

Portals.

Stateful Serverless Dataflow Streaming

Jonas Spenger (RISE, KTH), Paris Carbone (RISE, KTH), Philipp Haller (KTH)

This work was partially funded by the Swedish Foundation for Strategic Research (SSF grant no. BD15-0006) and by Digital Futures.

INTRODUCTION

All modern services are distributed collections of **microservices**. Current frameworks are ill-equipped for the complete job of building distributed services with: reliability, scalability, and privacy of complex business logic applications.


Portals aims to be the ultimate tool for building large ecosystems of microservices. It is a new distributed programming model for the next era of stateful serverless, with the goal to make it easy to build reliable, fault-tolerant, scalable, privacy-compliant distributed services.

HIGHLIGHTS

Portals provides built-in automated support for:

- **Scalability** through data-parallelism and task-parallelism
- **Exactly-once-processing**, end-to-end, across services, across deployments
- **Data-flow composition** with cycles
- **Request-reply** style communication + futures
- **Live consistent updates** to the running system

- Atomic Streams, Workflows



Atom:

- Sequence of events, transactional unit of computation.

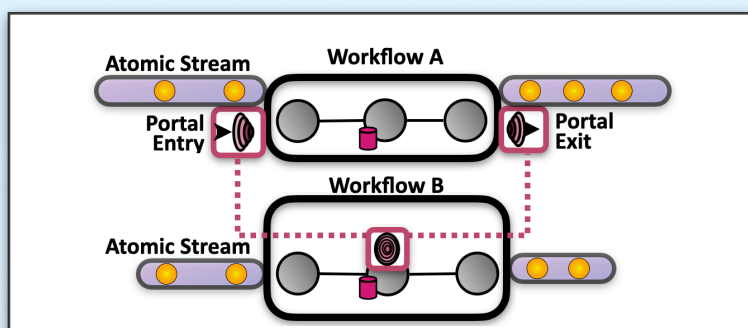
Workflow:

- Consumes atomic streams, produces atomic streams, DAG of tasks.

The Atomic Processing Contract:

- "The consumer/producer must always consume and process the whole atom, before consuming and processing the next atom."

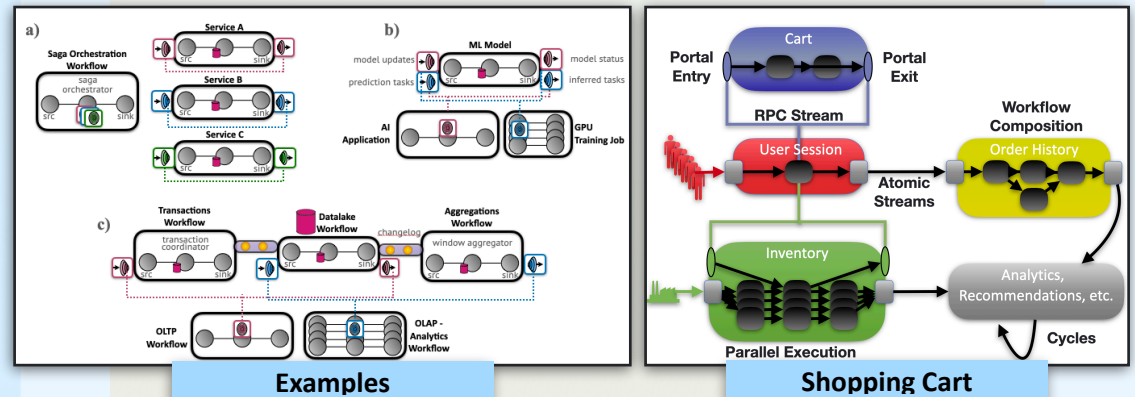
- "Portals"



Portal entry, portal exit:

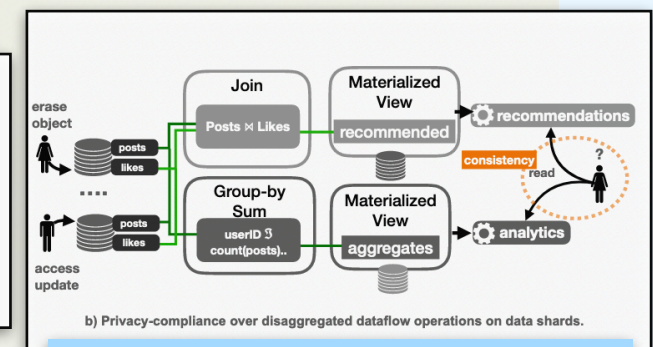
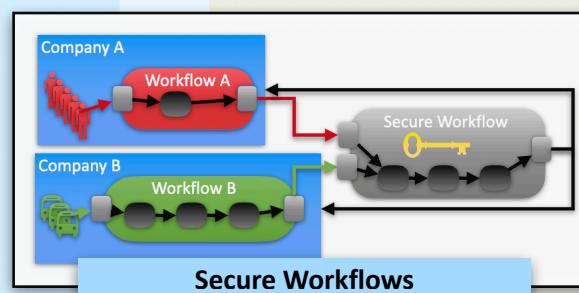
- Request-reply style programming with workflows, includes futures API

- Service Composition



- End-to-end guarantees across workflows, services, deployments, organisations. Complex compositional patterns, cycles.

- Privacy-by-Design



- Privacy compliance by design, GDPR, integration of secure multi-party computation.

THE PORTALS SYSTEM

- Implementation of the Portals programming model.
- Sign up for launch at www.portals-project.org
- <https://github.com/portals-project/portals>

ROADMAP

2022

- Release v1.0: distributed runtime, end-to-end exactly-once, workflows, atomic streams, etc.

2023

- Release v2.0: privacy, secure multi-party computation on workflows

CONCLUSIONS

Portals aims to be the best for complex service composition, with end-to-end guarantees across workflows, services, deployments, organisations.

[Poly'21] Spenger, J., Carbone, P., & Haller, P. (2021). WIP: Pods: Privacy Compliant Scalable Decentralized Data Services. In Heterogeneous Data Management, Polystores, and Analytics for Healthcare (pp. 70-82). Springer, Cham.

The presented work is under submission to Onward! at SPLASH 2022 Presented at CASTOR Software Days, 31 August 2022, KTH, Stockholm, Sweden

