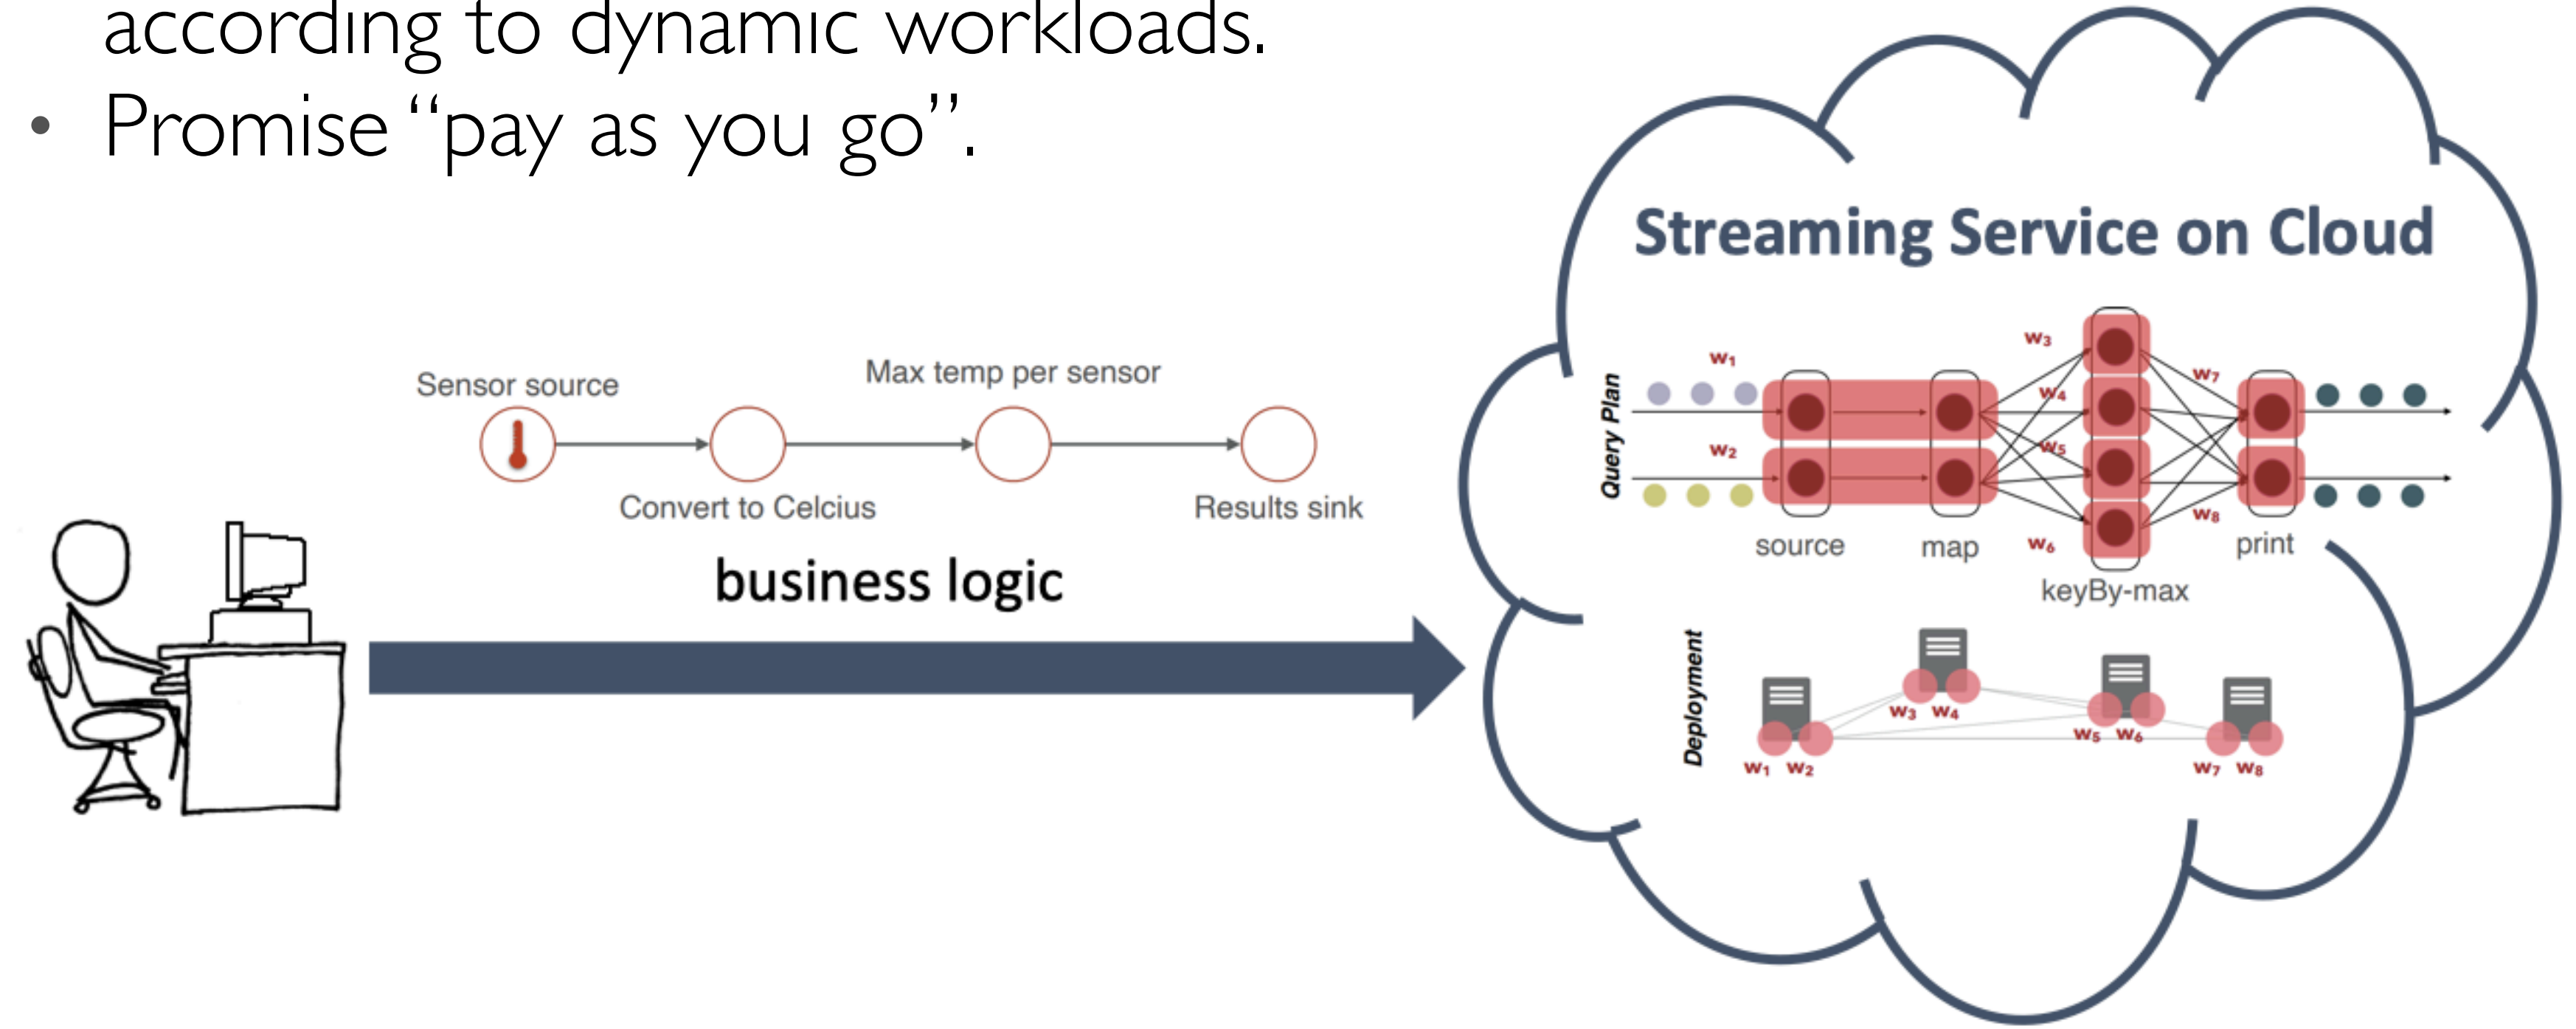


Streaming workloads are long-running, dynamic, and often unpredictable, eventually rendering any initial configuration out-of-tune.

**Expert oversight is required to achieve good performance.**

### Data Stream Analytics as managed services:

- Enable non-expert users to run stream processing jobs.
- Automatically deploy streaming job, provision and adjust resources, according to dynamic workloads.
- Promise "pay as you go".




**They over-charge users for under-utilized resources**

**They can not accurately identify bottlenecks**

**Workload:** Nexmark Queries with various operators and state characteristics. 1h high rate stage + 7h low-rate stage


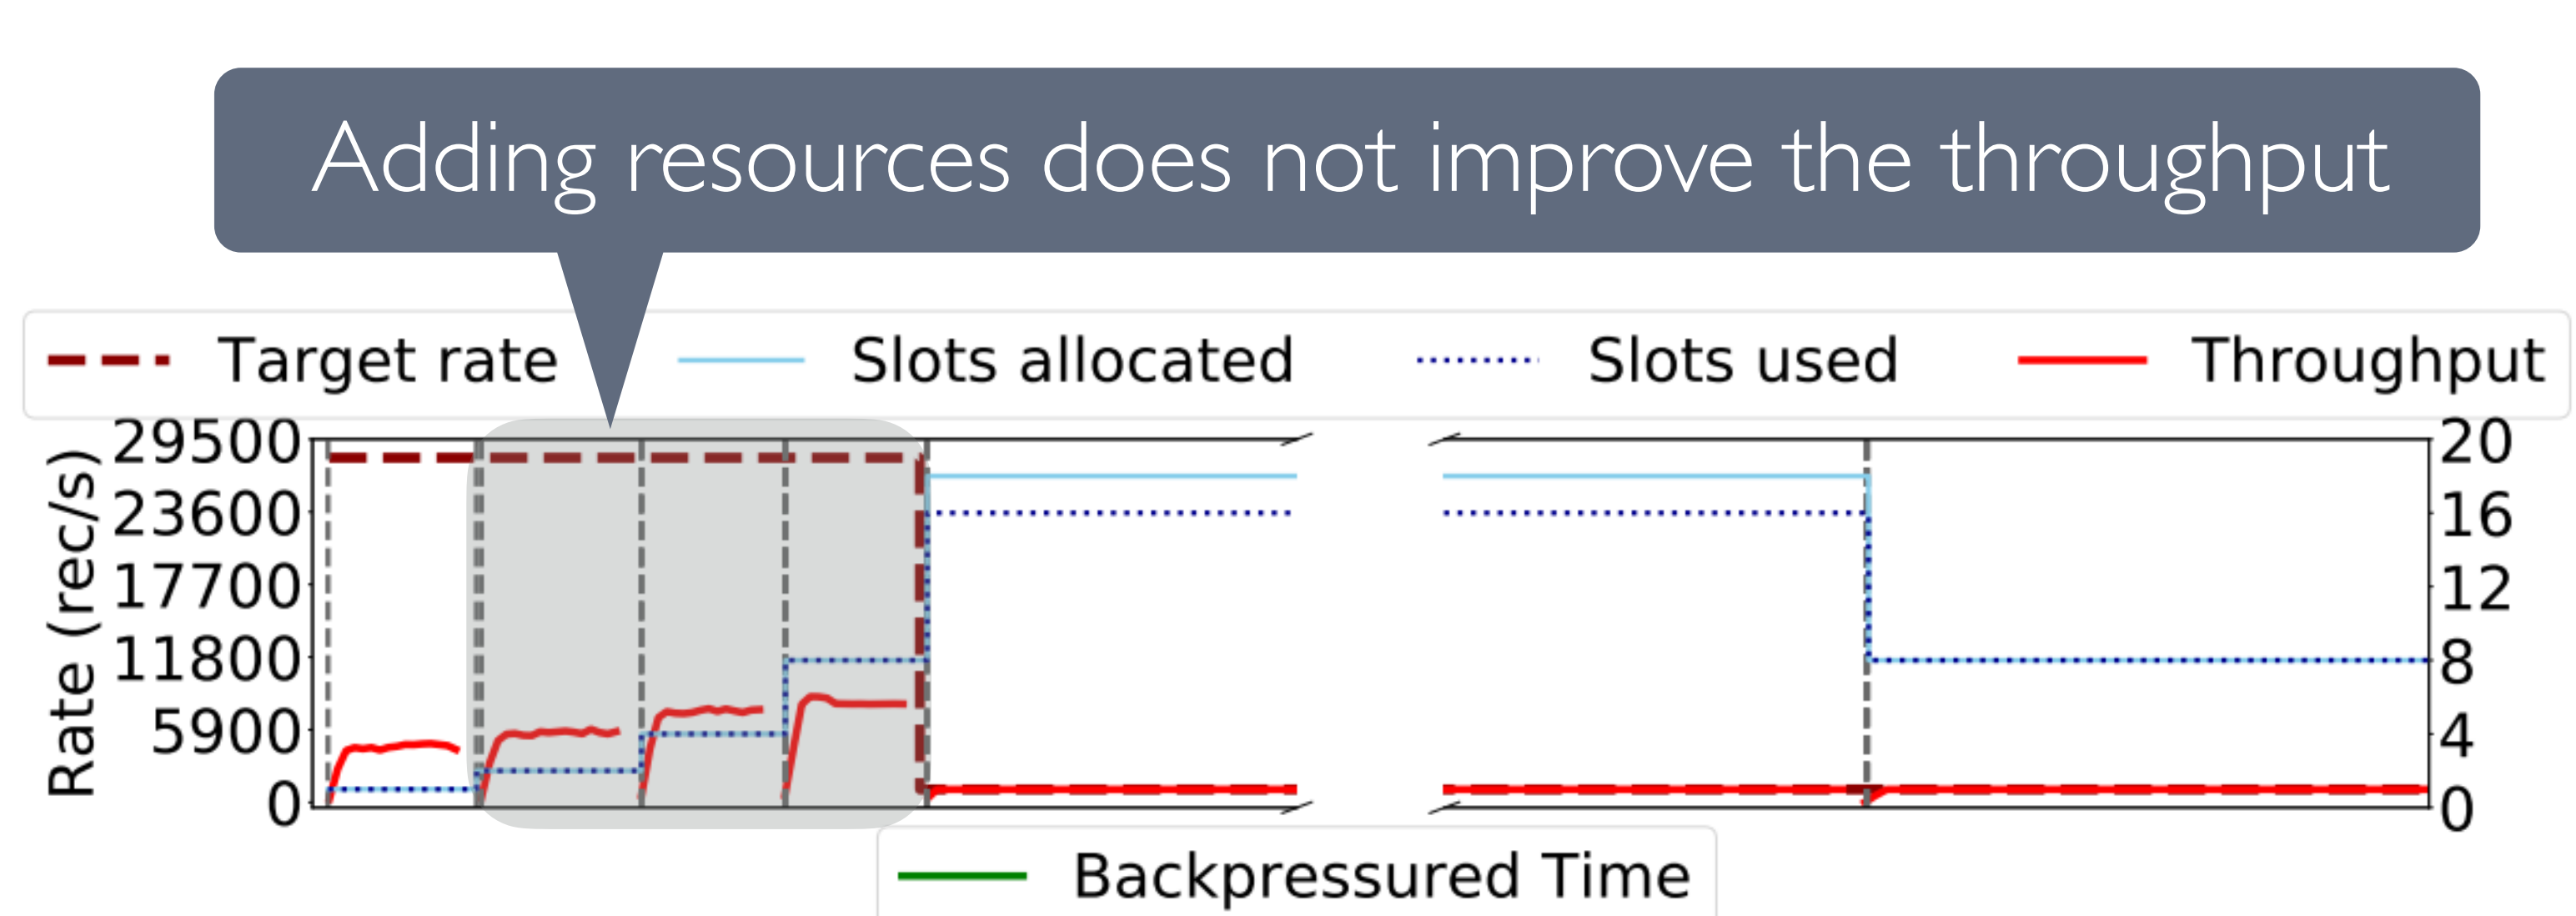
NEXMark Q1 on Amazon Kinesis Data Analytics



Aggressive scale-up, over-provisioning resources by 2x



Conservative scale-down: >24h to release underutilized resources

NEXMark Q8 on Amazon Kinesis Data Analytics



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

NEXMark Q5 on Google Dataflow



The scale-up and down policies are asymmetric

## 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.

- Over 300% Non-Expert-Tax per month for periodic workloads
- Over 500% Non-Expert-Tax for short-term jobs