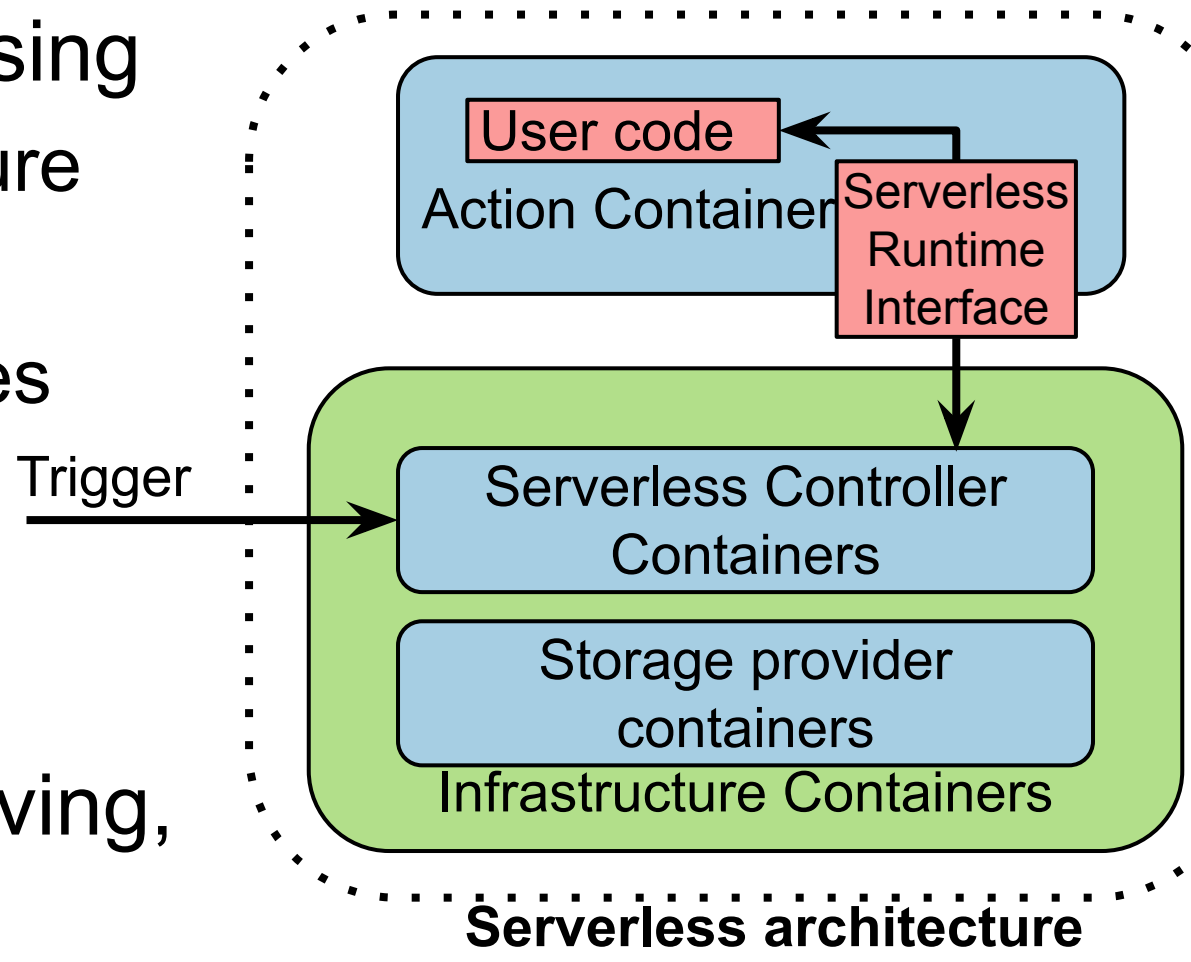


F3: Serving Files Efficiently in Serverless Computing

Alex Merenstein¹, Vasily Tarasov³, Ali Anwar², Scott Guthridge³, and Erez Zadok¹
 Stony Brook University; ²University of Minnesota; ³IBM Research - Almaden  

Background & Motivation

- Serverless adoption is increasing
 - Reduces burden of infrastructure management on developers
 - No longer pay for idle resources
 - Not just for new applications: existing systems can also benefit
- Serverless platforms are evolving, allowing for more generic computation
 - Restrictions being relaxed (e.g., longer run times, larger resource allocations allowed)
 - Arbitrary container images supported

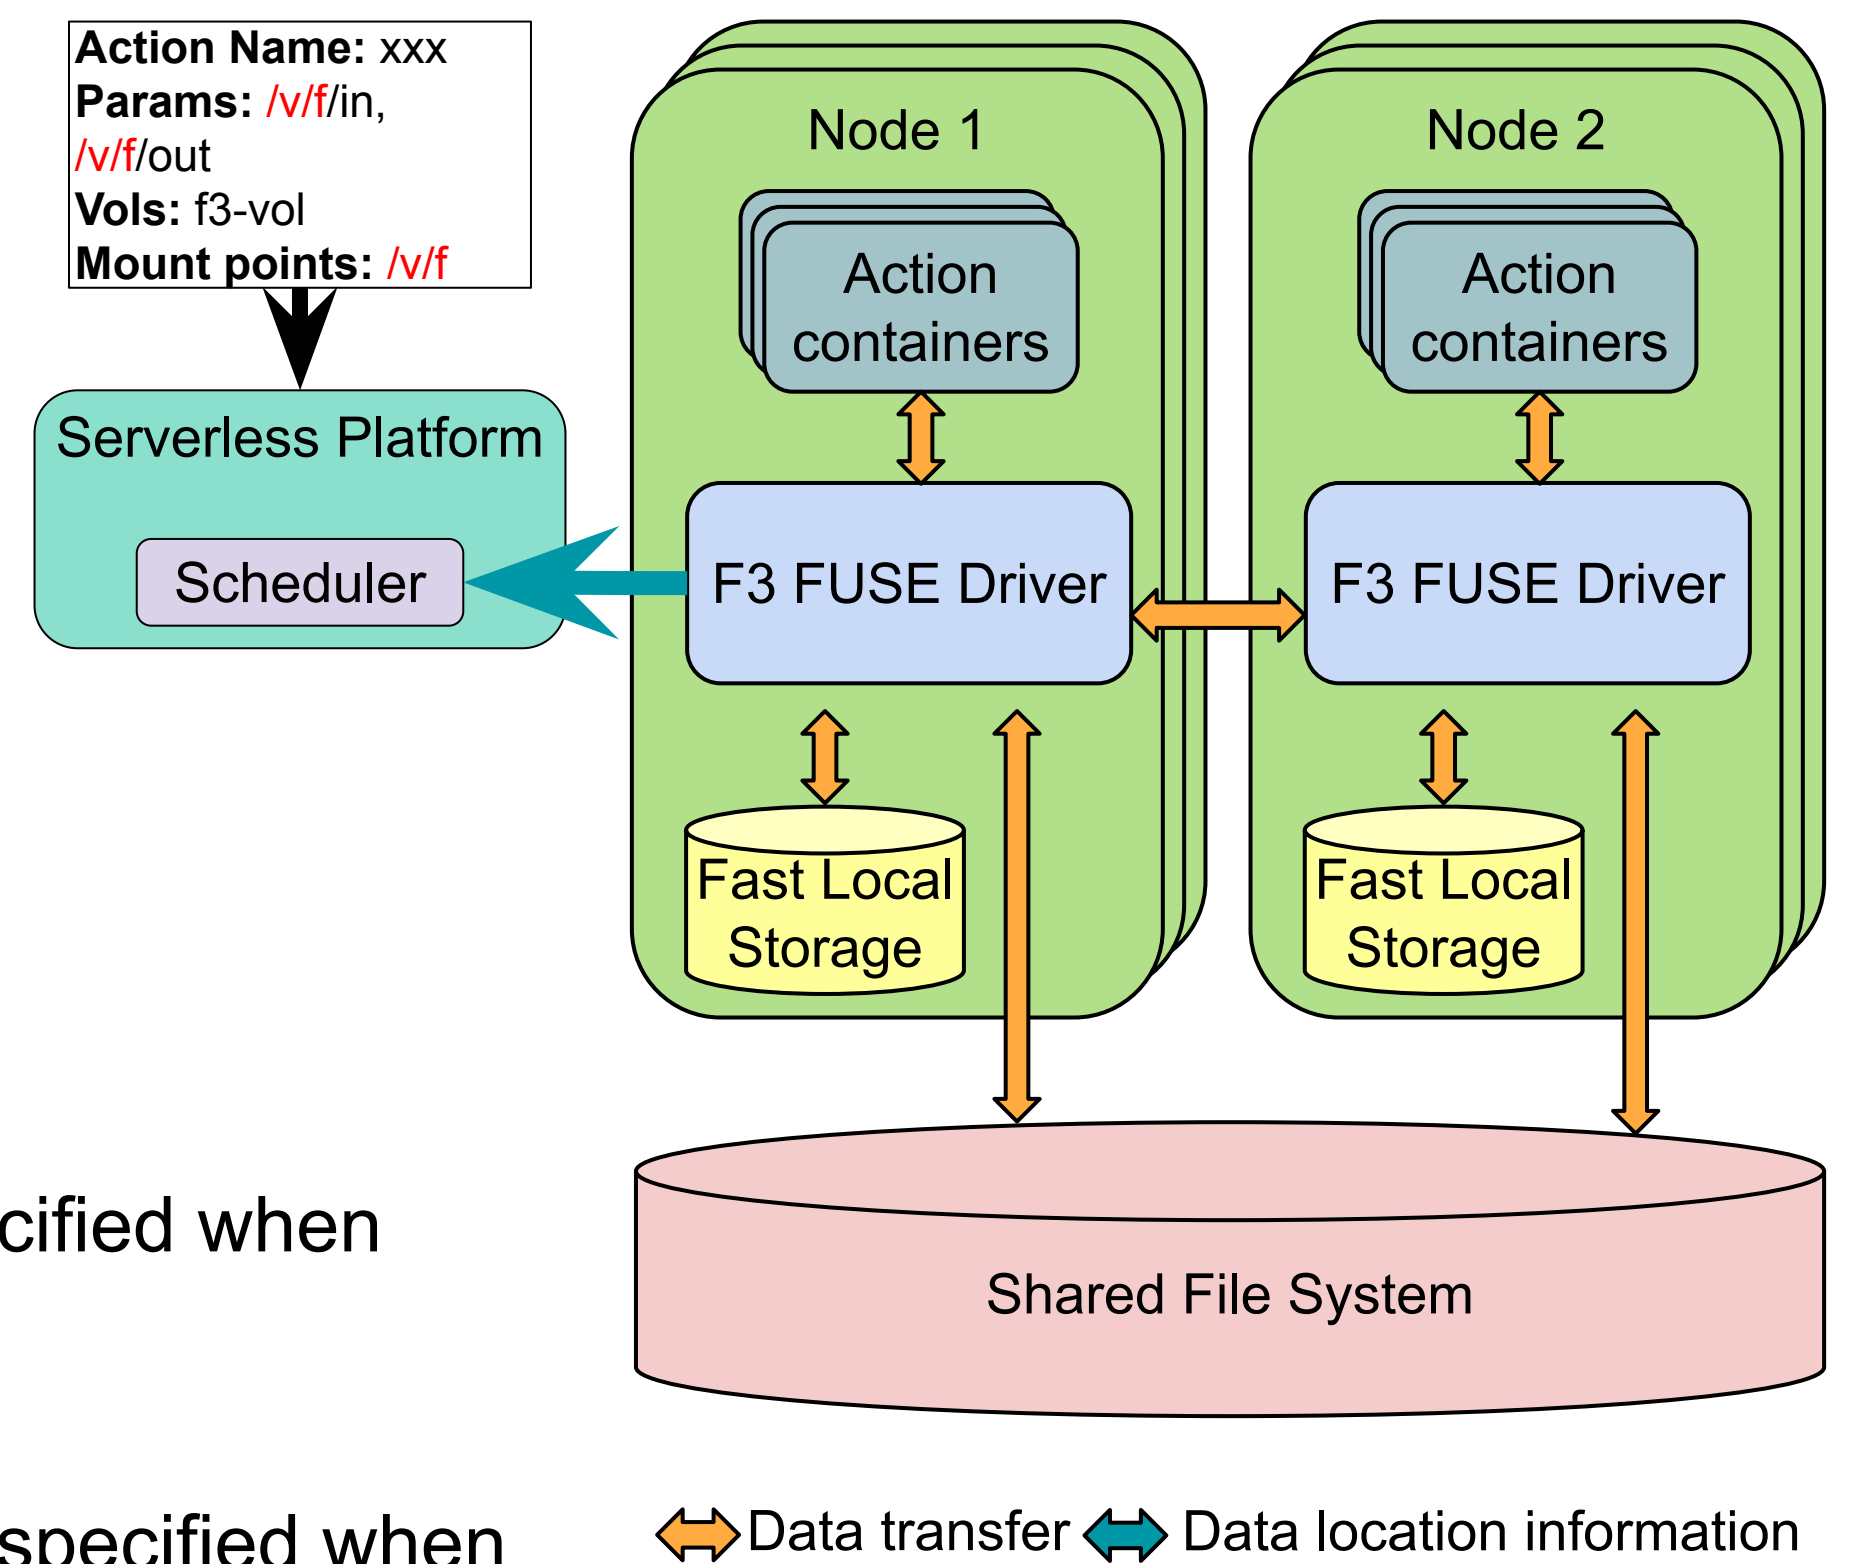


Storage for Serverless

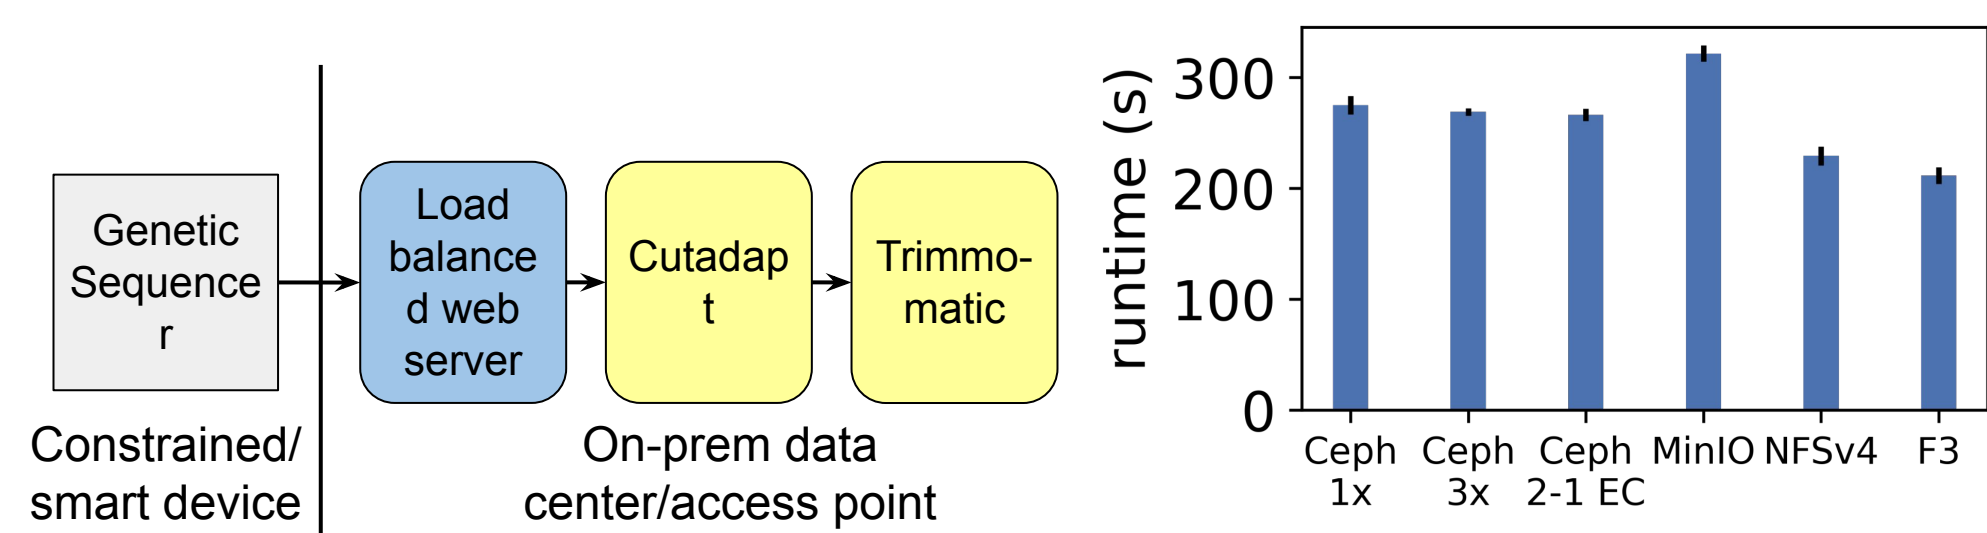
- Often object rather than file based
 - Many Existing applications already utilizing file interfaces
- No mechanism for providing data locality information to serverless schedulers
- Unnecessary durability features (some data is short lived, often easy to recreate)
 - Also need ability to store some data persistently
- In distributed file systems, overly cautious locking, not needed for write-once-read-many access patterns

F3 Design & Implementation

- Layers on top of shared file system
 - Uses fast local storage for ephemeral data
 - Uses shared file system for durable data and metadata
 - Special support for read-while-write access patterns
 - Interfaces with serverless scheduler to provide locality info
- Implementation
 - F3 FUSE based file system
 - CSI driver to interface with Kubernetes
 - OpenWhisk modifications
 - Enable attaching storage volumes to action containers (specified when action container is started)
 - Uses data locality information when scheduling actions
 - Enable using arbitrary container image for running actions (specified when action container is created)




Evaluation: Bioinformatics pipeline



8-34% shorter runtime vs. other storage systems

Evaluation: fio Benchmarks



Up to 6.46× higher throughput & 2.73× lower latency vs. other storage systems

Conclusion

- Problem
 - Serverless has potential for widespread adoption
 - Lack of file based storage options limiting this potential
- Solution
 - Implemented F3, a file system for serverless workloads
 - Modified open source serverless platform, OpenWhisk, to enable utilizing F3 volumes
- Evaluation/comparison to other storage systems
 - 1.40× to 6.46× higher throughput
 - 1.40× to 2.73× lower latency
- Ongoing work
 - Exploring other points along reliability-cost-performance continuum
 - Developing model to choose optimal placement of data