33. Narrowing the Gap Between Serverless and its State with Storage Functions



Utah Scalable Computer Systems UT University of Utah *Tian Zhang, Dong Xie, Feifei Li, Ryan Stutsman

Motivation - Pushing code to data. Flexibility of customized data mode at storage.

- Gap between serverless and it's states.
- Users pay for additional **idle time** waiting for data from external storage service.
- Pushing code to data greatly narrows the gap.
- V8 language runtime to isolate code, much lower isolation cost than process isolation.
- Users can leverage the flexibility of general programming language to implement customized data structures and logic.



Design - Data store and CSA builtin co-design to avoid runtime exits.

- Data store implemented C++.
- Embed V8 runtime to isolate functions.
- Functions access data through V8 builtin across language boundary.
- **Problem**: boundary crossing costs.
- **Optimization**: map data into V8 runtime.
- Allow functions to access data without exiting V8 runtime.
- Eliminate boundary crossing costs.
- Data are stored in hash tables in data store.
- Data store implements data lookup function in C++.
- CSA builtin implements the same data lookup logic.
- **CSA** (CodeStubAssembler) is the intermediate representation used by V8 runtime internally.
- JavaScript functions access data through CSA builtin



Results - Reduce data movements. High speed data access with CSA.

- Projection, query the first 4 bytes of a value.
- Without Shredder clients need to fetch the whole value.
- Bound by network bandwidth.
- Push projection logic to Shredder to reduce
- Graph query over Facebook social graph.
- Shredder **60X** better performance than remote get/put.
- CSA **3X** better performance than w/o CSA.
- A 1 hop query visits about 25 nodes.
- A 2 hop query visits about 700 nodes.
- A 4 hop query visits more than 400k nodes.
- CSA enables direct access to in-memory data (achieves **10s of GB/s** data access

data movements.

speed).





