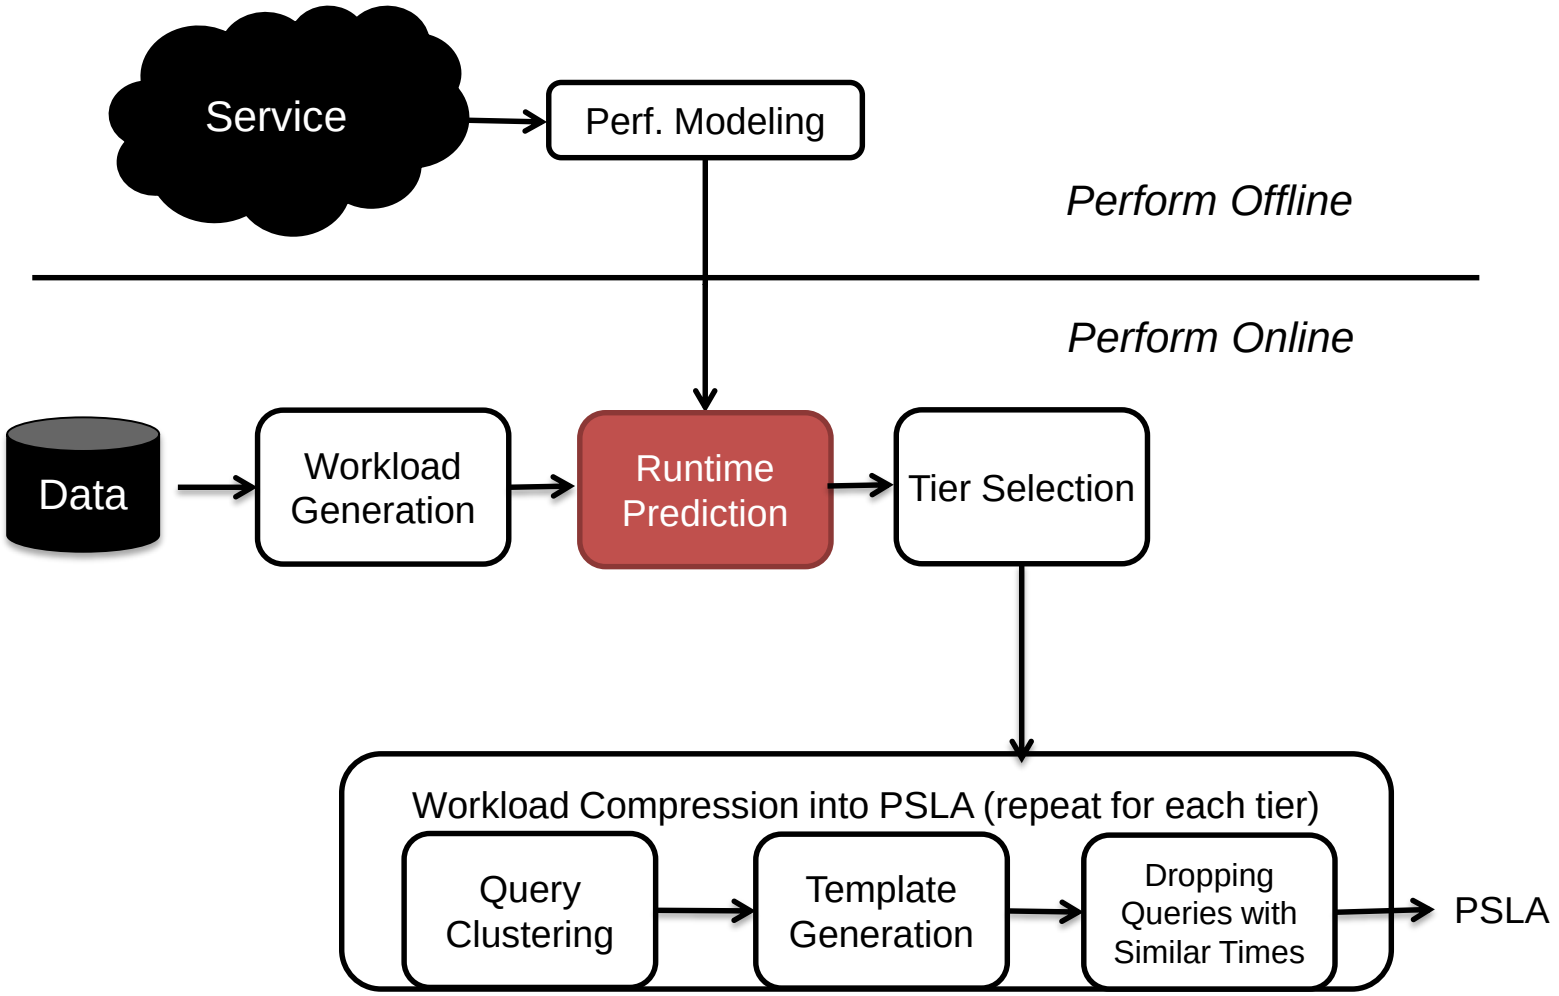

# Predicted Myria PSLA (Predicted Runtimes)

Tier #1

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (9 ATTR.) FROM (PART)<br>SELECT (9 ATTR.) FROM (CUSTOMER)<br>SELECT (17 ATTR.) FROM (DATE)<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 0.1% | 10 |
| SELECT (17 ATTR.) FROM (LINEITEM)<br>SELECT (9 ATTR.) FROM (2 TABLES)<br>SELECT (3 ATTR.) FROM (5 TABLES)<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 10% | 60 |
| SELECT (60 ATTR.) FROM (5 TABLES) | 300 |

Purchase @ $0.16/hour

Tier #2

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (27 ATTR.) FROM (5 TABLES) WHERE 10%<br>SELECT (60 ATTR.) FROM (5 TABLES) WHERE 1% | 10 |
| SELECT (11 ATTR.) FROM (2 TABLES)<br>SELECT (9 ATTR.) FROM (5 TABLES) | 60 |

Purchase @ $0.24/hour

Tier #3

| Query Template | Runtime (seconds) |
|---|---|

Purchase @ $0.32/hour

Tier #4

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (1 ATTR.) FROM (4 TABLES) | 10 |

Purchase @ $0.64/hour

# Looking Forward

- Direct extensions to the approach
  - Add support for indexes
  - Improve time predictions
- Longer-term future work
  - Can we guarantee the runtimes?
  - Can we update the PSLA as user queries data
    - Goal is to show increasingly more complex queries
- Usability testing

# Conclusion

- Many cloud DBMSs exist
- Require users to reason about resources
- We propose to re-think that interface
- Personalized Service Level Agreements
  - Service Tiers
  - Price/Capabilities/Performance
- Important direction for cloud DBMSs