

Expose

Experimental Performance Evaluation
of Stream Processing Engines Made Easy

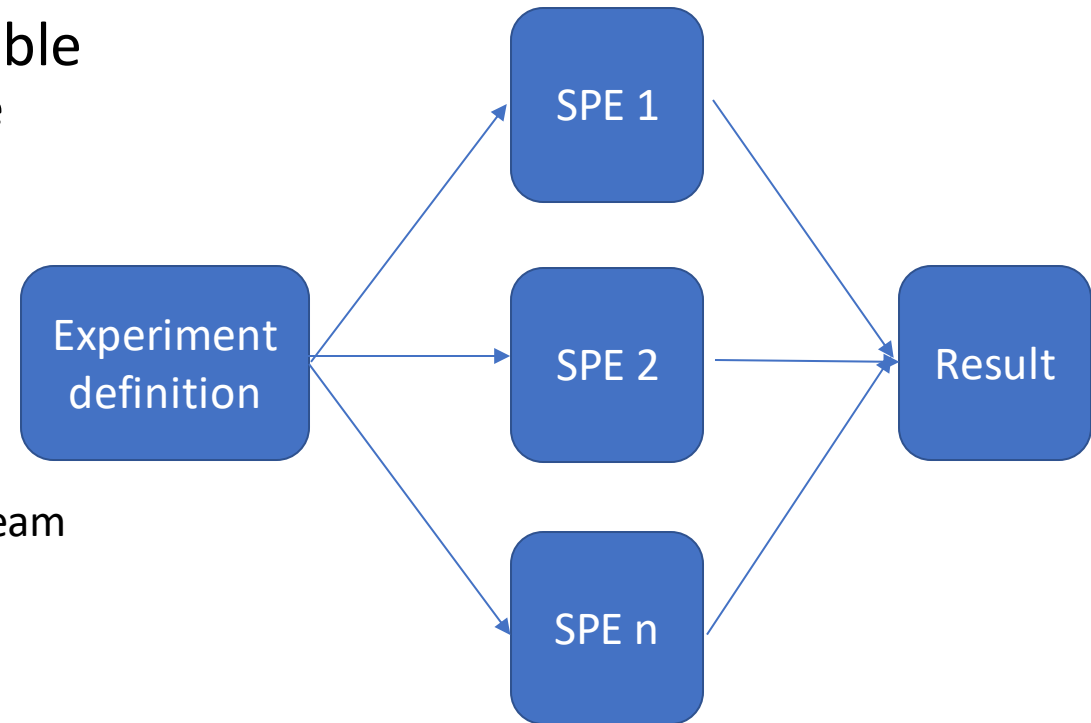
Authors: Espen Volnes, Thomas Plagemann, Vera Goebel, and Stein Kristiansen
University of Oslo

Challenges with benchmarking SPEs

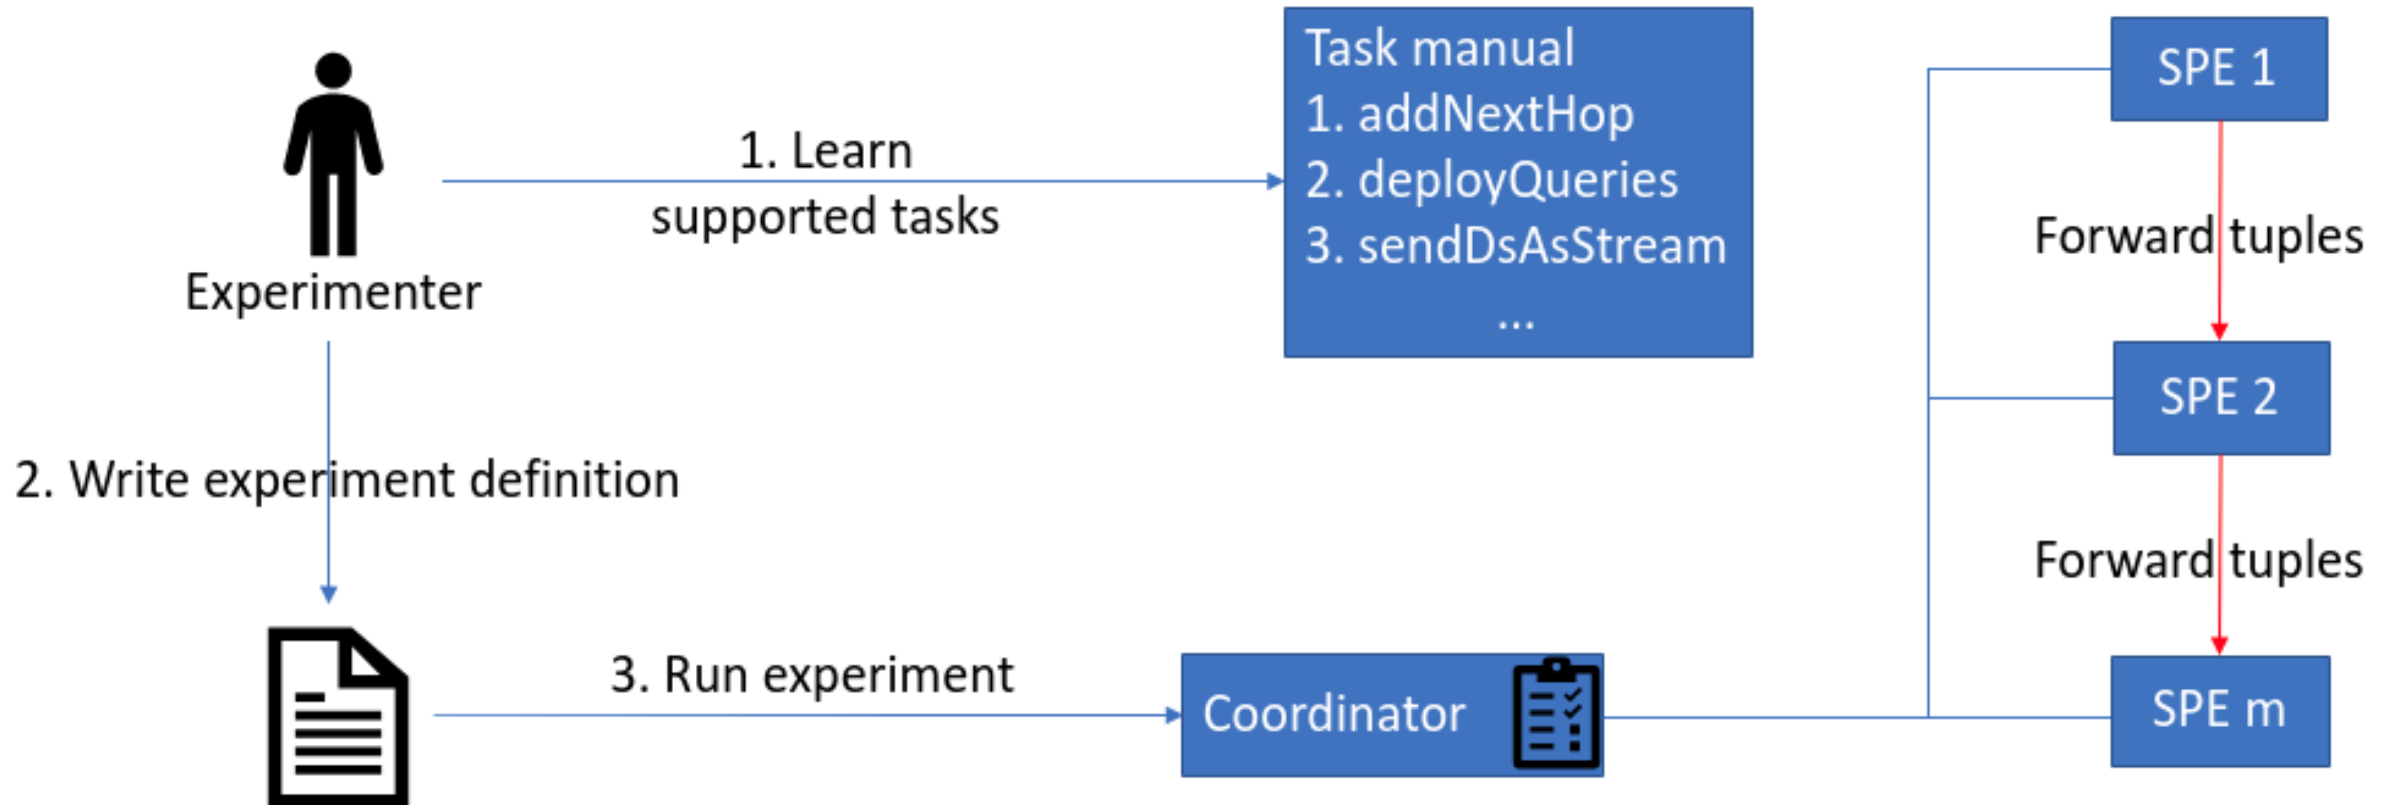
- Different implementations, similar features
- No standard for defining experiments
- Distribution demands coordination
- Fair comparisons of SPEs?

Design principles

- Experiments
 - Easy to define
 - Run on different SPEs
 - Sequence of high-level tasks
- Performance measurement must be comparable
 - Experiment definition determines what SPEs trace
- Tasks
 - Method calls
 - SPE experts implement API
 - Examples:
 - DeployQueries – Deploys query to node
 - AddNextHop – Forward stream to node
 - SendDsAsStream – Read dataset and convert it to stream
- Distribution
 - Coordinator manages experiment
 - Nodes register
 - Tasks can run on any node

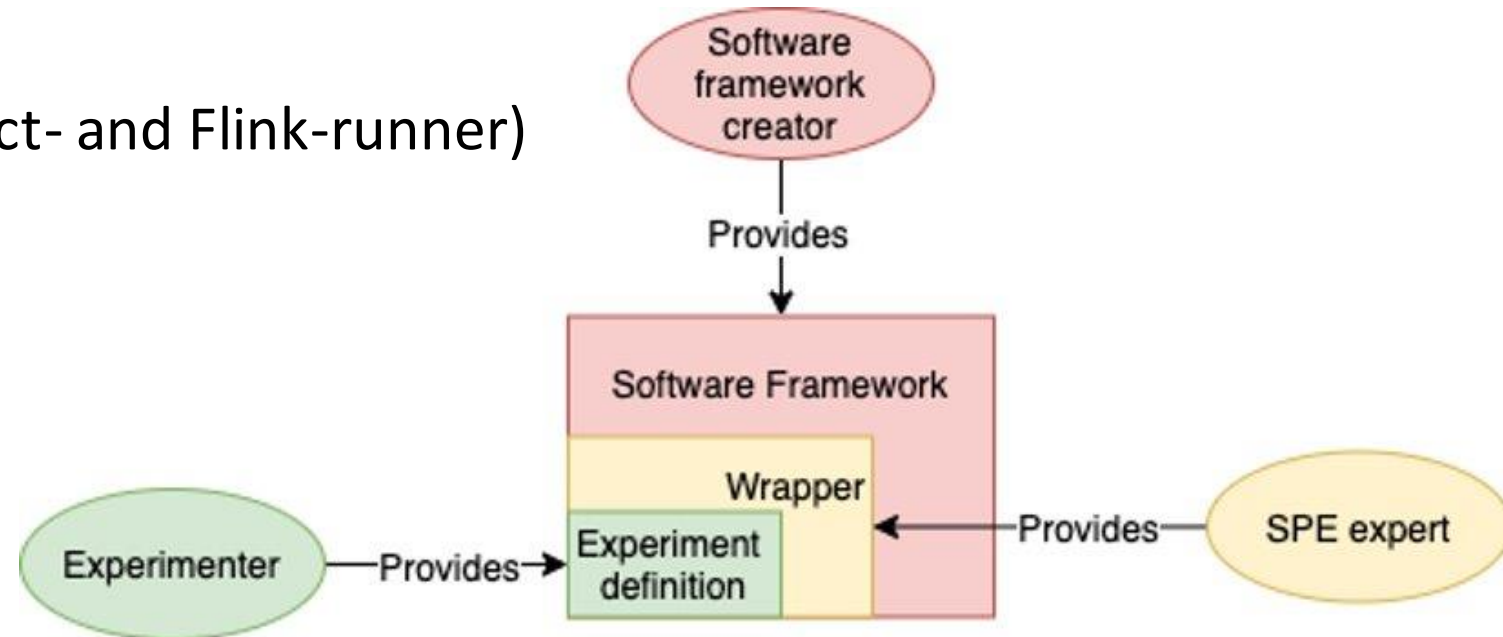


Workflow of using Expose



Stakeholders


- Three stakeholders:
 - Expose implementor
 - SPE expert
 - Experiment planner and executor
- Five wrappers:
 - Apache Beam (tested Direct- and Flink-runner)
 - Apache Flink
 - Siddhi
 - Esper
 - T-Rex



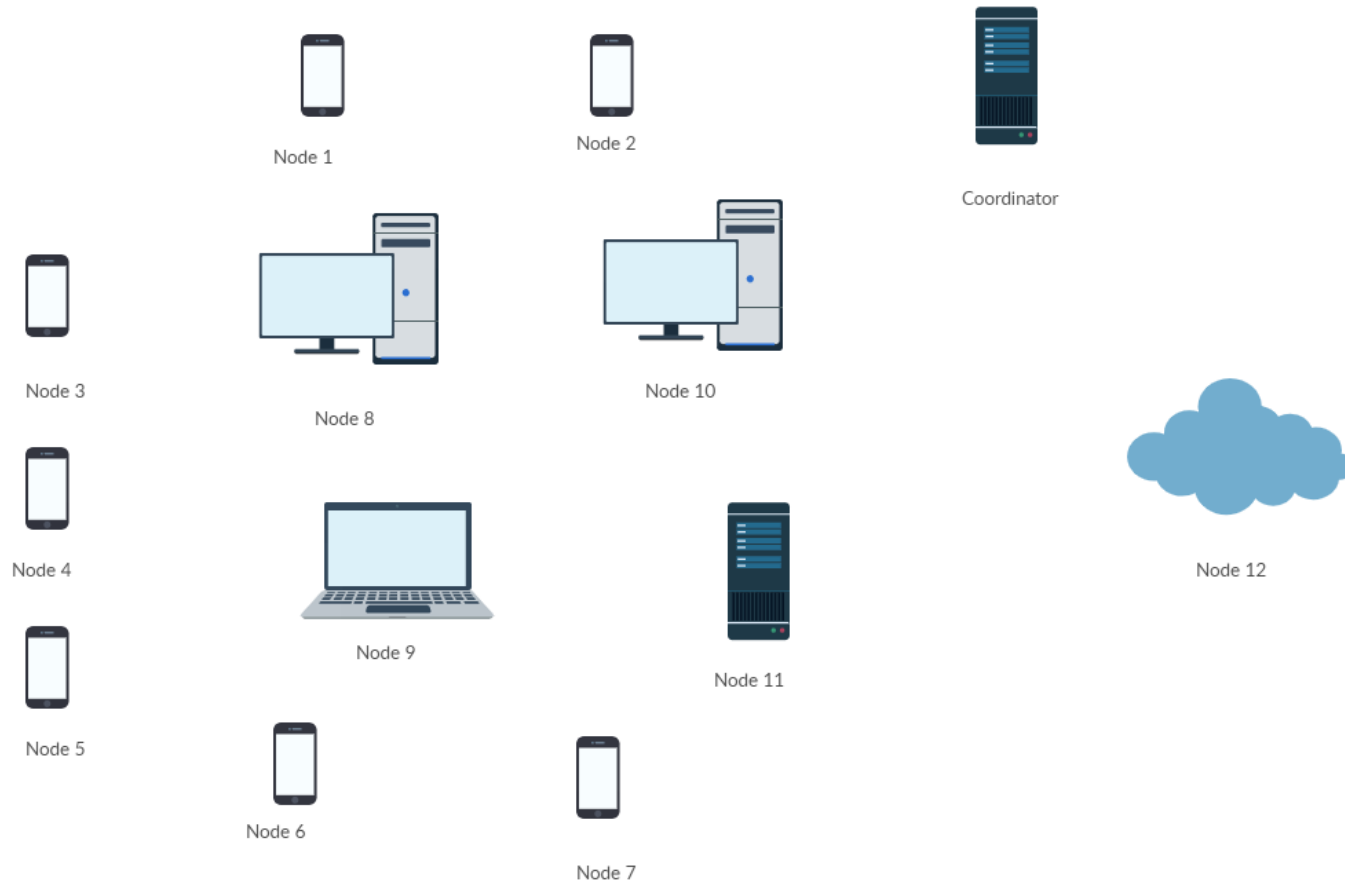
Use case: Example with many devices



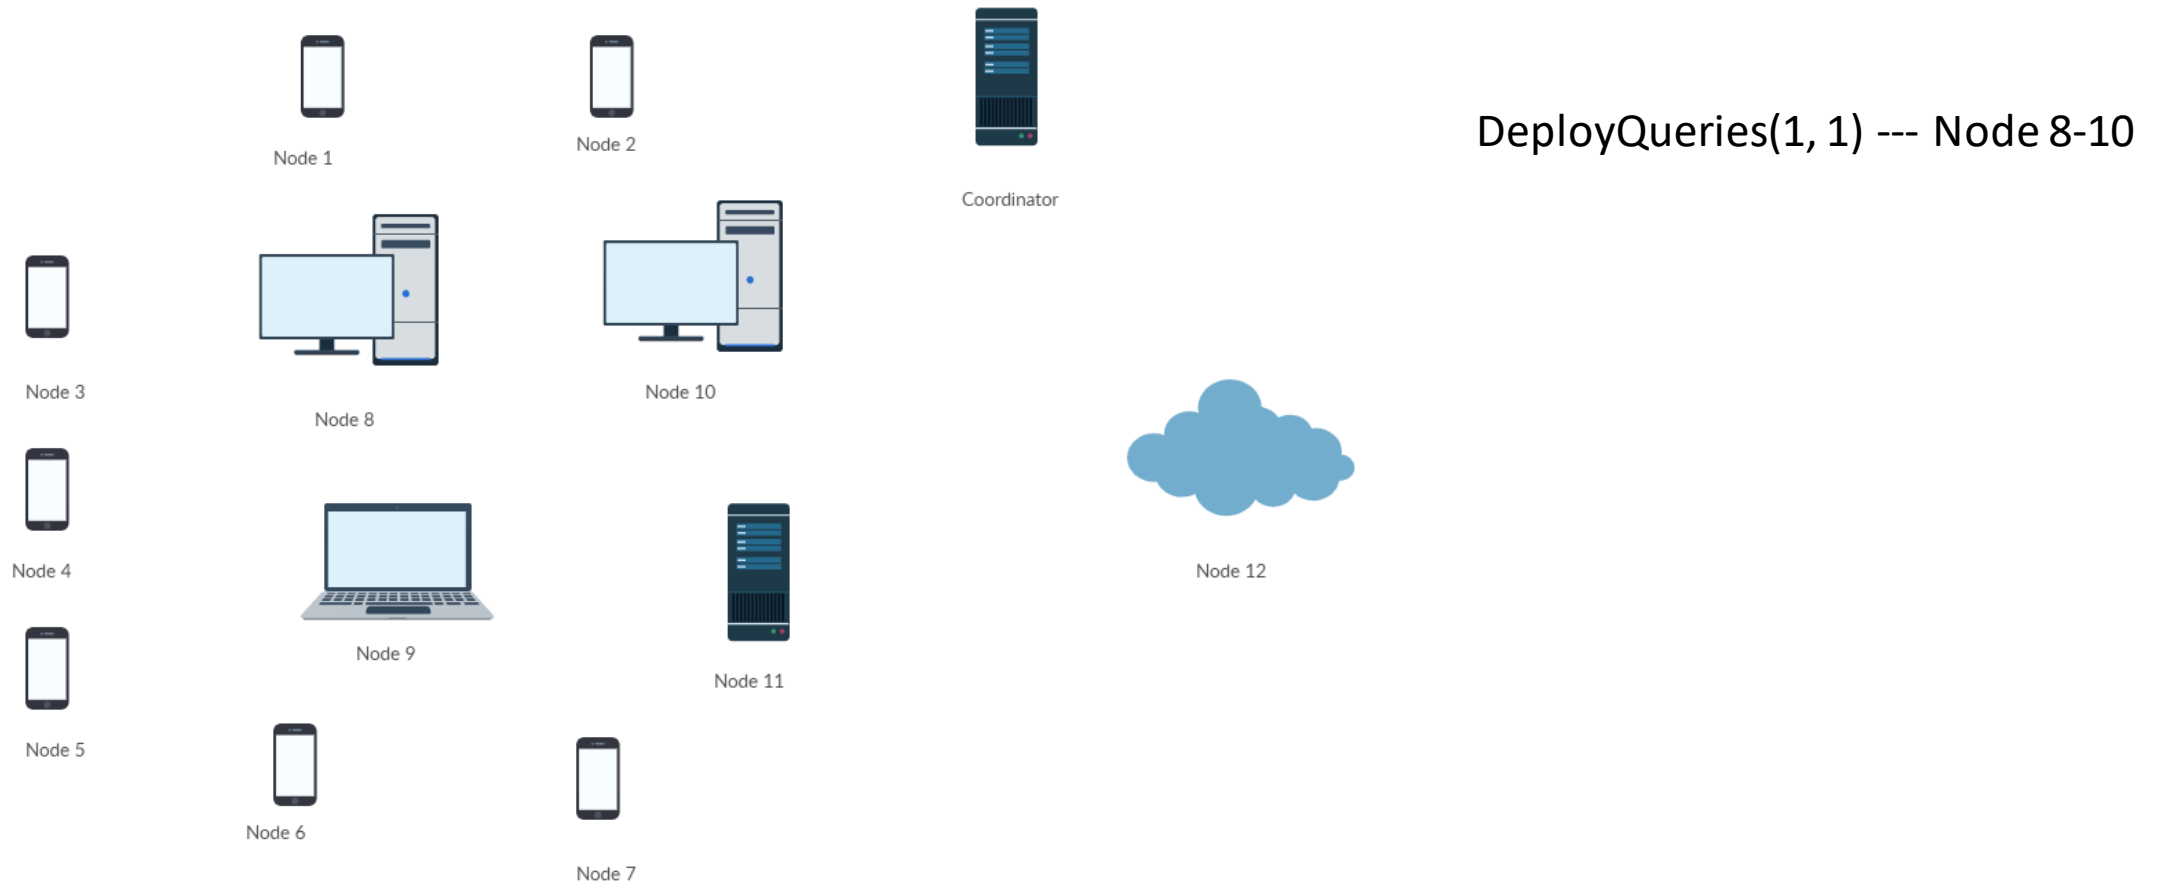
Use case: Example with many devices



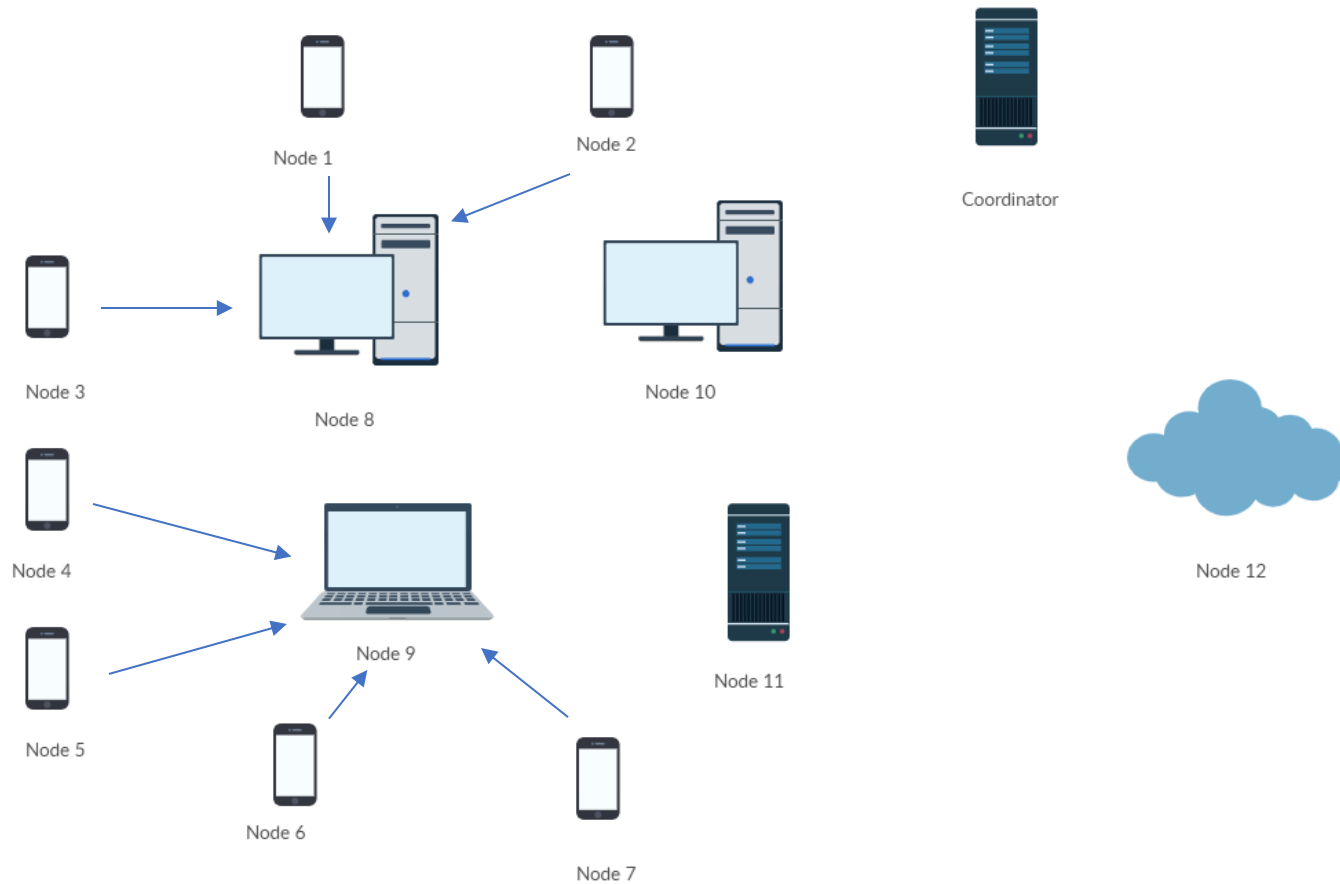
Use case: Example with many devices



Use case: Example with many devices



Use case: Example with many devices

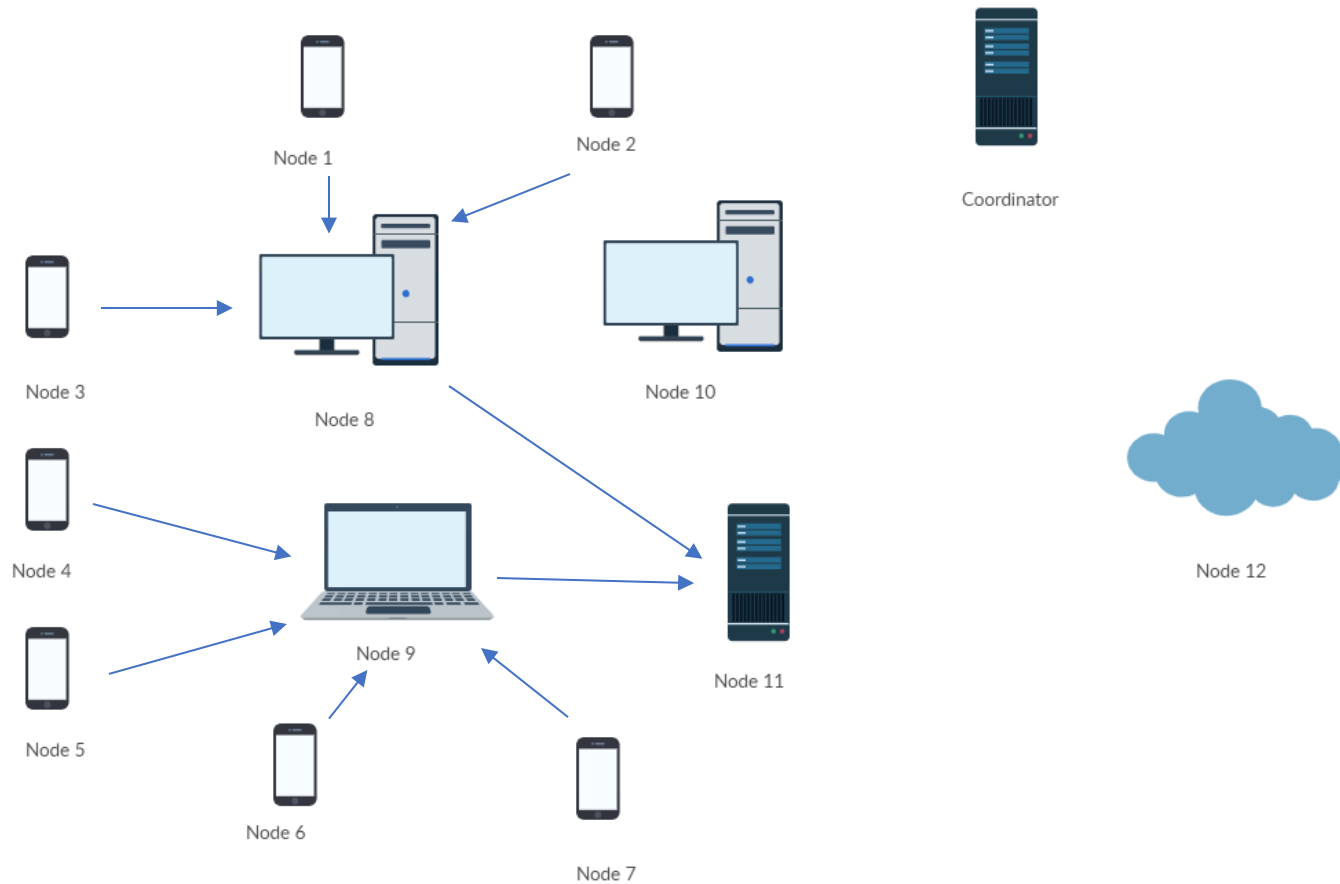


DeployQueries(1, 1) --- Node 8-10

AddNextHop(1, 8) --- Node 1-3

AddNextHop(1, 9) --- Node 4-7

Use case: Example with many devices



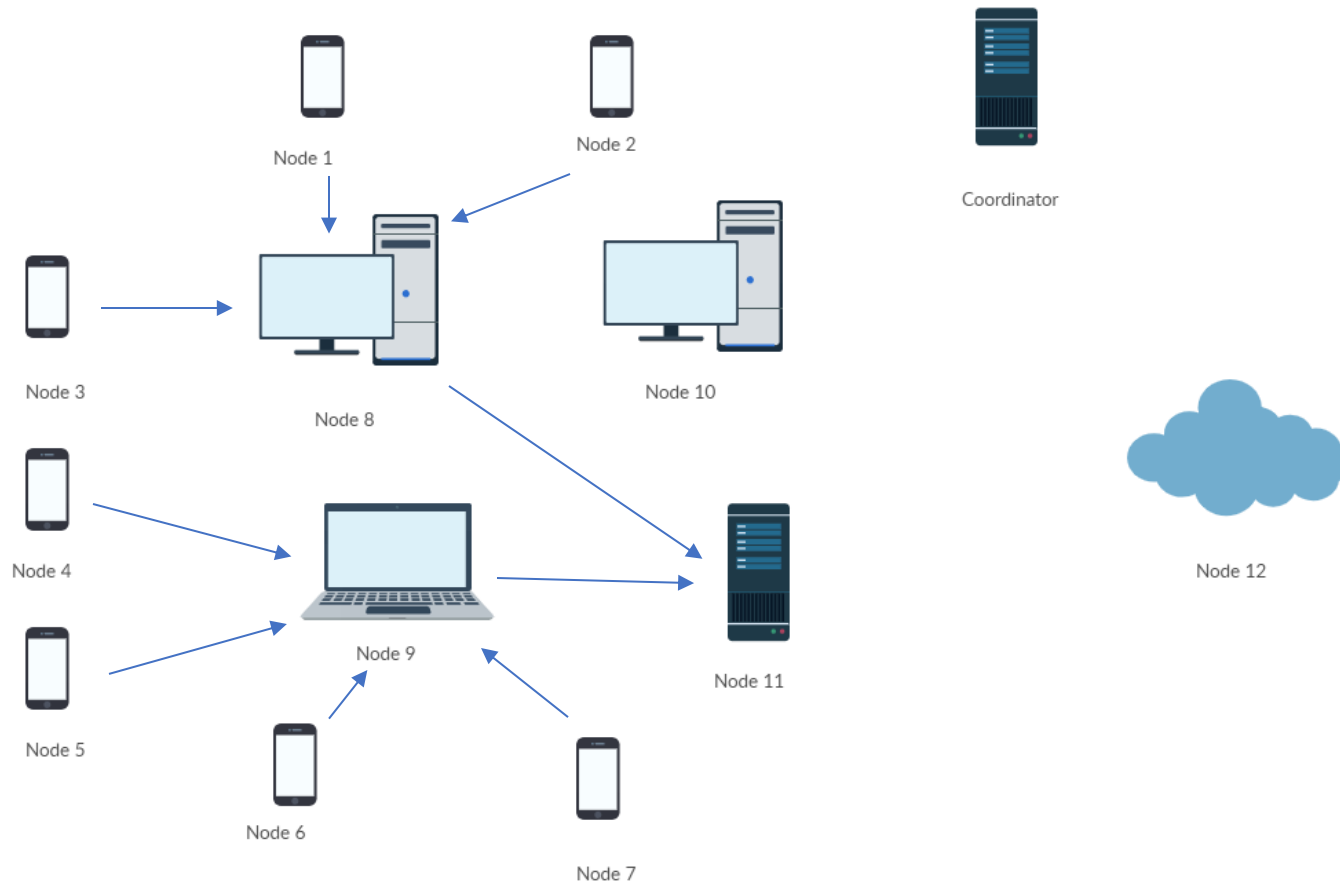
DeployQueries(1, 1) --- Node 8-10

AddNextHop(1, 8) --- Node 1-3

AddNextHop(1, 9) --- Node 4-7

AddNextHop(2, 11) --- Node 8-9

Use case: Example with many devices



DeployQueries(1, 1) --- Node 8-10


AddNextHop(1, 8) --- Node 1-3

AddNextHop(1, 9) --- Node 4-7

AddNextHop(2, 11) --- Node 8-9

DeployQueries(2, 1) --- Node 11

Use case: Example with many devices



DeployQueries(1, 1) --- Node 8-10

AddNextHop(1, 8) --- Node 1-3

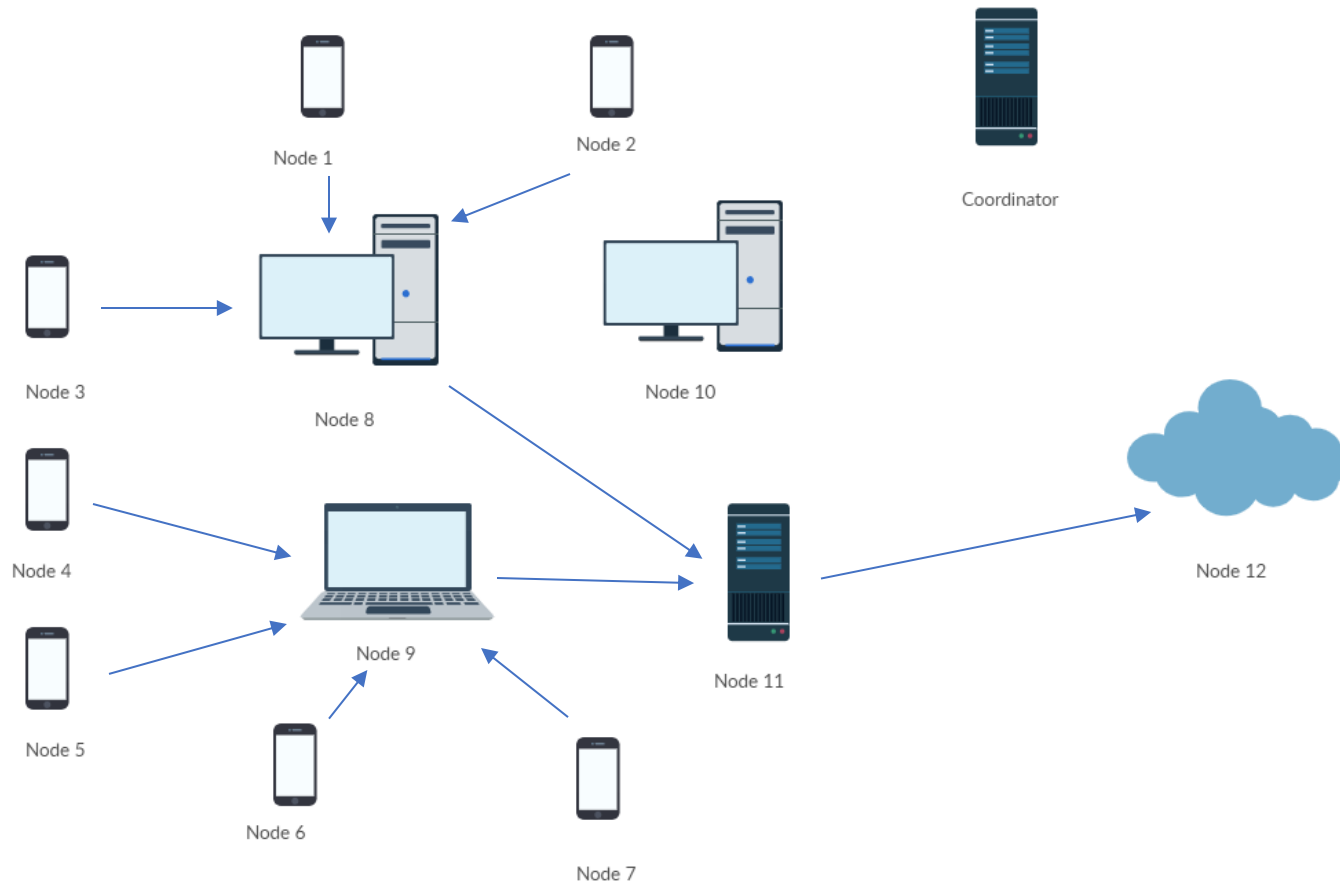
AddNextHop(1, 9) --- Node 4-7

AddNextHop(2, 11) --- Node 8-9

DeployQueries(2, 1) --- Node 11

AddNextHop(3, 12) --- Node 11

Use case: Example with many devices



DeployQueries(1, 1) --- Node 8-10

AddNextHop(1, 8) --- Node 1-3

AddNextHop(1, 9) --- Node 4-7


AddNextHop(2, 11) --- Node 8-9

DeployQueries(2, 1) --- Node 11

AddNextHop(3, 12) --- Node 11

WriteStreamToCsv(3) --- Node 12

Use case: Example with many devices



DeployQueries(1, 1) --- Node 8-10

AddNextHop(1, 8) --- Node 1-3

AddNextHop(1, 9) --- Node 4-7

AddNextHop(2, 11) --- Node 8-9

DeployQueries(2, 1) --- Node 11

AddNextHop(3, 12) --- Node 11

WriteStreamToCsv(3) --- Node 12

SendDsAsStream(15) --- Node 1-7

Use case: NEXMark benchmark

- Auction dataset and SQL queries
- Two machines: data center capable server, and edge device
 - Five SPEs:
 - Apache Beam with the Flink runner
 - Apache Flink
 - Siddhi
 - Esper
 - T-Rex
- SUT sends output tuples to the sink
- Throughput of the SUT is measured



Node 1
Node 3
Coordinator



Node 2 (SUT)
Raspberry Pi 4



Node 1
Node 3
Coordinator



Node 2 (SUT)
Intel Xeon

Thank you!

- Future work
 - Distributed coordination
 - Operator migration (distributed CEP)

- Expose is available on <https://github.com/espv/expose>