

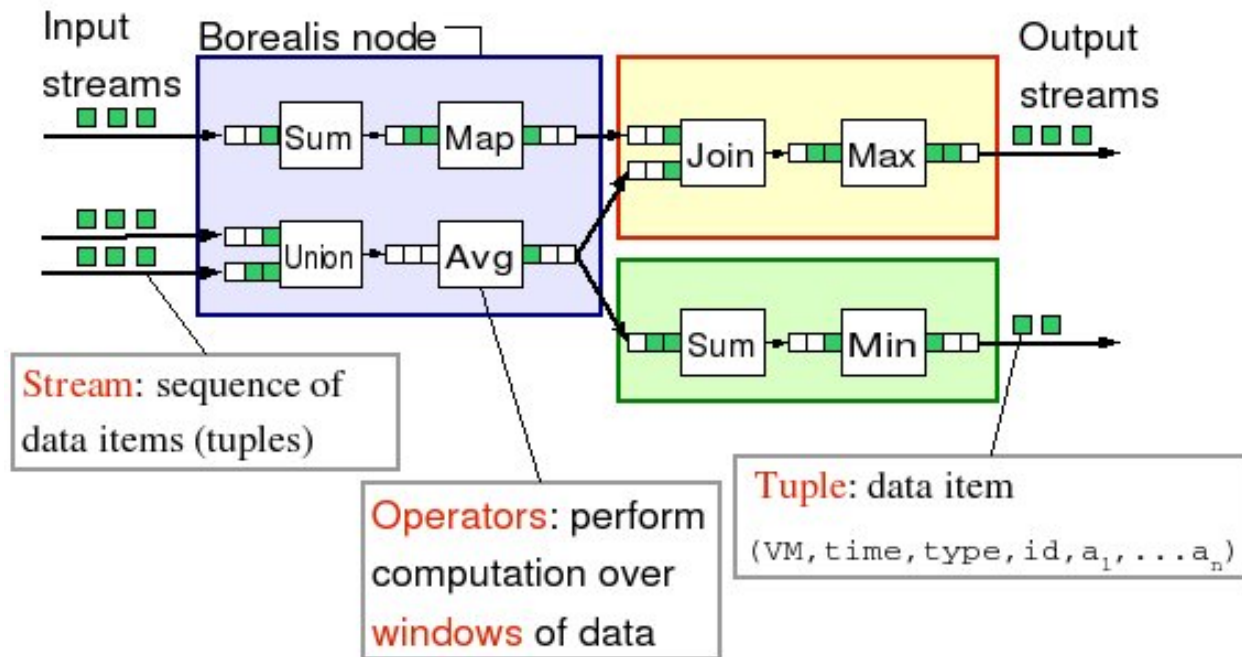
# Aurora Borealis



<http://www.cio.com/article/2370573/consumer-technology/aurora-borealis-through-a-circular-fisheye-lens.html>

# Borealis Stream Processing

“**Distributed** streaming processing engine with **more flexible stream processing model**” to meet the requirements of streaming applications.



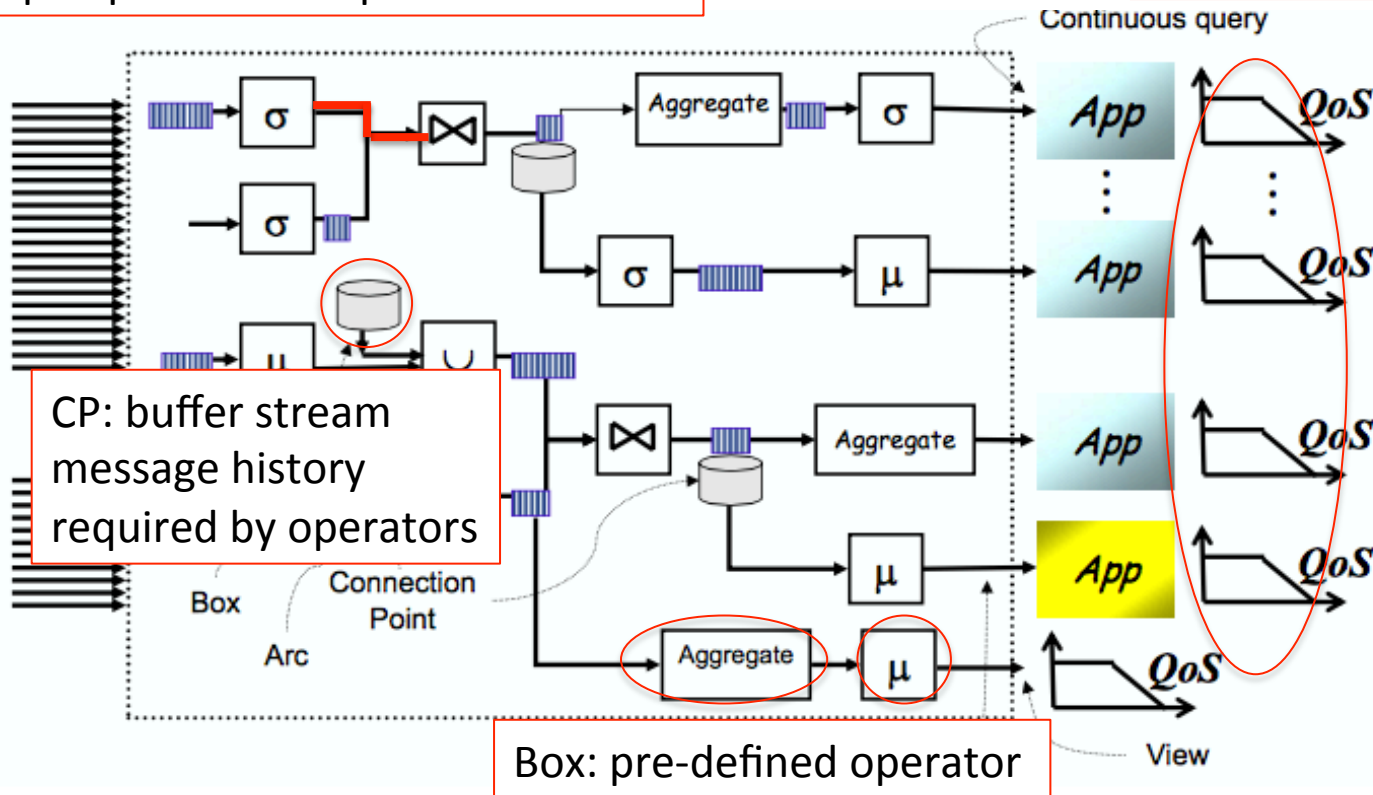
# Stream-Processing Engine

- Real-time processing of large volumes of streaming data.
  - Store then process → Process data before store
  - Relatively small set of operators (e.g. filter, aggregates, correlations) on windows of data that move with time
- **Aurora Stream-Processing Engine (SPE)**
  - Time-series ops. for streaming applications
  - Real-time latency (i.e. process before storing)
  - Gracefully deal with bursty message load

# Aurora Stream Processing

Arc: tuple queue that represents stream

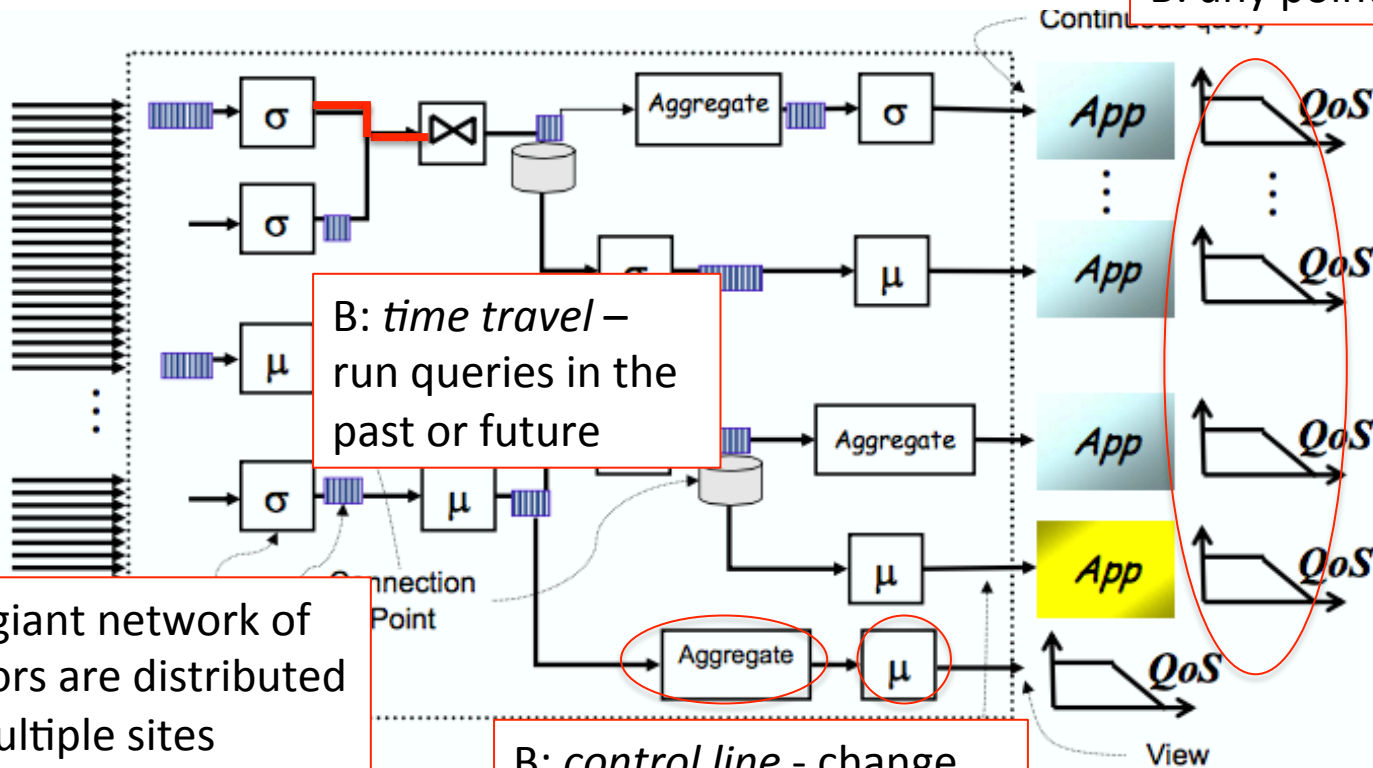
QoS: important metrics for resource management



# Extended in Borealis

A: tuples are appended to streams only.  
 B: revision – tuples can be inserted, deleted, or replaced

A: on the outputs  
 B: any point

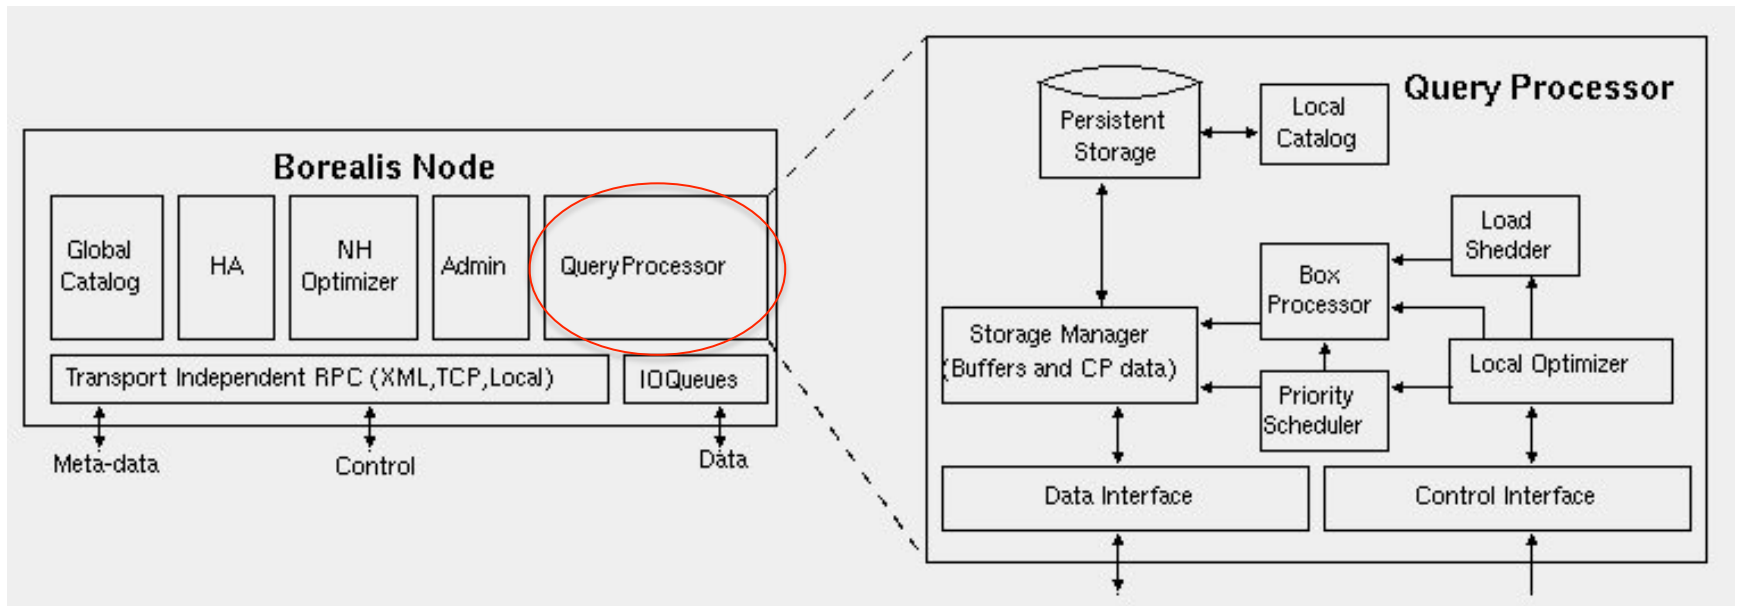


B: this giant network of operators are distributed over multiple sites

B: control line - change box semantics on the fly

# Borealis System Architecture

- Each site runs a Borealis server ...



# Borealis Stream Processing

“**Distributed** streaming processing engine with **more flexible stream processing model**” to meet the requirements of streaming applications.

- Built on top of Aurora + Medusa
- Advanced features for flexible stream proc.
  - Dynamic revision of query/results
  - Dynamic operator network optimization