

Non-orthodox applications of serverless computing

1. How far can we go beyond standard “reactive” or “event driven” scenarios:
 - High throughput scientific workflows, bioinformatics (sequence analysis)
 - More tightly coupled computations (PyWren?)
 - Spark on AWS Lambda, deep learning (GPU analogy)
2. Resource management challenges
 - User vs. provider perspective
 - How many “knobs” we can turn and how much can we gain
 - Granularity of billing vs. startup overheads, long running functions
 - Volunteer/p2p computing
3. What tooling is needed to make it all useful
 - Compilers, runtime supports
 - Deployment, transparency, “faasification”
 - Debugging, performance and cost profiling