

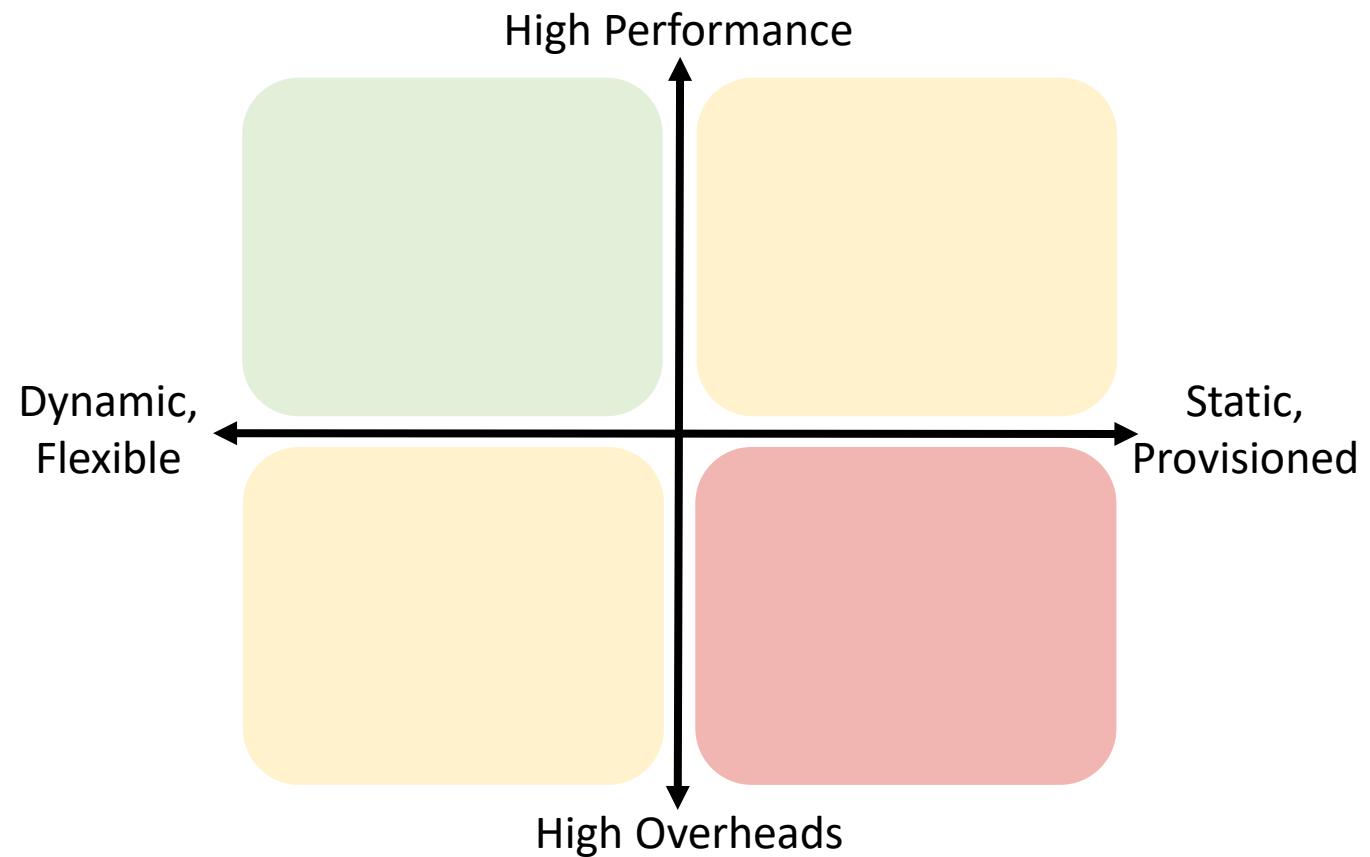


Interactive Computing with Serverless Functions in rFaaS

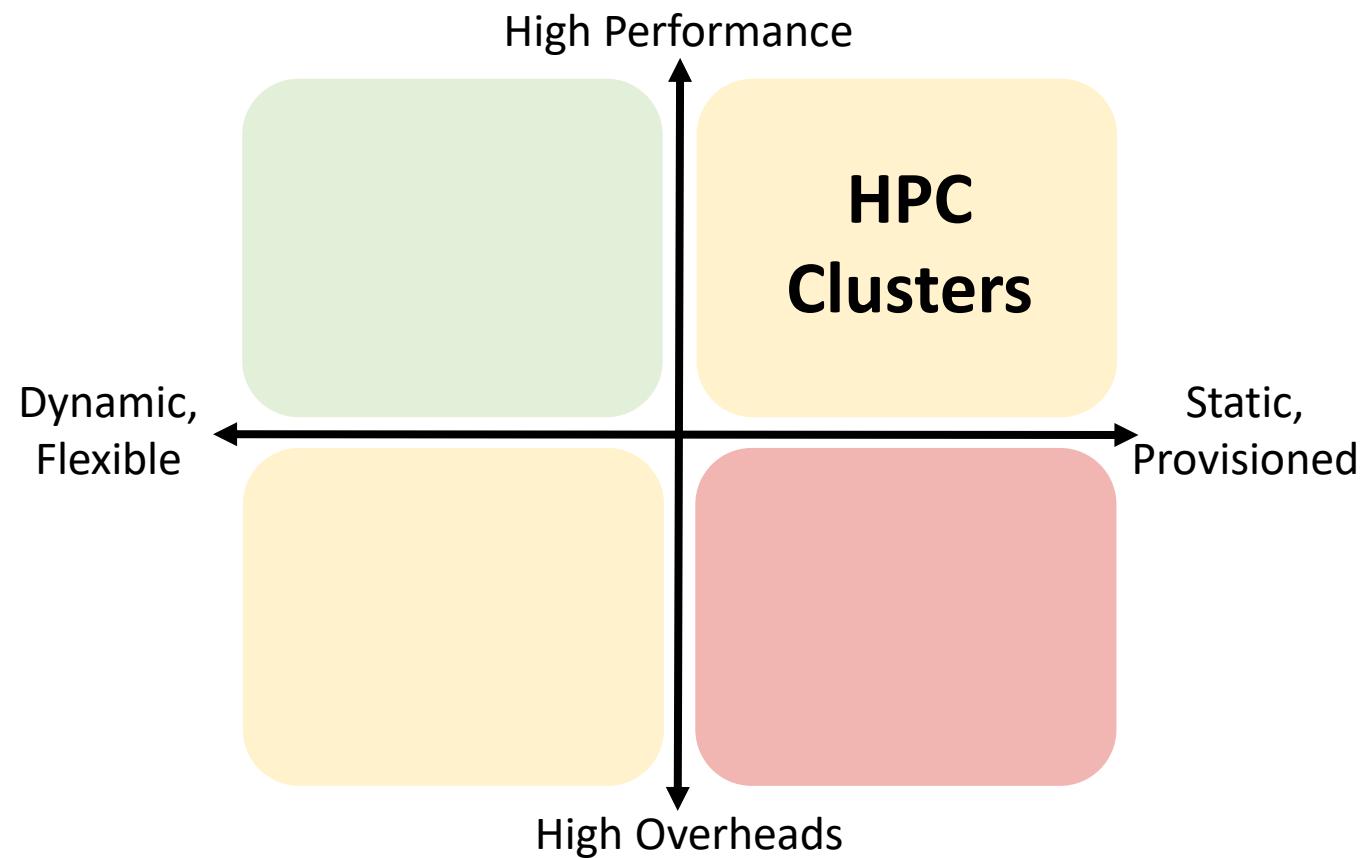
[Marcin Copik](#), Konstantin Taranov, Alexandru Calotoiu, Torsten Hoefler

ETH Zurich

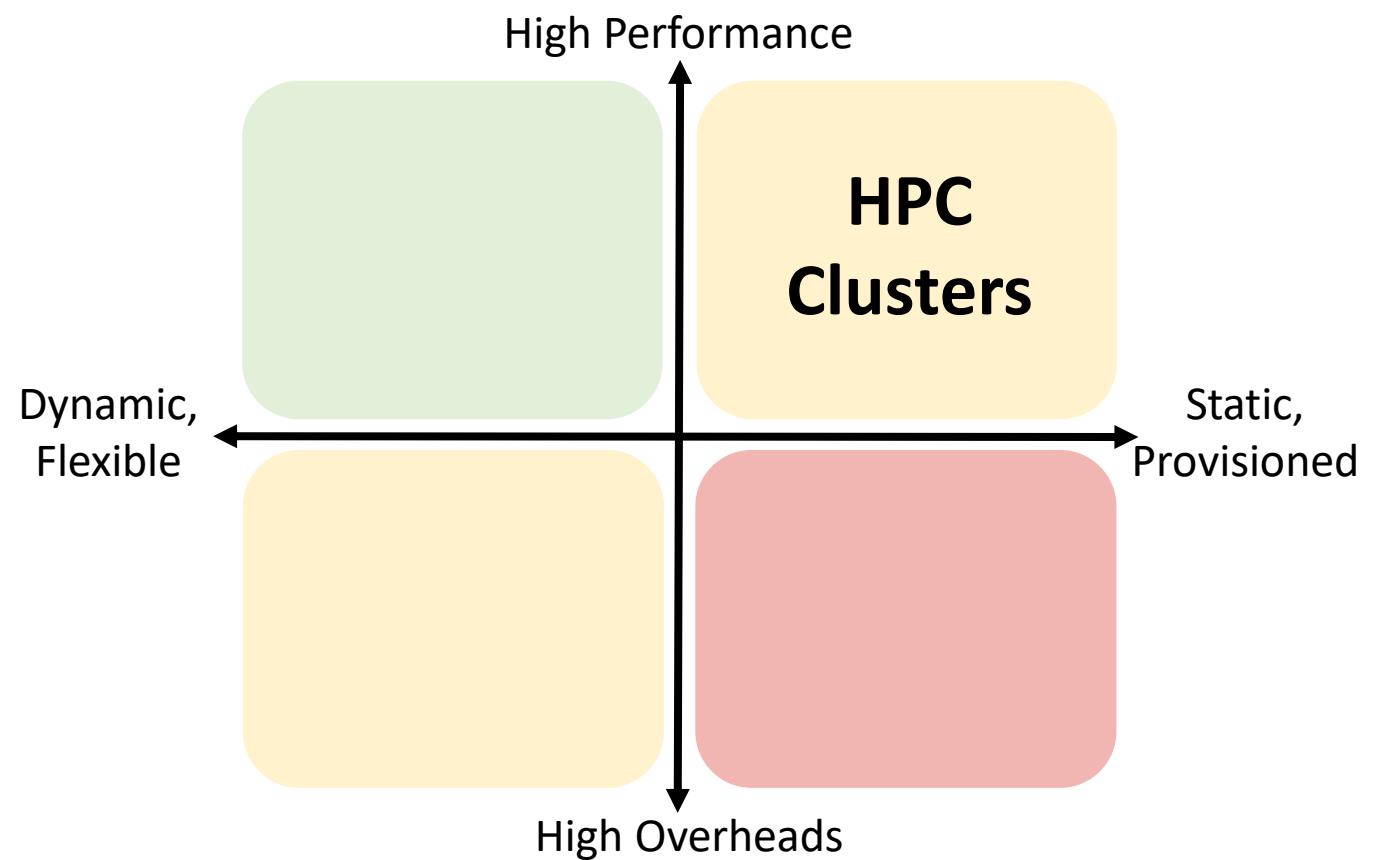
Function-as-a-Service for HPC



Function-as-a-Service for HPC



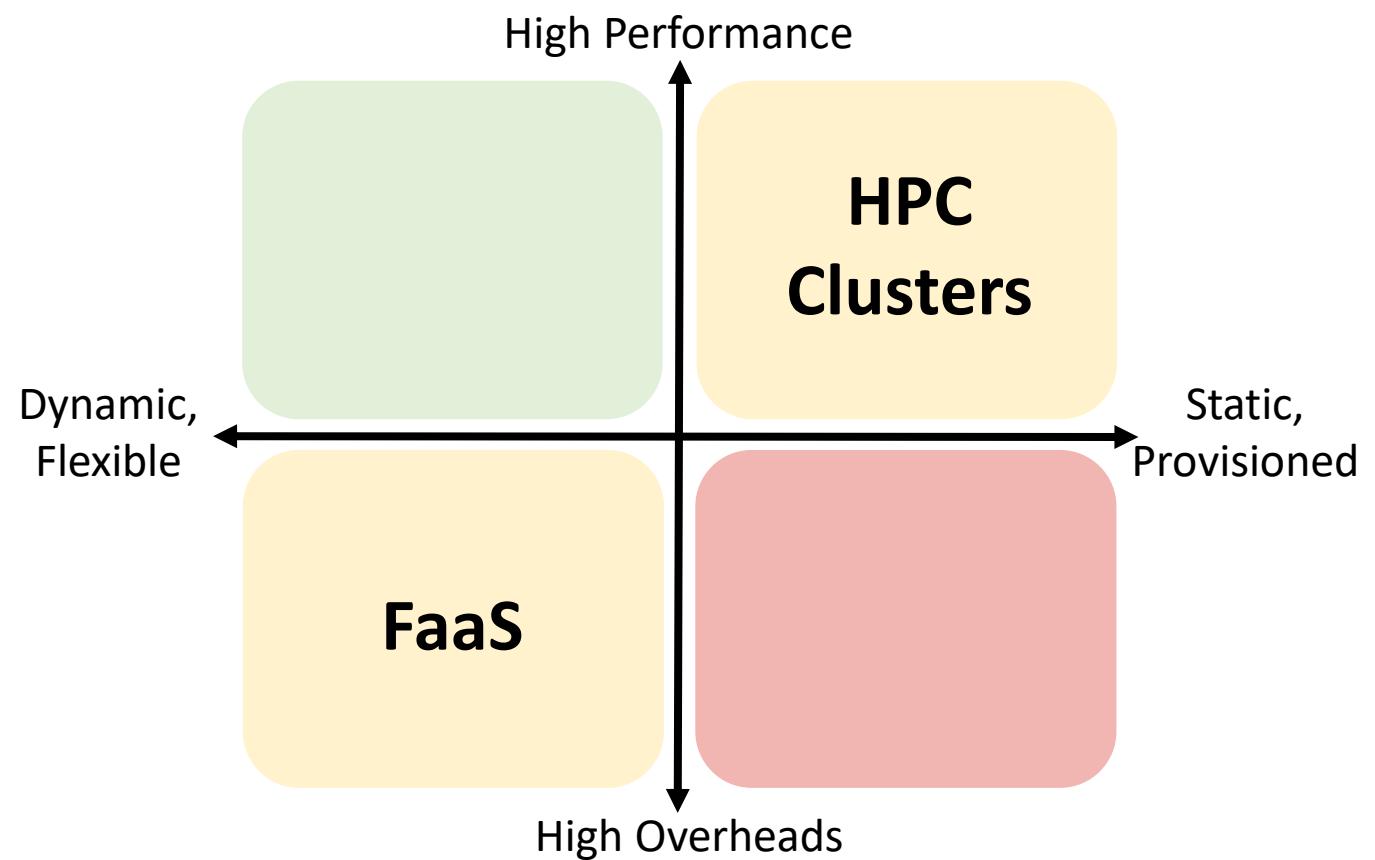
Function-as-a-Service for HPC



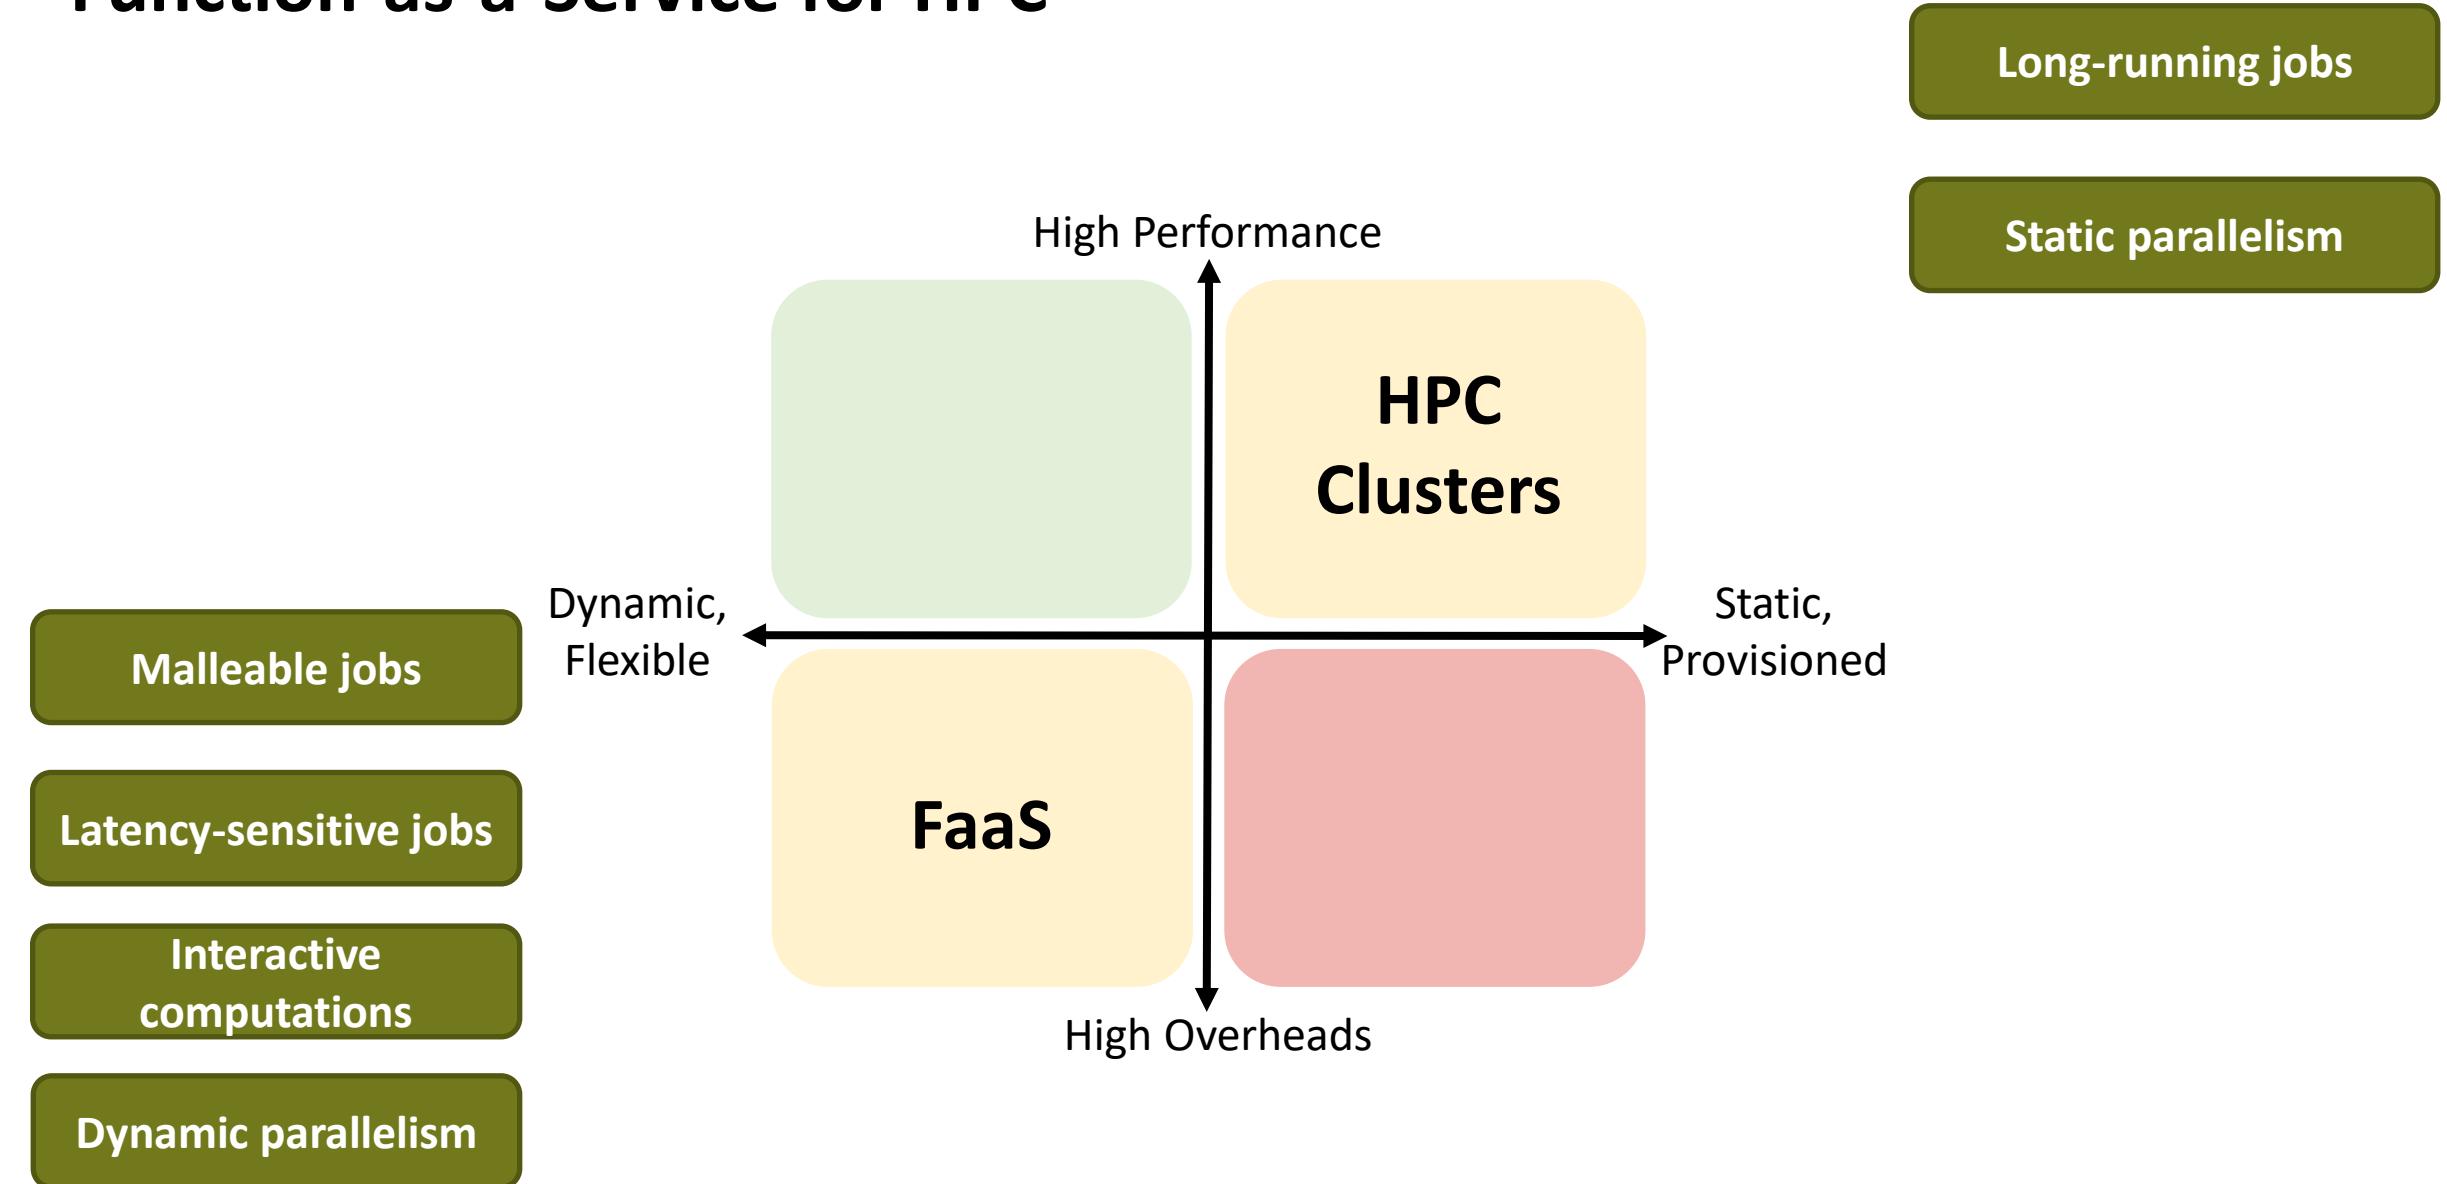
Long-running jobs

Static parallelism

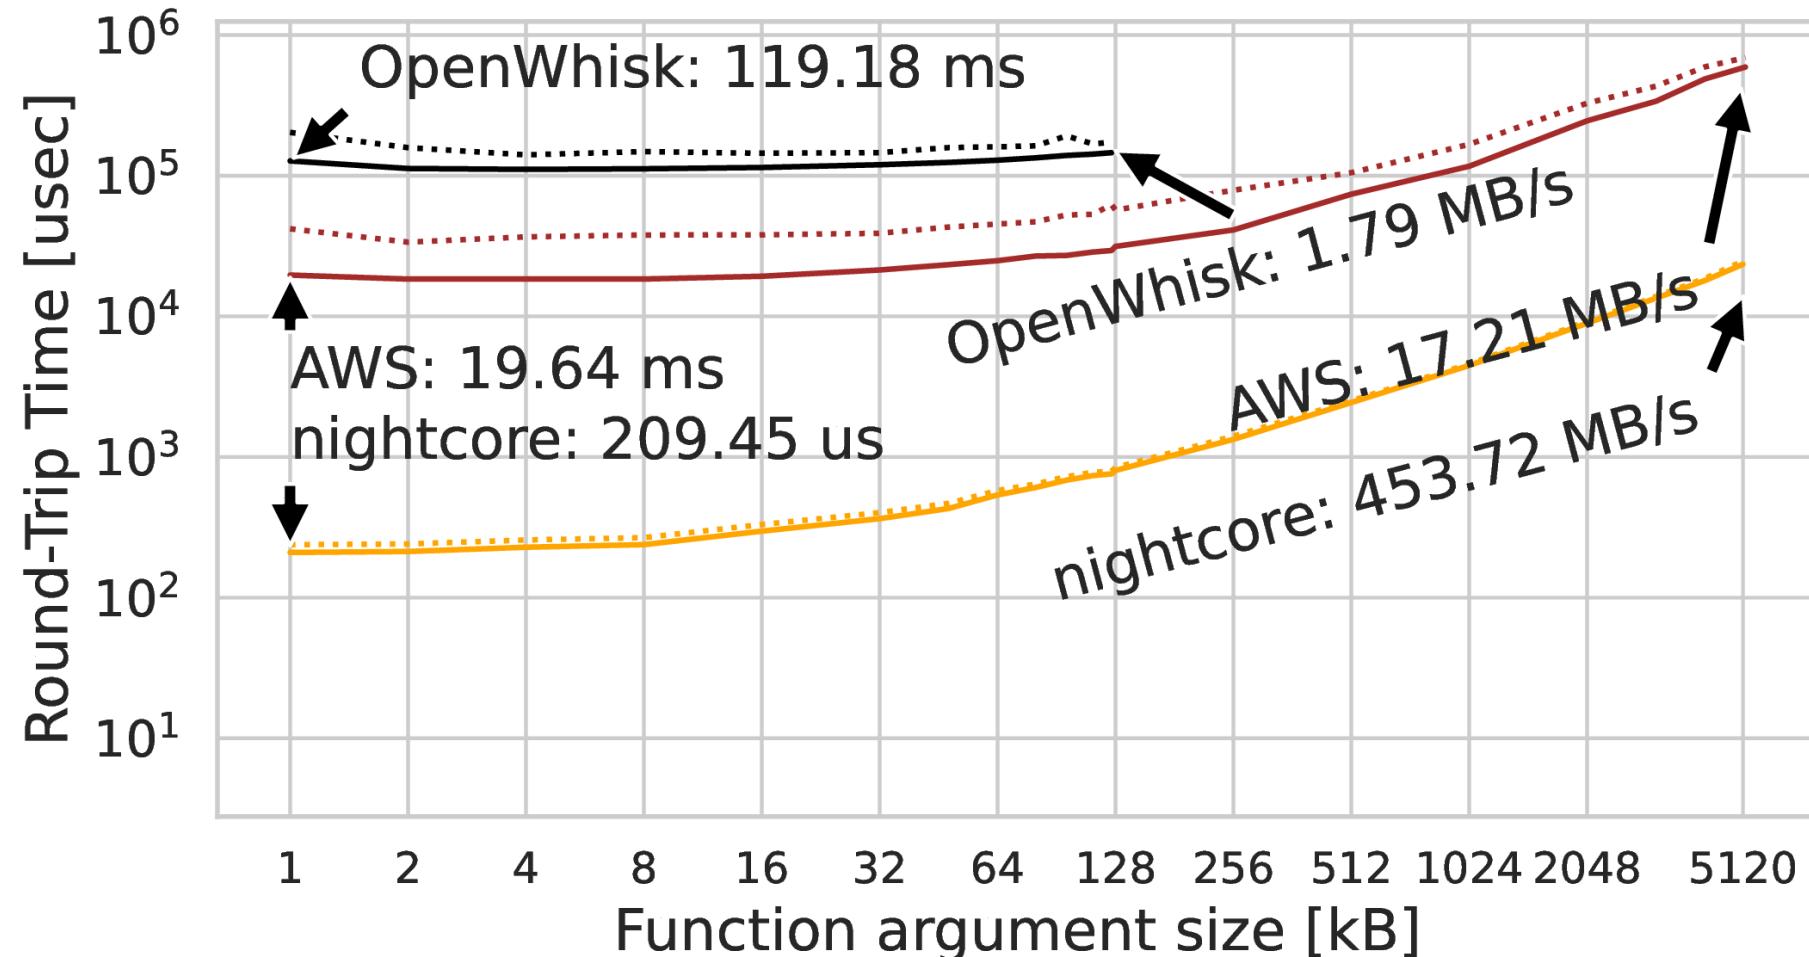
Function-as-a-Service for HPC



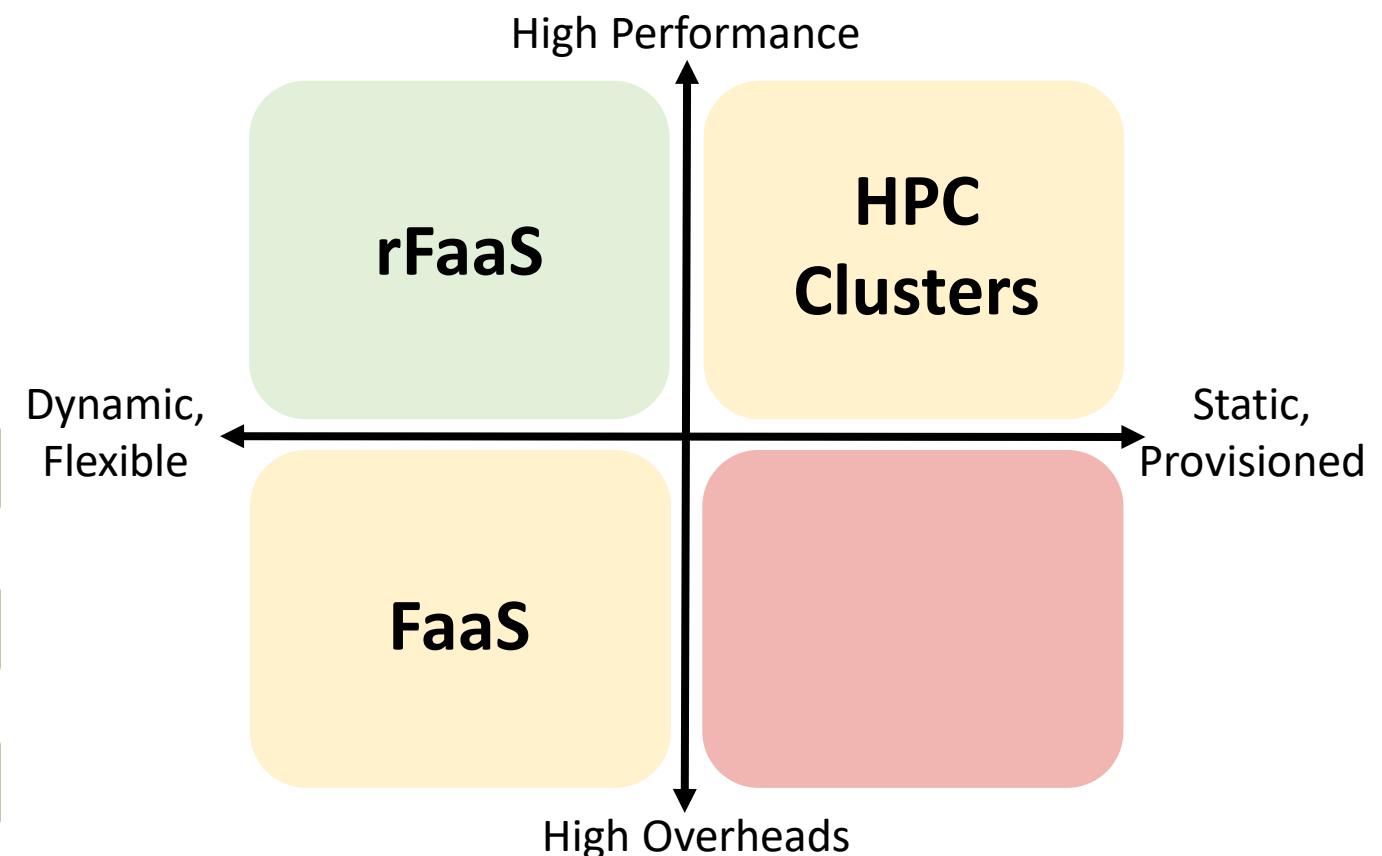
Function-as-a-Service for HPC



How fast FaaS invocations are?



Function-as-a-Service for HPC



Dynamic parallelism

Interactive computations

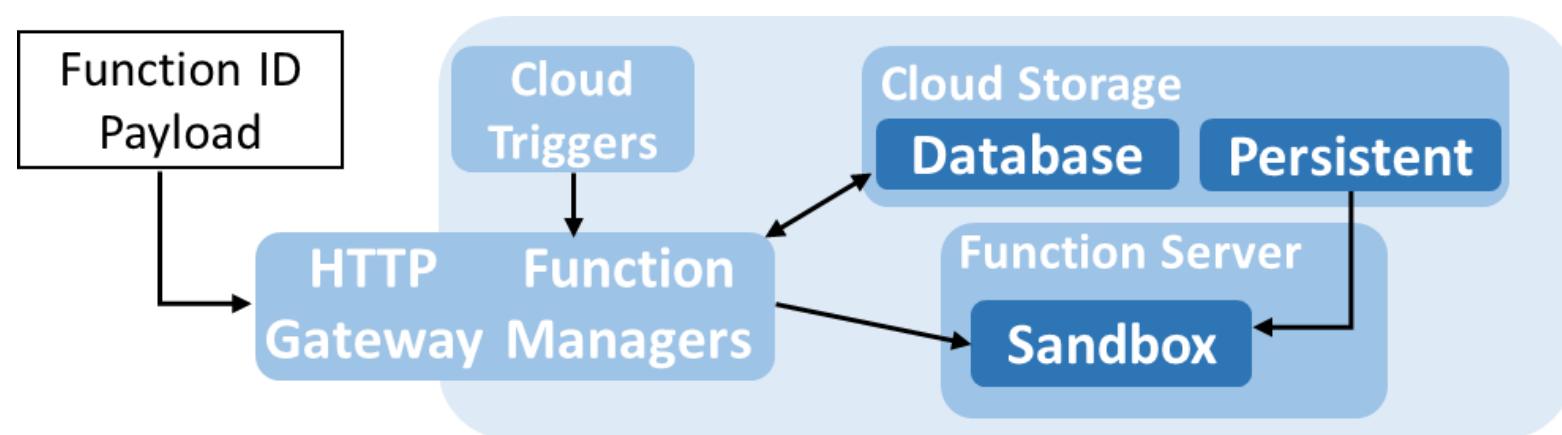
Latency-sensitive jobs

Malleable jobs

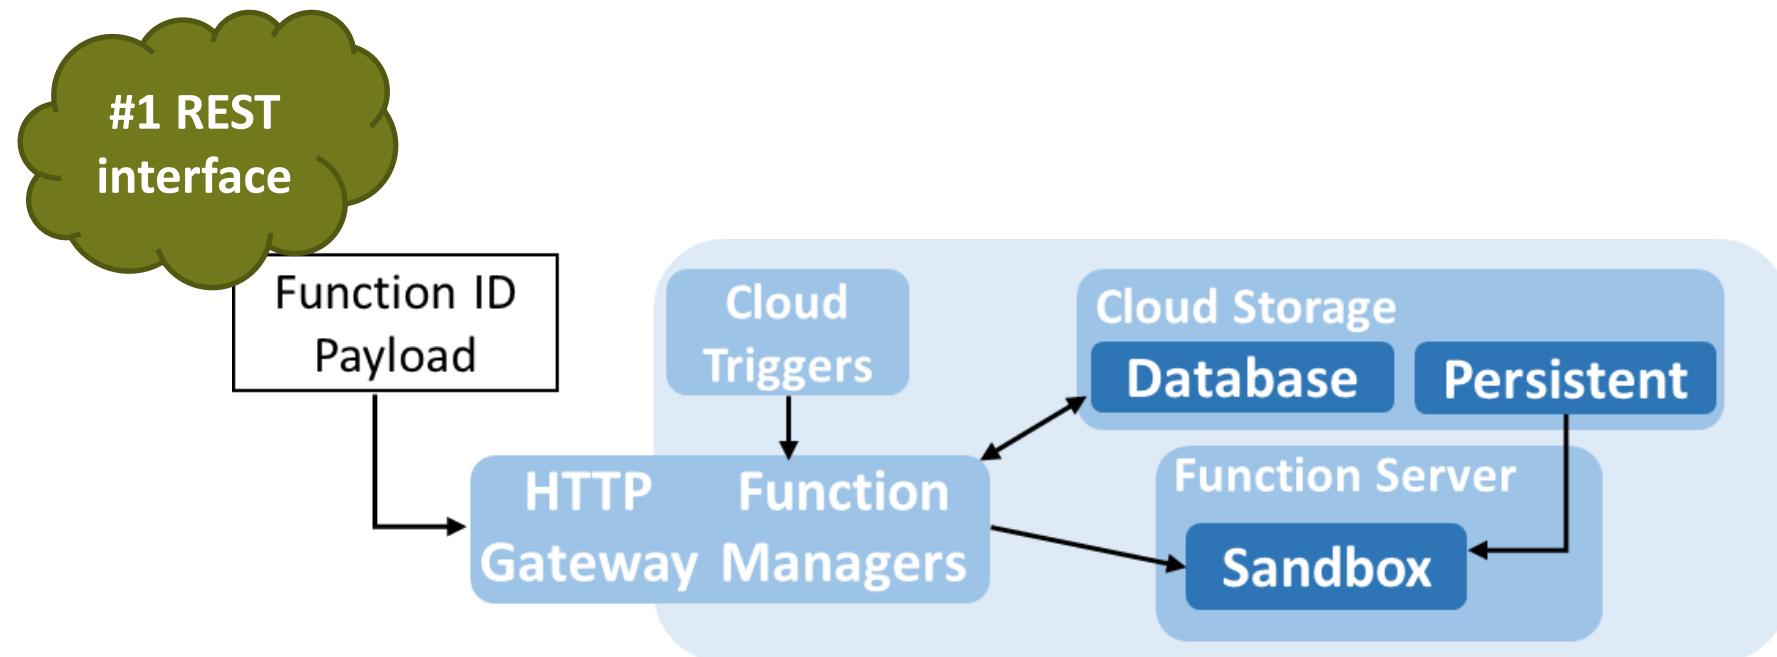
Long-running jobs

Static parallelism

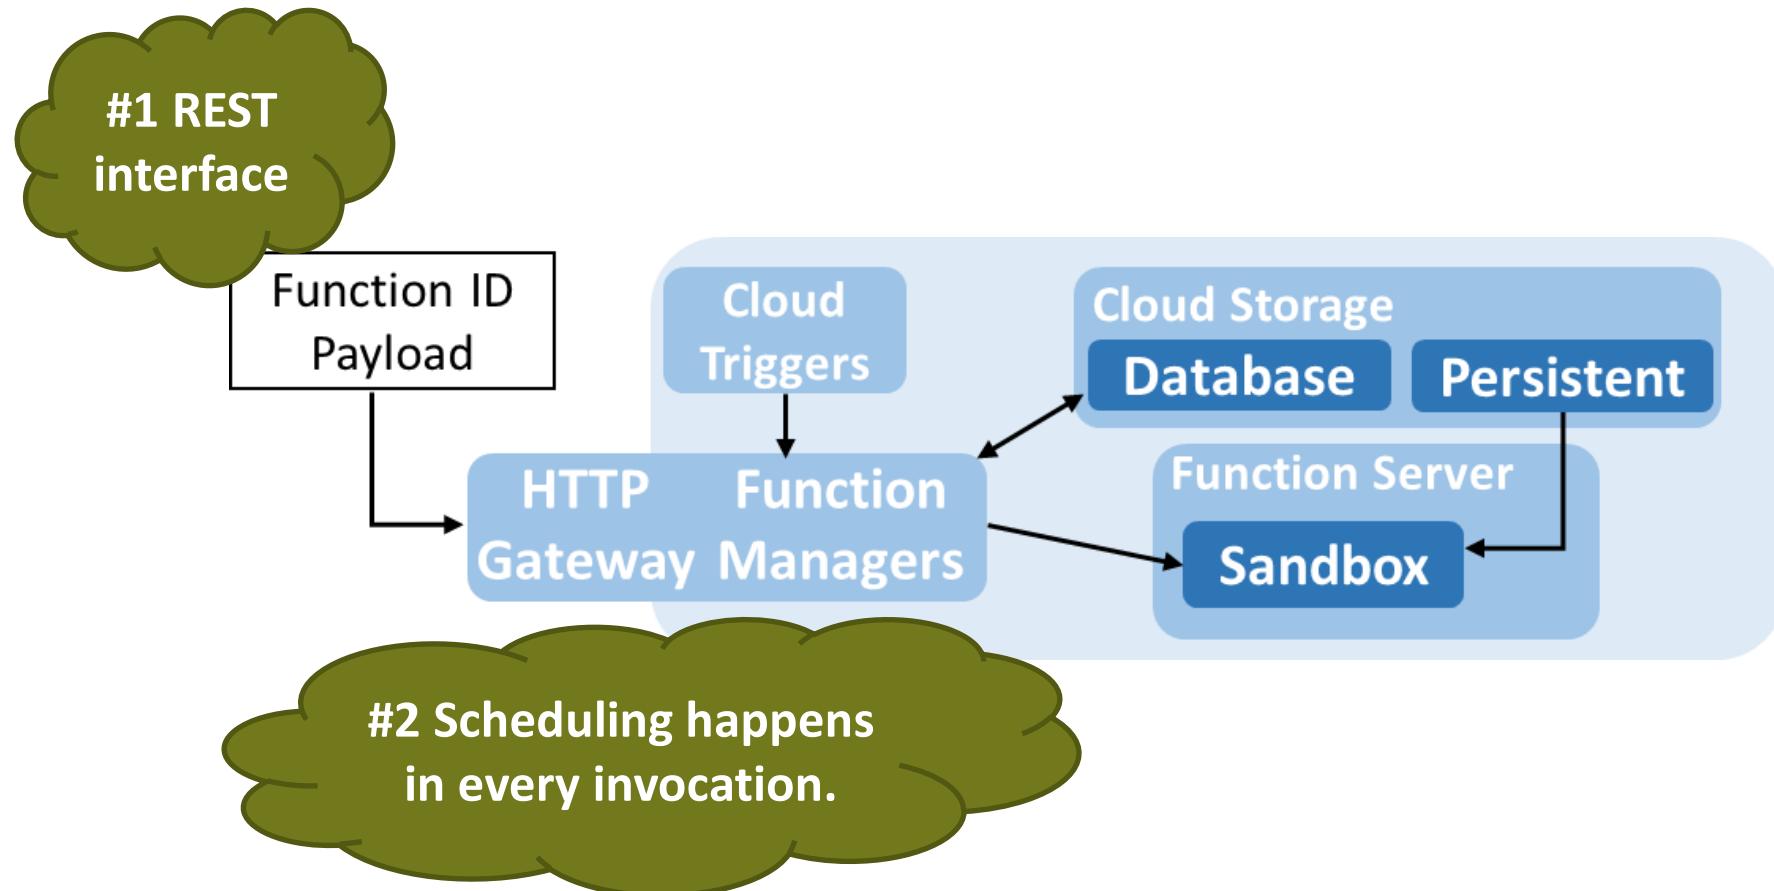
Function-as-a-Service – can it work in HPC?



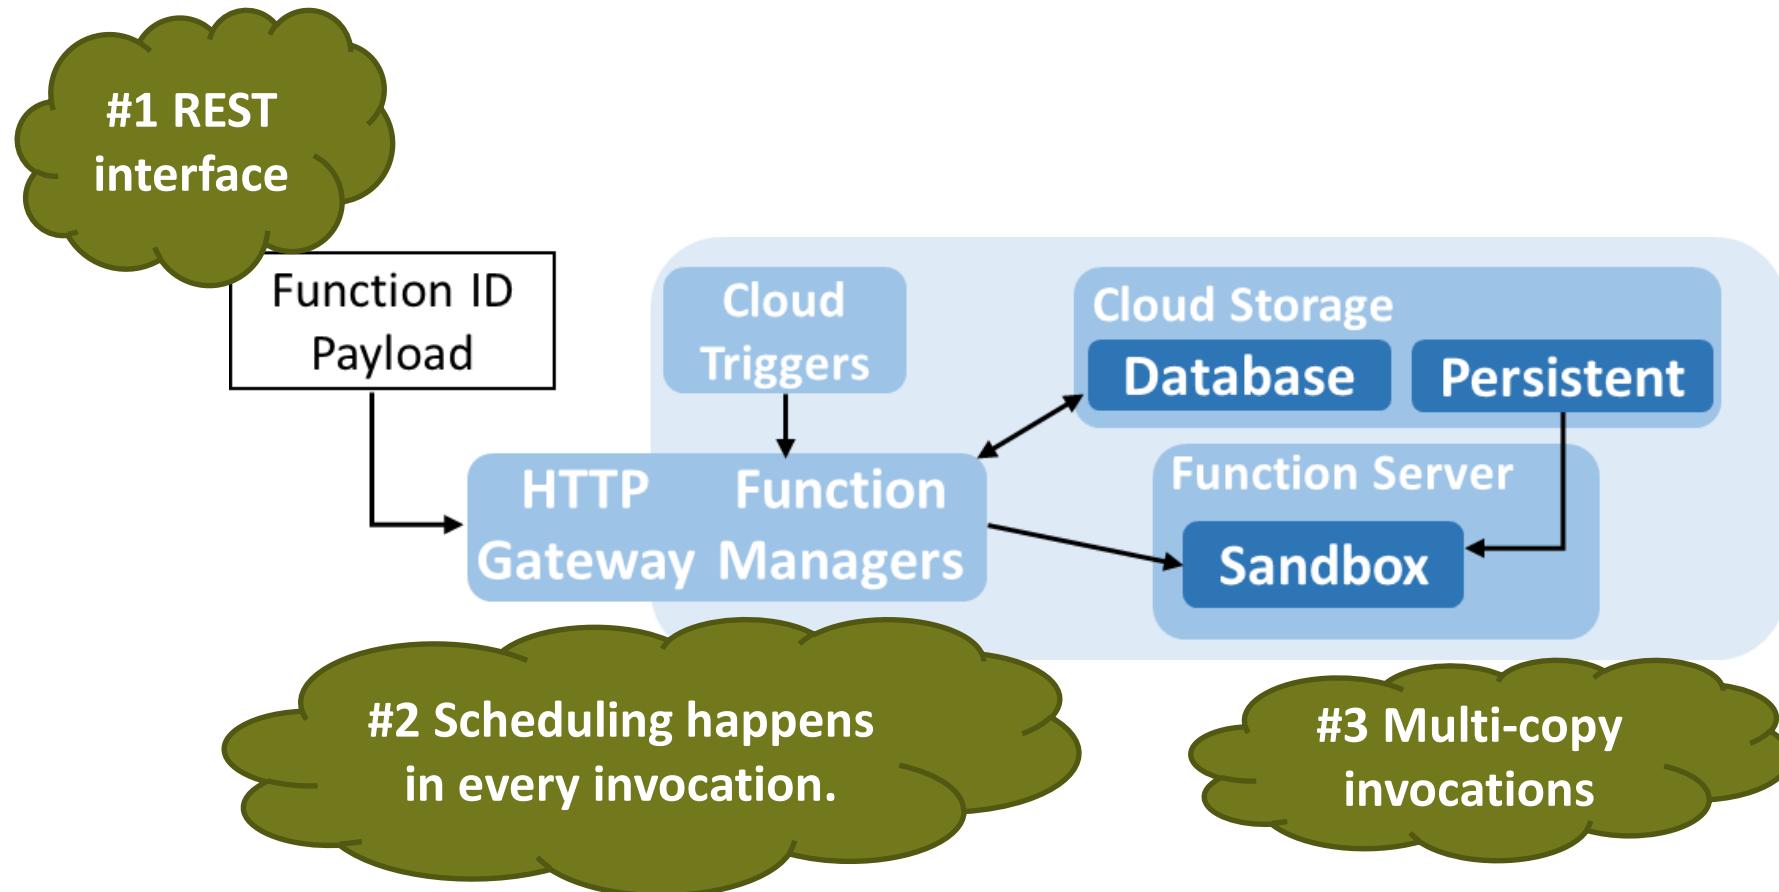
Function-as-a-Service – can it work in HPC?



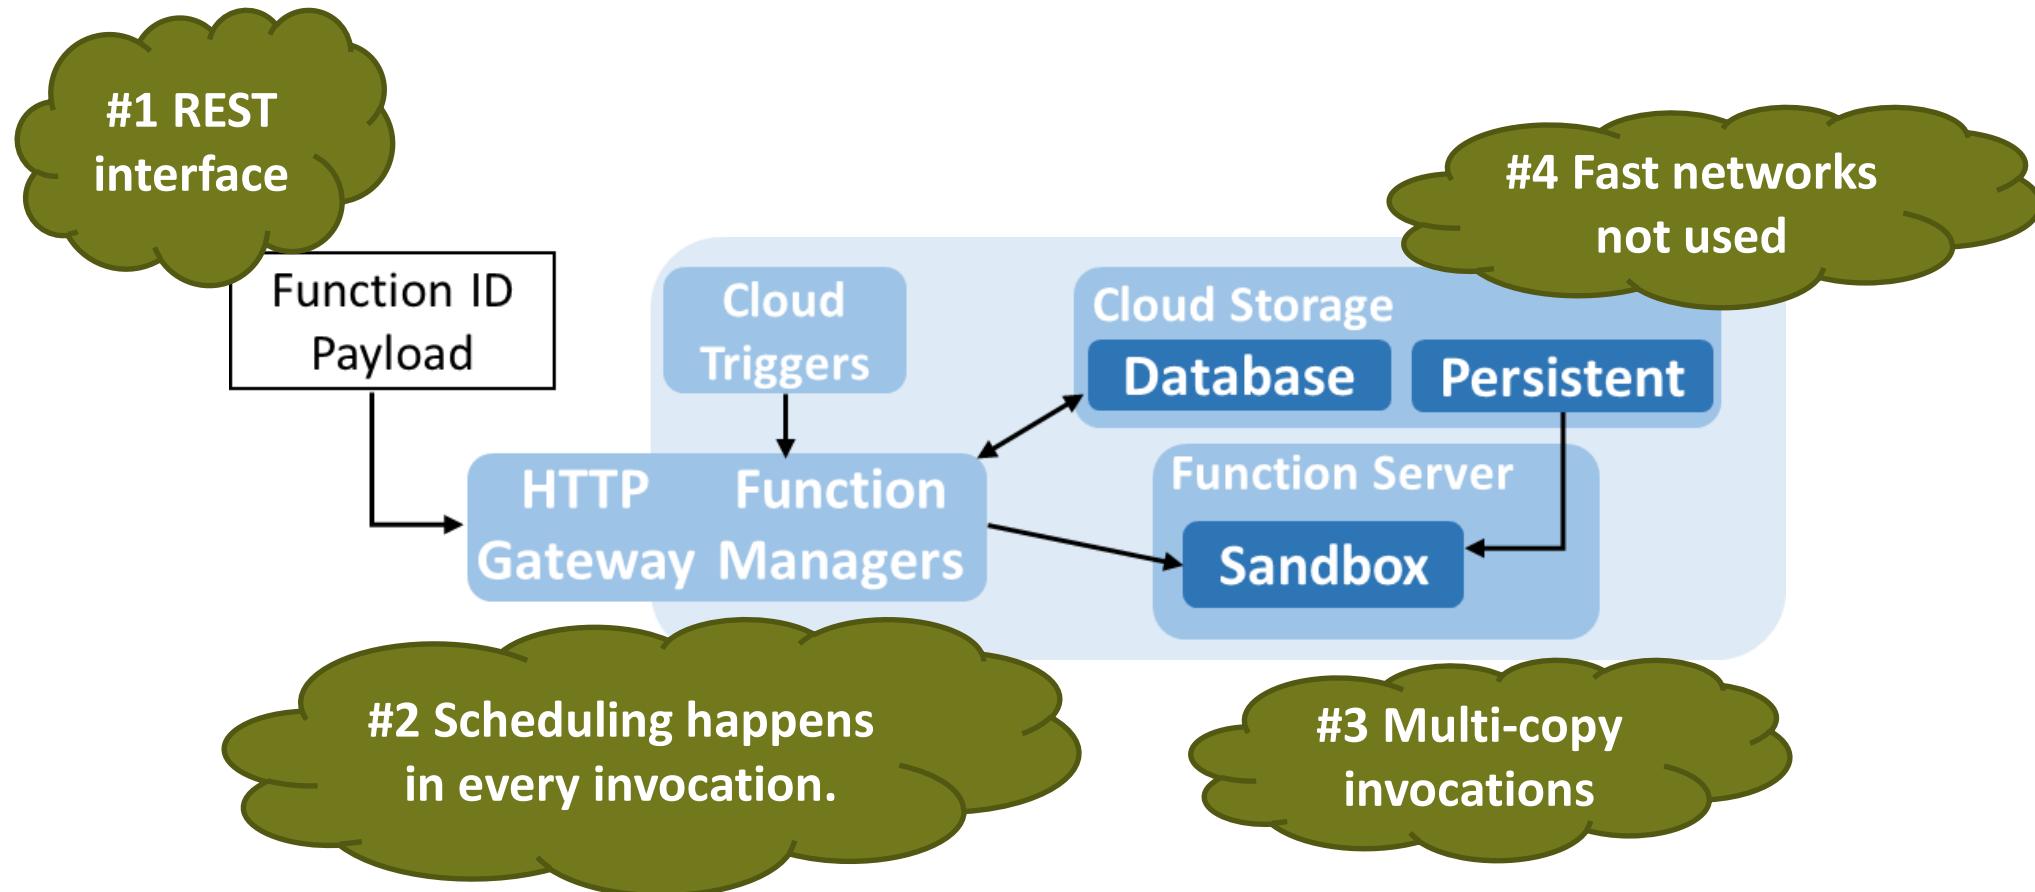
Function-as-a-Service – can it work in HPC?



Function-as-a-Service – can it work in HPC?



Function-as-a-Service – can it work in HPC?



Programming model

```
void compute(int size, options & opts) {
    rfaas::invoker invoker{opts.rnic_device};
    invoker.allocate(opts.lib, opts.size * sizeof(double),
        rfaas::invoker::ALWAYS_WARM_INVOCATIONS);

    auto alloc = invokerallocator<double>{};
    rfaas::buffer<double> in = alloc.input(2 * size);
    rfaas::buffer<double> out = alloc.output(2 * size);

    auto f = invoker.submit("task", in, size, out);
    local_task(in.data() + size, out.data() + size, size);
    f.get();

    invoker.deallocate();
}
```

Programming model

```
void compute(int size, options & opts) {
    rfaas::invoker invoker{opts.rnic_device};
    invoker.allocate(opts.lib, opts.size * sizeof(double),
        rfaas::invoker::ALWAYS_WARM_INVOCATIONS);

    auto alloc = invokerallocator<double>{};
    rfaas::buffer<double> in = alloc.input(2 * size);
    rfaas::buffer<double> out = alloc.output(2 * size);

    auto f = invoker.submit("task", in, size, out);
    local_task(in.data() + size, out.data() + size, size);
    f.get();

    invoker.deallocate();
}
```

□ Serverless leases

Programming model

```
void compute(int size, options & opts) {
    rfaas::invoker invoker{opts.rnic_device};
    invoker.allocate(opts.lib, opts.size * sizeof(double),
        rfaas::invoker::ALWAYS_WARM_INVOCATIONS);

    auto alloc = invokerallocator<double>{};
    rfaas::buffer<double> in = alloc.input(2 * size);
    rfaas::buffer<double> out = alloc.output(2 * size);

    auto f = invoker.submit("task", in, size, out);
    local_task(in.data() + size, out.data() + size, size);
    f.get();

    invoker.deallocate();
}
```

□ Serverless leases

□ RDMA abstractions

Programming model

```
void compute(int size, options & opts) {
    rfaas::invoker invoker{opts.rnic_device};
    invoker.allocate(opts.lib, opts.size * sizeof(double),
        rfaas::invoker::ALWAYS_WARM_INVOCATIONS);

    auto alloc = invokerallocator<double>{};
    rfaas::buffer<double> in = alloc.input(2 * size);
    rfaas::buffer<double> out = alloc.output(2 * size);

    auto f = invoker.submit("task", in, size, out);
    local_task(in.data() + size, out.data() + size, size);
    f.get();

    invoker.deallocate();
}
```

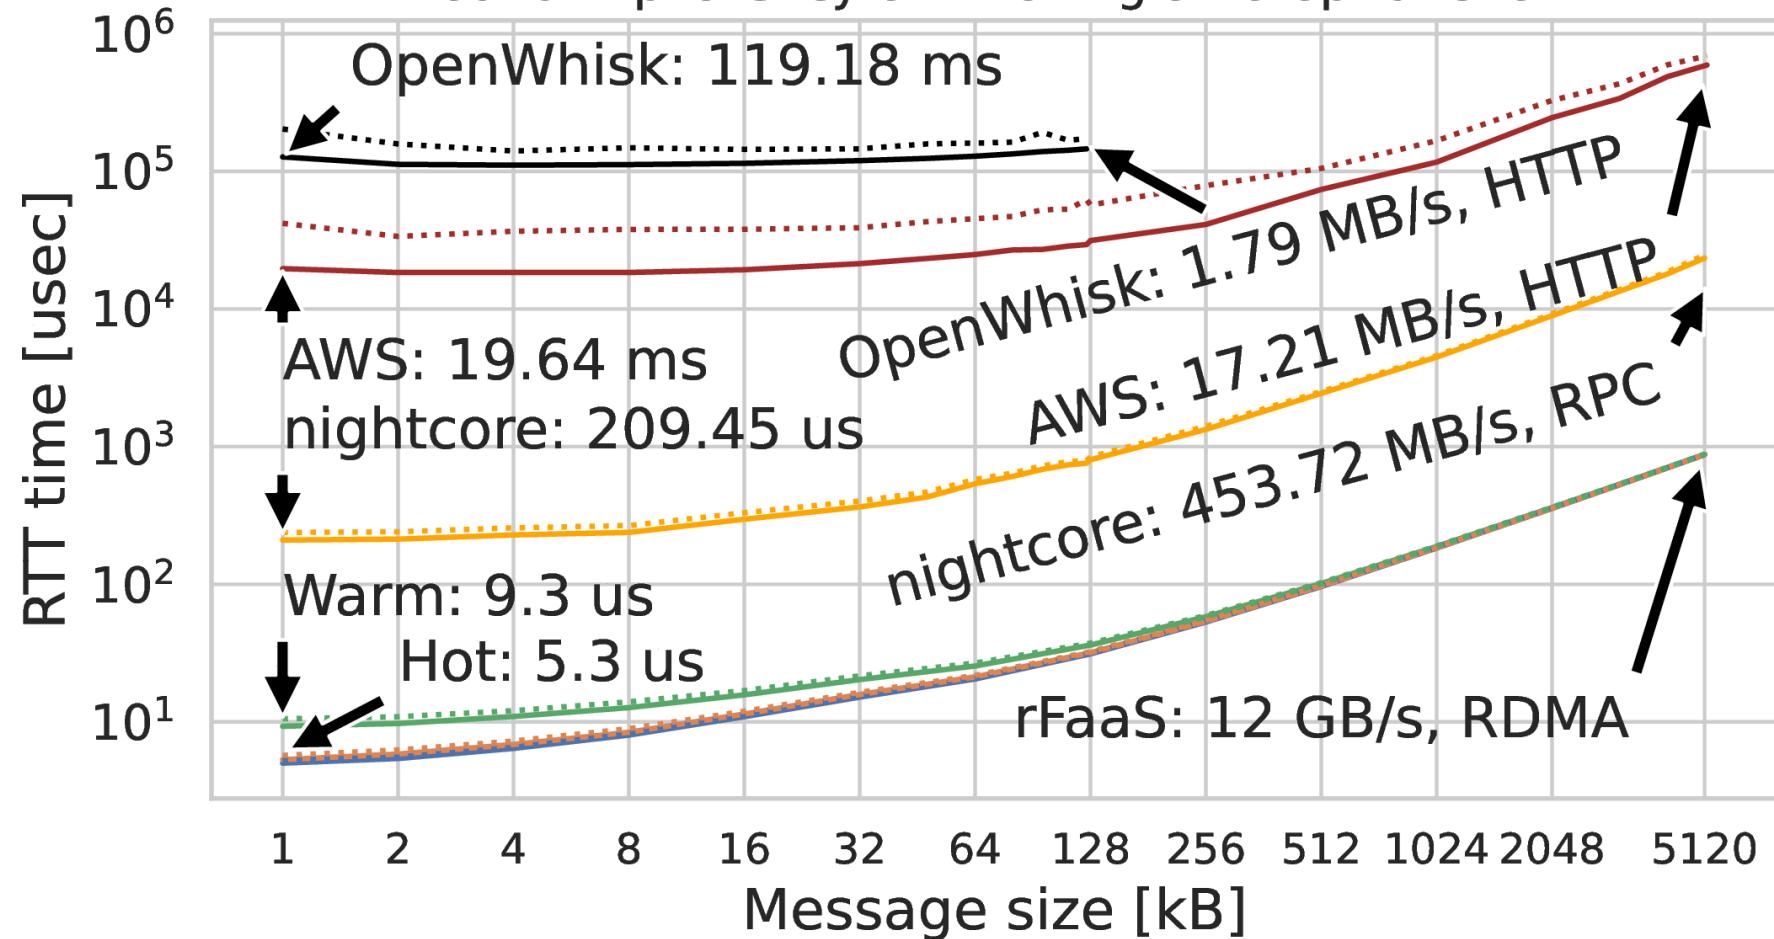
□ Serverless leases

□ RDMA abstractions

□ Zero-copy invocations

How fast rFaaS invocations are?

rFaaS versus OpenWhisk and nightcore (cluster) and AWS Lambda (cloud).
Round-trip latency of invoking a no-op function.

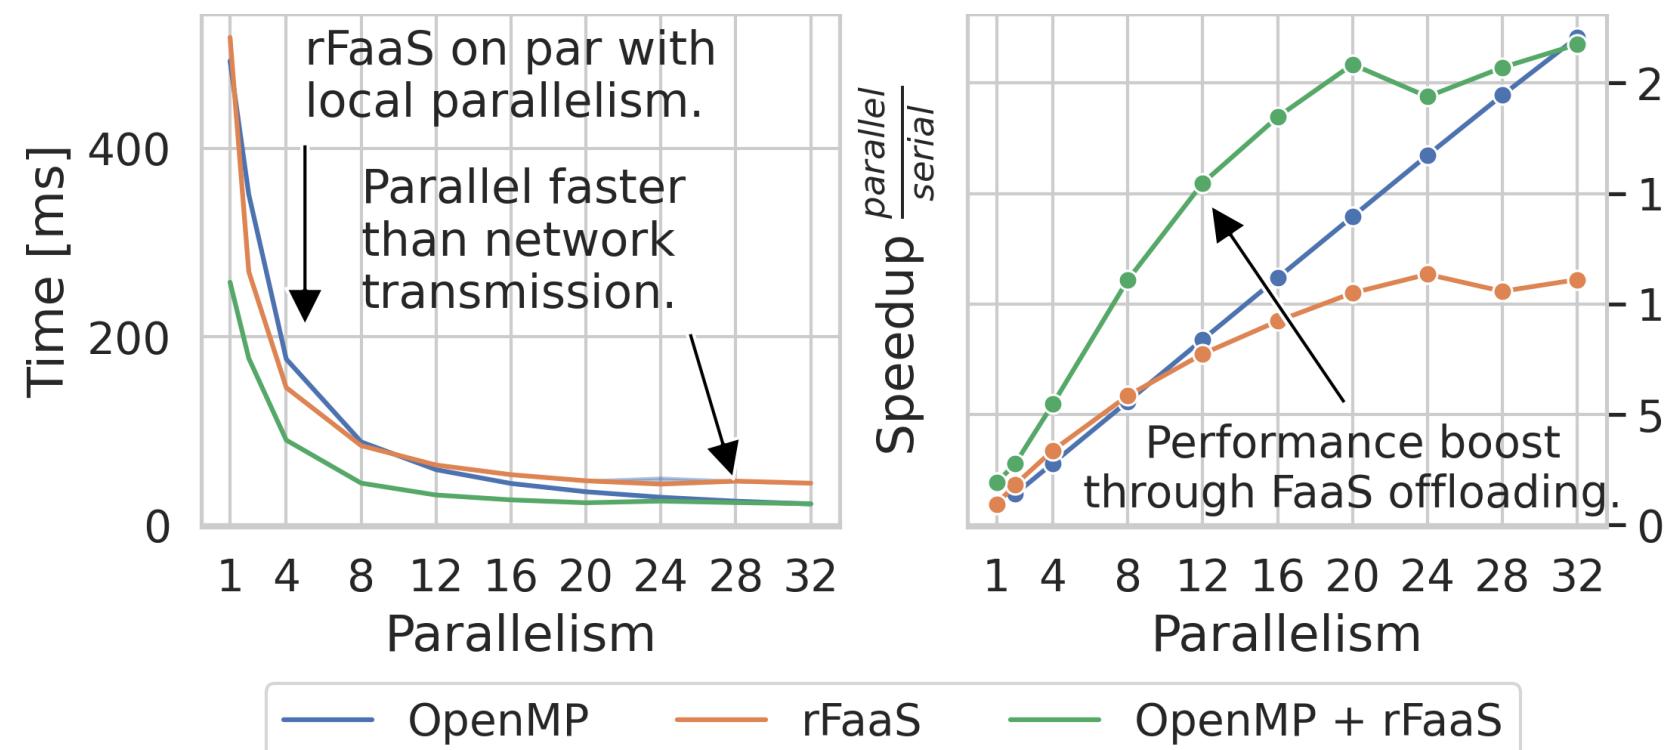


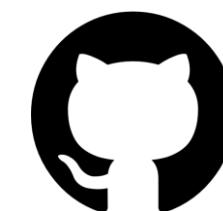
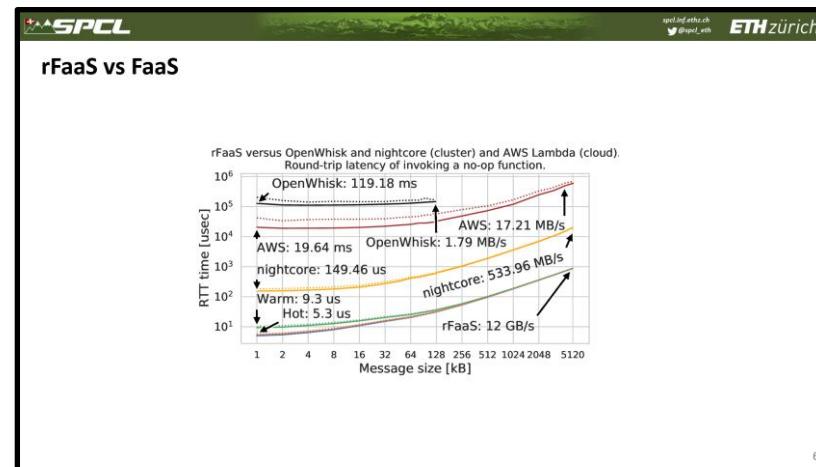
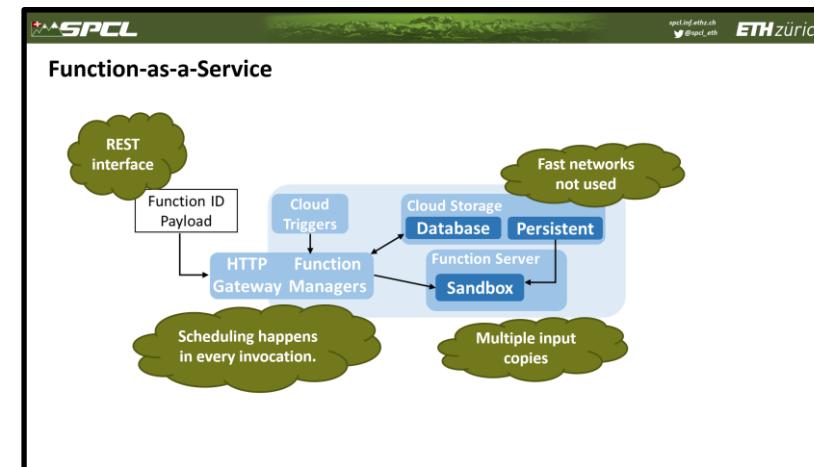
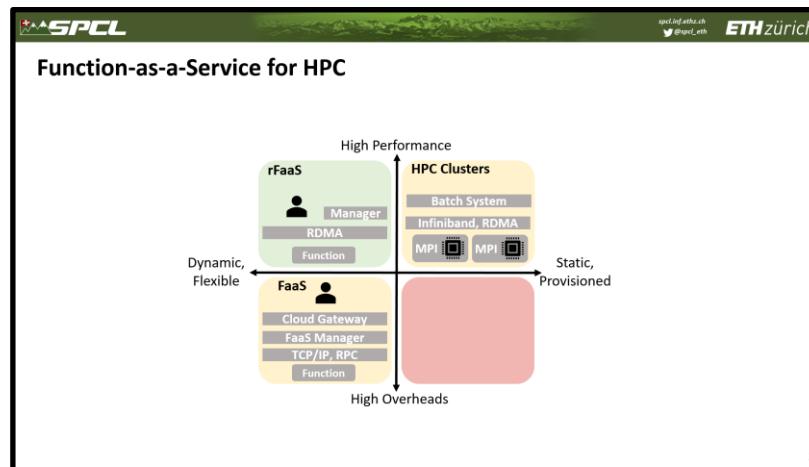
PARSEC: Black-Scholes

- Massively parallel computations
- Offload 50% of work to serverless functions.
- 10M equations, 229M input, 38M output.

PARSEC: Black-Scholes

- Massively parallel computations
- Offload 50% of work to serverless functions.
- 10M equations, 229M input, 38M output.





spcl/rFaaS



Paper preprint

<https://mcopik.github.io/projects/rfaas/>



[youtube.com/@spcl](https://www.youtube.com/@spcl)



twitter.com/spcl_eth



spcl.inf.ethz.ch



github.com/spcl