# The Non-Expert Tax: Quantifying the cost of auto-scaling in Cloud-based data stream analytics

**Yuanli Wang**, Baiqing Lyu, Vasiliki Kalavri

yuanliw@bu.edu

https://sites.bu.edu/casp/


# You are overpaying for your streaming services

- We conducted an empirical study of auto-scaling of major providers and found that…

# You are overpaying for your streaming services

- We conducted an empirical study of auto-scaling of major providers and found that…

- Users are over-charged on under-utilized resources
  - Up to 544% for short-term jobs
  - Up to 332% per month for periodic workloads
- Can not accurately identify bottlenecks even after more than 1h
  - Can not achieve target input rate

# Dataflow stream processing

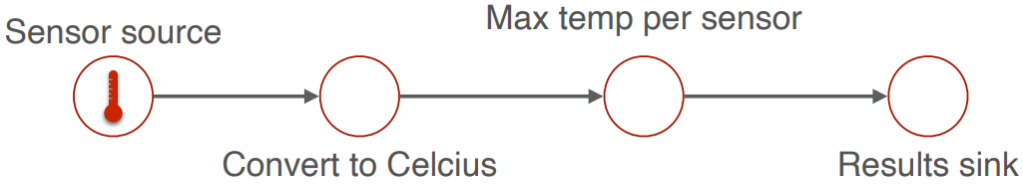


## Logical Dataflow

- Streaming dataflow programs describing business logic
- Vertices: operators
- Edges: data dependencies

# Dataflow stream processing

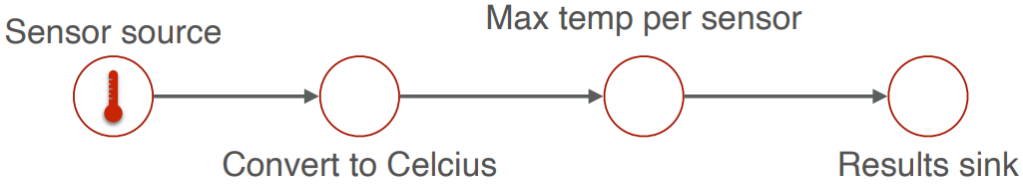


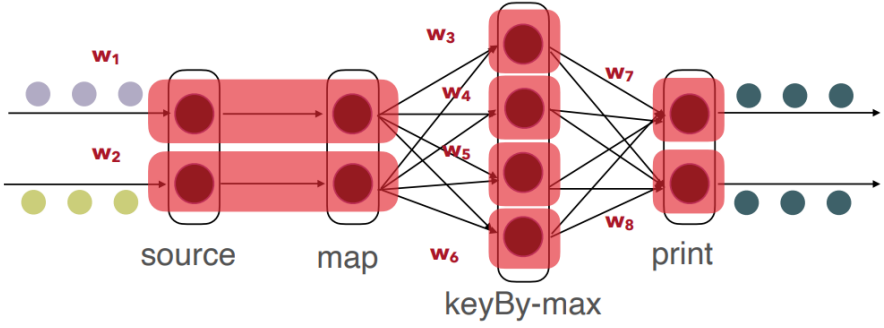


## Logical Dataflow

- Streaming dataflow programs describing business logic
- Vertices: operators
- Edges: data dependencies

## Physical Dataflow

- Decide parallelism of each operator

# Dataflow stream processing

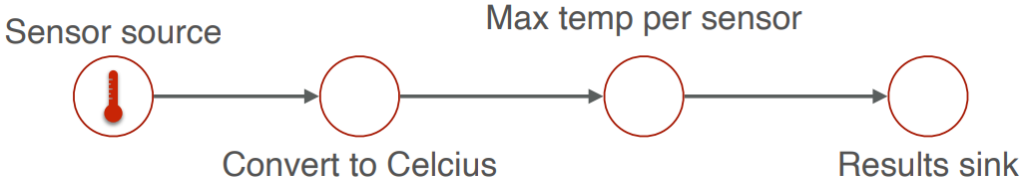


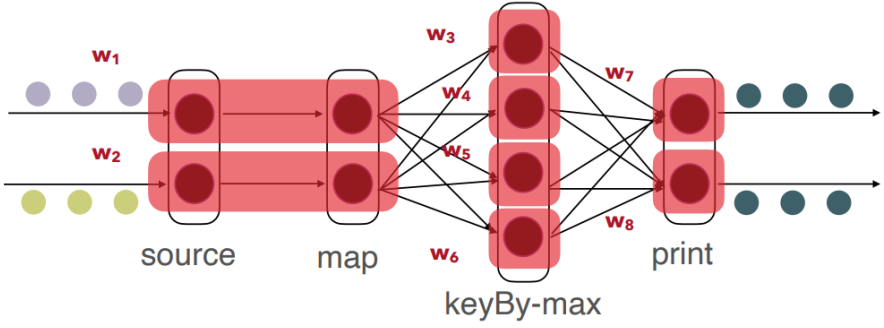


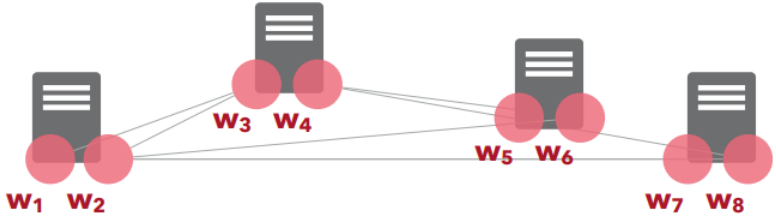


## Logical Dataflow

- Streaming dataflow programs describing business logic
- Vertices: operators
- Edges: data dependencies

## Physical Dataflow

- Decide parallelism of each operator

## Deployment

- Assign tasks to workers

# Main Challenge:
# Workloads are dynamic and often unpredictable

Varying workloads

**Dynamic conditions will eventually render any initial configuration out-of-tune**

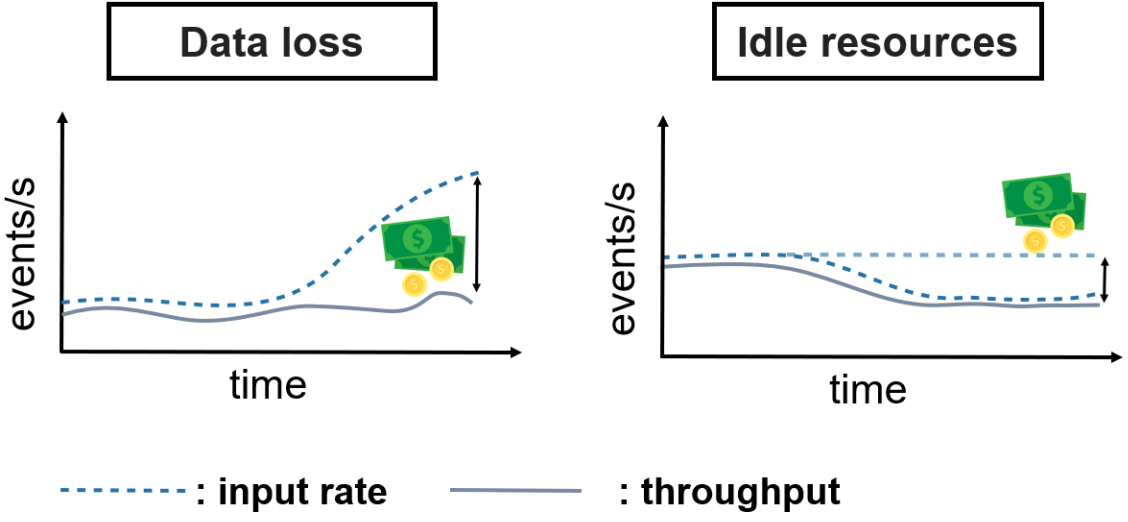

# Data Stream Analytics as managed services

- Enable non-expert users to run stream processing jobs
  - Users submit their business logic
  - The Cloud-hosted service automatically allocates and adjusts resources according to the dynamic workload.

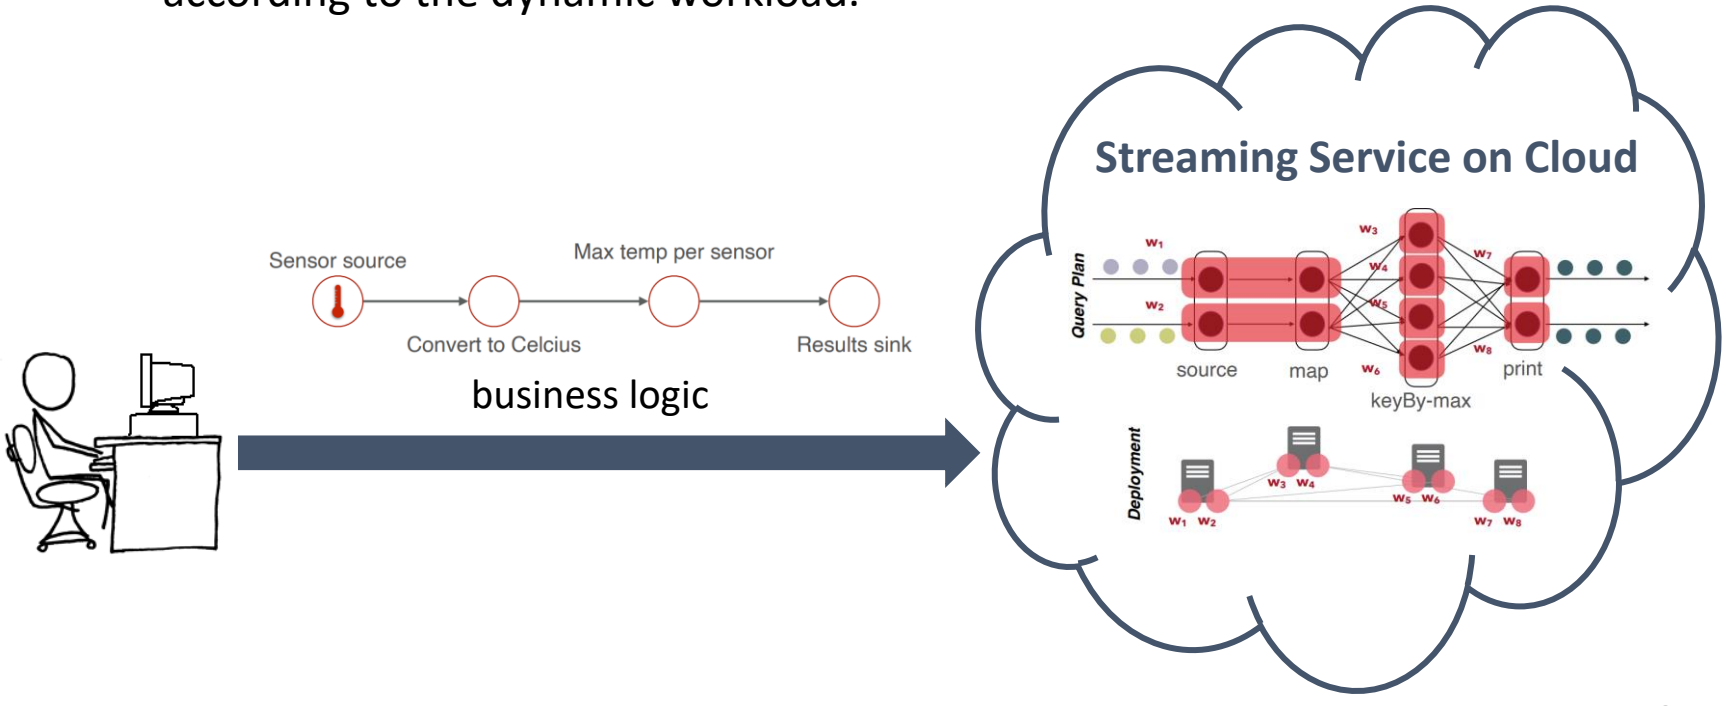

# Data Stream Analytics as managed services


Google Dataflow

Azure Stream Analytics

Alibaba Cloud

Realtime Compute for Apache Flink

# Data Stream Analytics as managed services


Google Dataflow

The cost-benefit trade-off of automatic scaling in cloud-hosted data stream services is not well understood

Azure Stream Analytics

Alibaba Cloud

Realtime Compute for Apache Flink

# Data Stream Analytics as managed services

| | AWS Kinesis Data Analytics | Google Cloud Dataflow | Microsoft Azure Stream Analytics | Alibaba Realtime Compute |
|---|---|---|---|---|
| Policy | Heuristic | Heuristic and predictive | Heuristic | Heuristic |
| Metrics | CPU utilization | CPU utilization and backlog | User-defined | CPU/memory utilization, Latency |
| Scale-out interval | 15 min | 10 minutes | User-defined | 6 min |
| Scale-down interval | 6 hours | 10 minutes | User-defined | 24 hours |

# Data Stream Analytics as managed services

| | AWS Kinesis Data Analytics | Google Cloud Dataflow | Microsoft Azure Stream Analytics | Alibaba Realtime Compute |
|---|---|---|---|---|
| Policy | Heuristic | Heuristic and predictive | Heuristic | Heuristic |
| Metrics | CPU utilization | CPU utilization and backlog | User-defined | CPU/memory utilization, Latency |
| Scale-out interval | 15 min | 10 minutes | User-defined | 6 min |
| Scale-down interval | 6 hours | 10 minutes | User-defined | 24 hours |

# Data Stream Analytics as managed services

| | AWS Kinesis Data Analytics | Google Cloud Dataflow | Microsoft Azure Stream Analytics | Alibaba Realtime Compute |
|---|---|---|---|---|
| Policy | Heuristic | Heuristic and predictive | Heuristic | Heuristic |
| Metrics | CPU utilization | CPU utilization and backlog | User-defined | CPU/memory utilization, Latency |
| Scale-out interval | 15 min | 10 minutes | User-defined | 6 min |
| Scale-down interval | 6 hours | 10 minutes | User-defined | 24 hours |

# Data Stream Analytics as managed services

| | AWS Kinesis Data Analytics | Google Cloud Dataflow | Microsoft Azure Stream Analytics | Alibaba Realtime Compute |
|---|---|---|---|---|
| Policy | Heuristic | Heuristic and predictive | Heuristic | Heuristic |
| Metrics | CPU utilization | CPU utilization and backlog | User-defined | CPU/memory utilization, Latency |
| Scale-out interval | 15 min | 10 minutes | User-defined | 6 min |
| Scale-down interval | 6 hours | 10 minutes | User-defined | 24 hours |

# Data Stream Analytics as managed services

| | AWS Kinesis Data Analytics | Google Cloud Dataflow | Microsoft Azure Stream Analytics | Alibaba Realtime Compute |
|---|---|---|---|---|
| Policy | Heuristic | Heuristic and predictive | Heuristic | Heuristic |
| Metrics | CPU utilization | CPU utilization and backlog | User-defined | CPU/memory utilization, Latency |
| Scale-out interval | 15 min | 10 minutes | User-defined | 6 min |
| Scale-down interval | 6 hours | 10 minutes | User-defined | 24 hours |

# Experimental Methodology

- Workload
  - Nexmark Queries with various operators and state characteristics
  - Dynamic input rate

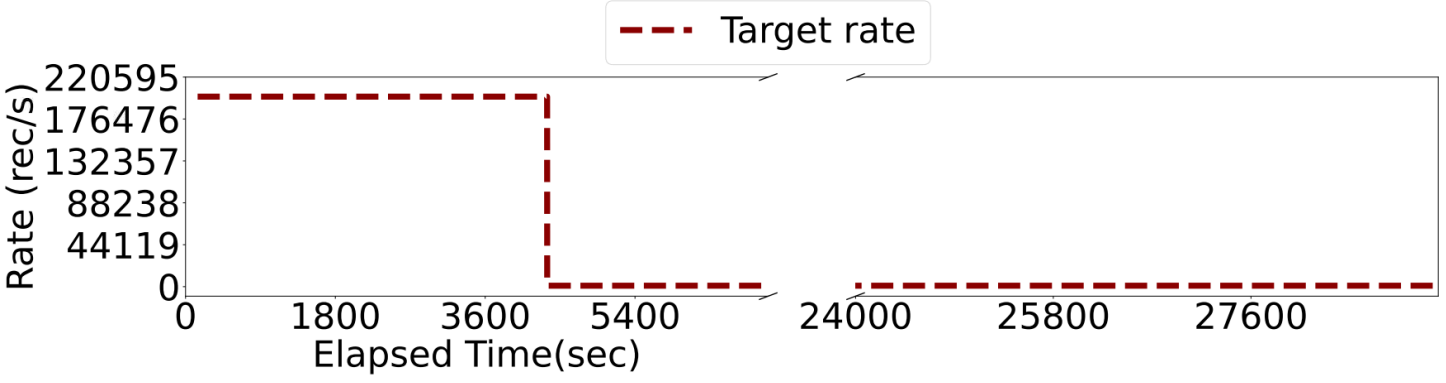

# Experimental Methodology

- Workload
  - Nexmark Queries with various operators and state characteristics
  - Dynamic input rate

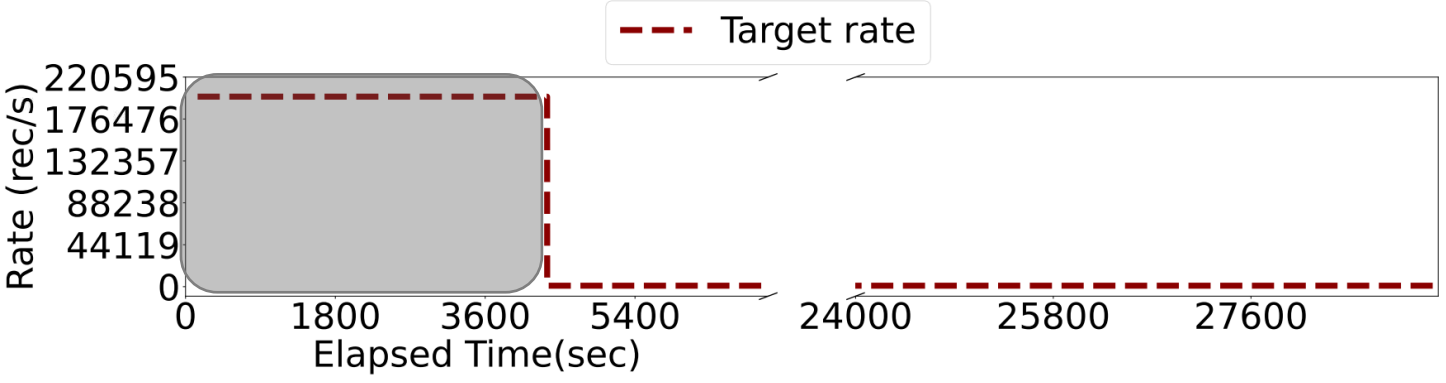

# Experimental Methodology

- Workload
  - Nexmark Queries with various operators and state characteristics
  - Dynamic input rate

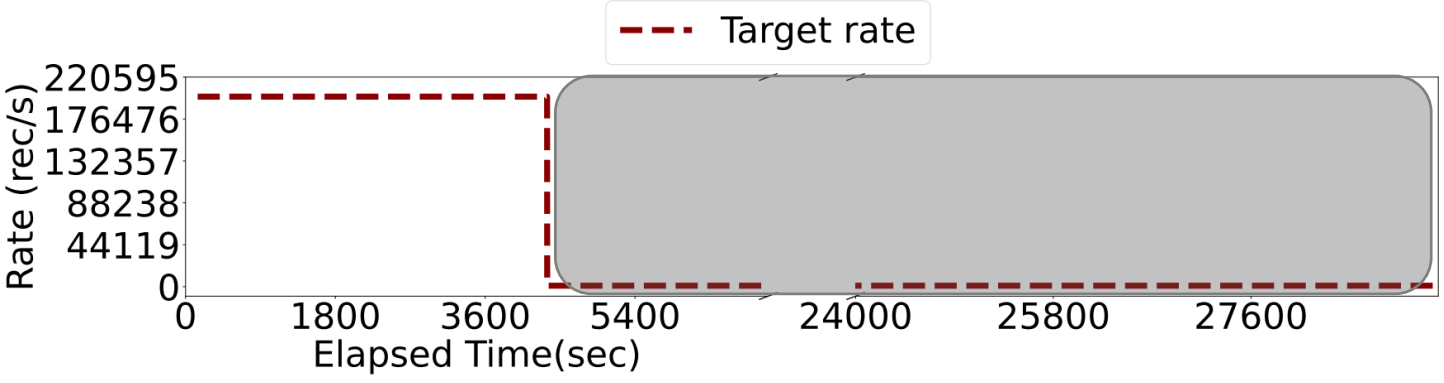

# Experimental Methodology

- Metrics
  - Throughput
  - backpressured time
  - CPU utilization
  - backlog (for Google Dataflow)

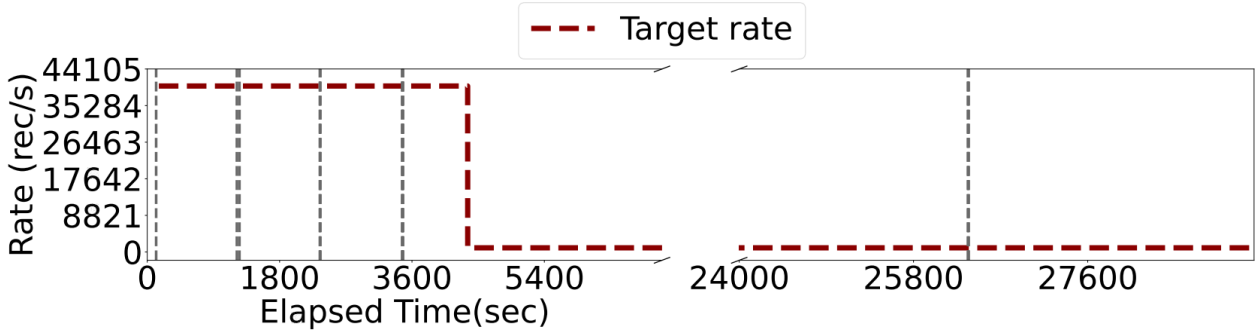


NEXMark Q5 on Amazon Kinesis Data Analytics with auto-scaling enabled
1h10min high input rate, 7h low input rate

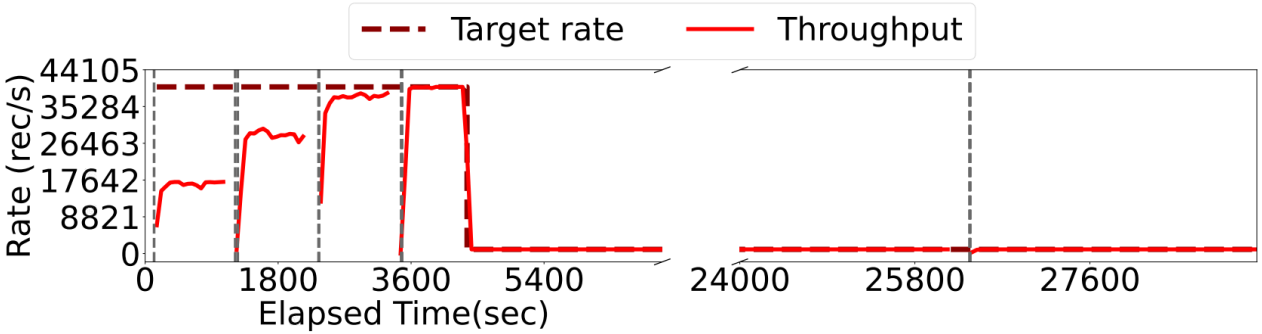


NEXMark Q5 on Amazon Kinesis Data Analytics with auto-scaling enabled
1h10min high input rate, 7h low input rate

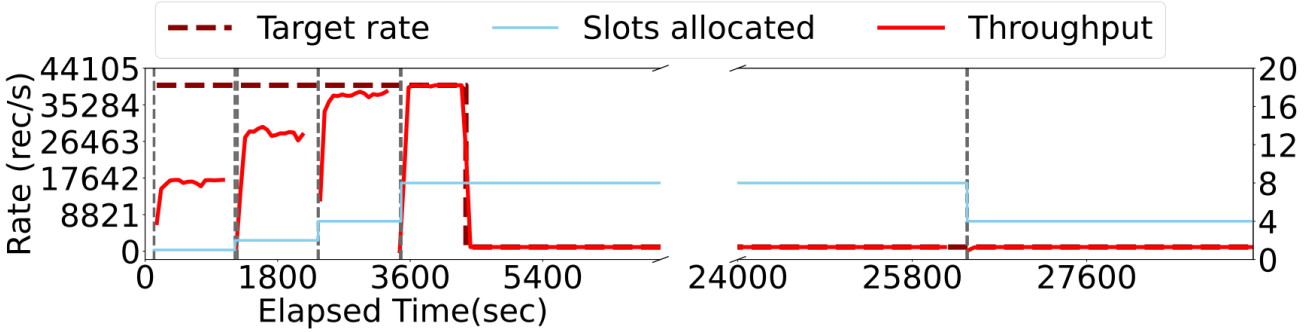


NEXMark Q5 on Amazon Kinesis Data Analytics with auto-scaling enabled
1h10min high input rate, 7h low input rate

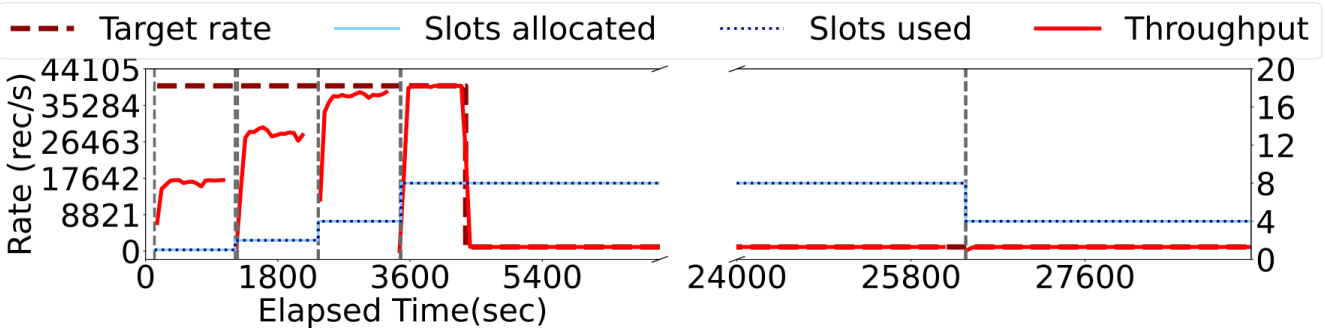


NEXMark Q5 on Amazon Kinesis Data Analytics with auto-scaling enabled
1h10min high input rate, 7h low input rate

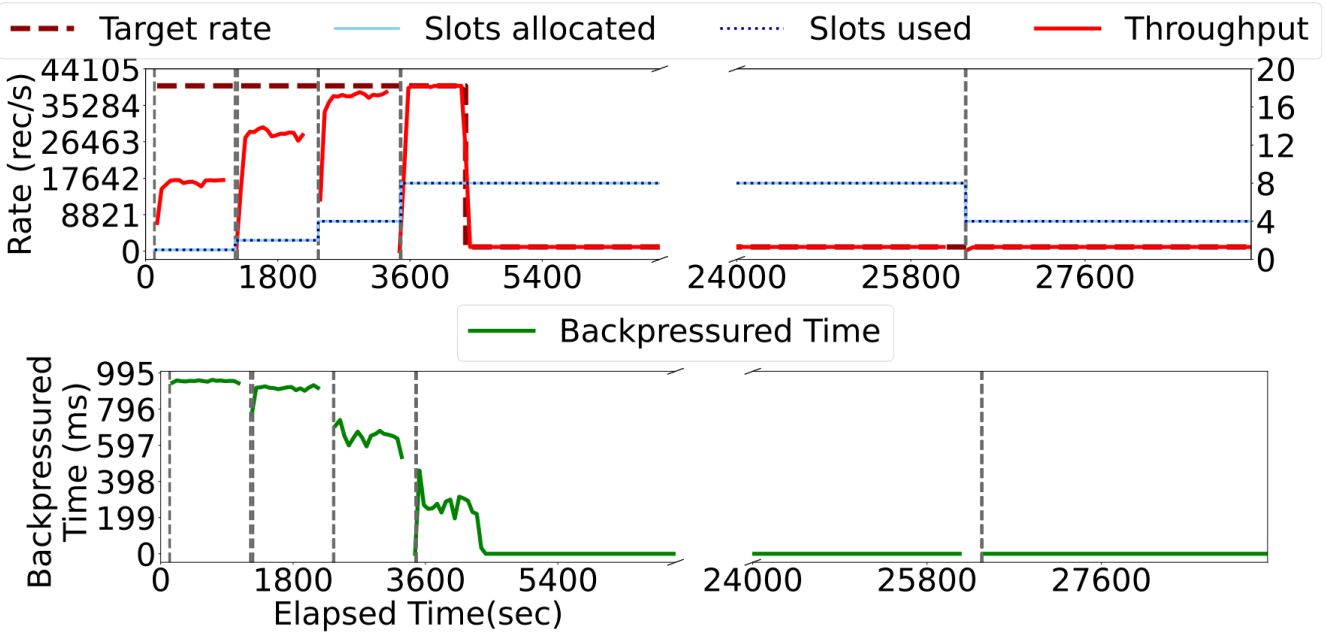


NEXMark Q5 on Amazon Kinesis Data Analytics with auto-scaling enabled
1h10min high input rate, 7h low input rate

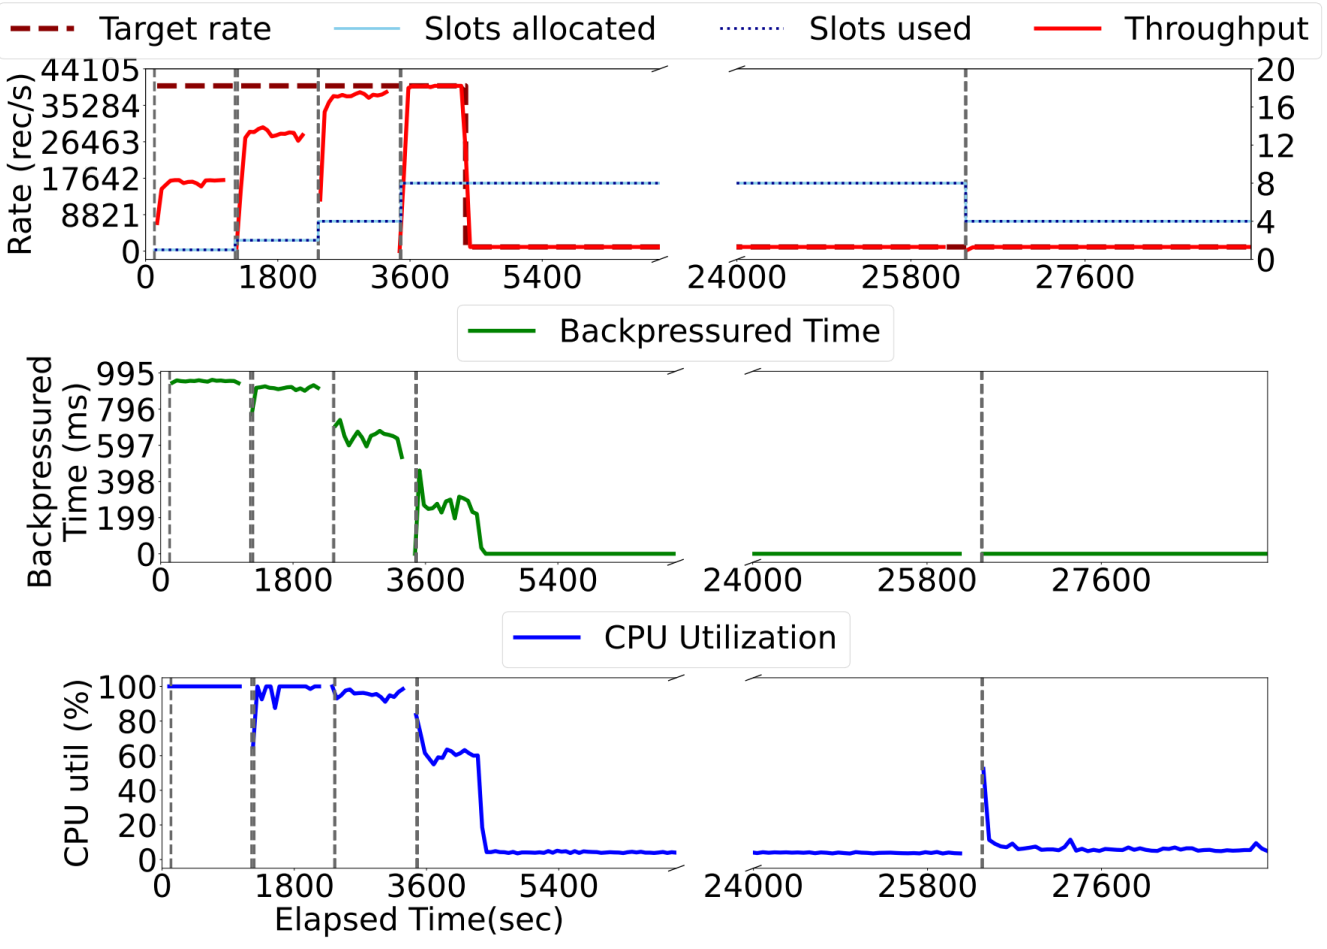


NEXMark Q5 on Amazon Kinesis Data Analytics with auto-scaling enabled
1h10min high input rate, 7h low input rate

Scale up interval is 15 minutes

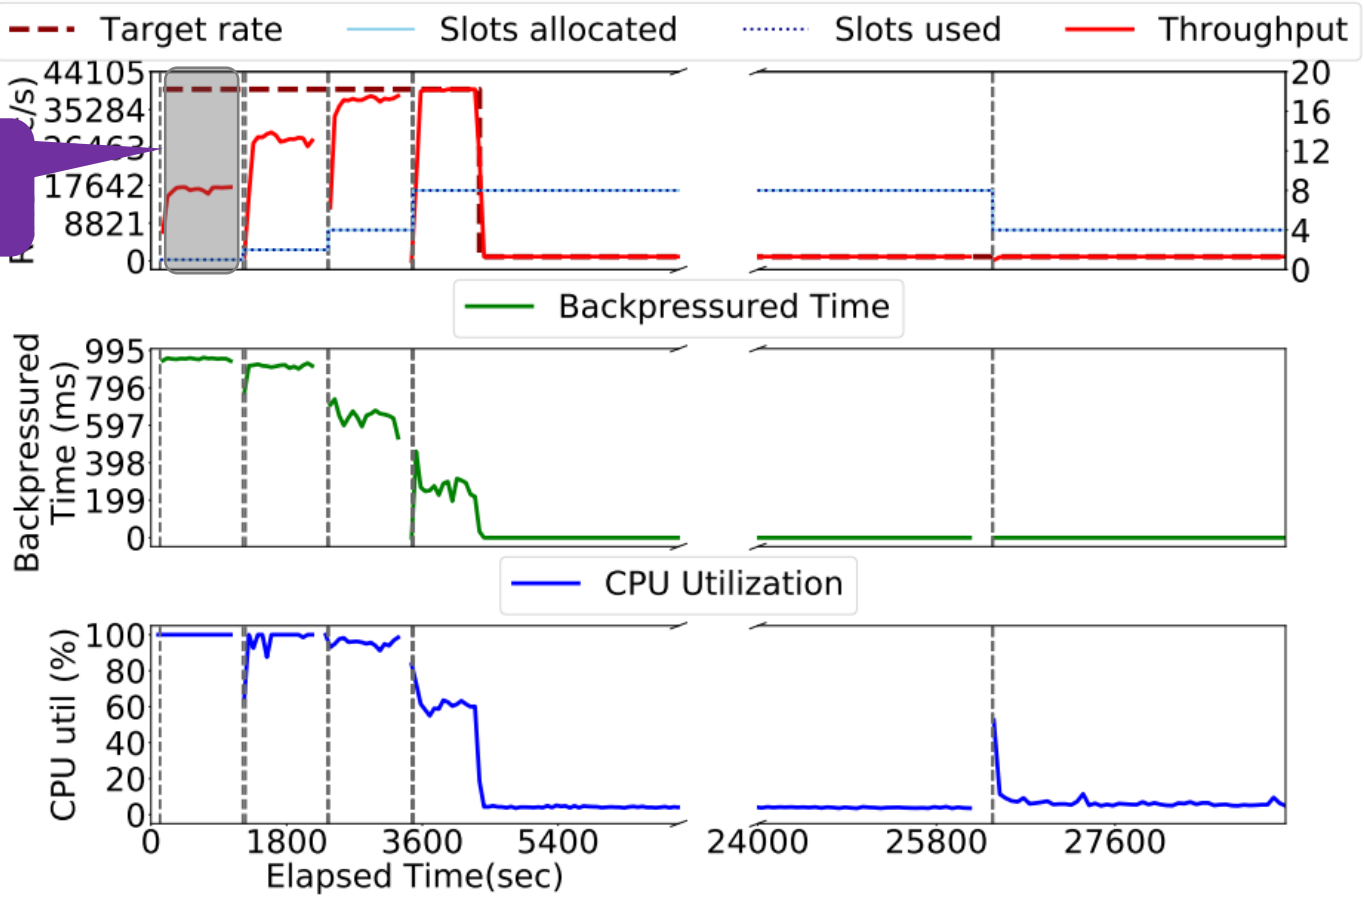


Aggressive scale-up. Conservative and slow scale-down
Users might be considerably over-charged for underutilized resources

Scale down interval is 6 hours

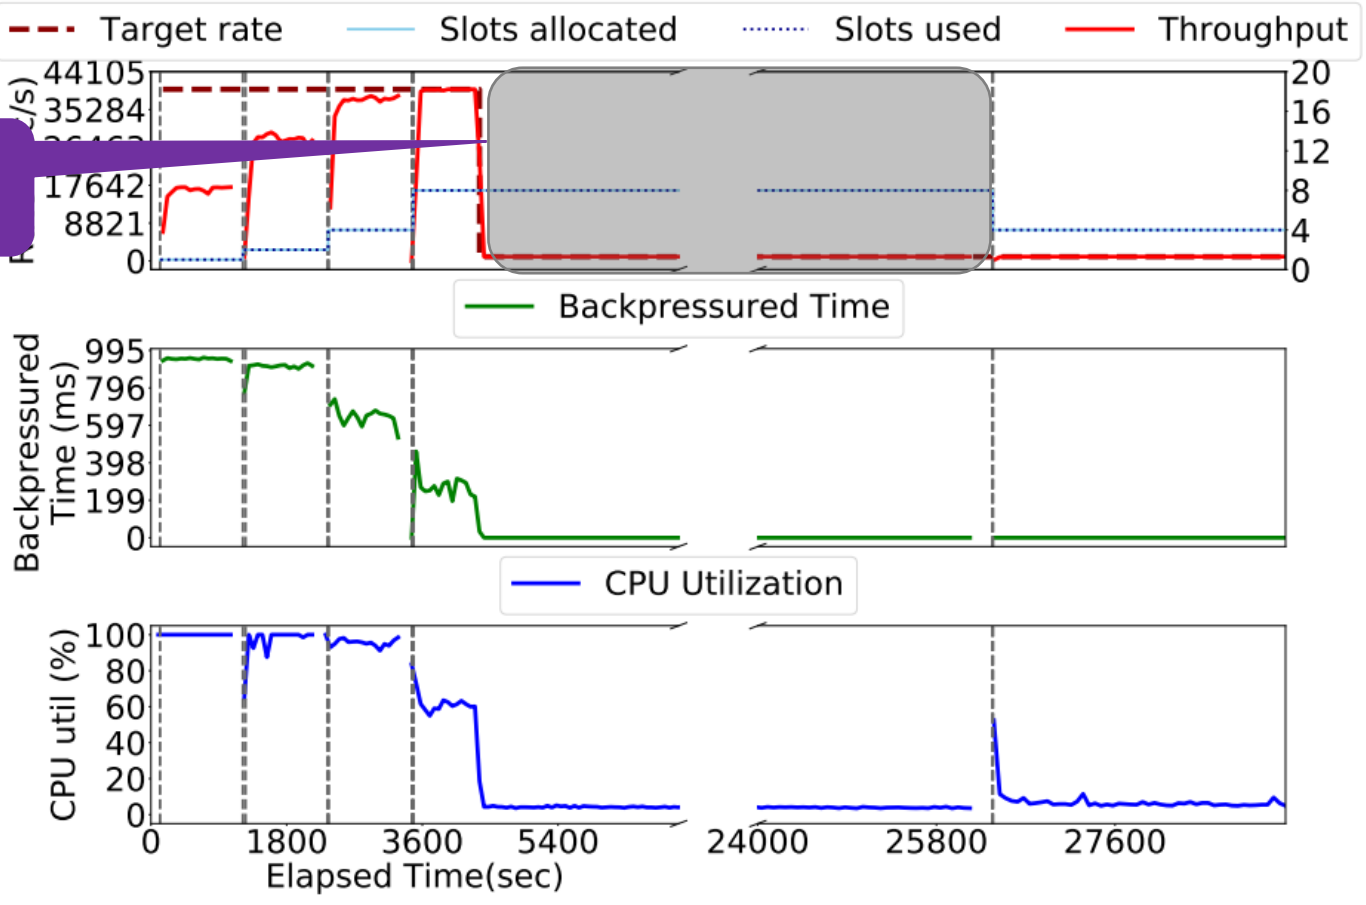


Aggressive scale-up. Conservative and slow scale-down
Users might be considerably over-charged for underutilized resources

Need 18h to arrive ideal configuration (1 slot)

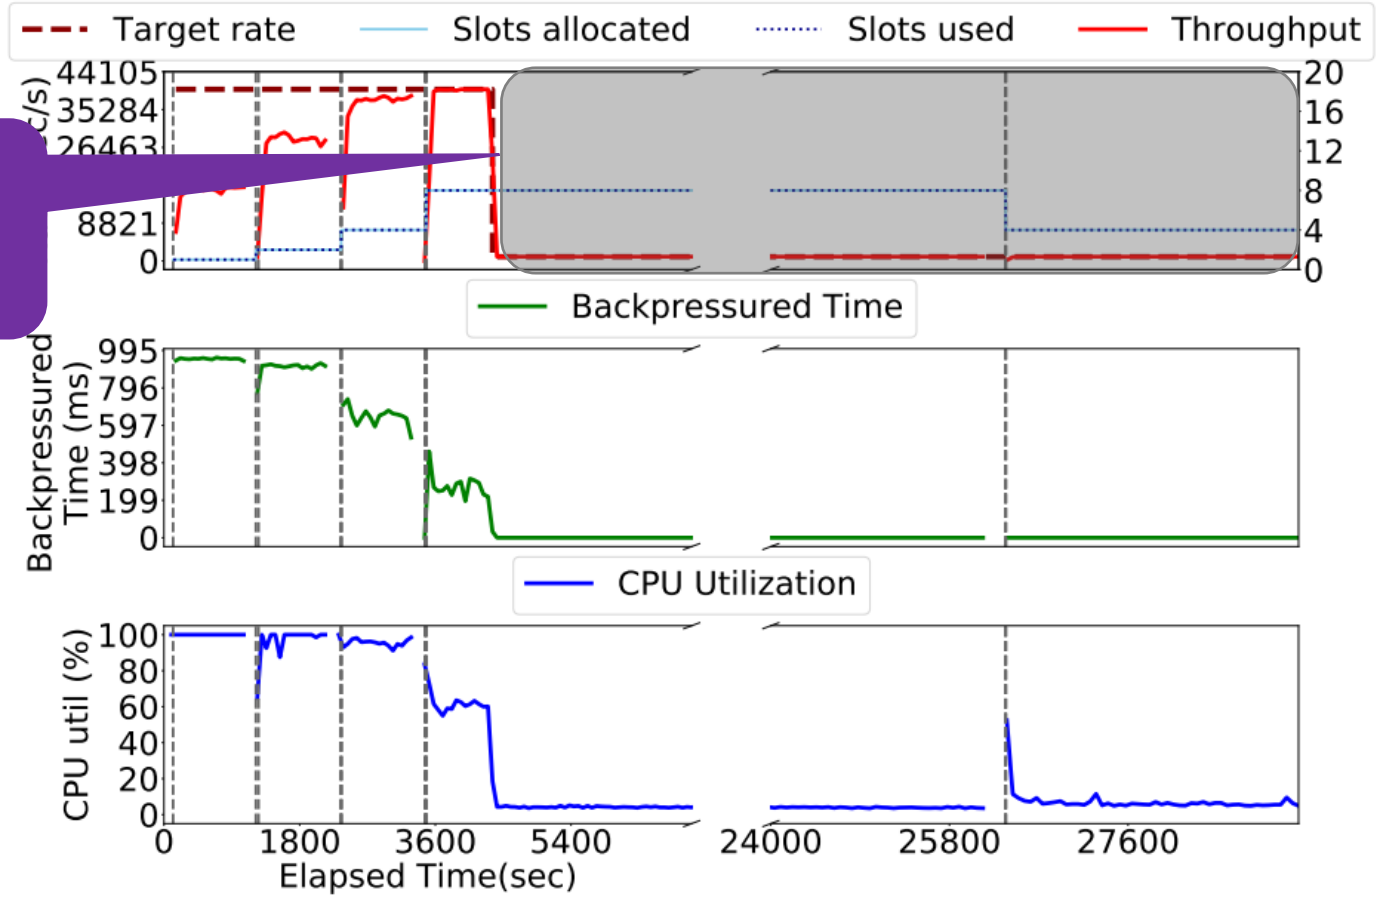


Aggressive scale-up. Conservative and slow scale-down
Users might be considerably over-charged for underutilized resources

Adding resources does not improve the throughput

The job remains back-pressured while resources are under-utilized

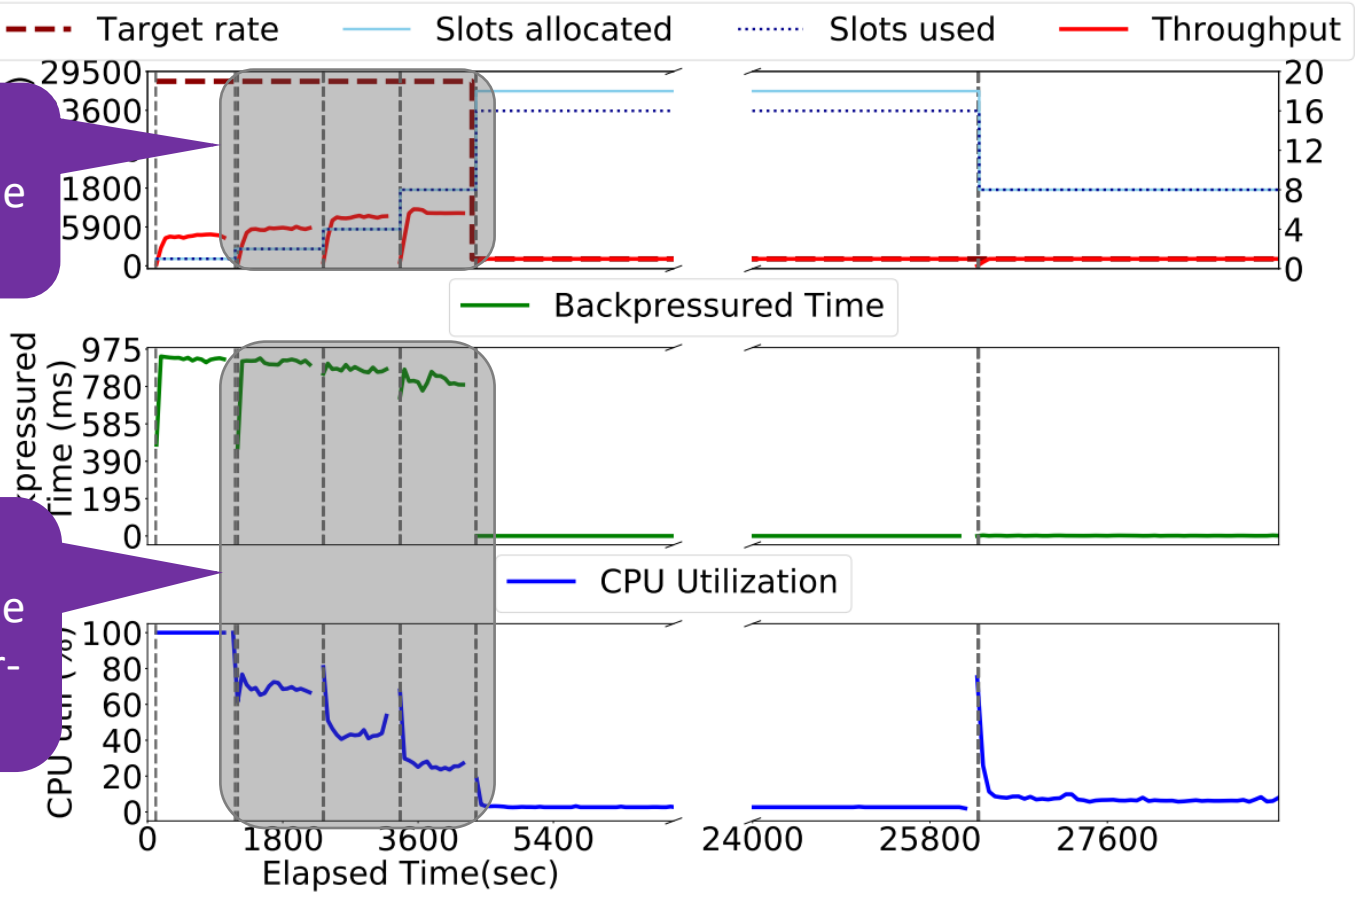


NEXMark Q8 on Amazon Kinesis Data Analytics with auto-scaling enabled
Kinesis is not always capable of accurately identifying bottlenecks

Not all allocated slots are using

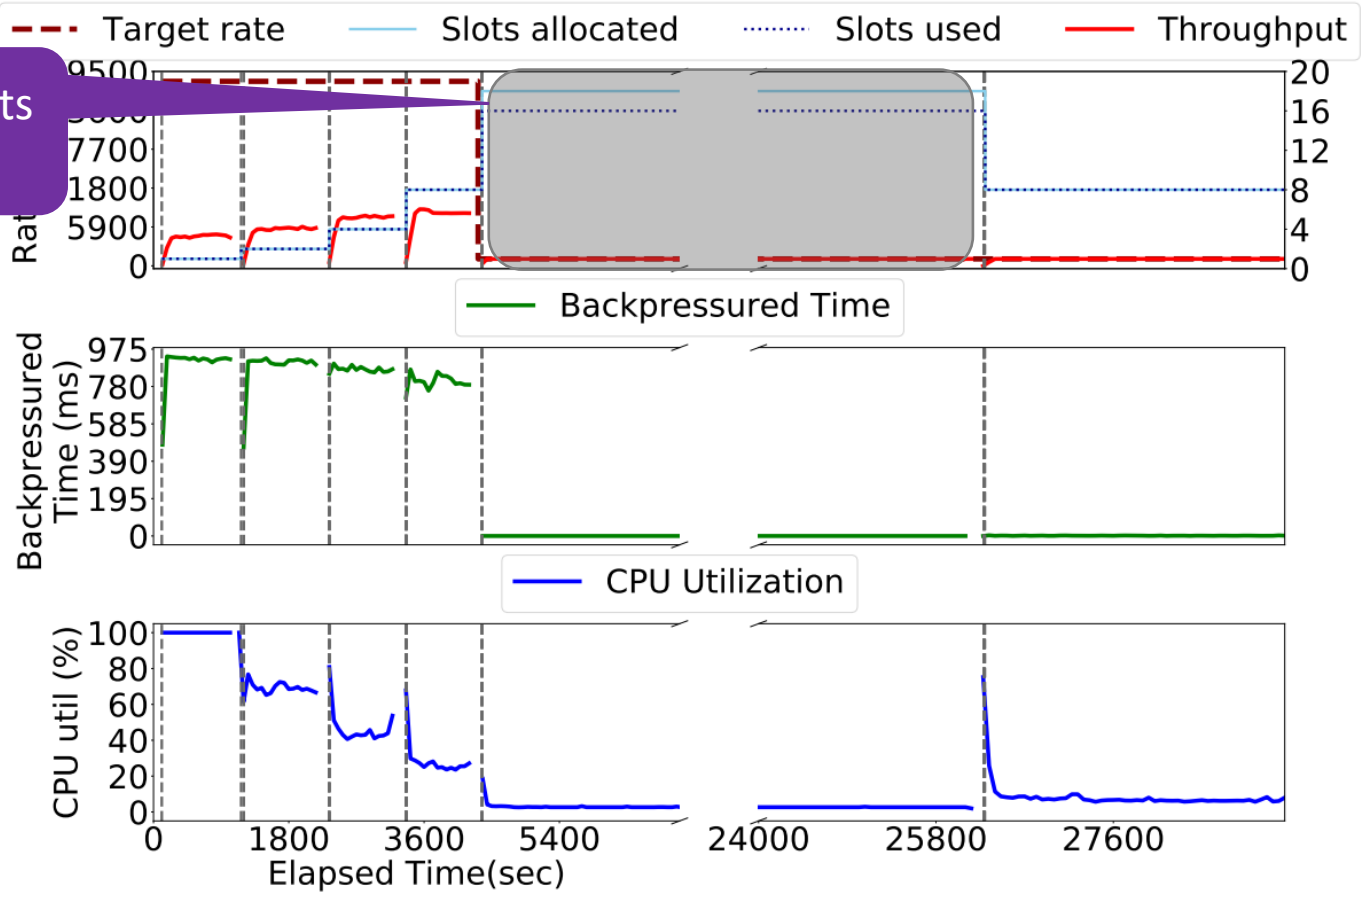


NEXMark Q8 on Amazon Kinesis Data Analytics with auto-scaling enabled
Kinesis is not always capable of accurately identifying bottlenecks

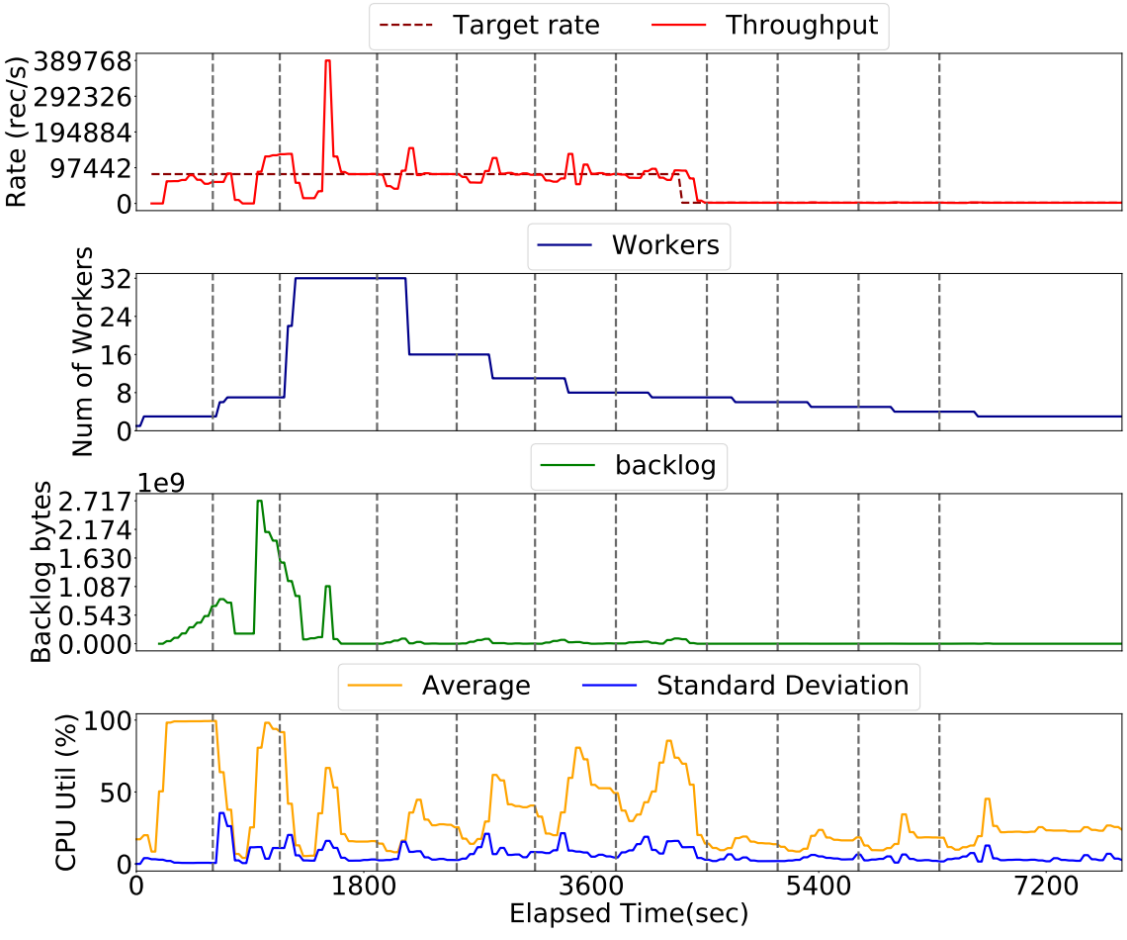


NEXMark Q5 on Google Dataflow with maximum number of workers==32
1h10min high input rate, 7h low input rate

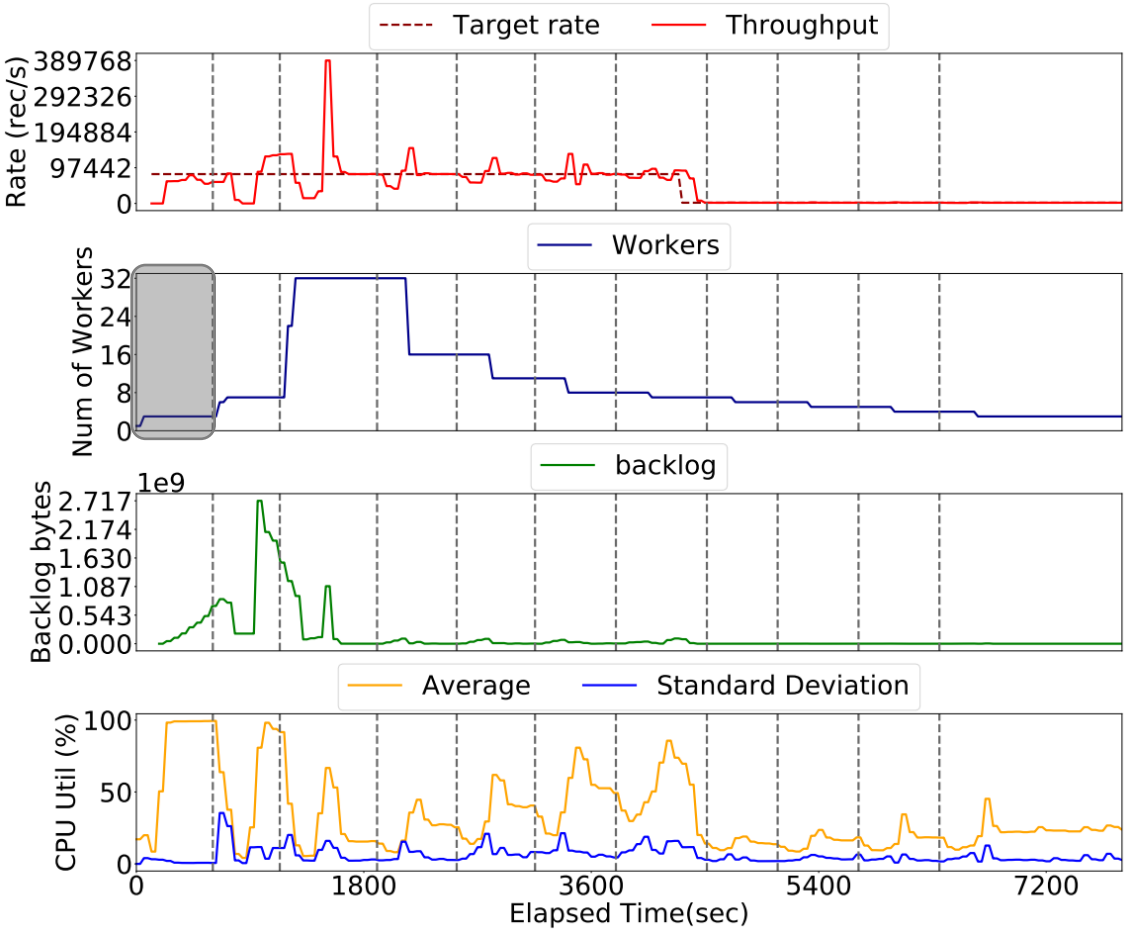


Dataflow makes scaling decisions about every 10min
Dataflow remembers previous scale-up decisions

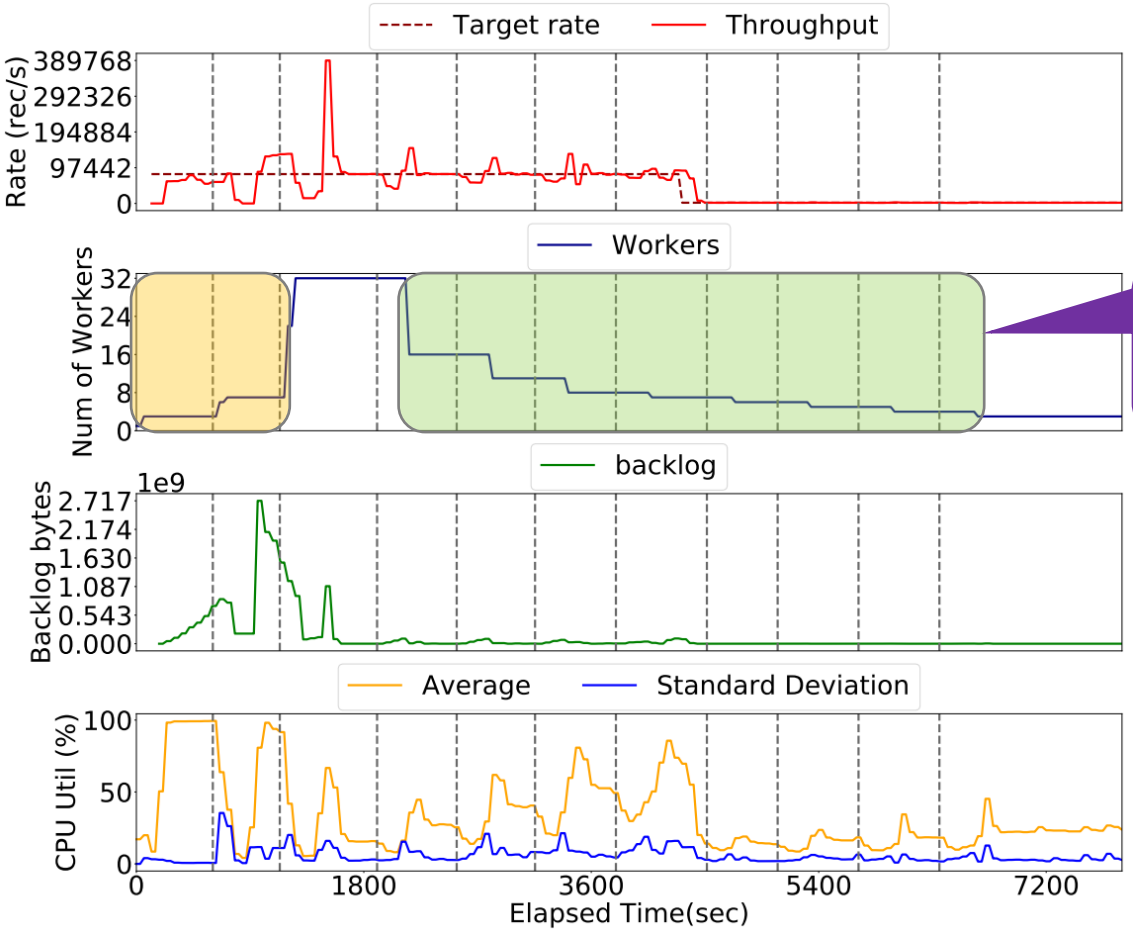


Scale up aggressively
Scale down conservative

The scale-up and down policies are asymmetric

# The Non-Expert Tax

- A configuration is **ideal** if:
  - it can keep up with the target rate without inducing backpressure or backlog, while using the minimum amount of resources.
- Non-Expert Tax: relative error between **the expected cost under ideal configuration** and **the actual cost charged by the Cloud provider**

$$Non\_Expert\ Tax = \frac{actual\_cost - expected\_cost}{expected\_cost}$$

# Quantifying the Non-Expert Tax

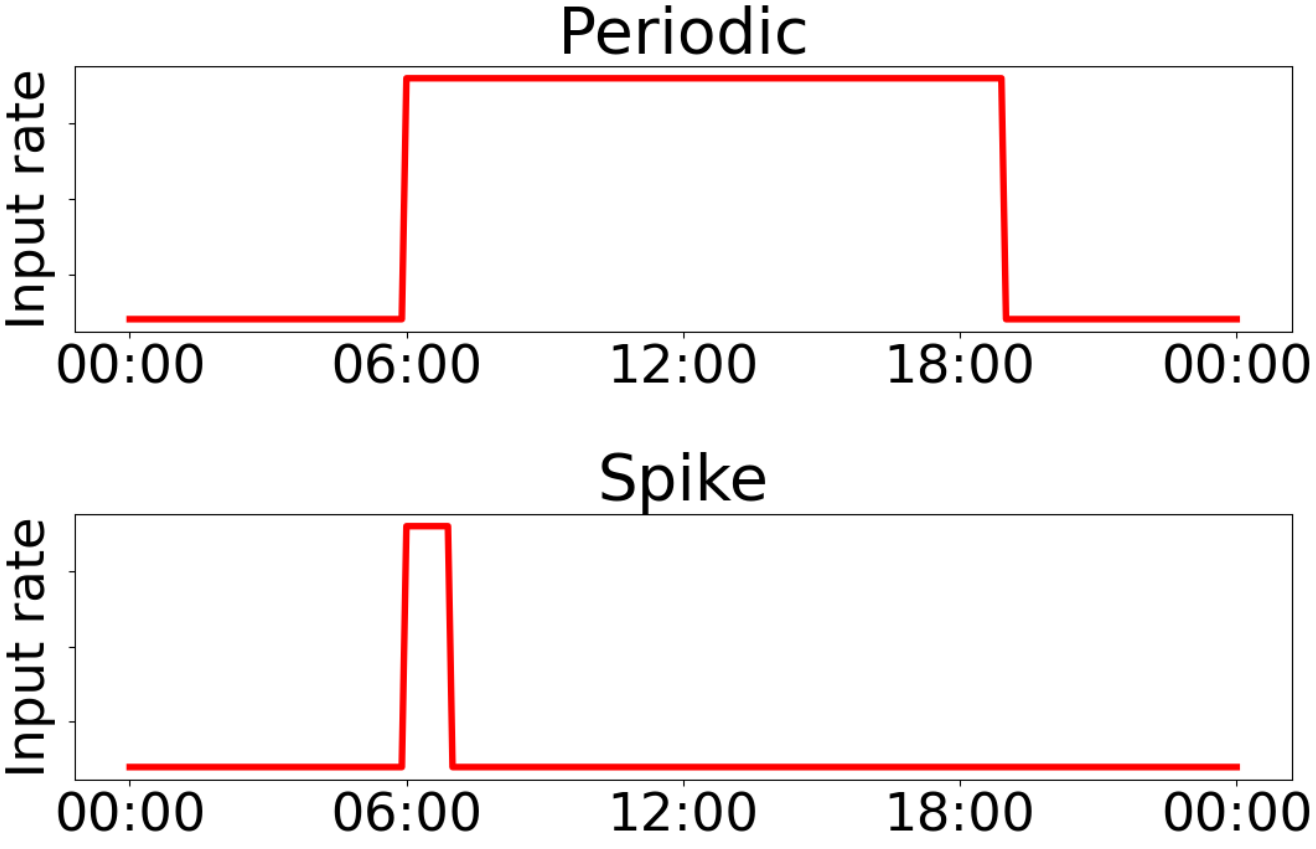

# Quantifying the Non-Expert Tax

| | | Experiment workload (per query) | | | Periodic workload for a month | | | Spike workload for a month | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | | Actual Cost ($) | Ideal Cost ($) | Non-Expert Tax (%) | Actual Cost ($) | Ideal Cost ($) | Non-Expert Tax (%) | Actual Cost ($) | Ideal Cost ($) | Non-Expert Tax (%) |
| Kinesis | Q1 | 13.98 | 2.17 | 544.23 | 1252.7 | 289.8 | 332.26 | 419.55 | 173.97 | 141.16 |
| | Q3 | 13.98 | 2.87 | 387.1 | 1252.7 | 542.52 | 130.90 | 419.55 | 195.03 | 115.12 |
| | Q5 | 7.9 | 2.64 | 199.24 | 666.77 | 458.28 | 45.49 | 419.55 | 188.01 | 123.15 |
| | Q8 | 13.98 | 2.87 | 387.1 | 1252.7 | 542.52 | 130.90 | 419.55 | 195.03 | 115.12 |
| Dataflow | Q1 | 2.68 | 1.06 | 152.83 | 275.03 | 190.08 | 44.69 | 206.3 | 73.92 | 179.09 |
| | Q3 | 6.68 | 2.21 | 202.26 | 642.8 | 428.4 | 50.05 | 425.88 | 147.9 | 187.95 |
| | Q5 | 6.16 | 2.04 | 201.96 | 586.9 | 367.2 | 59.83 | 425.85 | 142.8 | 198.21 |
| | Q8 | 8.04 | 3.91 | 105.63 | 1209.55 | 1040.4 | 16.26 | 464.24 | 198.9 | 133.40 |

**These results do not provide a reliable way to compare Kinesis with Dataflow in terms of cost. The services were evaluated using different input rates and instance types.*

**Table 3: Actual cost, ideal cost, and incurred Non-Expert Tax for the Nexmark queries in three workload scenarios.**

# Quantifying the Non-Expert Tax

| | | Experiment workload (per query) | | | Periodic workload for a month | | | Spike workload for a month | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | | Actual Cost ($) | Ideal Cost ($) | Non-Expert Tax (%) | Actual Cost ($) | Ideal Cost ($) | Non-Expert Tax (%) | Actual Cost ($) | Ideal Cost ($) | Non-Expert Tax (%) |
| Kinesis | Q1 | 13.98 | 2.17 | 544.23 | 1252.7 | 289.8 | 332.26 | 419.55 | 173.97 | 141.16 |
| | Q3 | 13.98 | 2.87 | 387.1 | 1252.7 | 542.52 | 130.90 | 419.55 | 195.03 | 115.12 |
| | Q5 | 7.9 | 2.64 | 199.24 | 666.77 | 458.28 | 45.49 | 419.55 | 188.01 | 123.15 |
| | Q8 | 13.98 | 2.87 | 387.1 | 1252.7 | 542.52 | 130.90 | 419.55 | 195.03 | 115.12 |
| Dataflow | Q1 | 2.68 | 1.06 | 152.83 | 275.03 | 190.08 | 44.69 | 206.3 | 73.92 | 179.09 |
| | Q3 | 6.68 | 2.21 | 202.26 | 642.8 | 428.4 | 50.05 | 425.88 | 147.9 | 187.95 |
| | Q5 | 6.16 | 2.04 | 201.96 | 586.9 | 367.2 | 59.83 | 425.85 | 142.8 | 198.21 |
| | Q8 | 8.04 | 3.91 | 105.63 | 1209.55 | 1040.4 | 16.26 | 464.24 | 198.9 | 133.40 |

*These results do not provide a reliable way to compare Kinesis with Dataflow in terms of cost. The services were evaluated using different input rates and instance types.

**Table 3: Actual cost, ideal cost, and incurred Non-Expert Tax for the Nexmark queries in three workload scenarios.**

Over 300% Non-Expert-Tax per month for periodic workloads

Over 500% Non-Expert-Tax for short-term jobs

# Quantifying the Non-Expert Tax

| | | Experiment workload (per query) | | | Periodic workload for a month | | | Spike workload for a month | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | | Actual Cost ($) | Ideal Cost ($) | Non-Expert Tax (%) | Actual Cost ($) | Ideal Cost ($) | Non-Expert Tax (%) | Actual Cost ($) | Ideal Cost ($) | Non-Expert Tax (%) |
| Kinesis | Q1 | 13.98 | 2.17 | 544.23 | 1252.7 | 289.8 | 332.26 | 419.55 | 173.97 | 141.16 |
| | Q3 | 13.98 | 2.87 | 387.1 | 1252.7 | 542.52 | 130.90 | 419.55 | 195.03 | 115.12 |
| | Q5 | 7.9 | 2.64 | 199.24 | 666.77 | 458.28 | 45.49 | 419.55 | 188.01 | 123.15 |
| | Q8 | 13.98 | 2.87 | 387.1 | 1252.7 | 542.52 | 130.90 | 419.55 | 195.03 | 115.12 |
| Dataflow | Q1 | 2.68 | 1.06 | 152.83 | 275.03 | 190.08 | 44.69 | 206.3 | 73.92 | 179.09 |
| | Q3 | 6.68 | 2.21 | 202.26 | 642.8 | 428.4 | 50.05 | 425.88 | 147.9 | 187.95 |
| | Q5 | 6.16 | 2.04 | 201.96 | 586.9 | 367.2 | 59.83 | 425.85 | 142.8 | 198.21 |
| | Q8 | 8.04 | 3.91 | 105.63 | 1209.55 | 1040.4 | 16.26 | 464.24 | 198.9 | 133.40 |

*These results do not provide a reliable way to compare Kinesis with Dataflow in terms of cost. The services were evaluated using different input rates and instance types.

**Table 3: Actual cost, ideal cost, and incurred Non-Expert Tax for the Nexmark queries in three workload scenarios.**

The Kinesis Non-Expert Tax is higher for periodic workload

# Quantifying the Non-Expert Tax

| | | Experiment workload (per query) | | | Periodic workload for a month | | | Spike workload for a month | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | | Actual Cost ($) | Ideal Cost ($) | Non-Expert Tax (%) | Actual Cost ($) | Ideal Cost ($) | Non-Expert Tax (%) | Actual Cost ($) | Ideal Cost ($) | Non-Expert Tax (%) |
| Kinesis | Q1 | 13.98 | 2.17 | 544.23 | 1252.7 | 289.8 | 332.26 | 419.55 | 173.97 | 141.16 |
| | Q3 | 13.98 | 2.87 | 387.1 | 1252.7 | 542.52 | 130.90 | 419.55 | 195.03 | 115.12 |
| | Q5 | 7.9 | 2.64 | 199.24 | 666.77 | 458.28 | 45.49 | 419.55 | 188.01 | 123.15 |
| | Q8 | 13.98 | 2.87 | 387.1 | 1252.7 | 542.52 | 130.90 | 419.55 | 195.03 | 115.12 |
| Dataflow | Q1 | 2.68 | 1.06 | 152.83 | 275.03 | 190.08 | 44.69 | 206.3 | 73.92 | 179.09 |
| | Q3 | 6.68 | 2.21 | 202.26 | 642.8 | 428.4 | 50.05 | 425.88 | 147.9 | 187.95 |
| | Q5 | 6.16 | 2.04 | 201.96 | 586.9 | 367.2 | 59.83 | 425.85 | 142.8 | 198.21 |
| | Q8 | 8.04 | 3.91 | 105.63 | 1209.55 | 1040.4 | 16.26 | 464.24 | 198.9 | 133.40 |

*These results do not provide a reliable way to compare Kinesis with Dataflow in terms of cost. The services were evaluated using different input rates and instance types.

**Table 3: Actual cost, ideal cost, and incurred Non-Expert Tax for the Nexmark queries in three workload scenarios.**

Up to 200% Non-Expert-Tax

# These services are not easy to use, even by experts

- User input is required to achieve good performance or avoid failures
- User settings might conflict with policy decisions and silently cause unexpected behavior.
- Sometimes documents are inaccurate/incomplete.

# Conclusion

There is large room for improving the **cost** and **resource efficiency** of Cloud-hosted streaming analytics services

- Over-charge users on under-utilized resources.
- Cannot accurately identify bottlenecks.
- Not easy to use.

# Conclusion

There is large room for improving the **cost** and **resource efficiency** of Cloud-hosted streaming analytics services

- Over-charge users on under-utilized resources.
- Cannot accurately identify bottlenecks.
- Not easy to use.

**Yuanli Wang**, Baiqing Lyu, Vasiliki Kalavri
yuanliw@bu.edu
https://sites.bu.edu/casp/

BOSTON UNIVERSITY

CASP Systems