


FAAS ORCHESTRATION OF PARALLEL WORKLOADS

Daniel Barcelona Pons, Pedro García López, Álvaro Ruiz
Ollobarren, Amanda Gómez-Gómez, Gerard París, Marc
Sánchez Artigas



H2020 CLOUDBUTTON: SERVERLESS DATA ANALYTICS

- Large Serverless European Research Project (4.3M)
- The main goal is to create CloudButton: a Serverless Data Analytics Platform. CloudButton will “democratize big data” by overly simplifying the overall life cycle and programming model thanks to serverless technologies
- To demonstrate the impact of the project, we target two settings with large data volumes: bioinformatics (genomics, metabolomics) and geospatial data (LiDAR, satellital)



The Serverless Supercomputer Stack

Applications: Big data, analytics, ML/AI, simulations, r/t video transcoding, ...

Video

ML

LinPack

Monte
Carlo

M-R

Analytics

CRISPR

Serverless Dataflow Graphs (SDGs)

Real-time Scheduler

Low-latency serverless storage

Networking

The ExCamera
stack



Cloud Vendors: Functions, APIs, Managed Service Portfolios

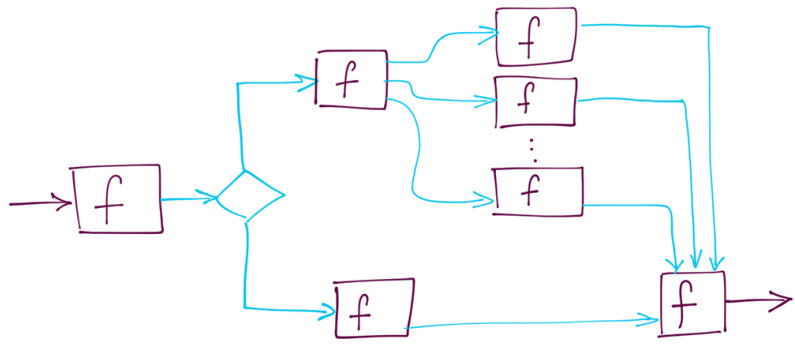
Systems	Components	
	<i>Serverful</i>	<i>Serverless</i>
Locus	Scheduler, Redis	Lambda Functions, S3
PyWren	Scheduler	Lambda Functions, S3
IBM PyWren	Scheduler	IBM Cloud Functions, COS, RabbitMQ
ExCamera	Coordinator, Rendezvous	Lambda Functions, S3
Flint	Scheduler	Lambda Functions, S3, SQS
NumPyWren	Provisioner	Lambda Functions, S3
Cirrus	Scheduler, Parameter Servers	Lambda Functions, S3

SERVERMIX, SERVERLESS, SERVERFUL


ServerMix: Tradeoffs and Challenges of Serverless Data Analytics




COMPARISON (WOSC4)



■ USER FUNCTIONS
■ WORKFLOW FRAMEWORK



AWS Step Functions



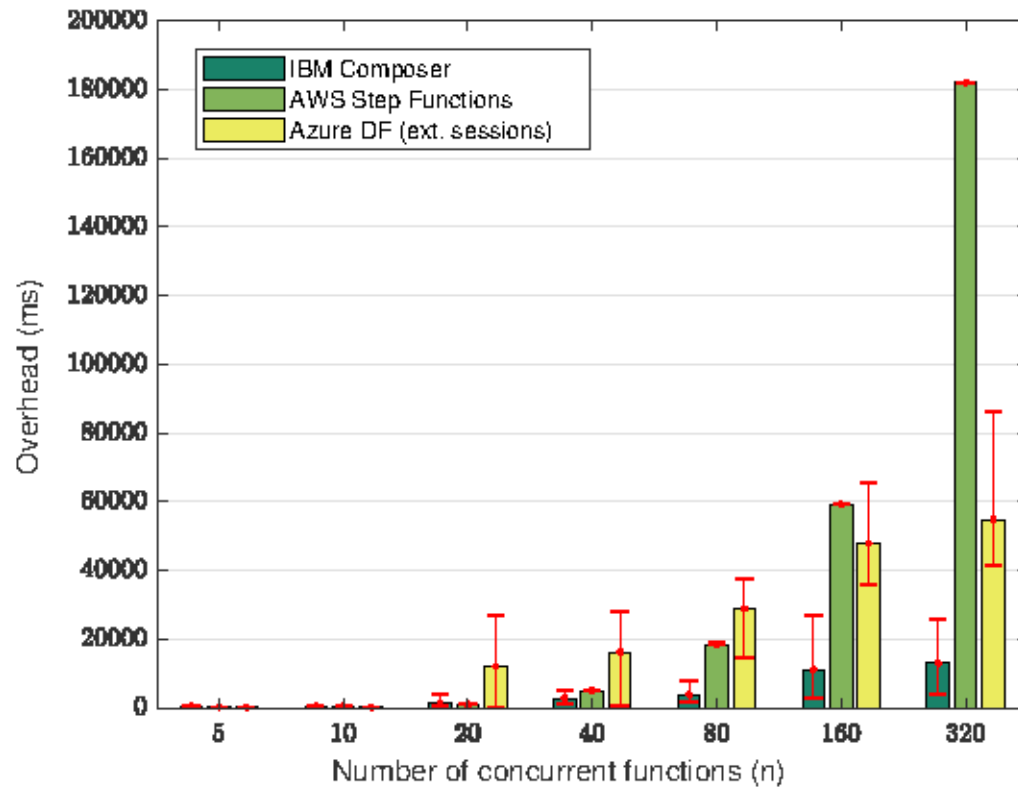
Azure Durable Functions



IBM Function Composer



COMPARISON (WOSC5)




Is IBM Composer the answer now ?



AWS Step Functions

State Machine



Event Sourcing

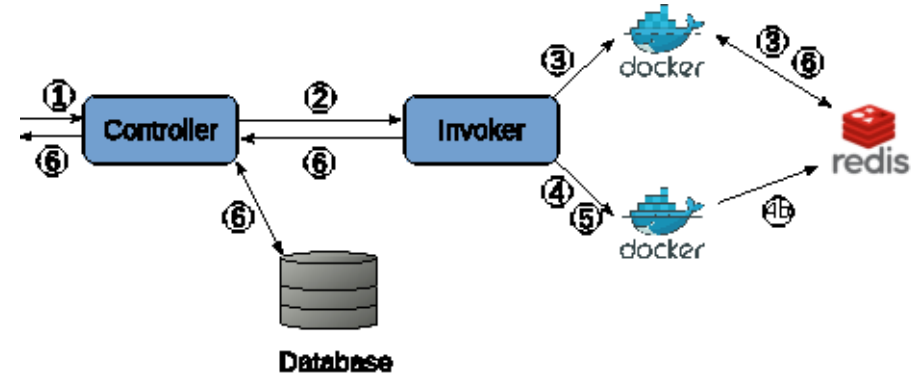
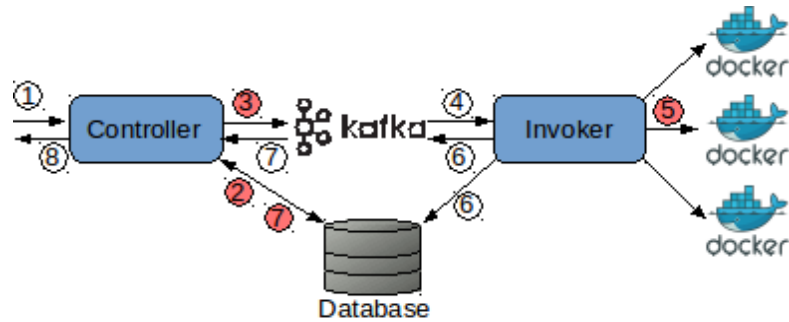
Azure Durable Functions



AST
continuations



OPENWHISK AND COMPOSER

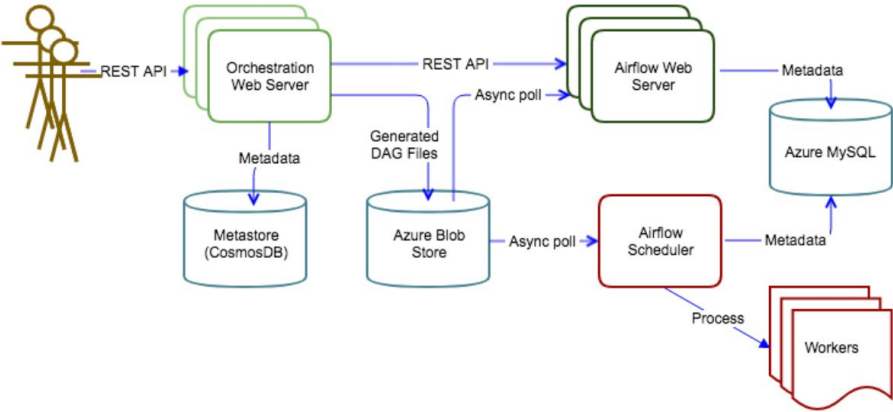
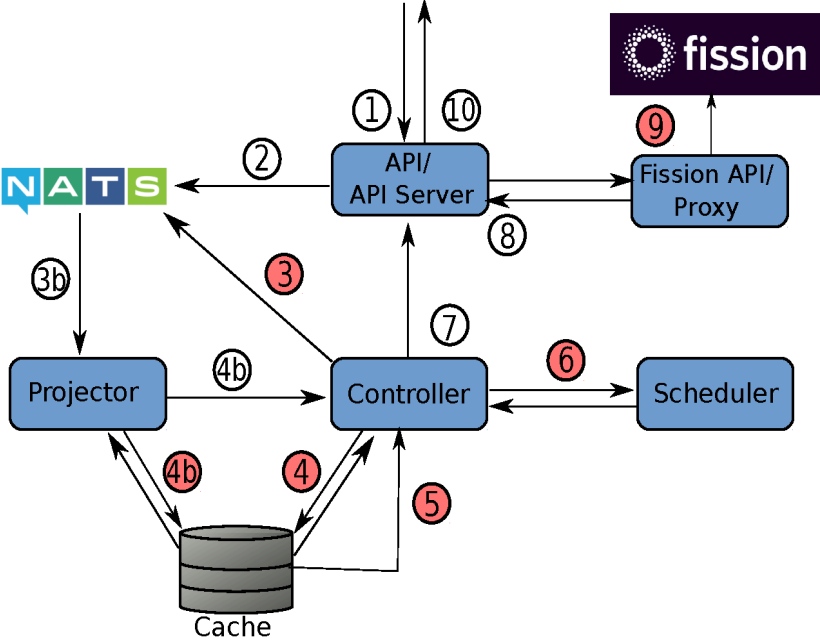


COMPOSER LIMITATIONS

- Coupled Design with OpenWhisk Conductor Actions
- Resource inefficiencies when launching parallel functions
 - High Variability in experiments
- Difficult billing of conductor and secondary actions
- Designed for short-running workflows
- Limited fault-tolerance
- It relies on a user-provided **serverful** Redis Service for fork-join aggregation



FISSION WORKFLOWS, AIRFLOW



AWS EXPRESS WORKFLOWS (BONUS)

- Amazing performance for 320 workers, overheads below 1 second
- Valid for short-running workflows (5 mins), function working all time
- It is relaxing fault-tolerance and checkpointing
 - At-least-once workflow execution vs Exactly-once workflow execution.
 - Message logs, not current state
- Different pricing model
 - Express Workflows: Priced by the number of executions you run, their duration, and memory consumption.
 - Standard Workflows Priced per state transition. A state transition is counted each time a step in your execution is completed.



CONCLUSIONS

- **Innovation in the Cloud providers is needed !!!**
- Existing FaaS Orchestration systems are not currently designed for supporting parallel Big Data pipelines
 - Both IBM Composer and Amazon Express Functions have now good performance
 - But they are both designed for short-running workflows
- Messages to Cloud providers:
 - Cloud providers must offer the required tools to orchestrate parallel Data Analytics tasks
 - Cloud providers must offer event-based building blocks to enable the construction of third-party schedulers
 - Multi-tenancy, cost, and resource efficiency will be of paramount importance for these services

