Self-Driving Database Management Systems

CIDR 2017

@andy_pavlo

CARNEGIE MELLON
DATABASE GROUP

1980s

Timothy Pavlo

1950s

Joseph Pavlo

1920s

Cornelius Von Pavlo

2015 Median DBA Salary

$81,710

[Source]

Possible

» Physical Database Design

» Resource Allocation

» Query Optimization & Tuning

» Knob Configuration

What's Different?

» Previous tools only dealt with handling problems in the past.

» Humans still make final decisions.

» Hardware & algorithm advancements.

Peloton

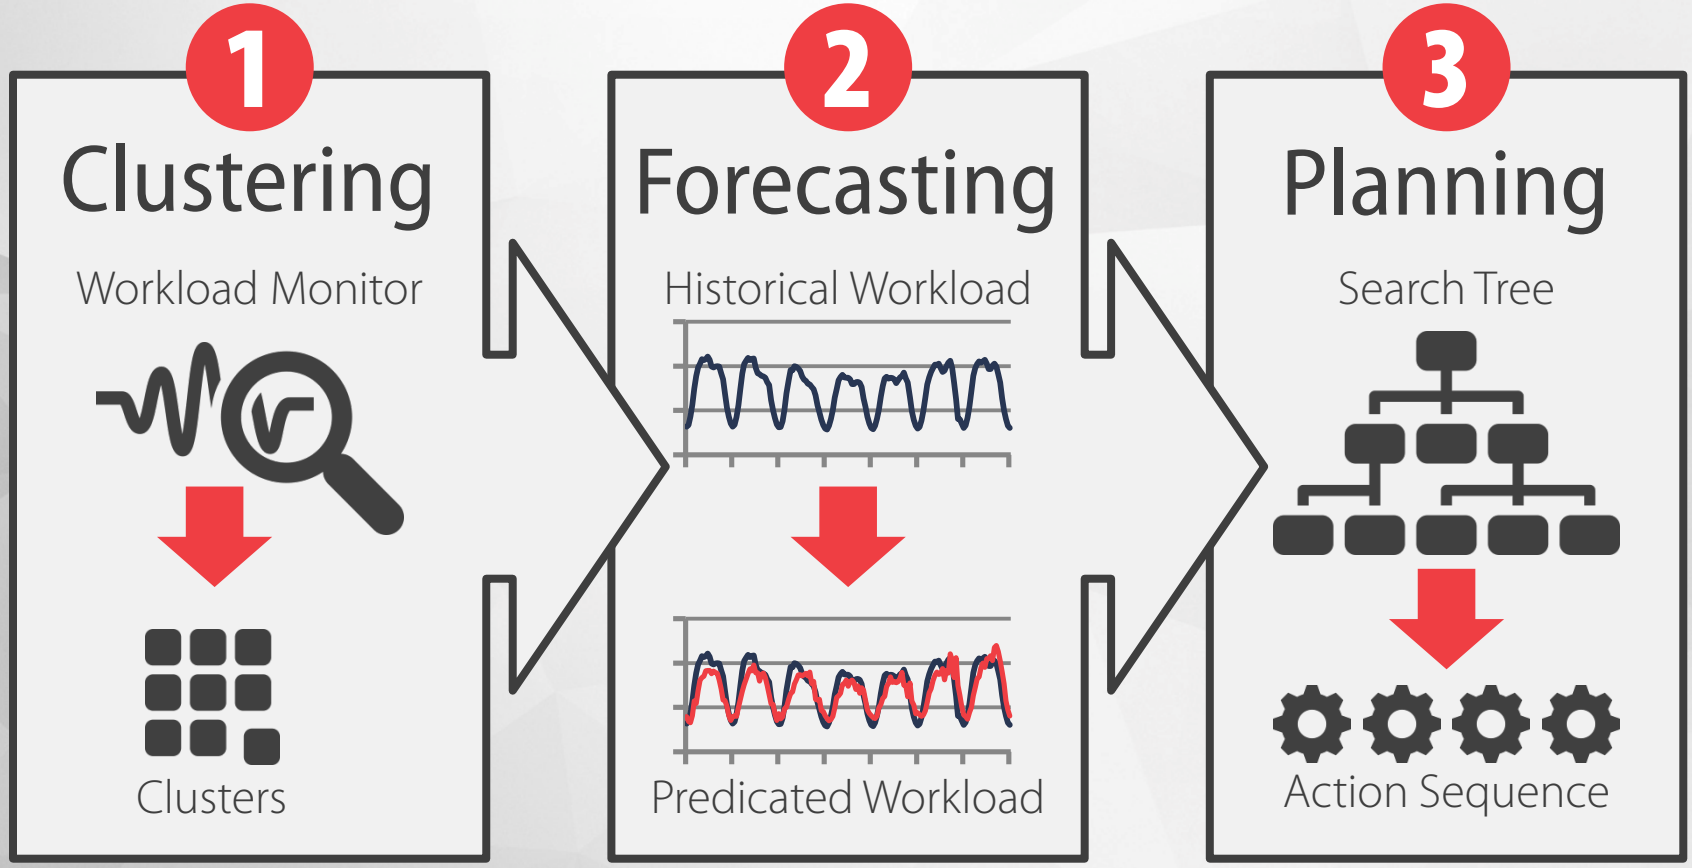
1
Clustering
Workload Monitor
Clusters
2
Forecasting
Historical Workload
Predicated Workload
3
Planning
Search Tree
Action Sequence

#1 – Clustering

» Group similar queries together to improve the forecasting models.

» Logical vs. Physical Features

```
SELECT C_ID
  FROM CUSTOMER
 WHERE C_W_ID = ?
   AND C_D_ID = ?
   AND C_LAST = ?
 ORDER BY C_FIRST
```

Logical Features

```
table={CUSTOMER}
attributes={C_ID,C_W_ID,C_D_ID,C_LAST}
orderby={C_FIRST}
aggregate={∅}
```

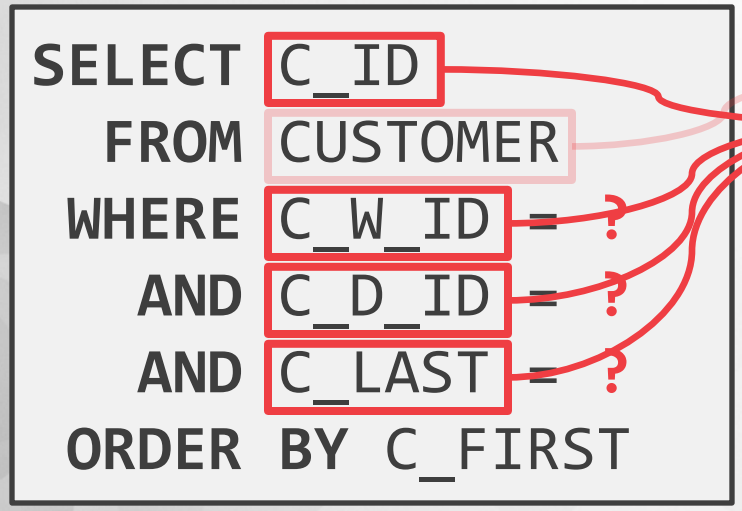

Logical Features

```
table={CUSTOMER}
attributes={C_ID,C_W_ID,C_D_ID,C_LAST}
orderby={C_FIRST}
aggregate={∅}
```

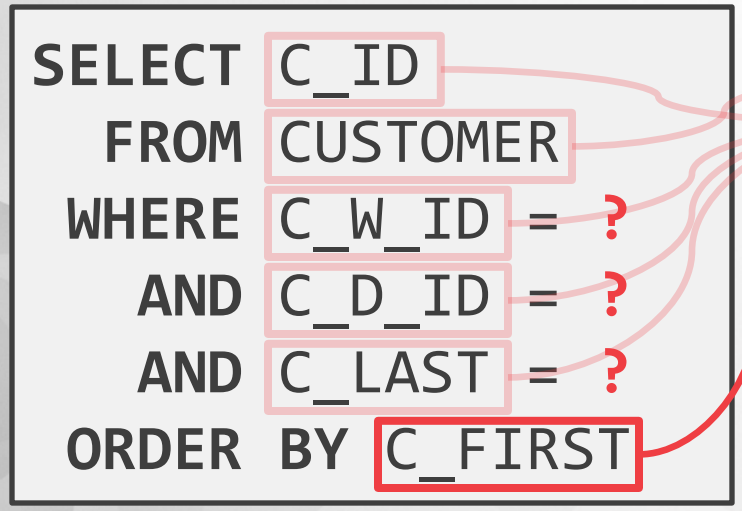

Logical Features

```
table={CUSTOMER}
attributes={C_ID,C_W_ID,C_D_ID,C_LAST}
orderby={C_FIRST}
aggregate={∅}
```

```
SELECT C_ID
  FROM CUSTOMER
 WHERE C_W_ID = ?
   AND C_D_ID = ?
   AND C_LAST = ?
 ORDER BY C_FIRST
```

Logical Features

```
table={CUSTOMER}
attributes={C_ID,C_W_ID,C_D_ID,C_LAST}
orderby={C_FIRST}
aggregate={∅}
```

Physical Features

```
tuplesRead={##}
tuplesWritten={##}
cpu={##}
memory={##}
```

```
lockWait={##}
indexPages={##}
networkRead={##}
networkWritten={##}
```

- ➕ Fixed/Immutable
- ➕ Cheap to Compute
- ➖ Lacks Execution Info

Logical Features

```
table={CUSTOMER}
attributes={C_ID,C_W_ID,C_D_ID,C_LAST}
orderby={C_FIRST}
aggregate={∅}
```

- ➕ Descriptive
- ➕ Identifies Problems
- ➖ Unstable/Changes

Physical Features

tuplesRead={##}	lockWait={##}
tuplesWritten={##}	indexPages={##}
cpu={##}	networkRead={##}
memory={##}	networkWritten={##}

#2 – Forecasting

» Generate forecasting models for each cluster to predict future arrival rate.

» Multiple horizons & intervals.

TensorFlow

- Linear Regression
- LSTM RNN

Real Workload

Predicted Workload

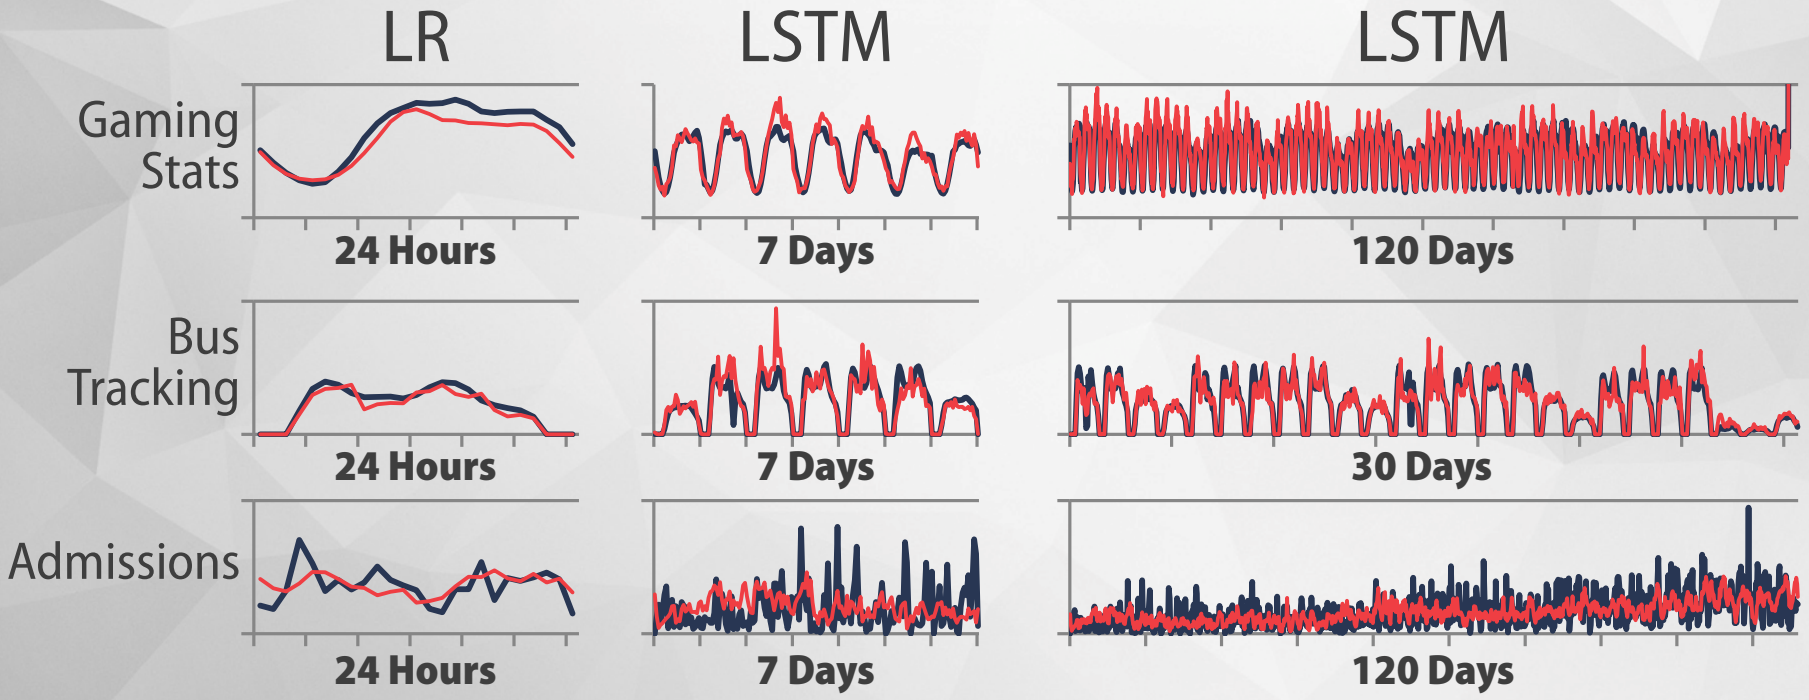

#3 – Planning

» Generate optimization actions for the DBMS based on the workload forecasts.

» Select a sequence of actions that optimize the target metric.

Action Catalog

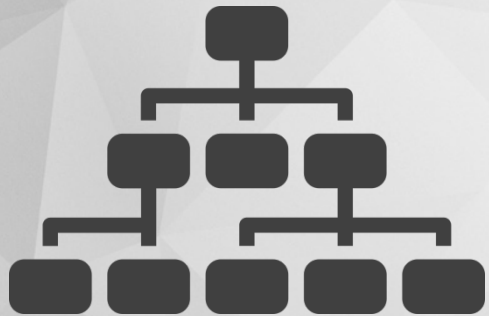

Search Tree

Action Sequence

Action Catalog

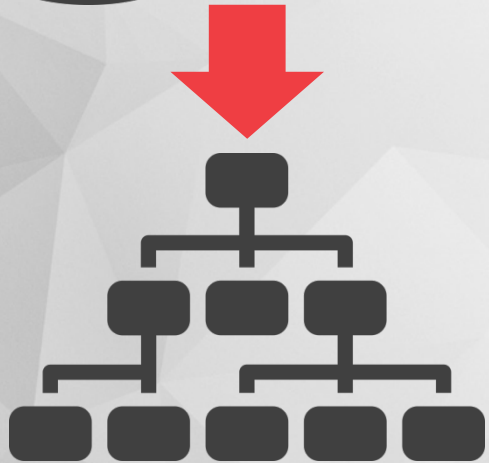

Search Tree

Benefit

Cost

Action Sequence

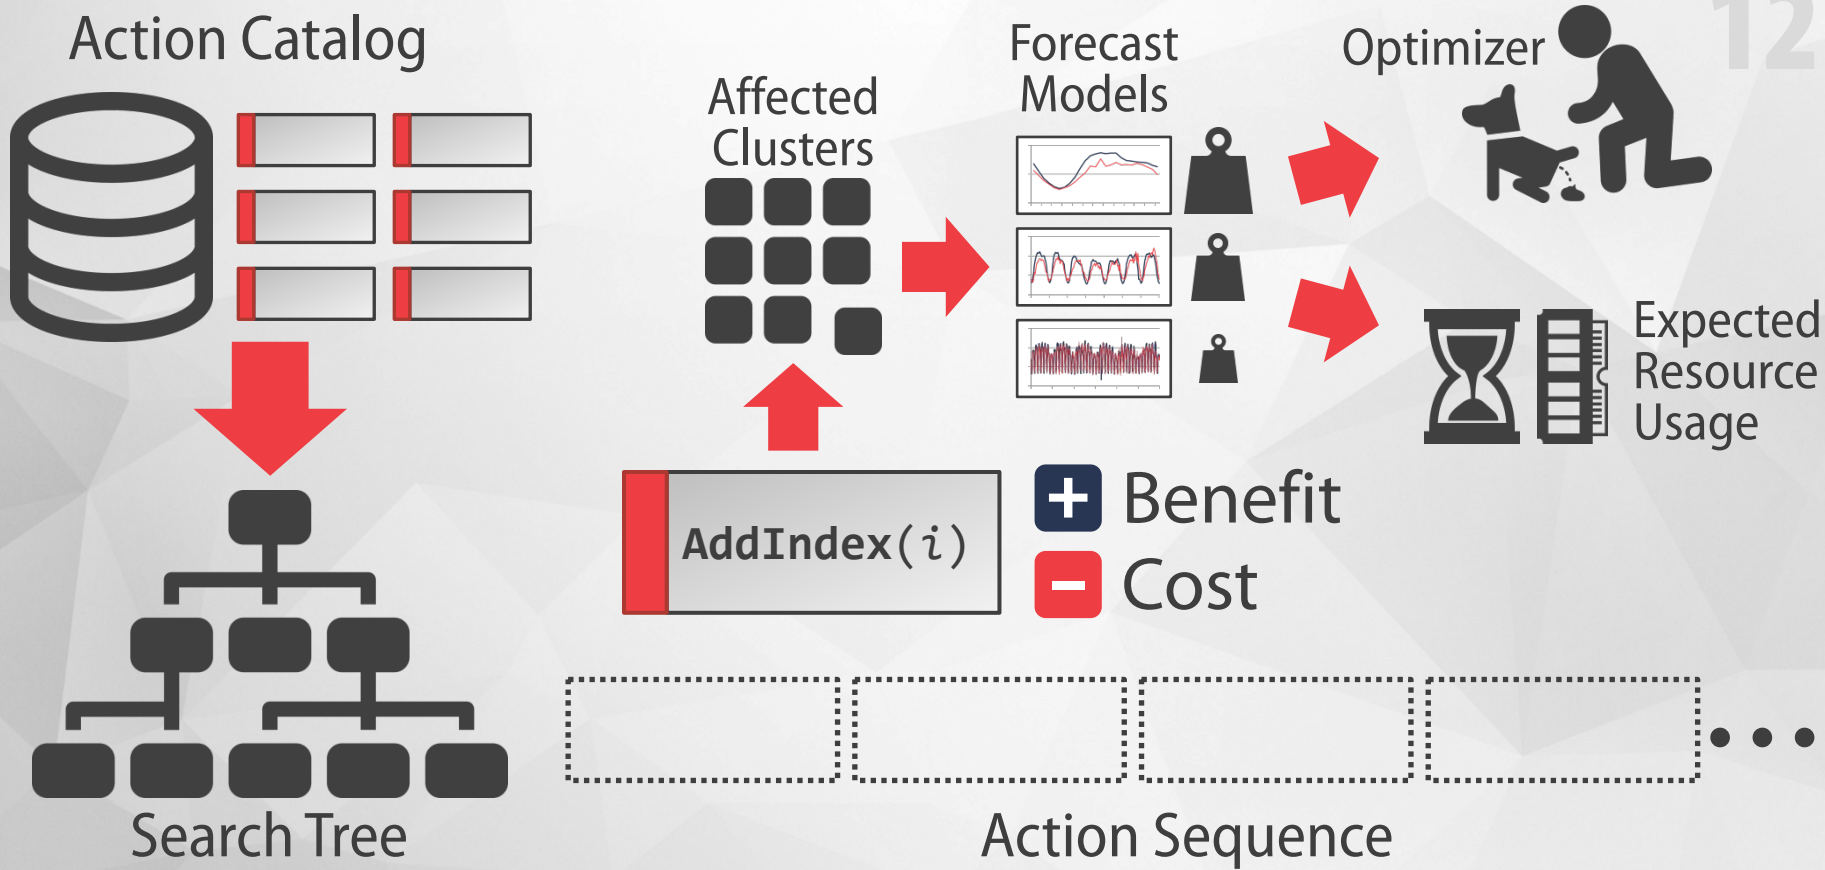
Action Catalog
Affected Clusters
Forecast Models
Optimizer
Expected Resource Usage
AddIndex(i)
Benefit
Cost
Search Tree
Action Sequence

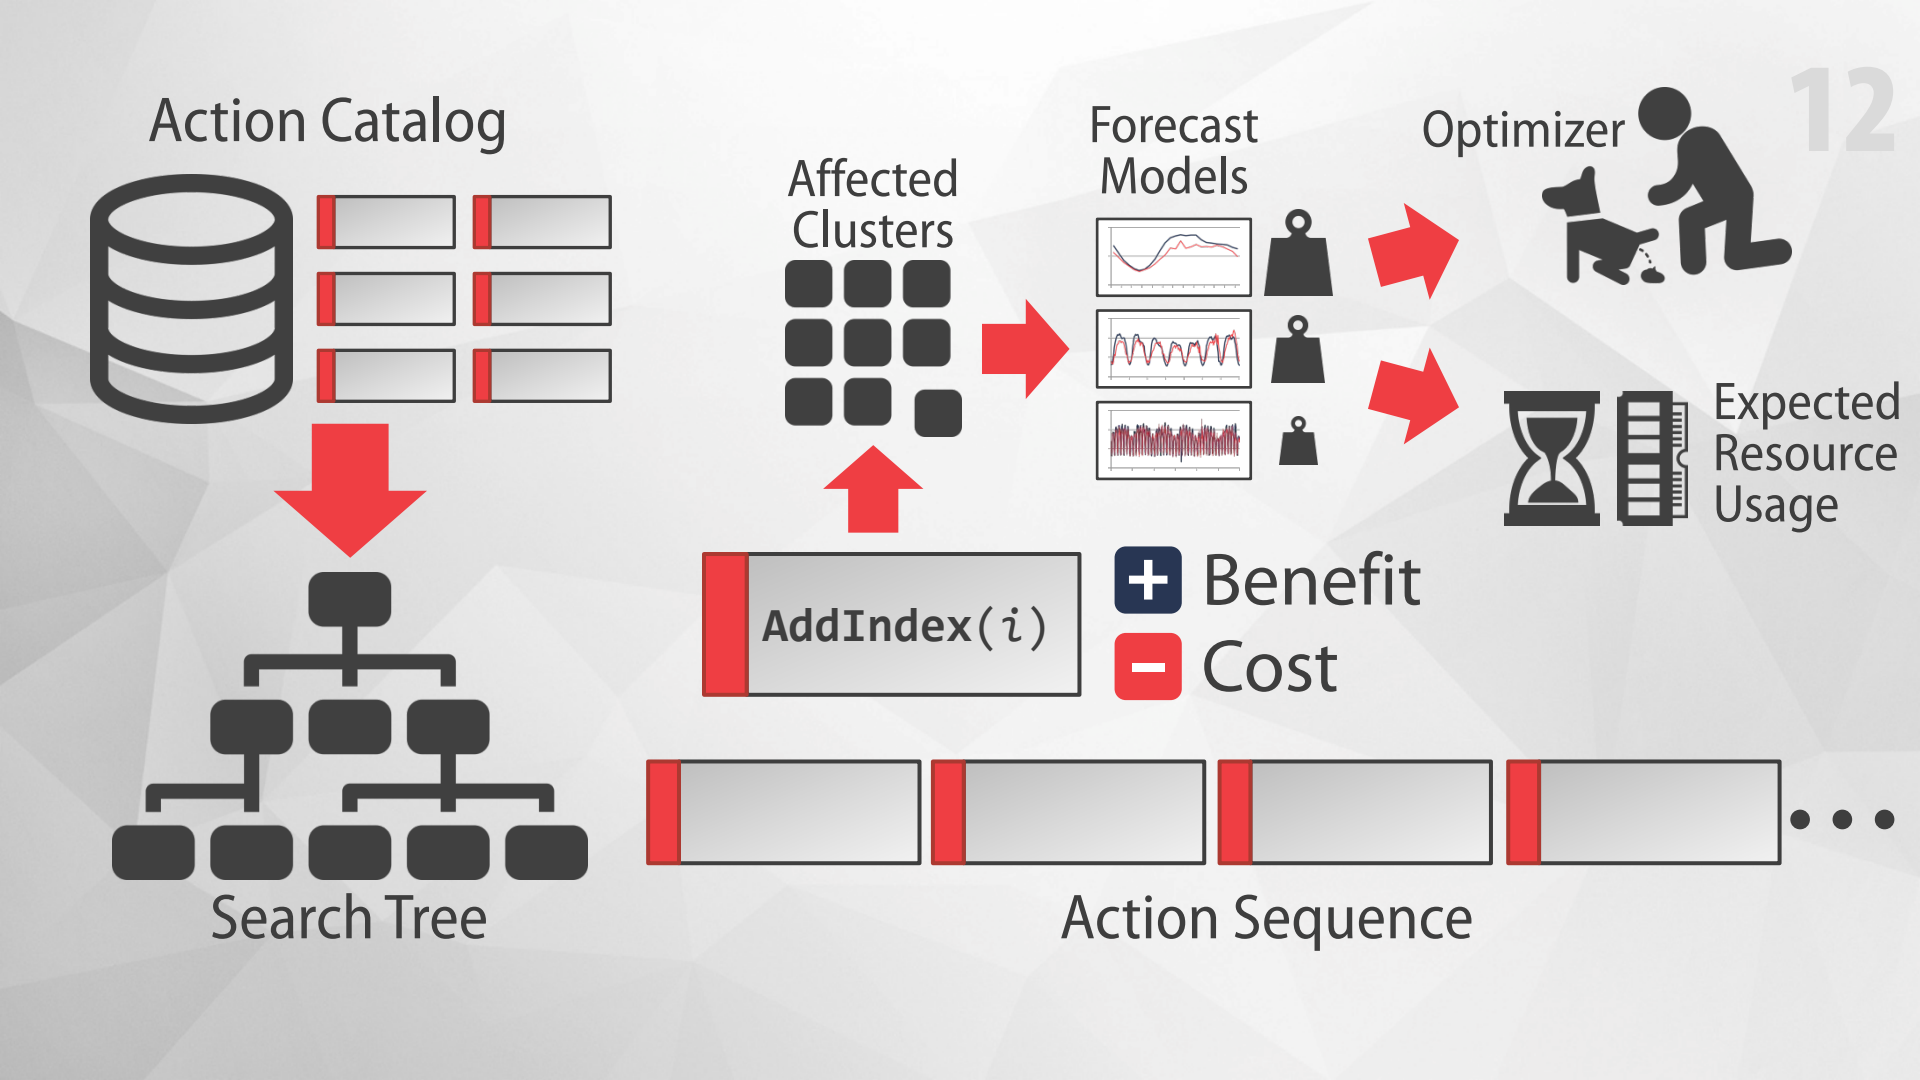
Action Catalog
Affected Clusters
Forecast Models
Optimizer
Expected Resource Usage
AddIndex(i)
Benefit
Cost
Search Tree
Action Sequence

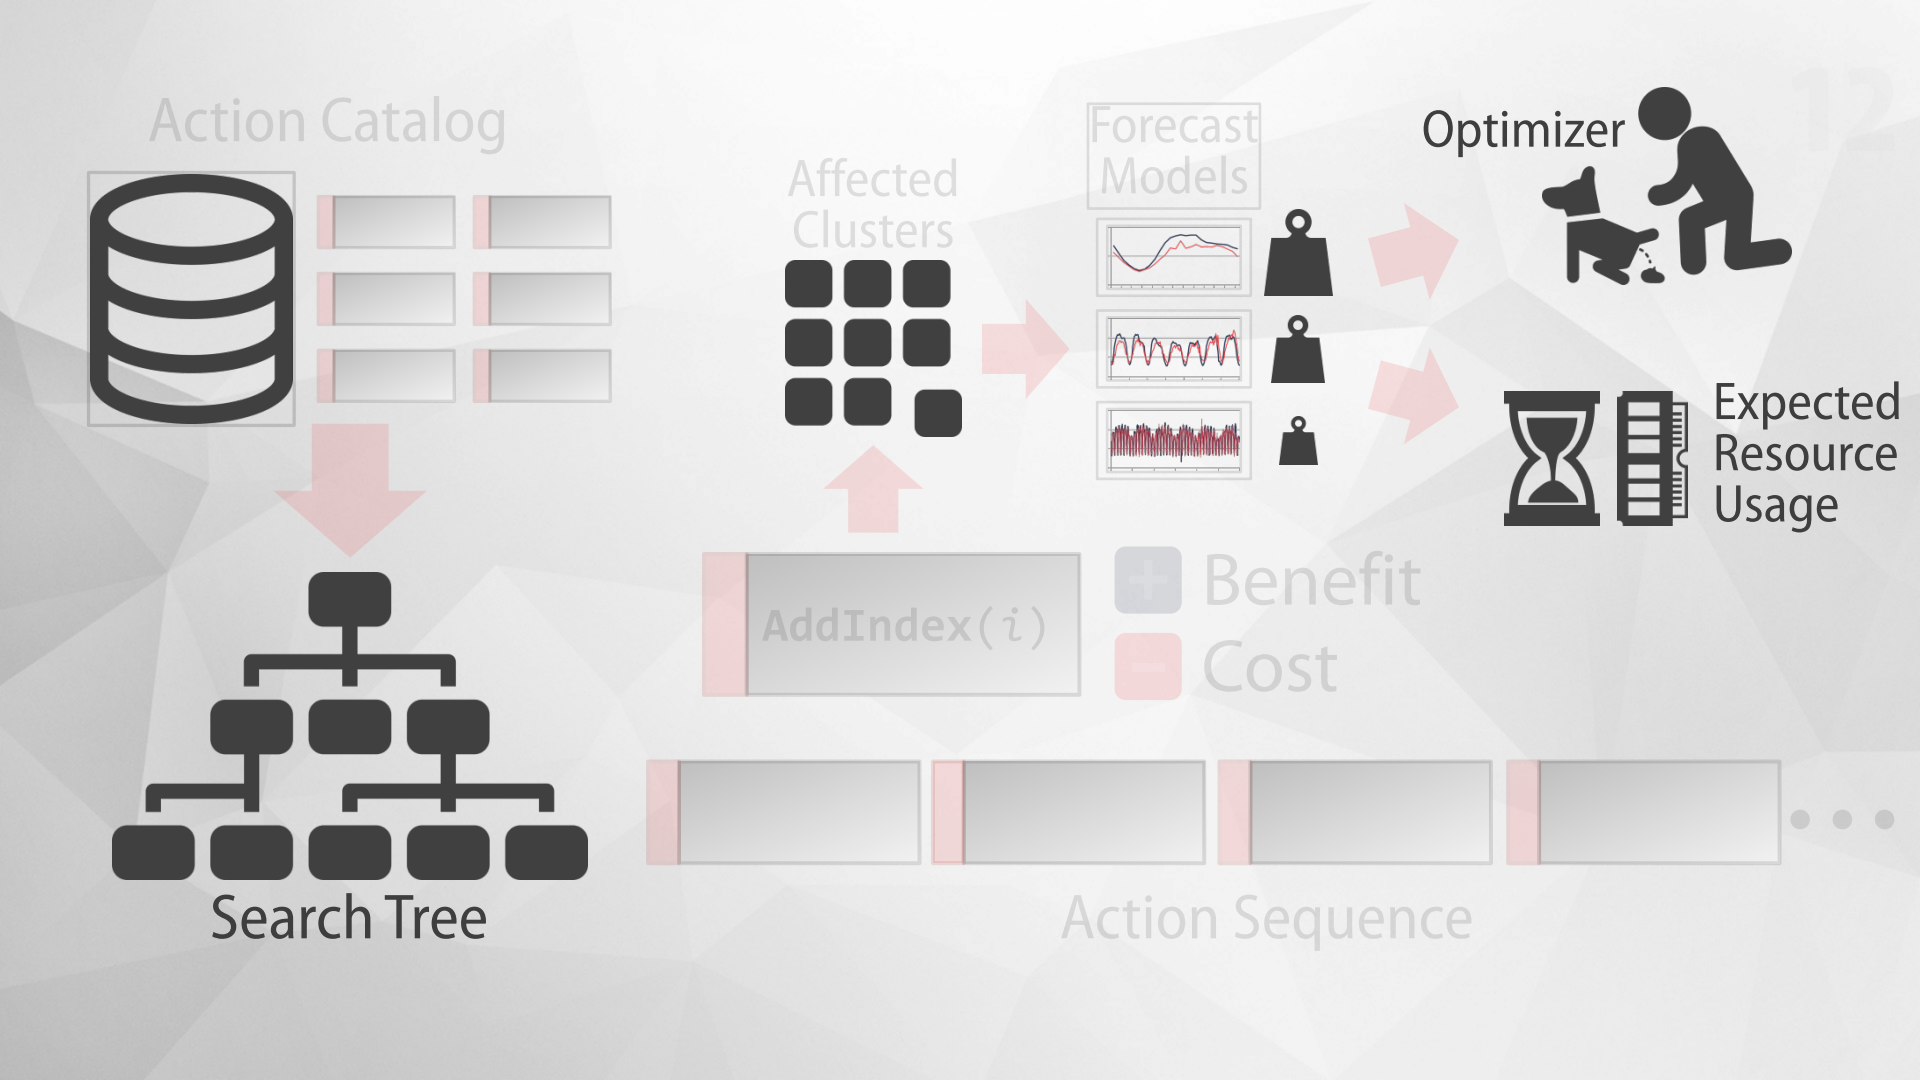
Action Catalog
Affected Clusters
Forecast Models
Optimizer
Expected Resource Usage
AddIndex(i)
Benefit
Cost
Search Tree
Action Sequence

Demo

» **Peloton** (v2017-01)

» TPC-C with 100 warehouses

» Database loaded without indexes

Current Status

» Clusters/forecasts computed off-line.

» No universal planning algorithm.

» We lost our catalog, planner, and optimizer in the “purge”.

2016

In-Memory / NVM Storage

Open Bw-Tree

WAL (SSD) / WBL (NVM)

Index / Layout Tuning

Apache v2.0 License

2017

More Self-Driving

TensorFlow Integration

LLVM Execution Engine

Cascades Optimizer

Intra-Query Parallelism

Unsolved Problems

» Cluster Prioritization (OLTP vs. OLAP)

» Self-Driving Components Interference

» Human Interactions

» “Traditional” ML Problems


SOLID STATE
EST. 2011
UNIVERSITY

Peloton

http://pelotondb.io