rFaaS: Enabling High Performance Serverless with RDMA and Leases

Marcin Copik, Konstantin Taranov, Alexandru Calotoiu, Torsten Hoefler

IPDPS
2023 • St. Petersburg, Florida USA

Function-as-a-Service for HPC

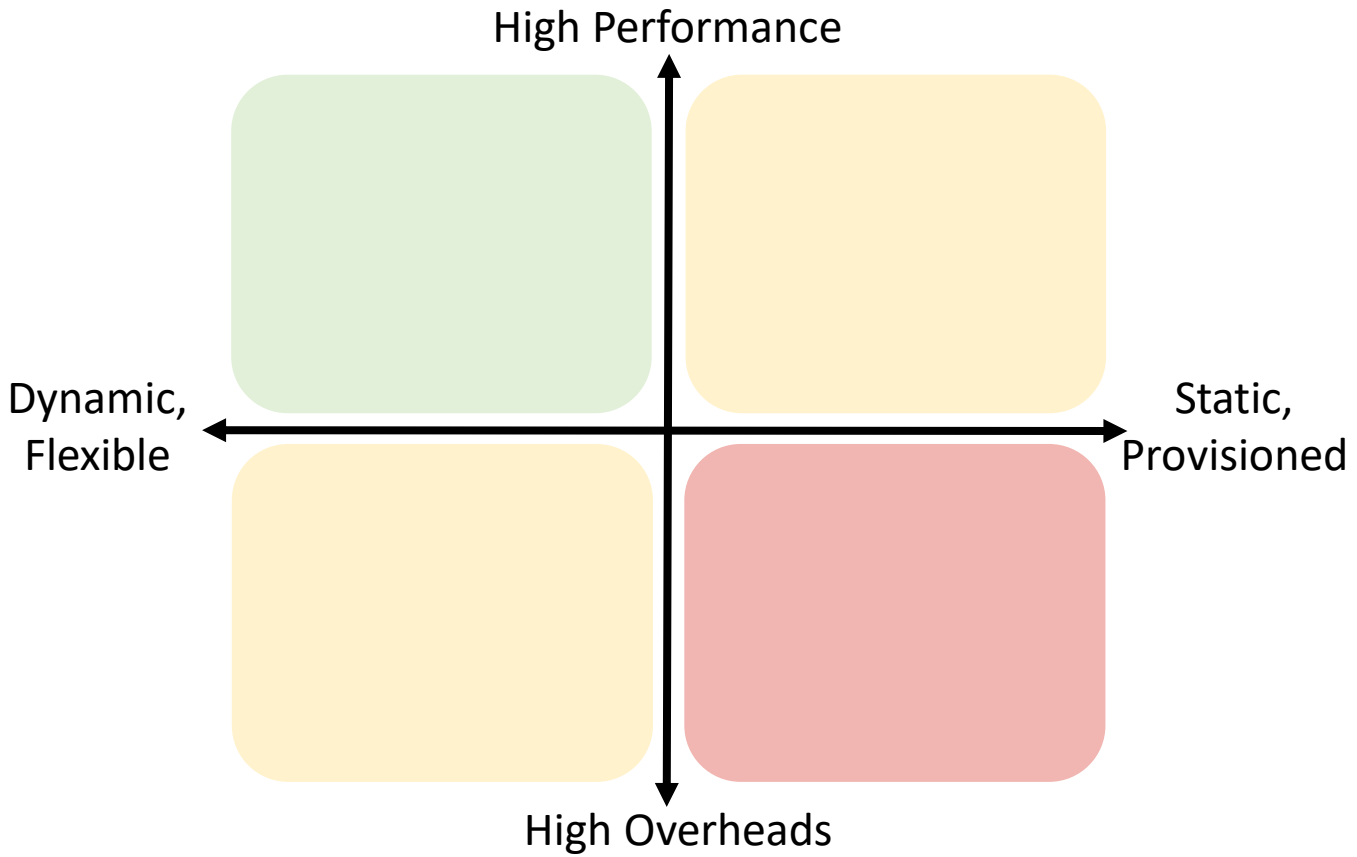

Function-as-a-Service for HPC

Long-running jobs

Static parallelism

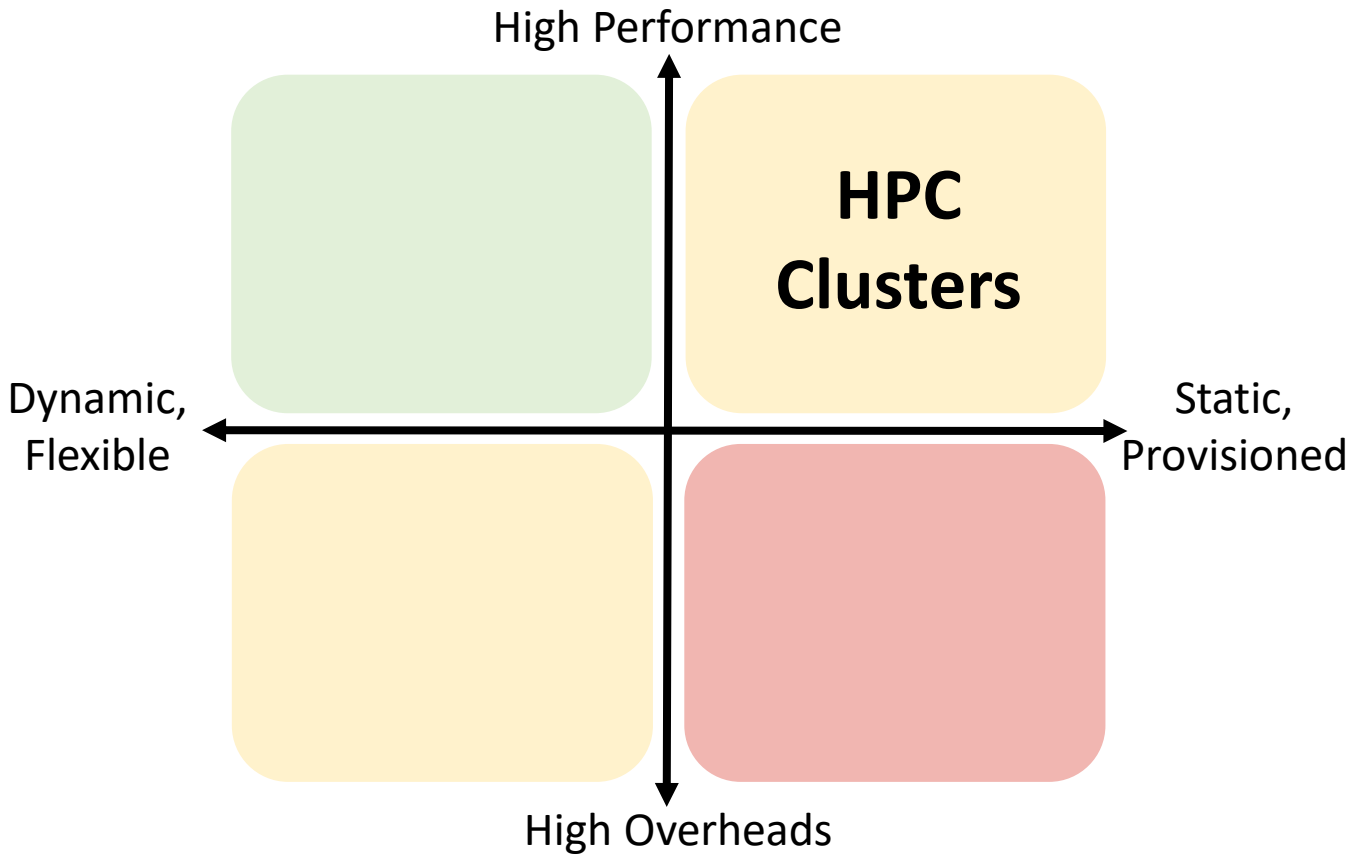

Function-as-a-Service for HPC

Long-running jobs

Static parallelism

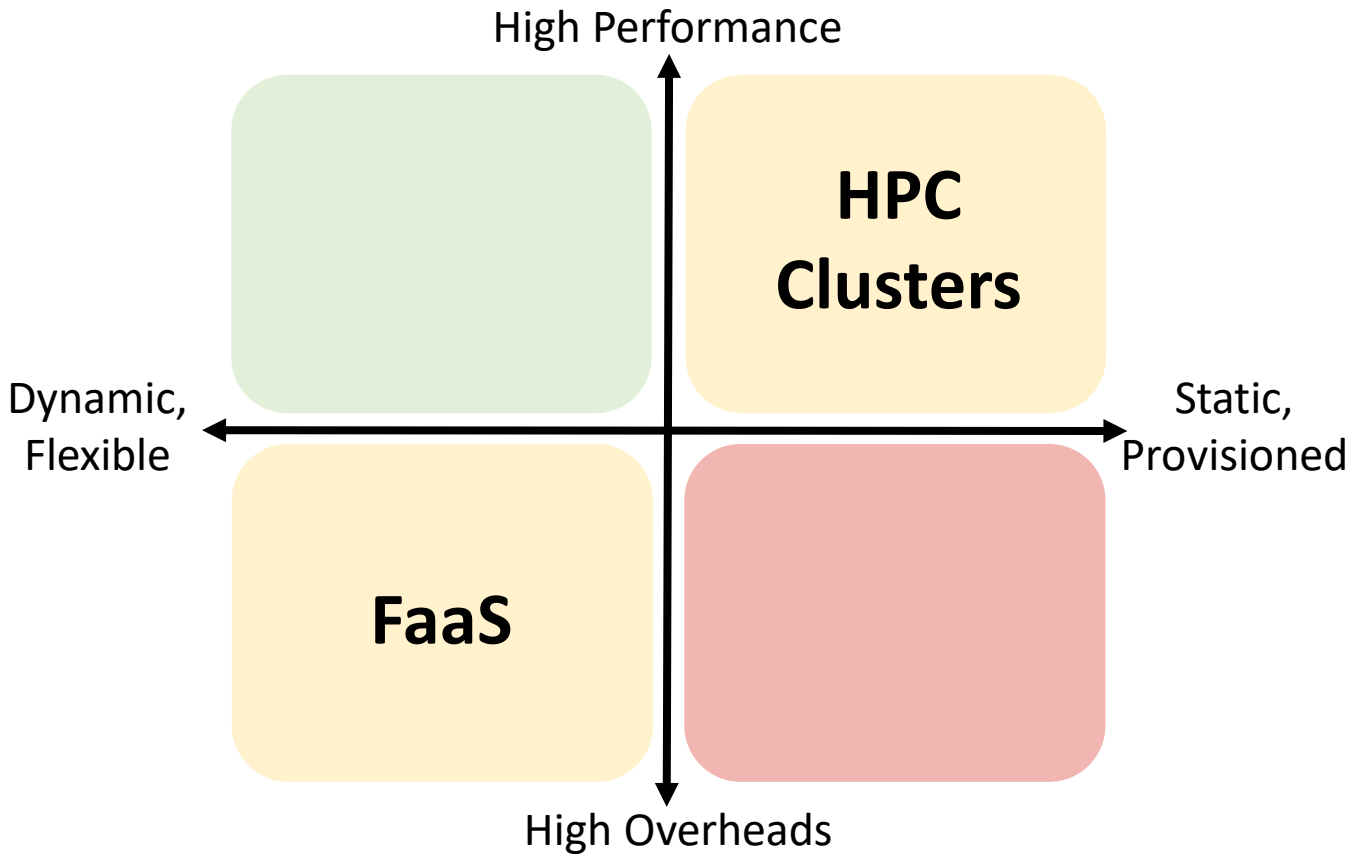

Malleable jobs

Latency-sensitive jobs

Interactive computations

Dynamic parallelism

How does FaaS work?

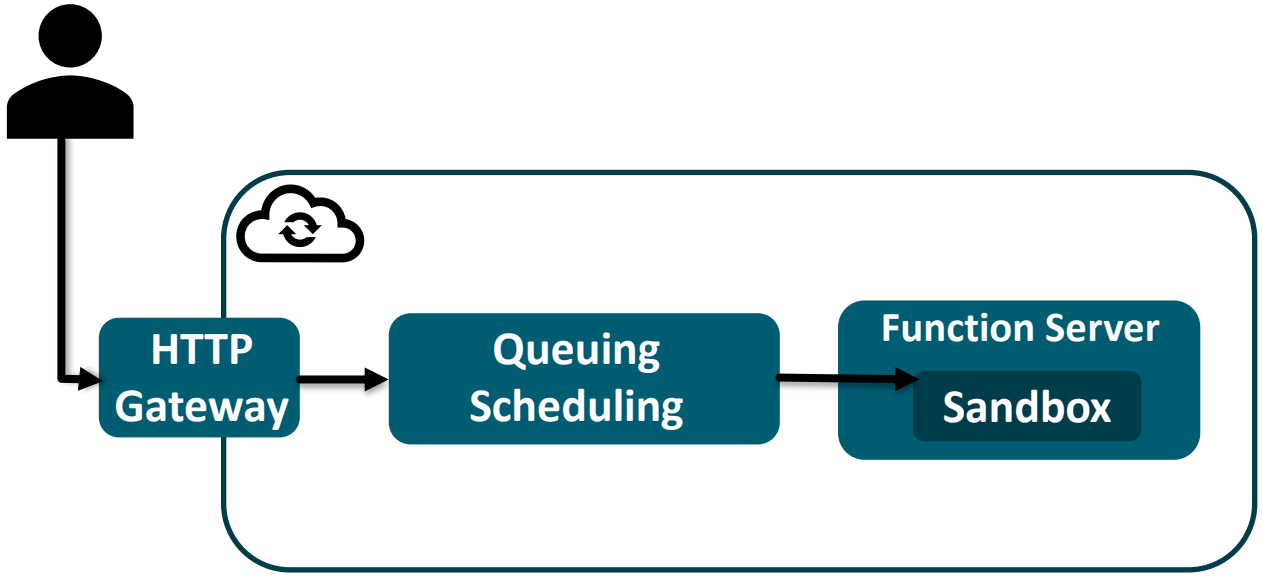

How does FaaS work?

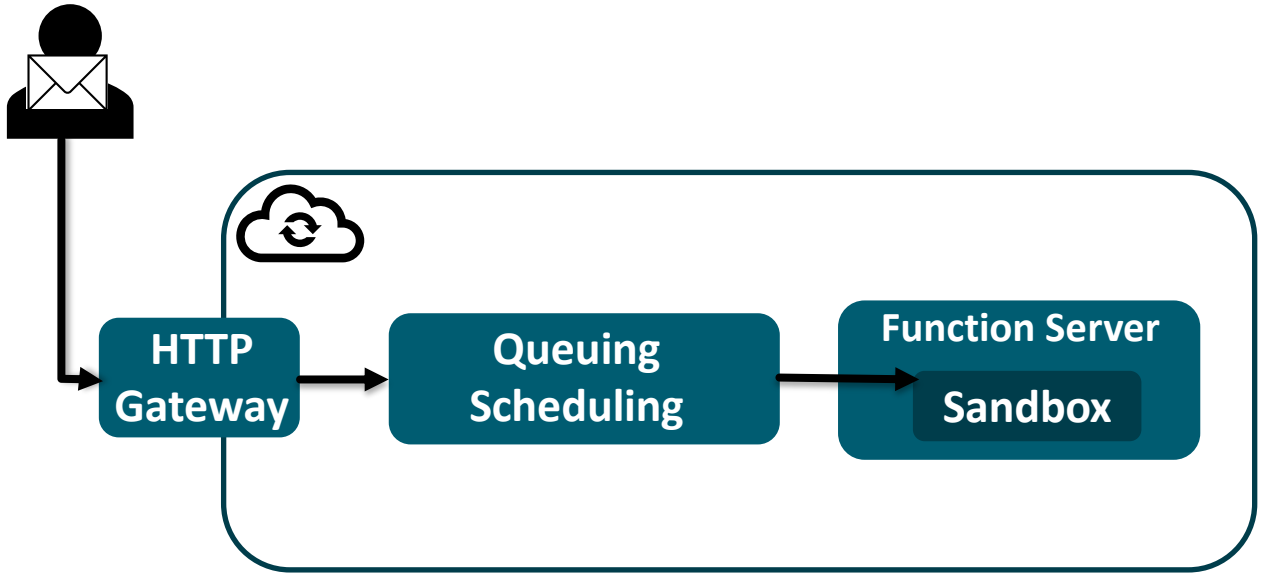

How does FaaS work?

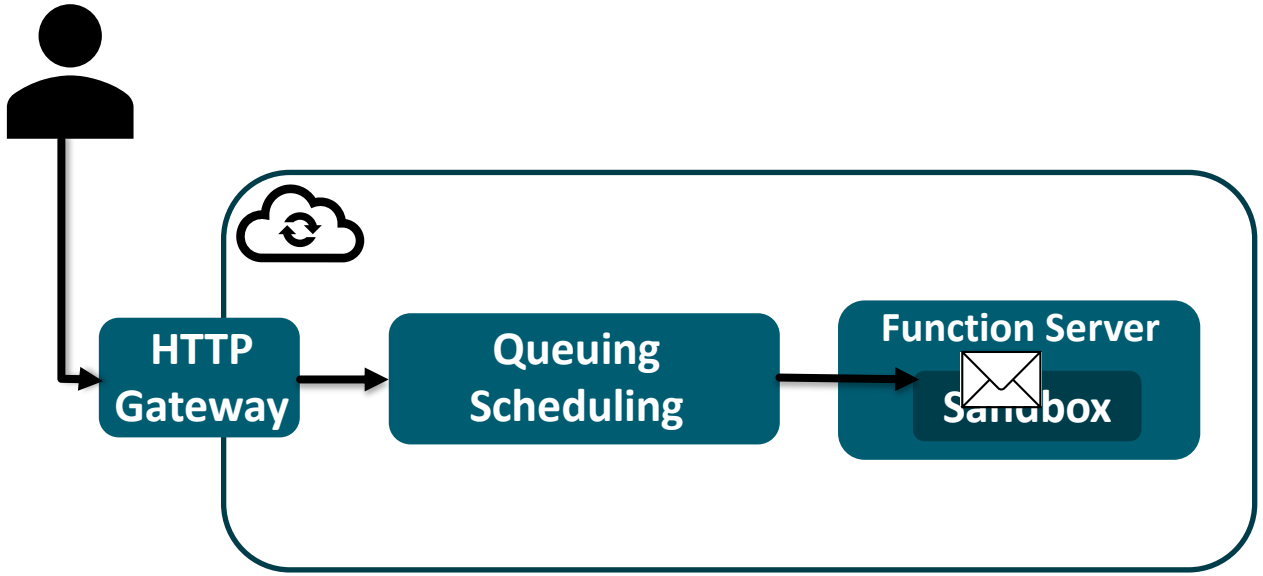

How fast are invocations in FaaS?

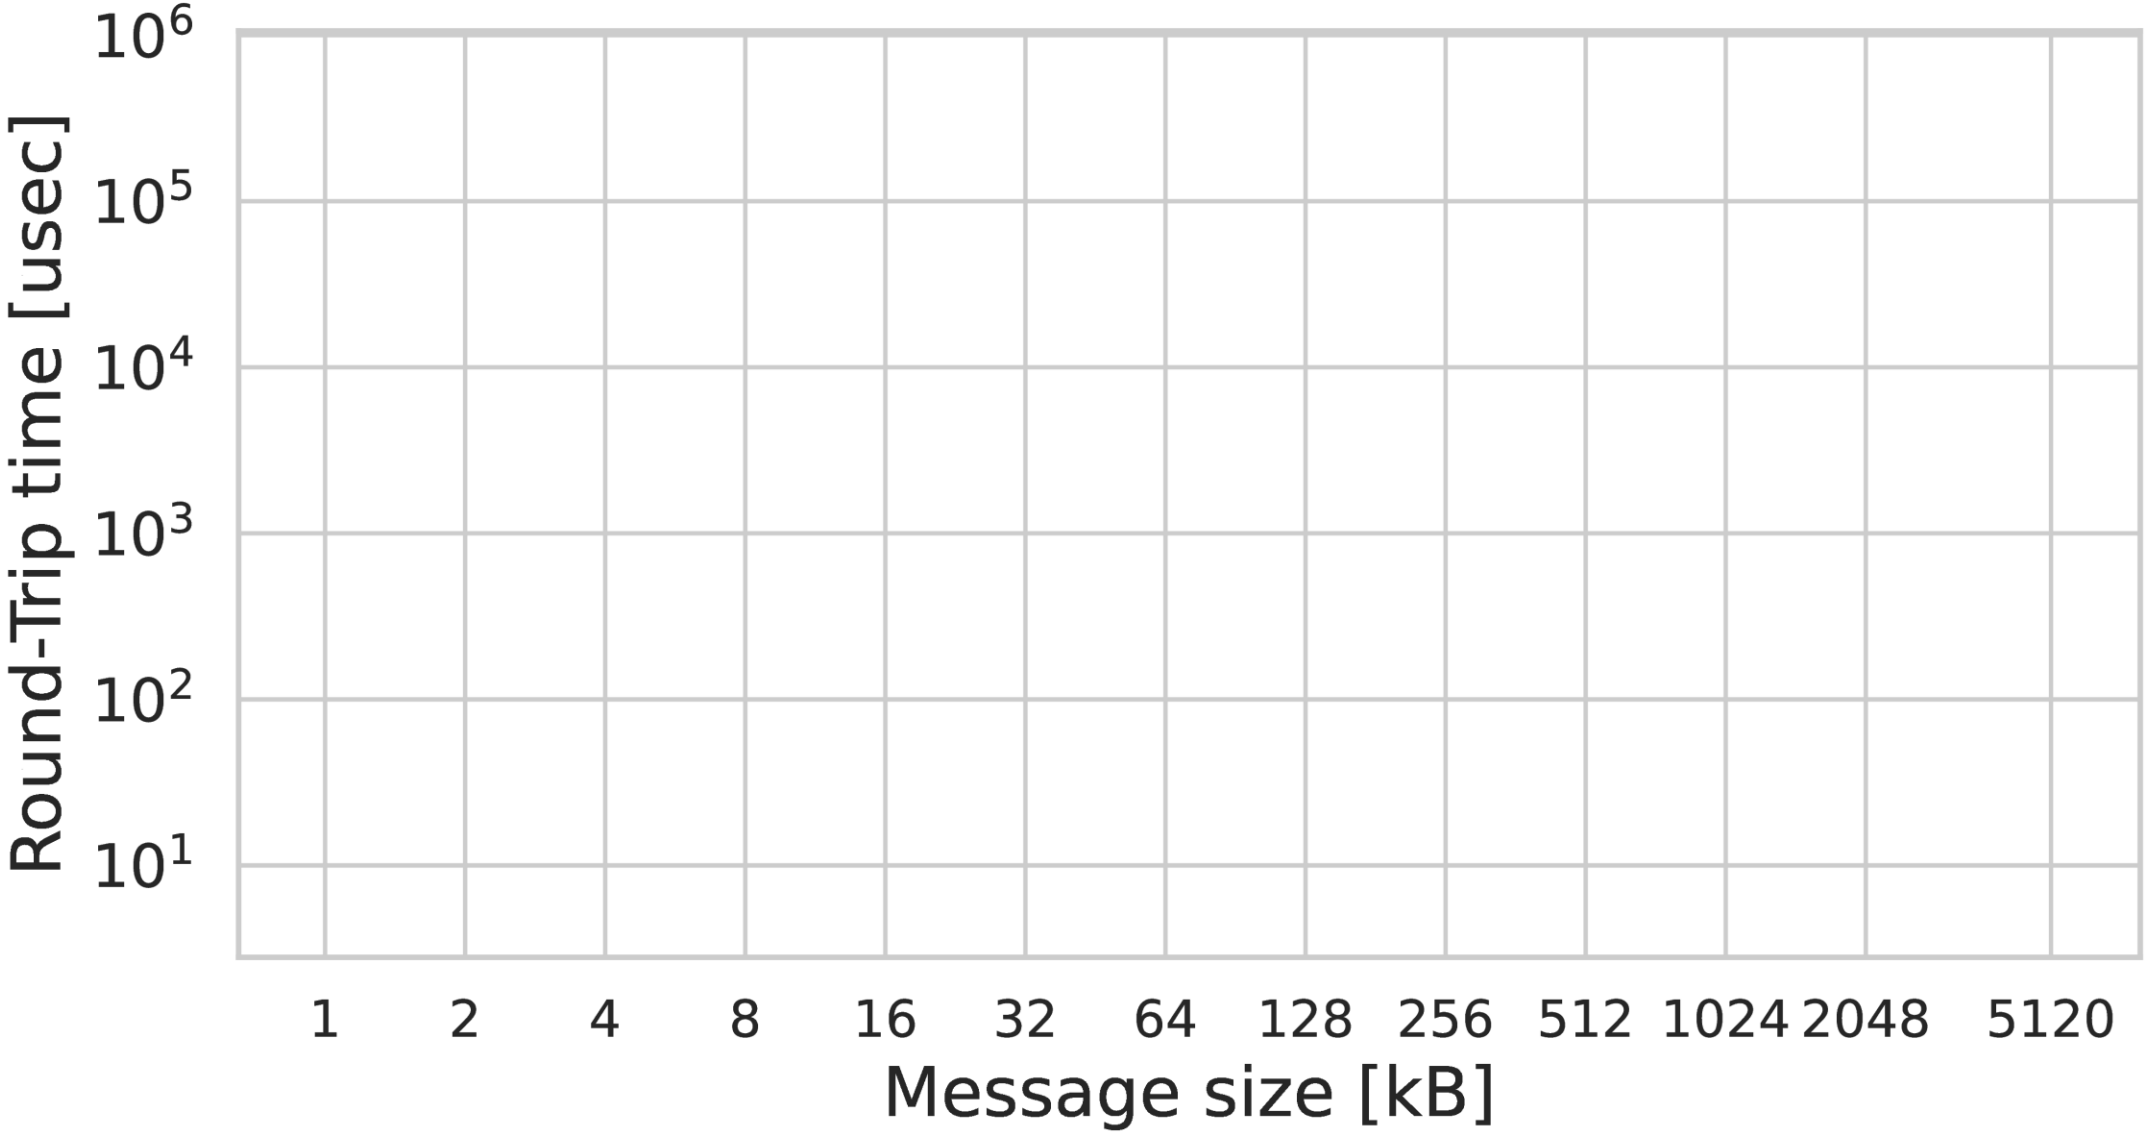

How fast are invocations in FaaS?

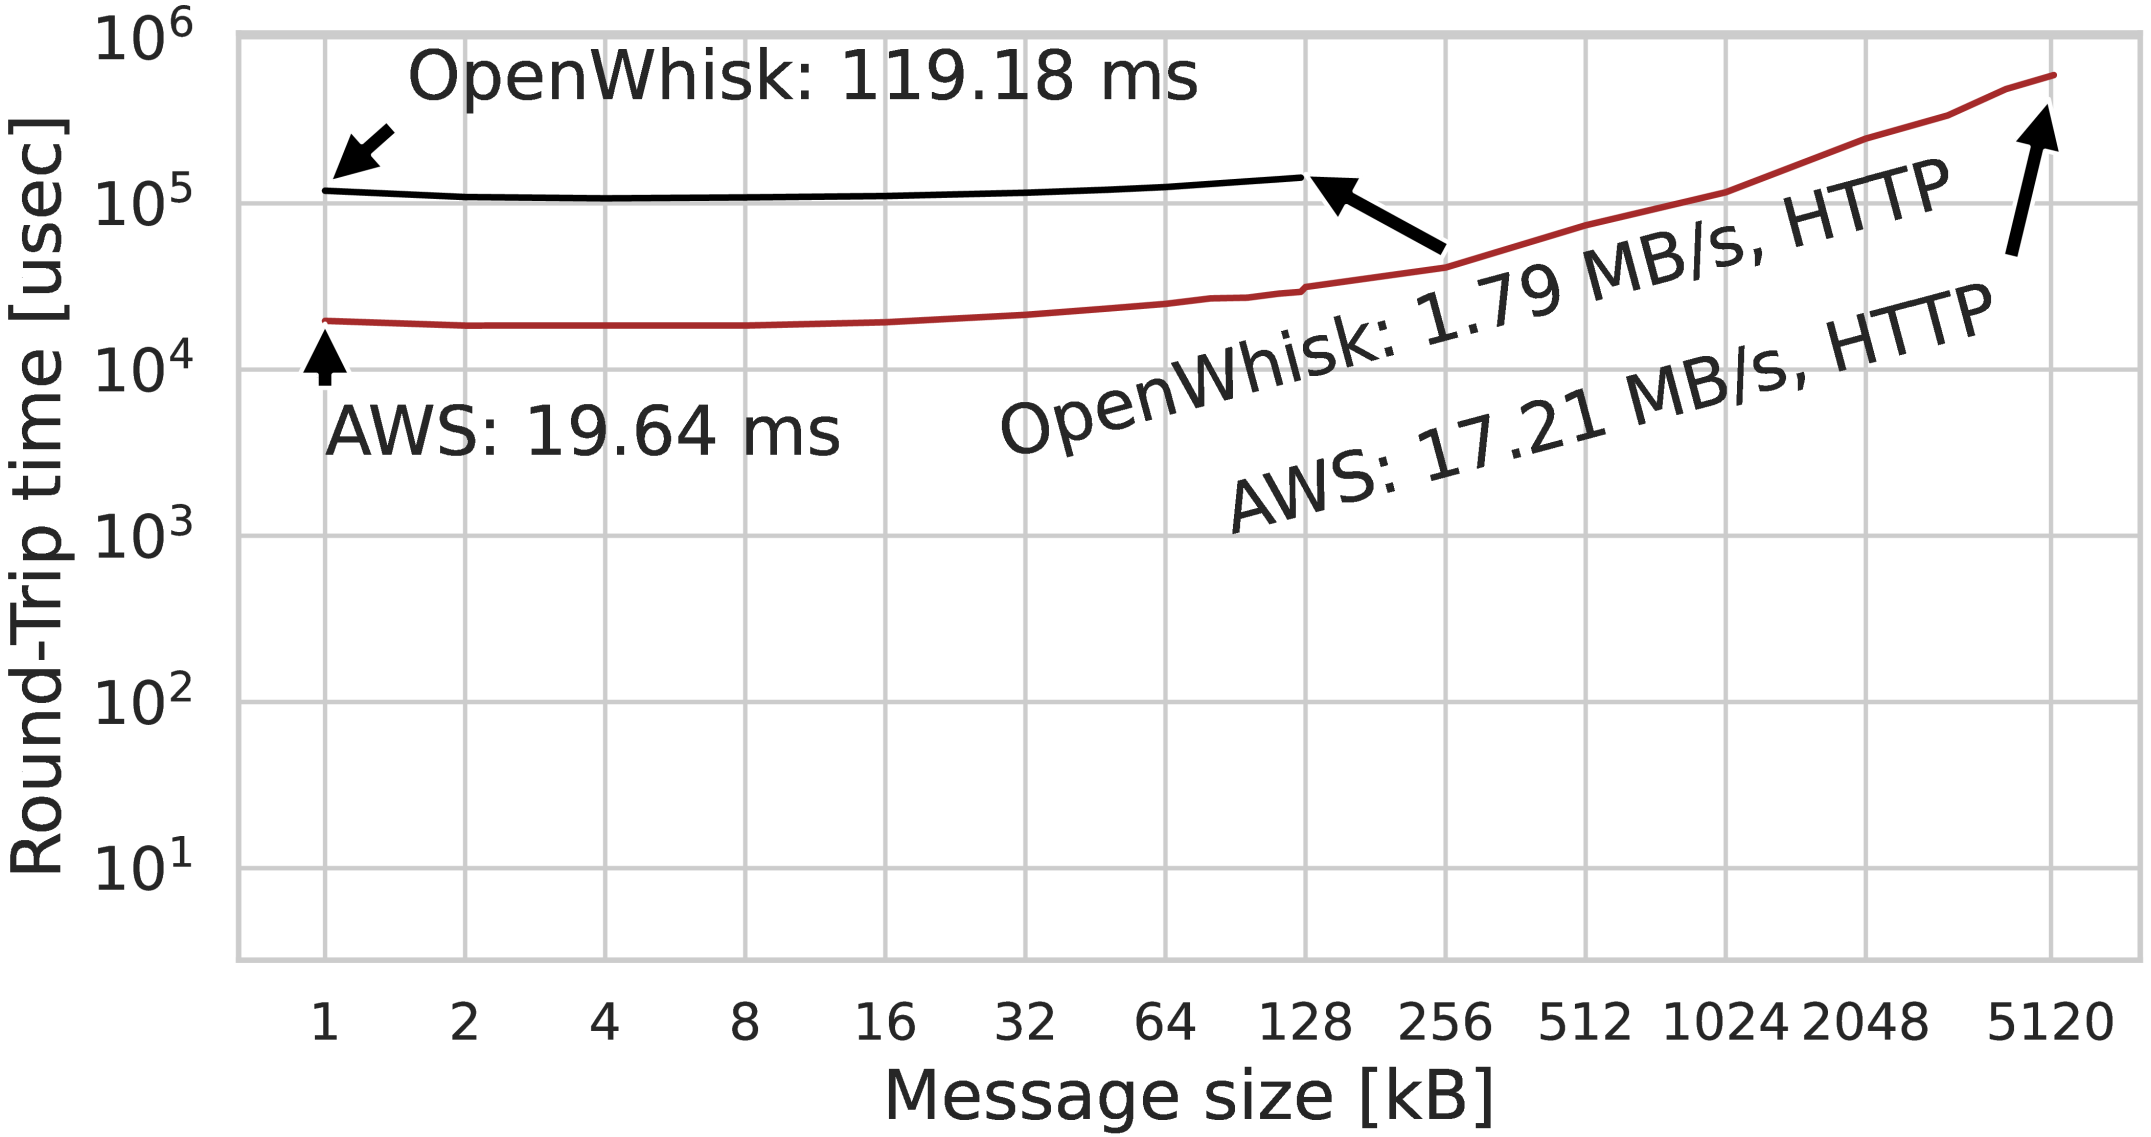

How fast are invocations in FaaS?

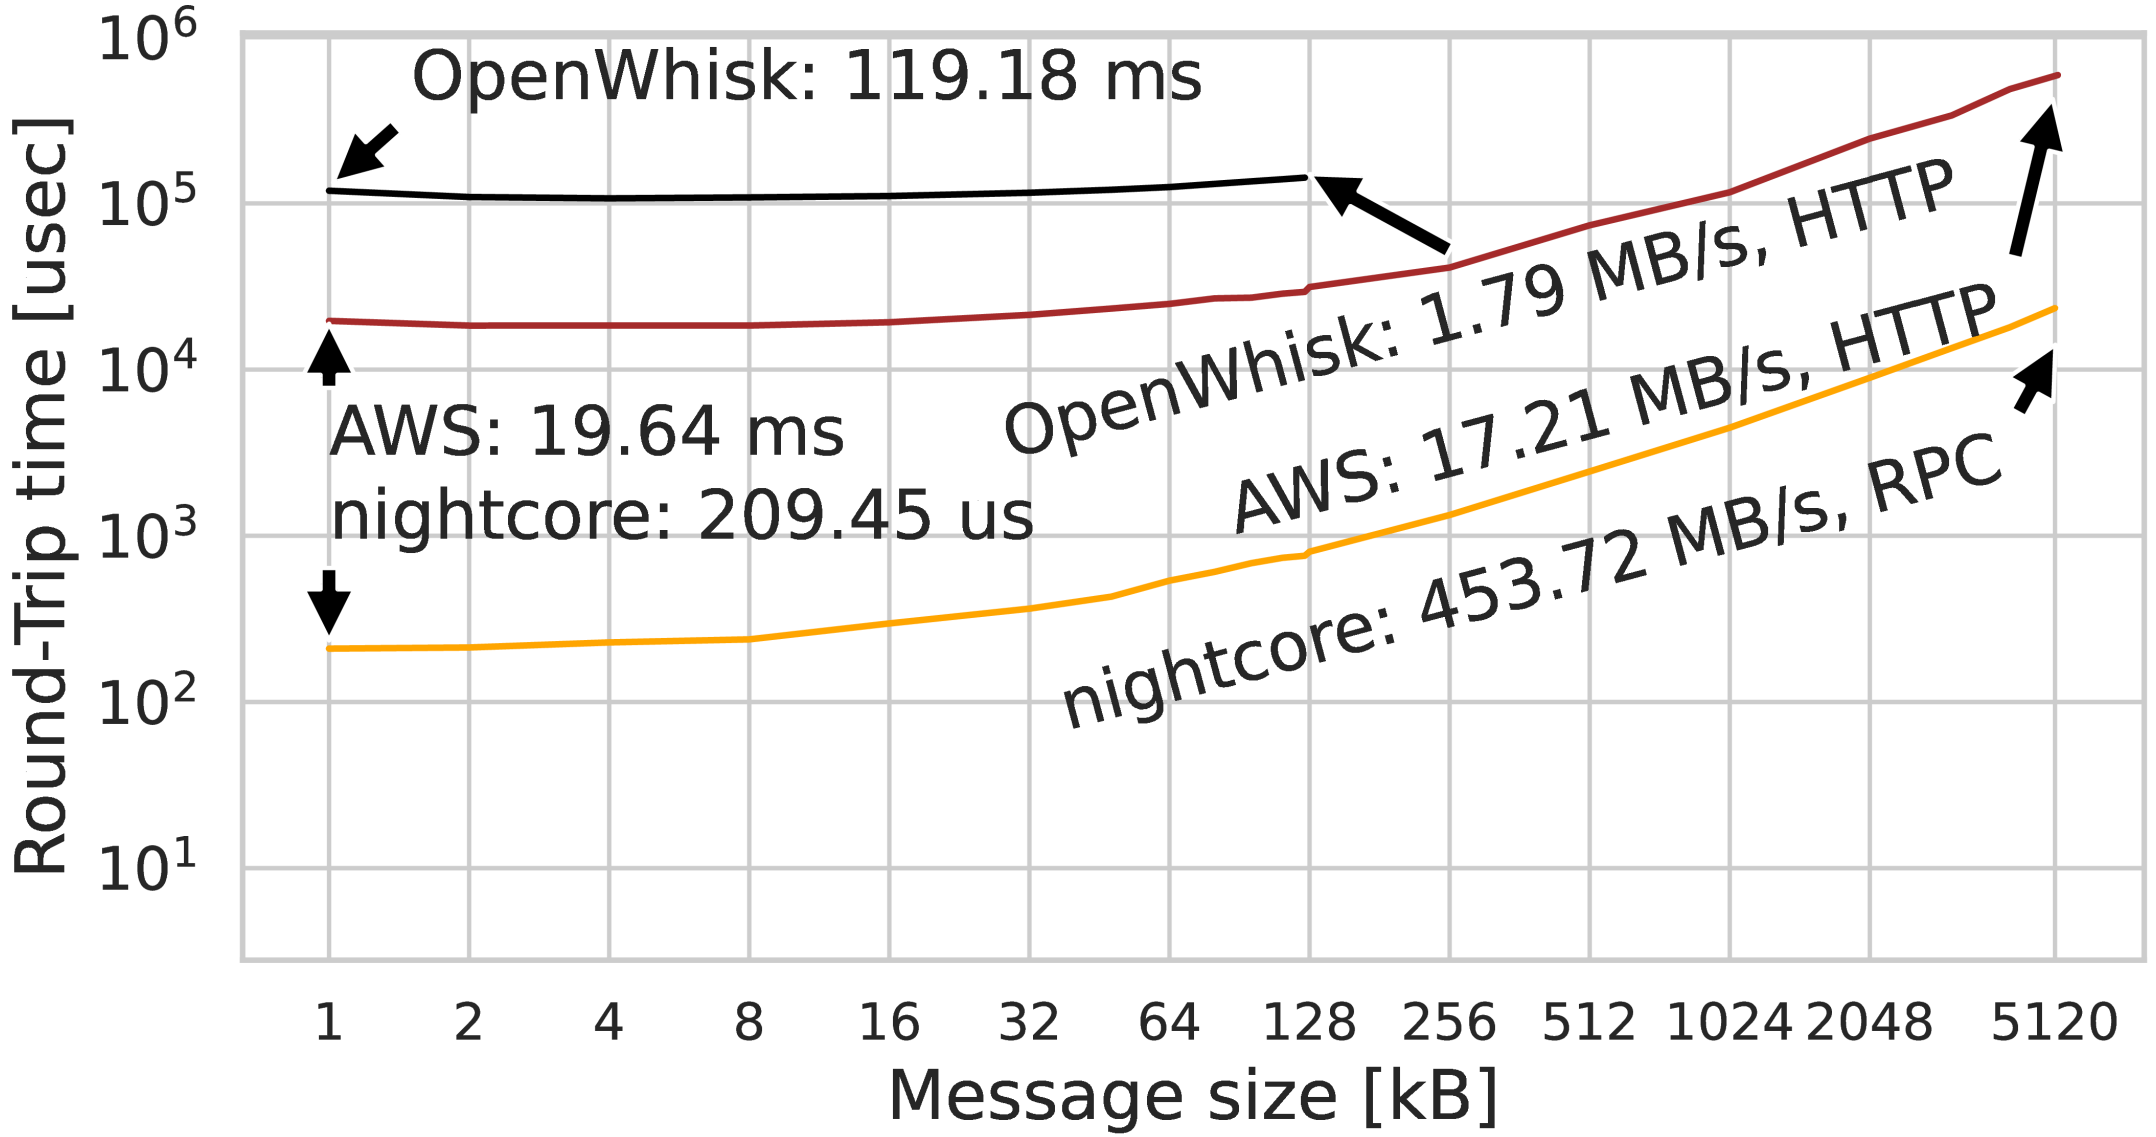

Why is FaaS slow?

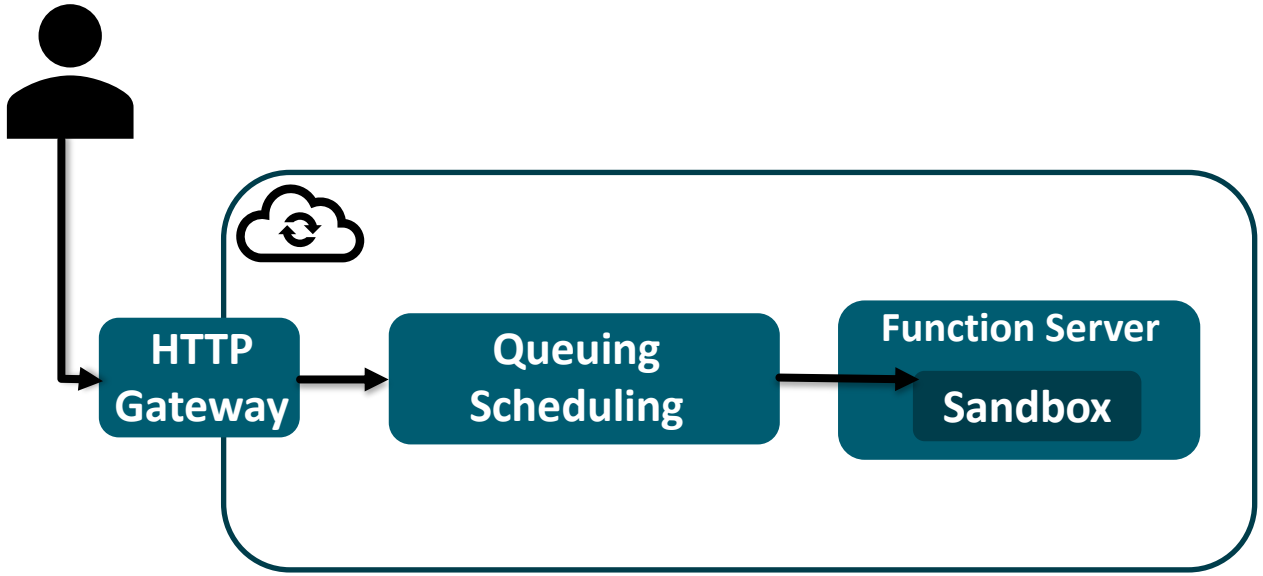

Why is FaaS slow?

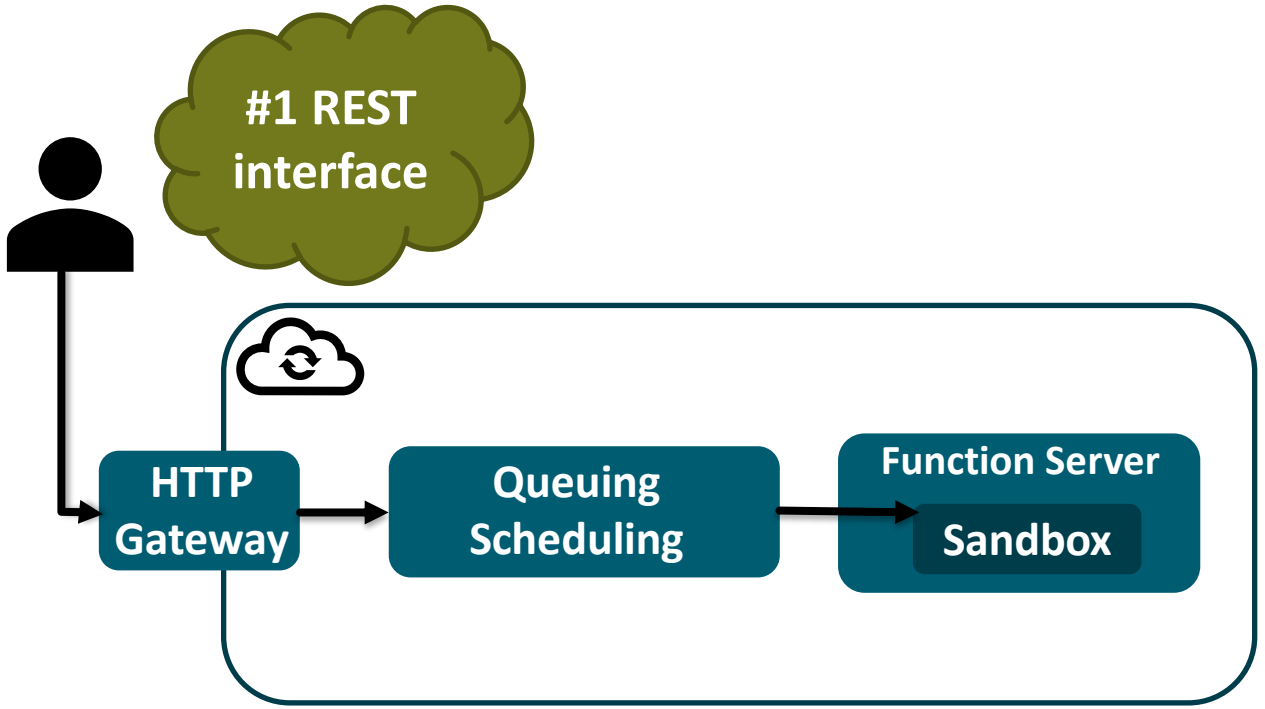

Why is FaaS slow?

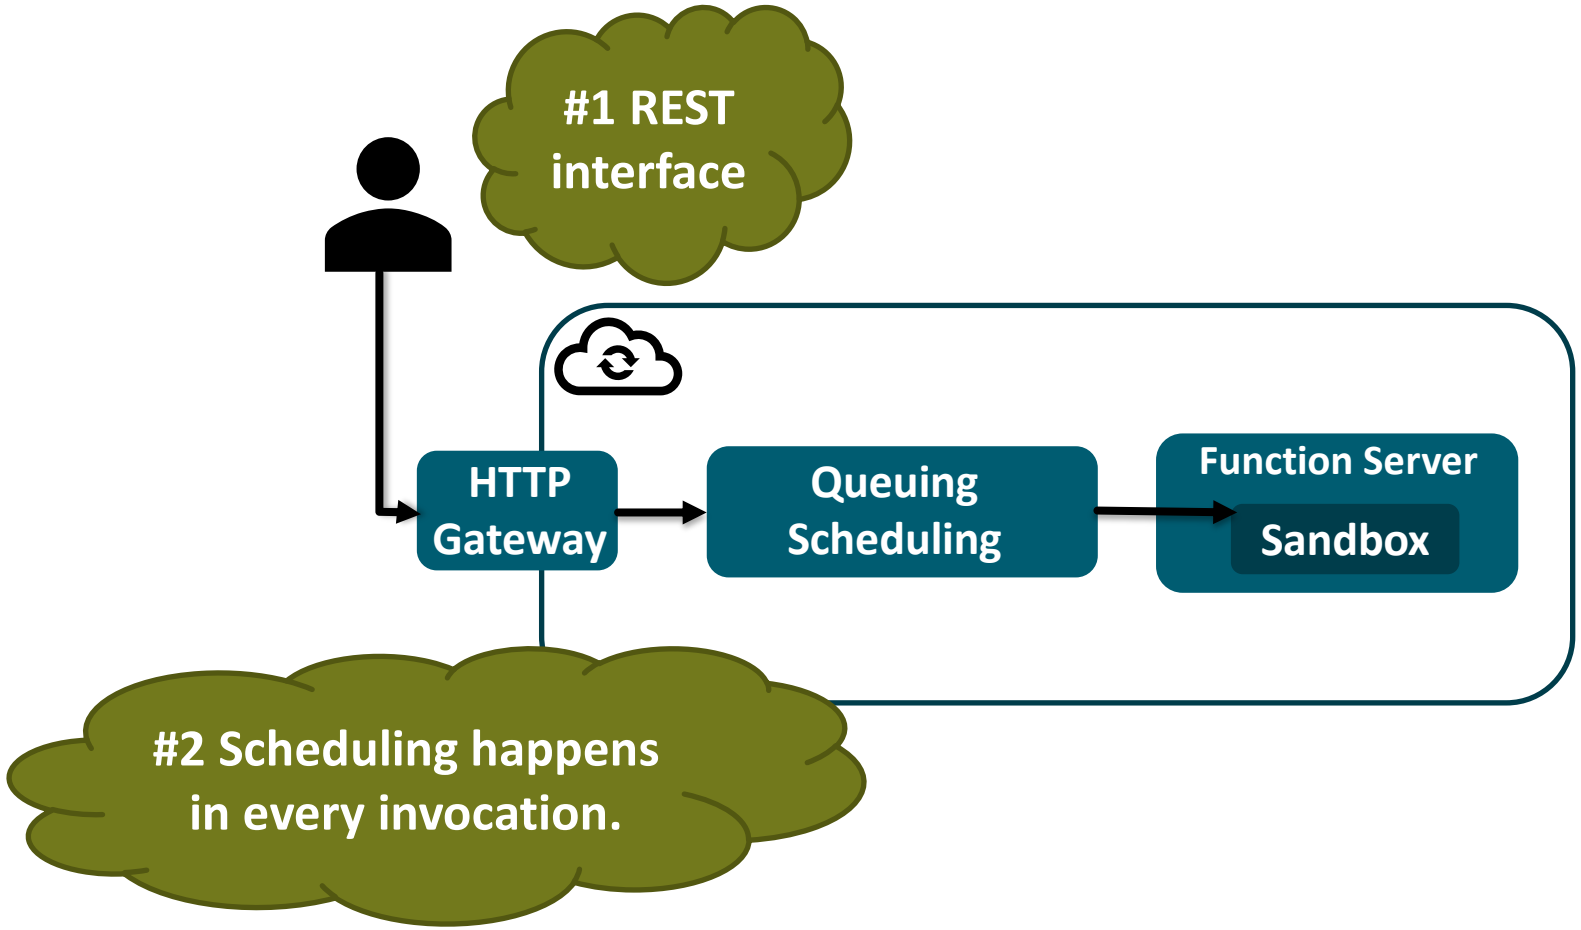

Why is FaaS slow?

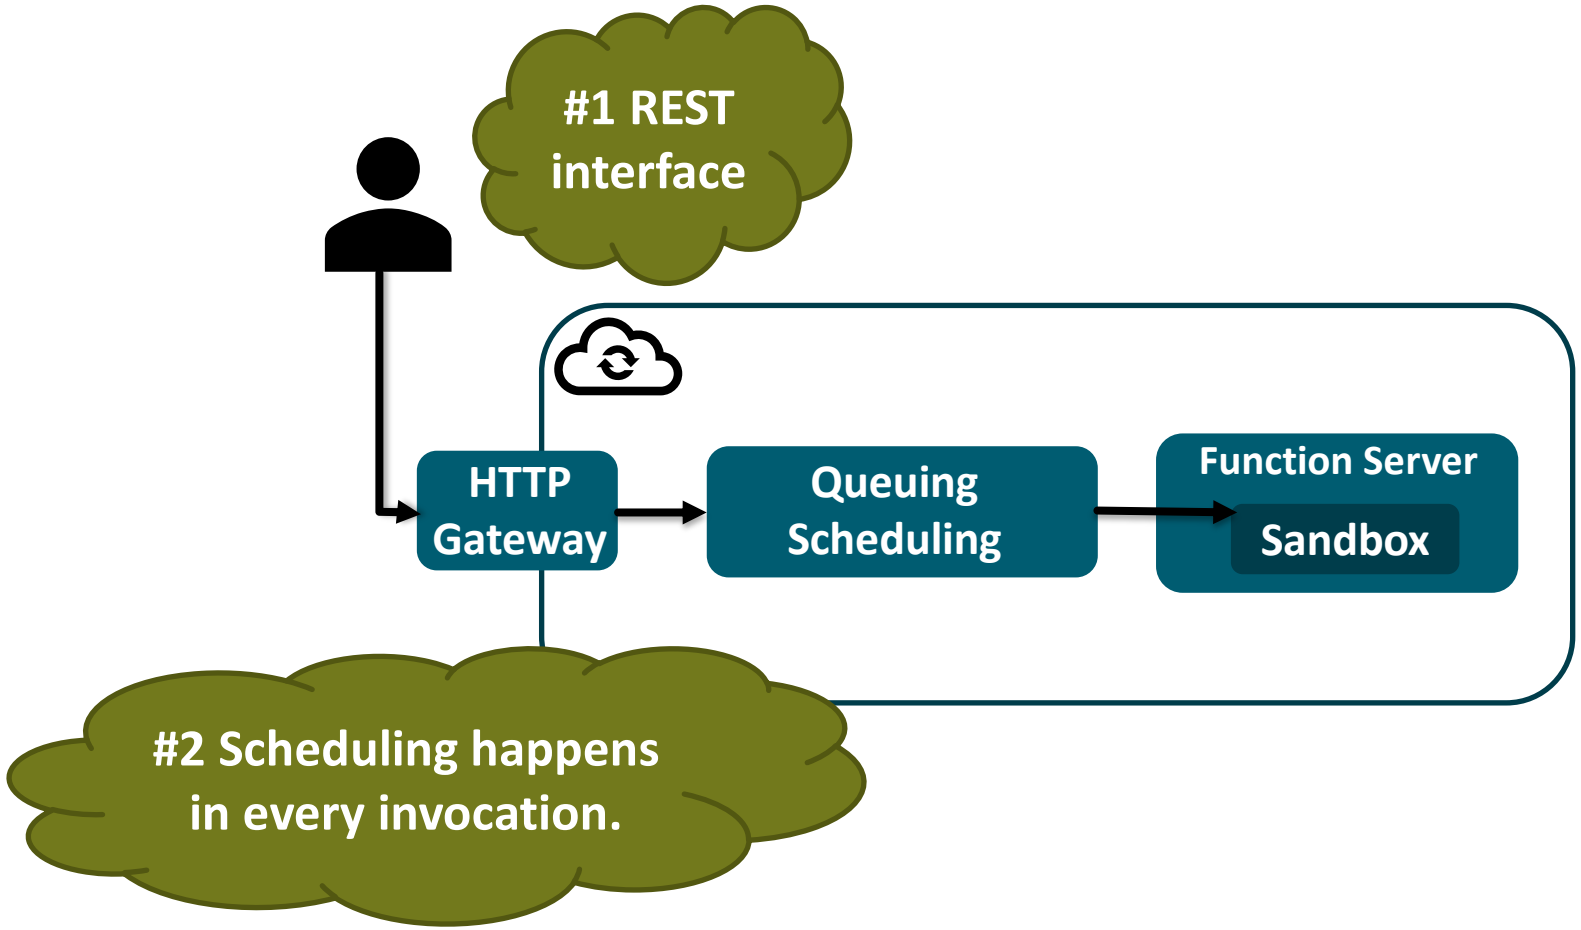

"SeBS: a Serverless Benchmark Suite for Function-as-a-Service Computing"

Why is FaaS slow?

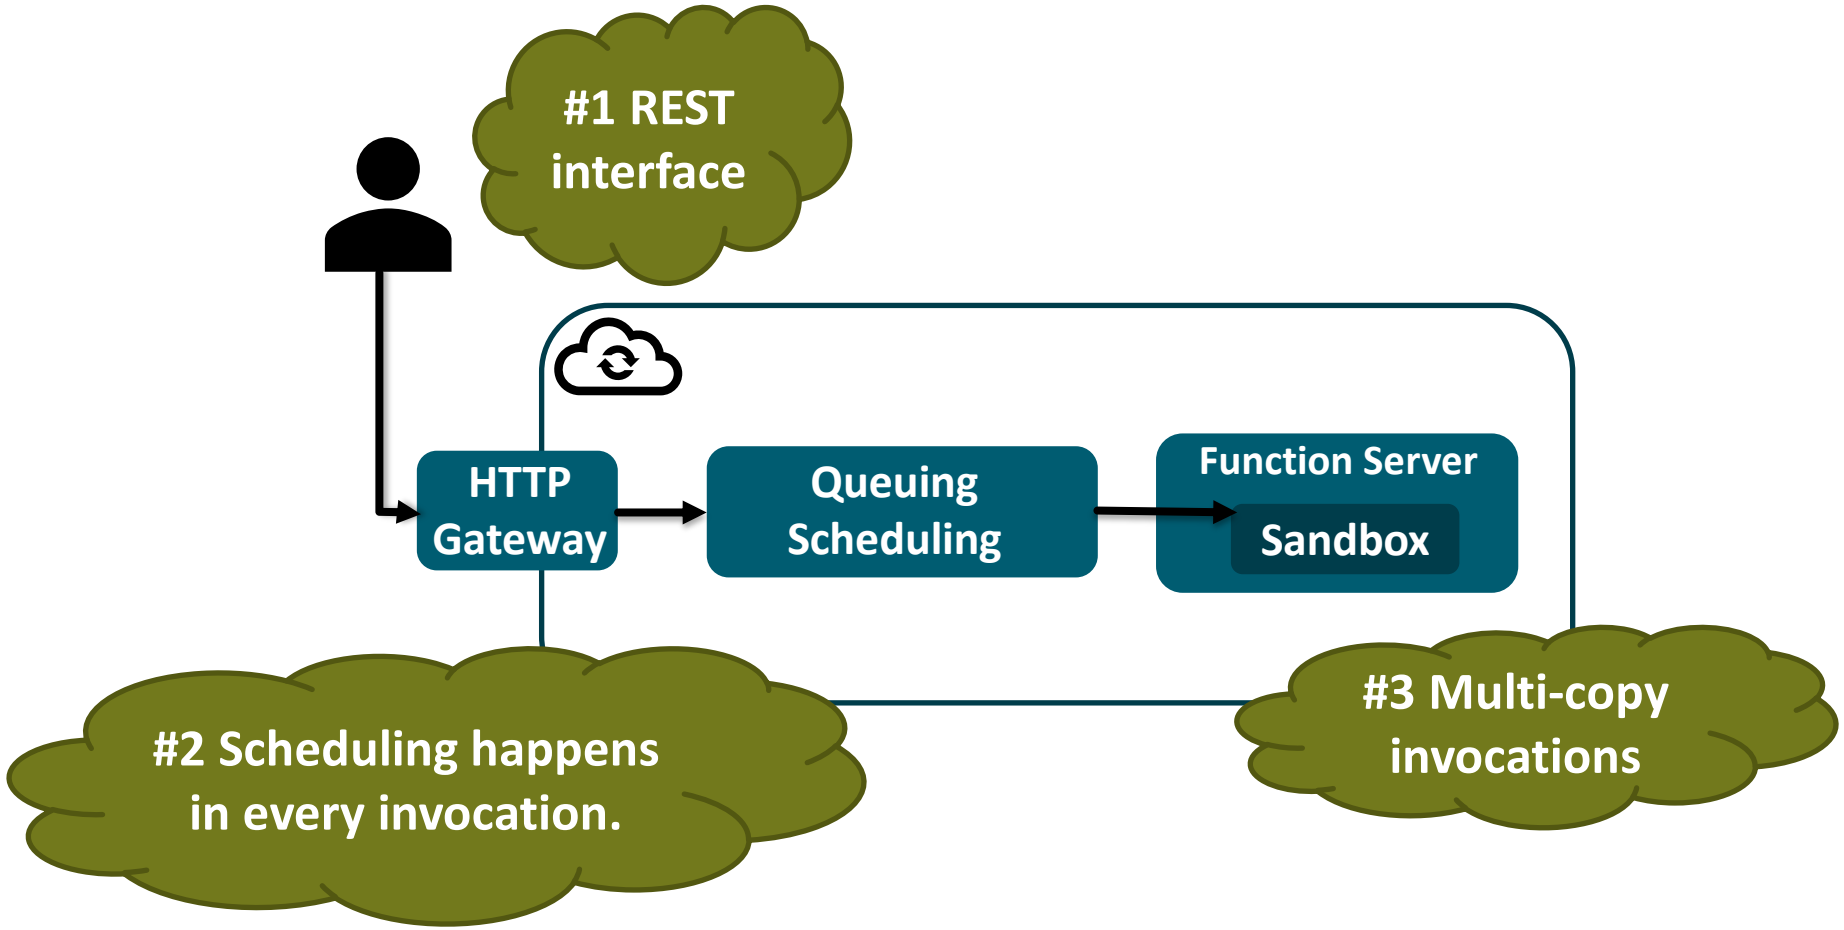

"SeBS: a Serverless Benchmark Suite for Function-as-a-Service Computing"

Why is FaaS slow?

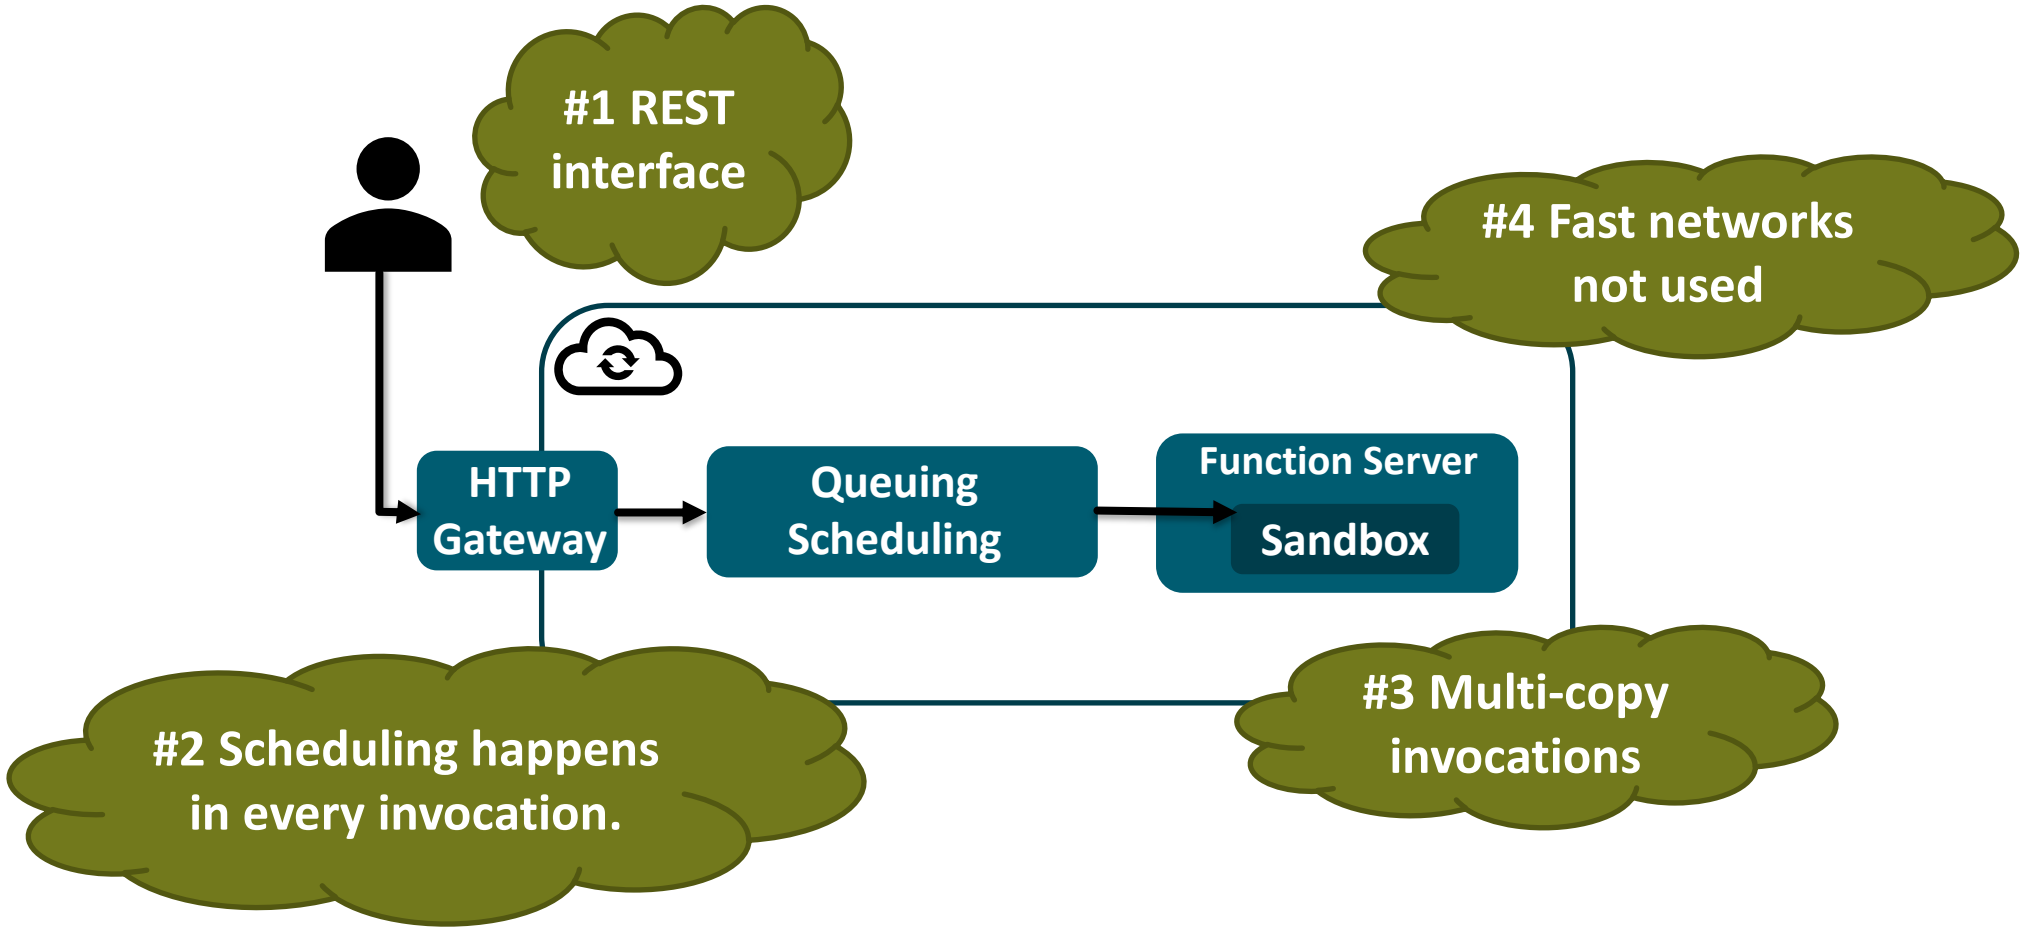

"SeBS: a Serverless Benchmark Suite for Function-as-a-Service Computing"

Function-as-a-Service for HPC

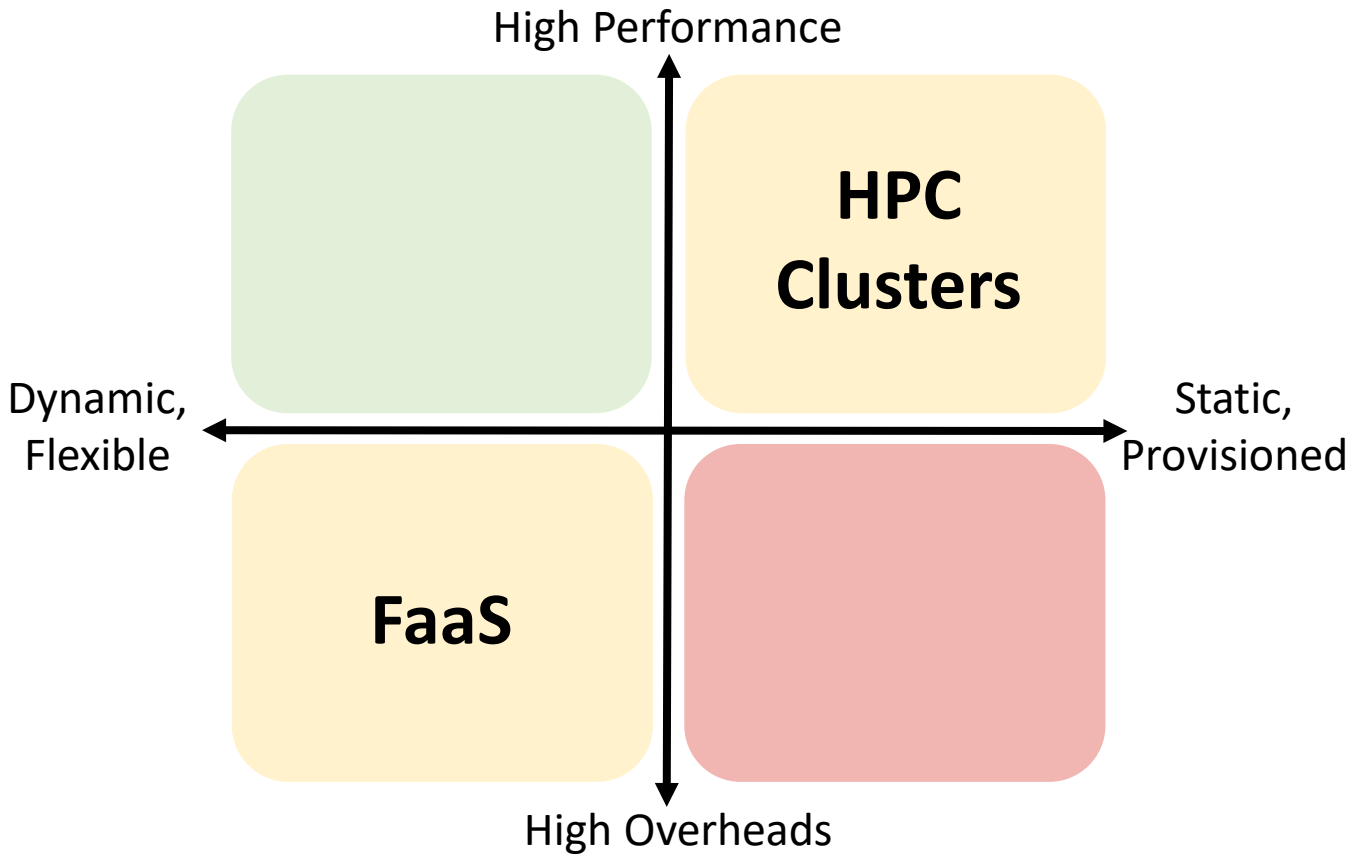

Function-as-a-Service for HPC

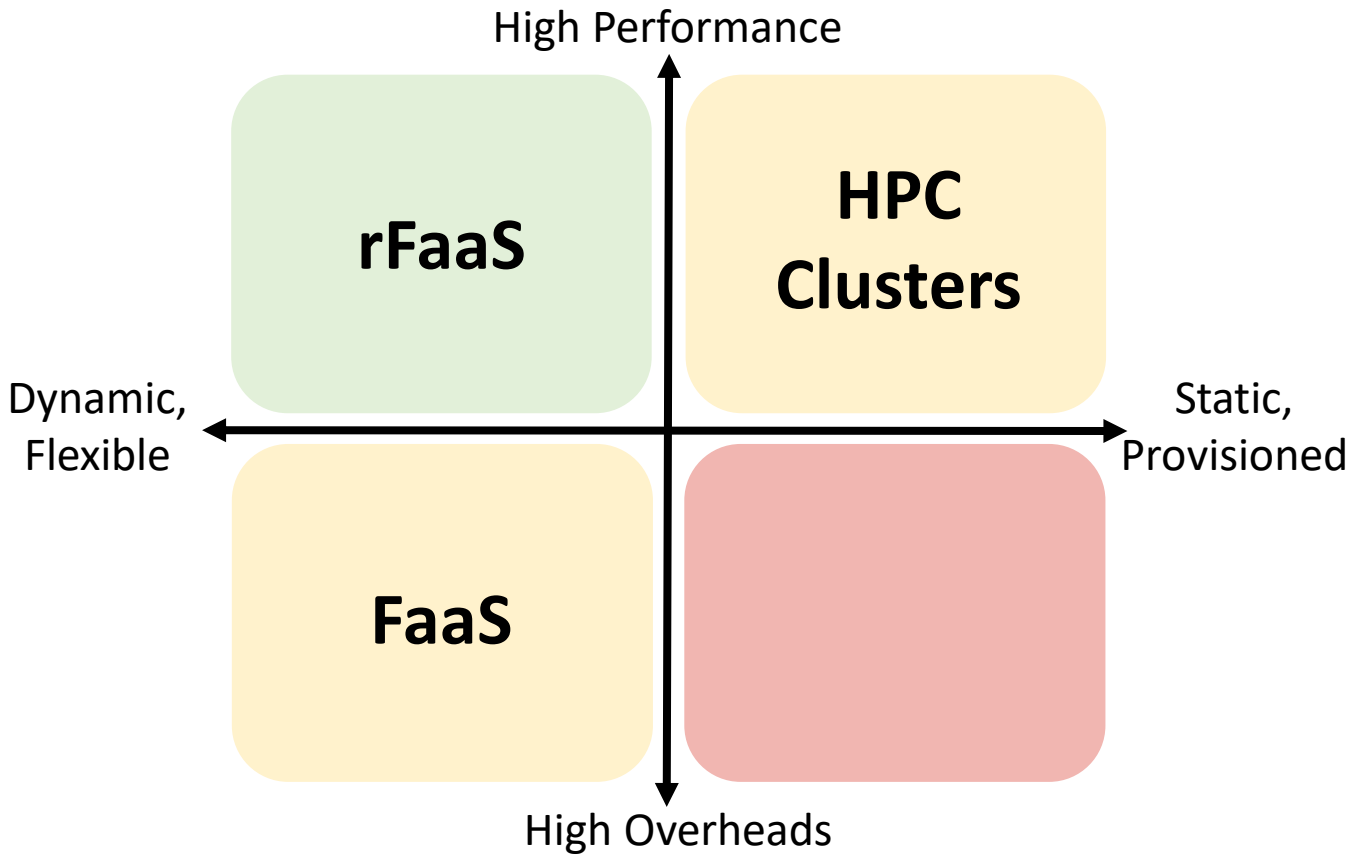

Function-as-a-Service for HPC

Reduced invocation critical path

Zero-copy RDMA

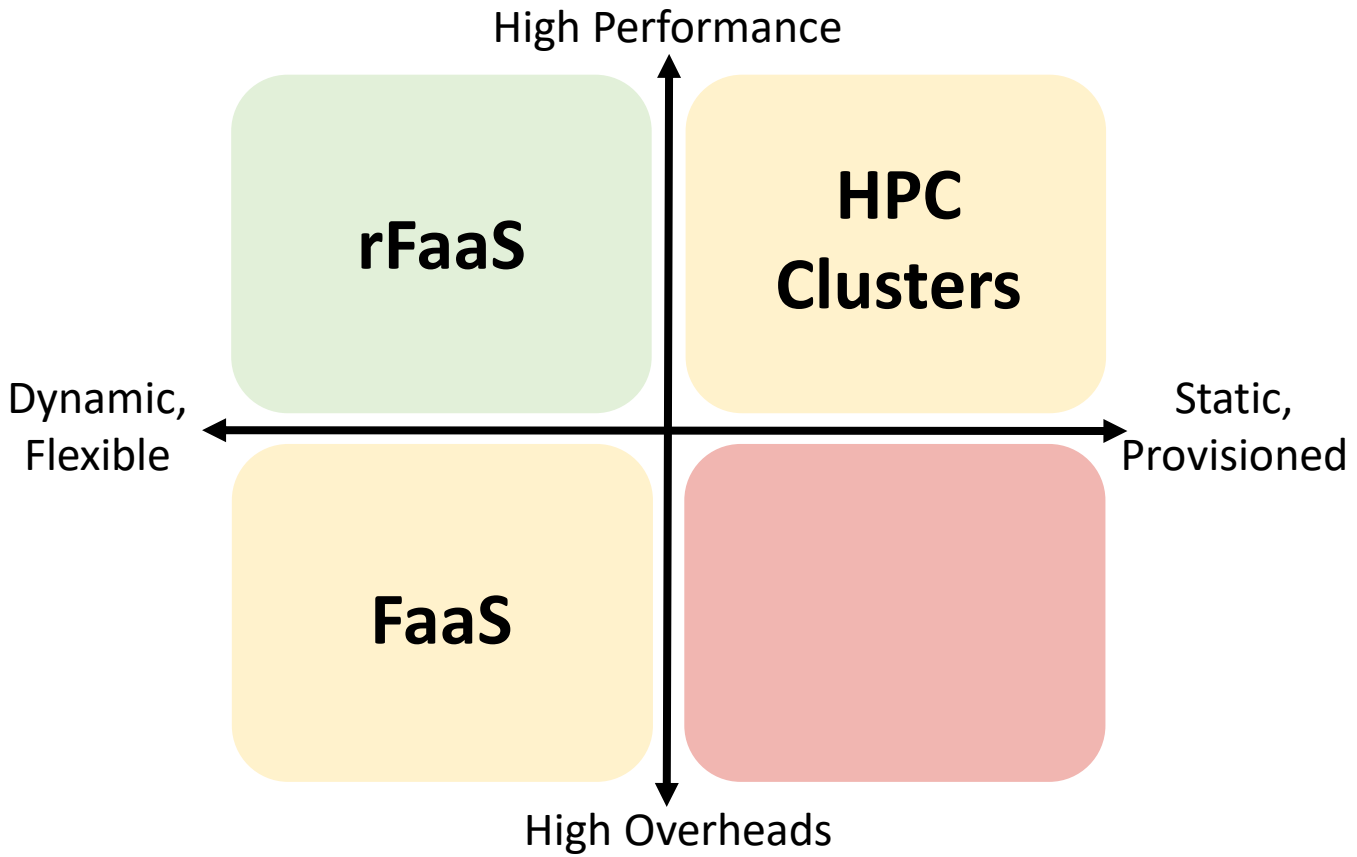

Invocations in FaaS and rFaaS

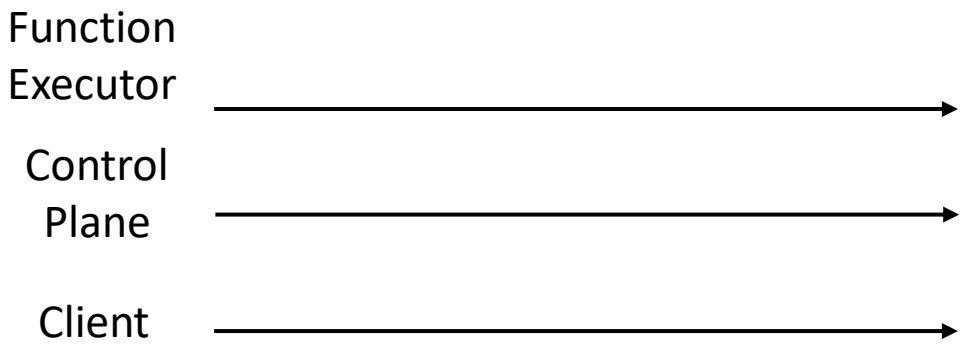

Invocations in FaaS and rFaaS

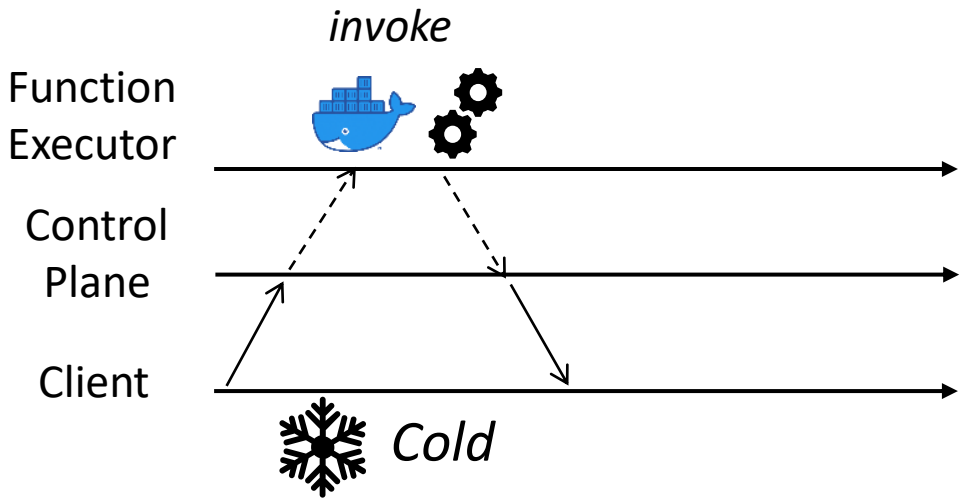

Invocations in FaaS and rFaaS

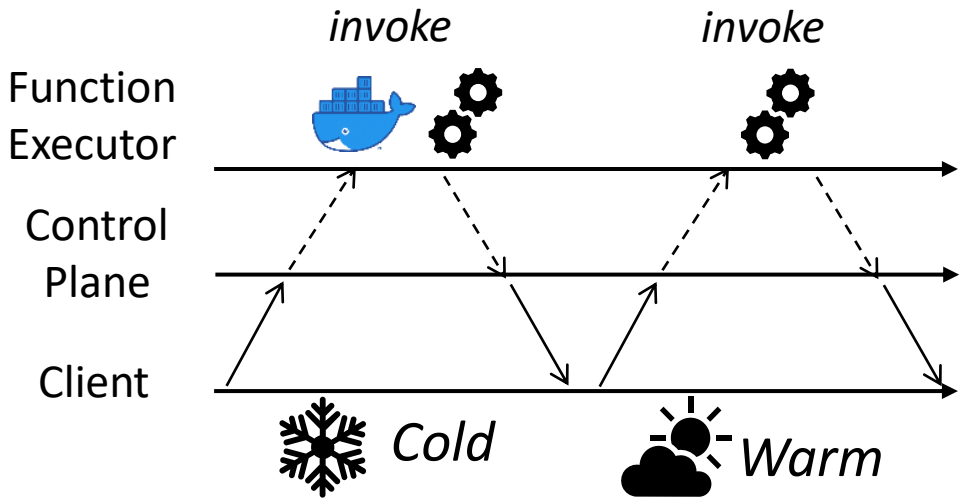

Invocations in FaaS and rFaaS

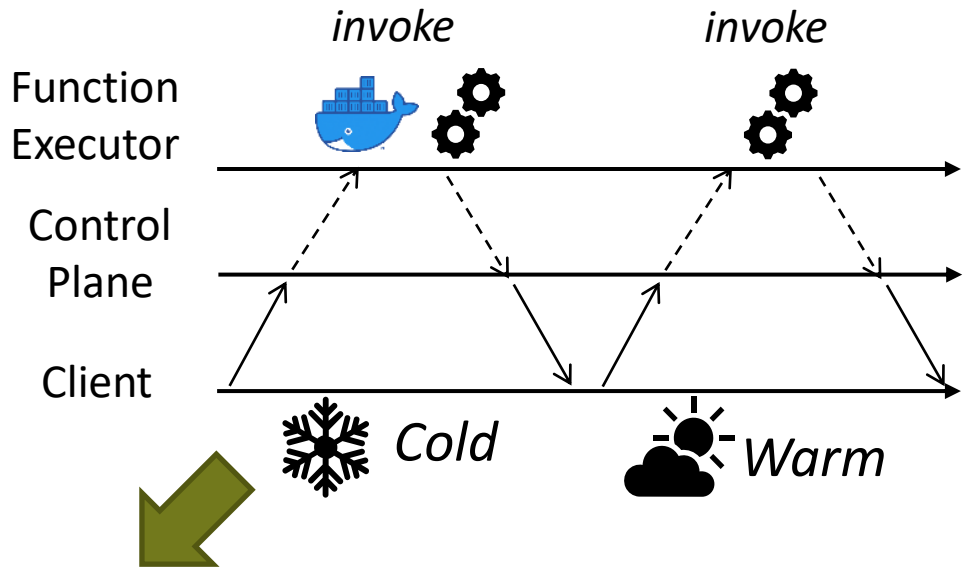

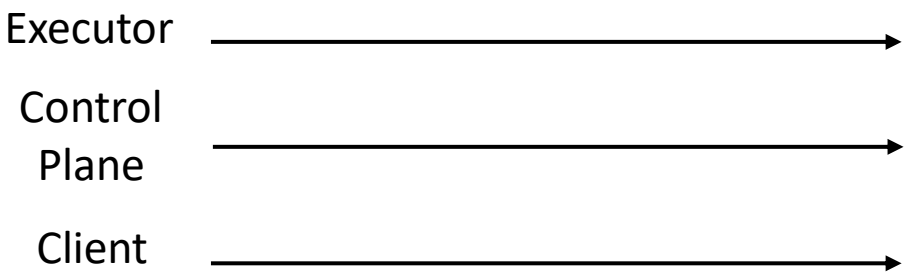

Invocations in FaaS and rFaaS

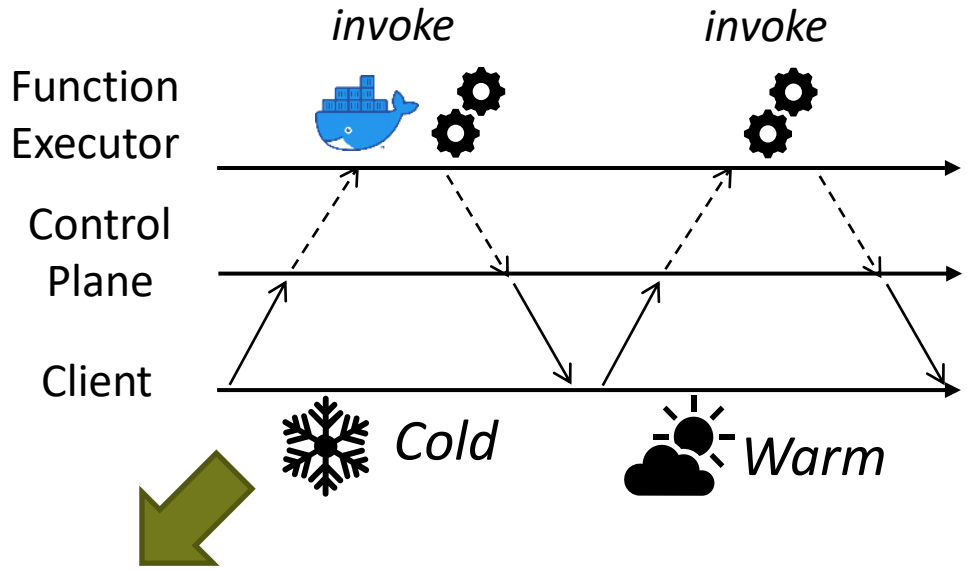

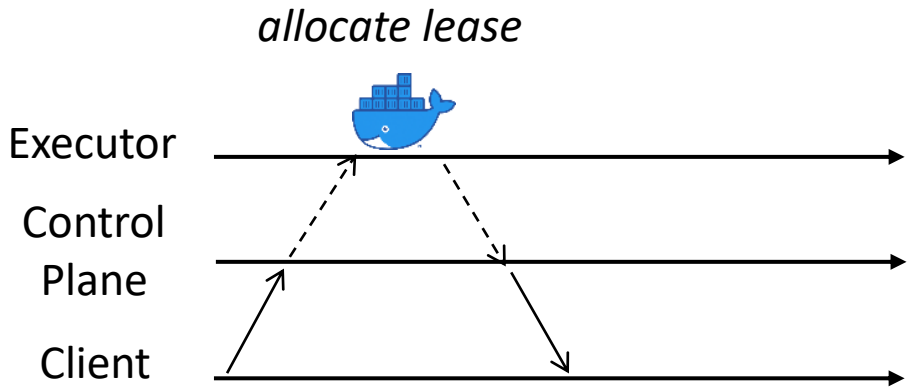

Invocations in FaaS and rFaaS

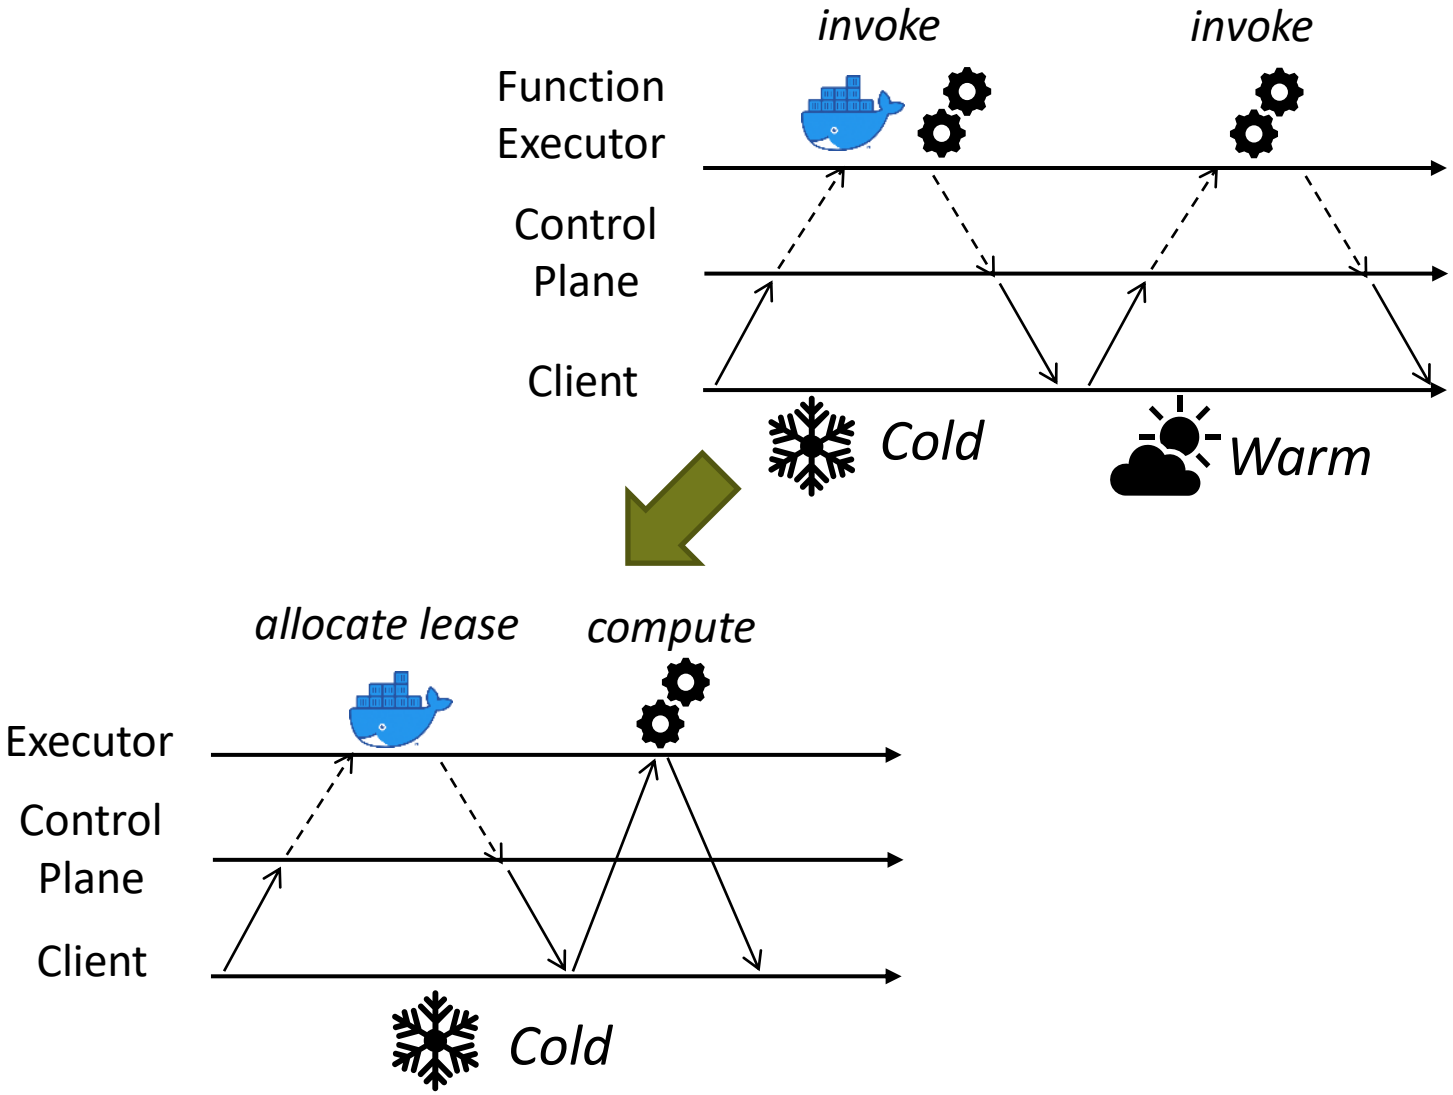

Invocations in FaaS and rFaaS

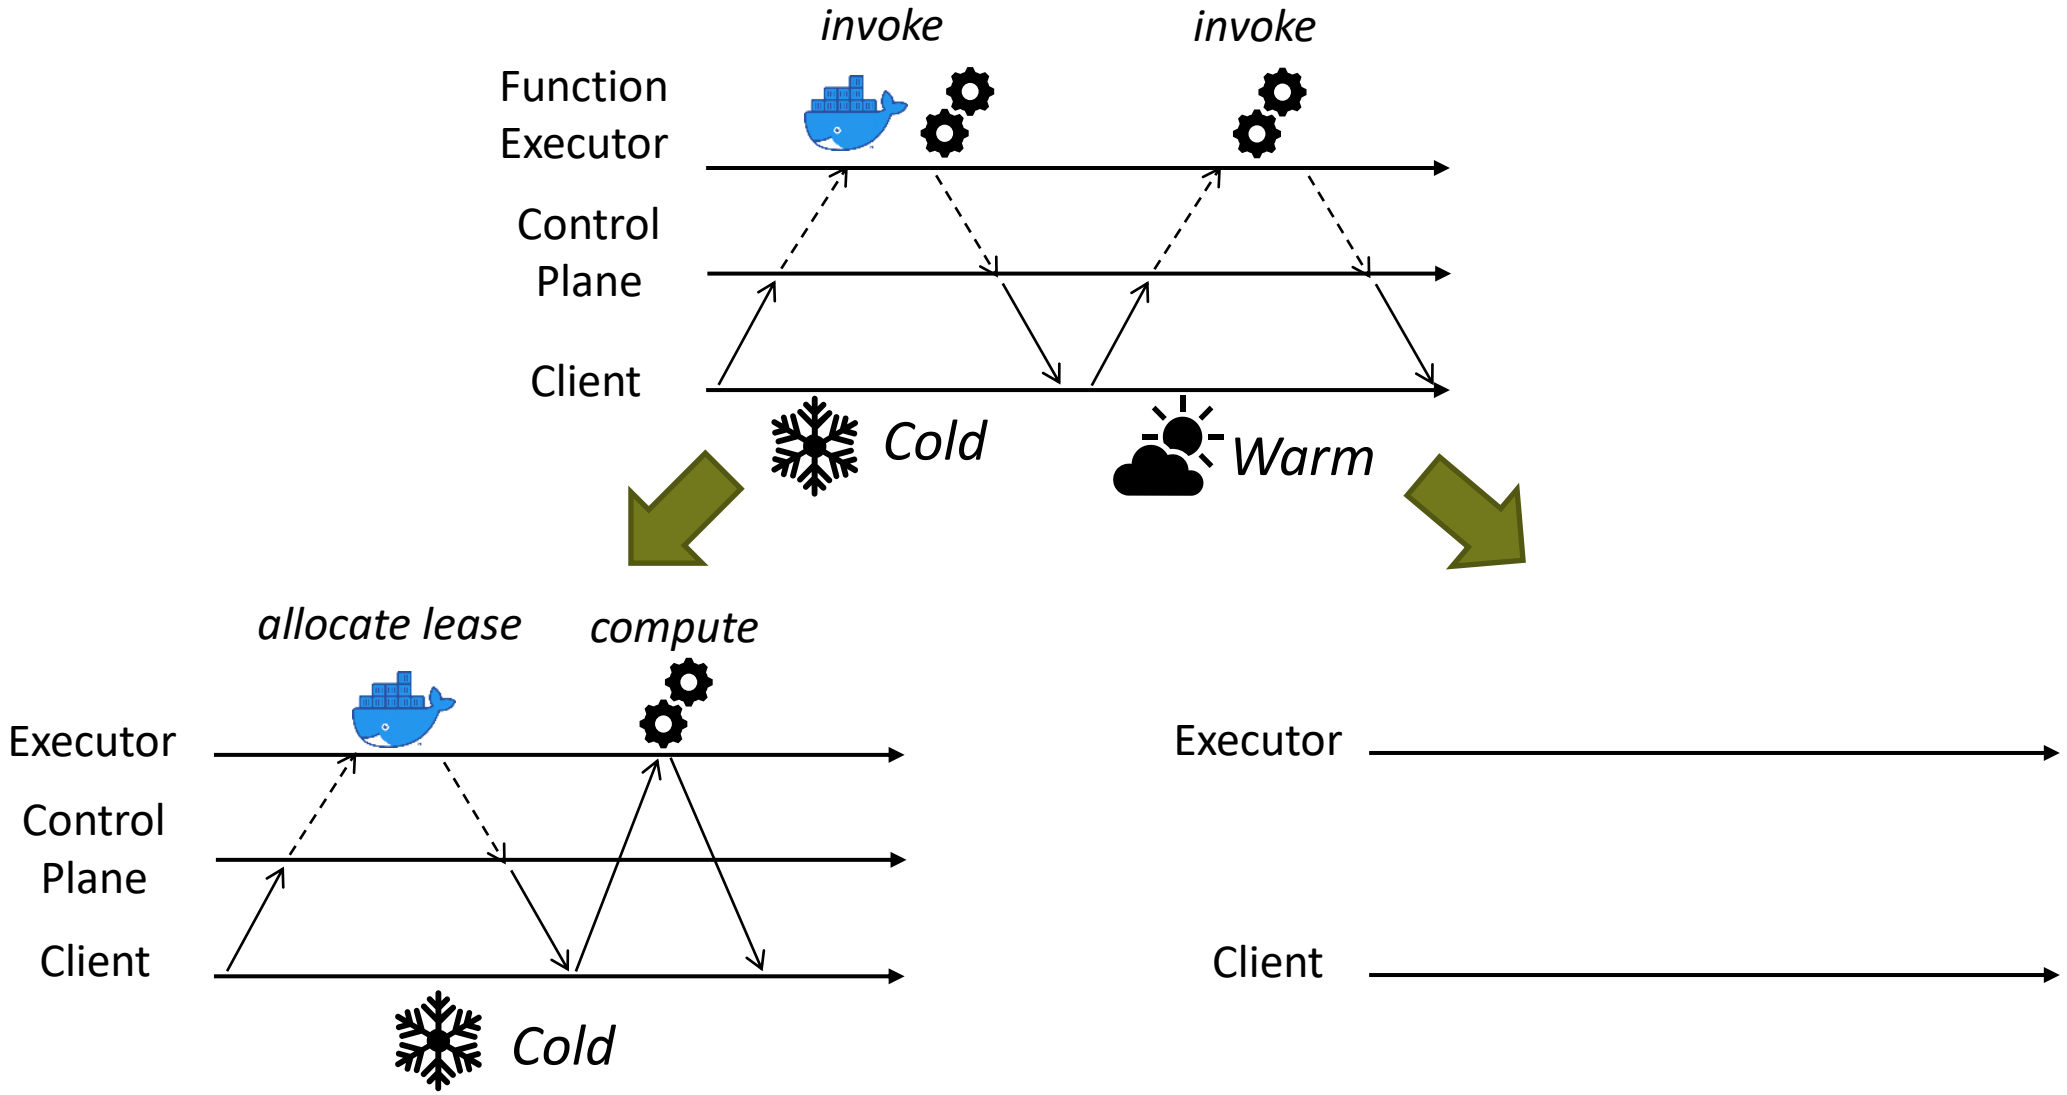

Invocations in FaaS and rFaaS

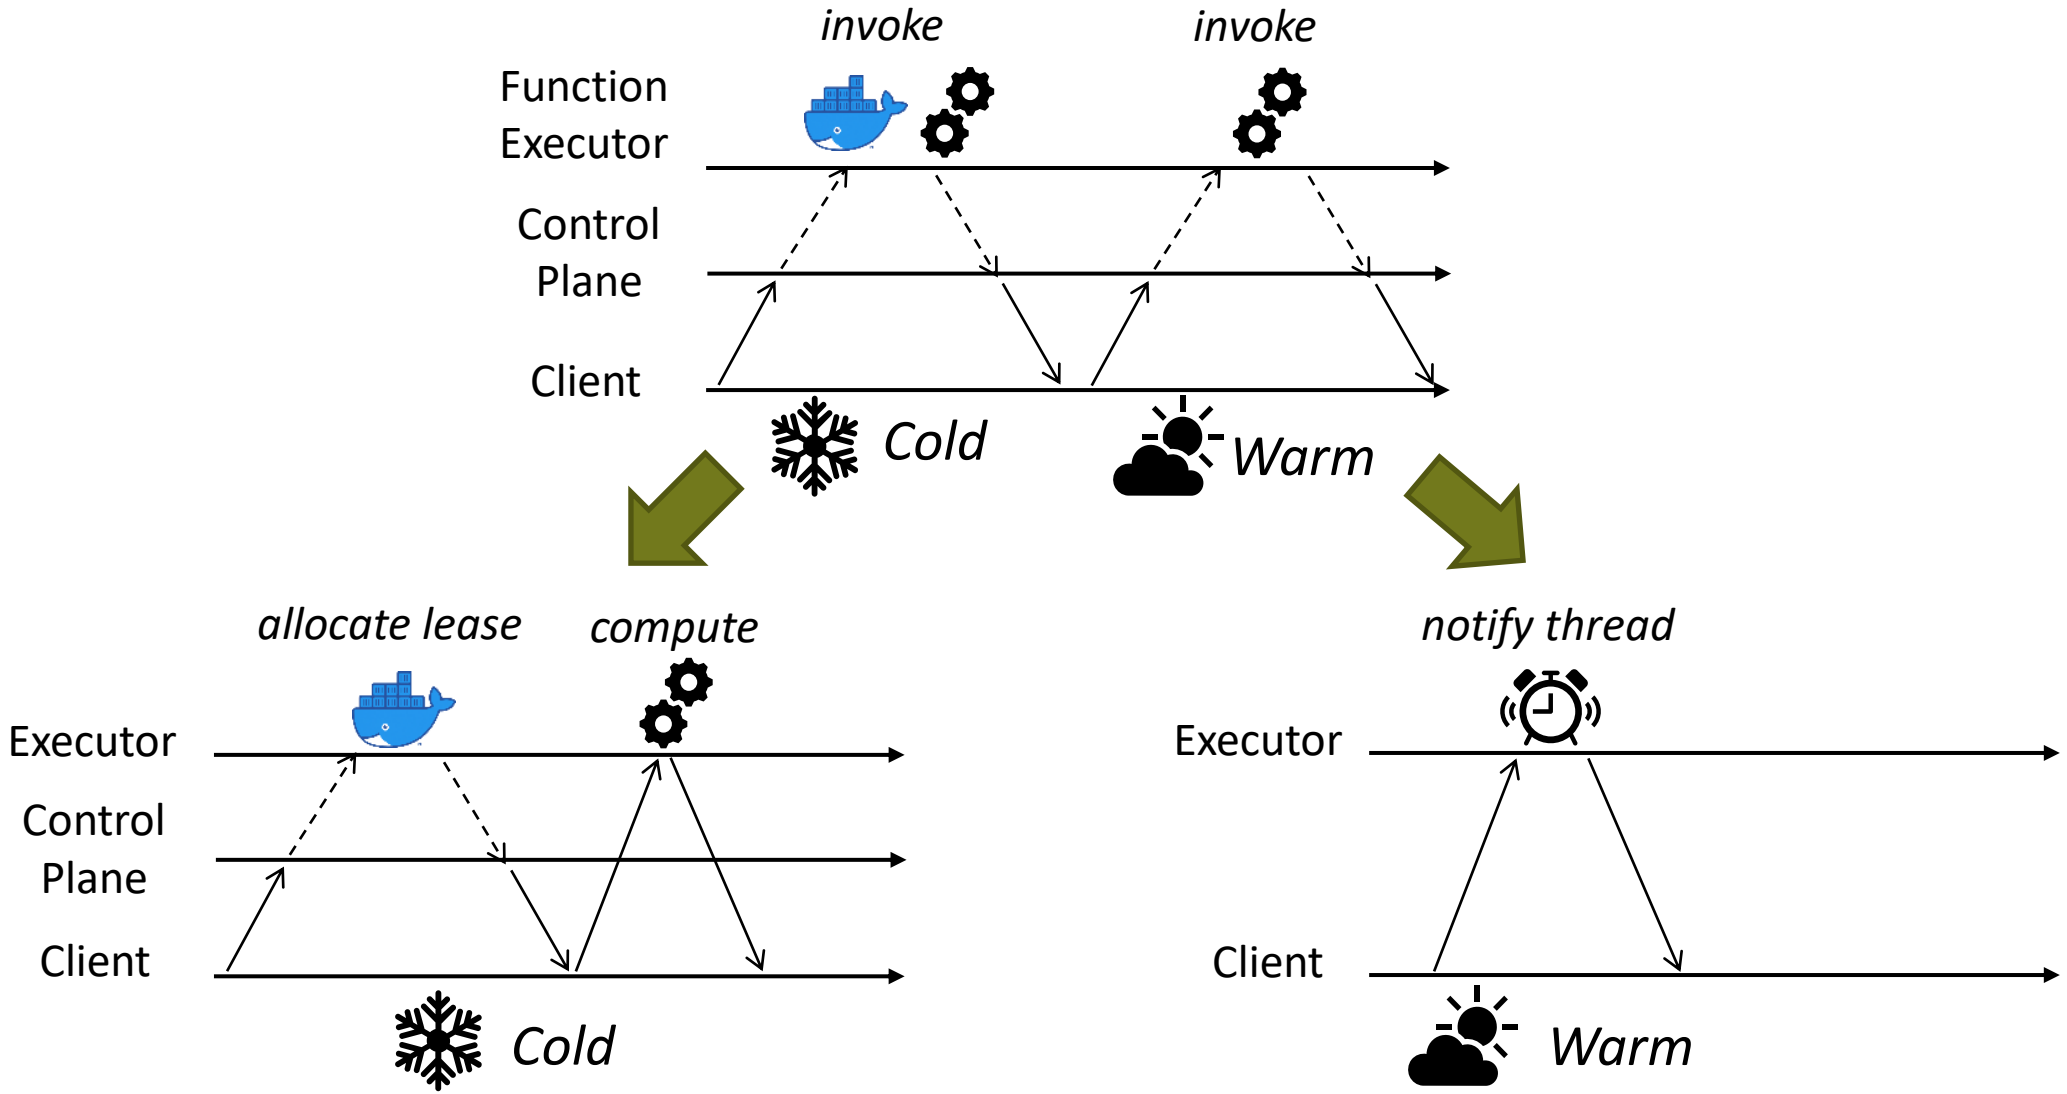

Invocations in FaaS and rFaaS

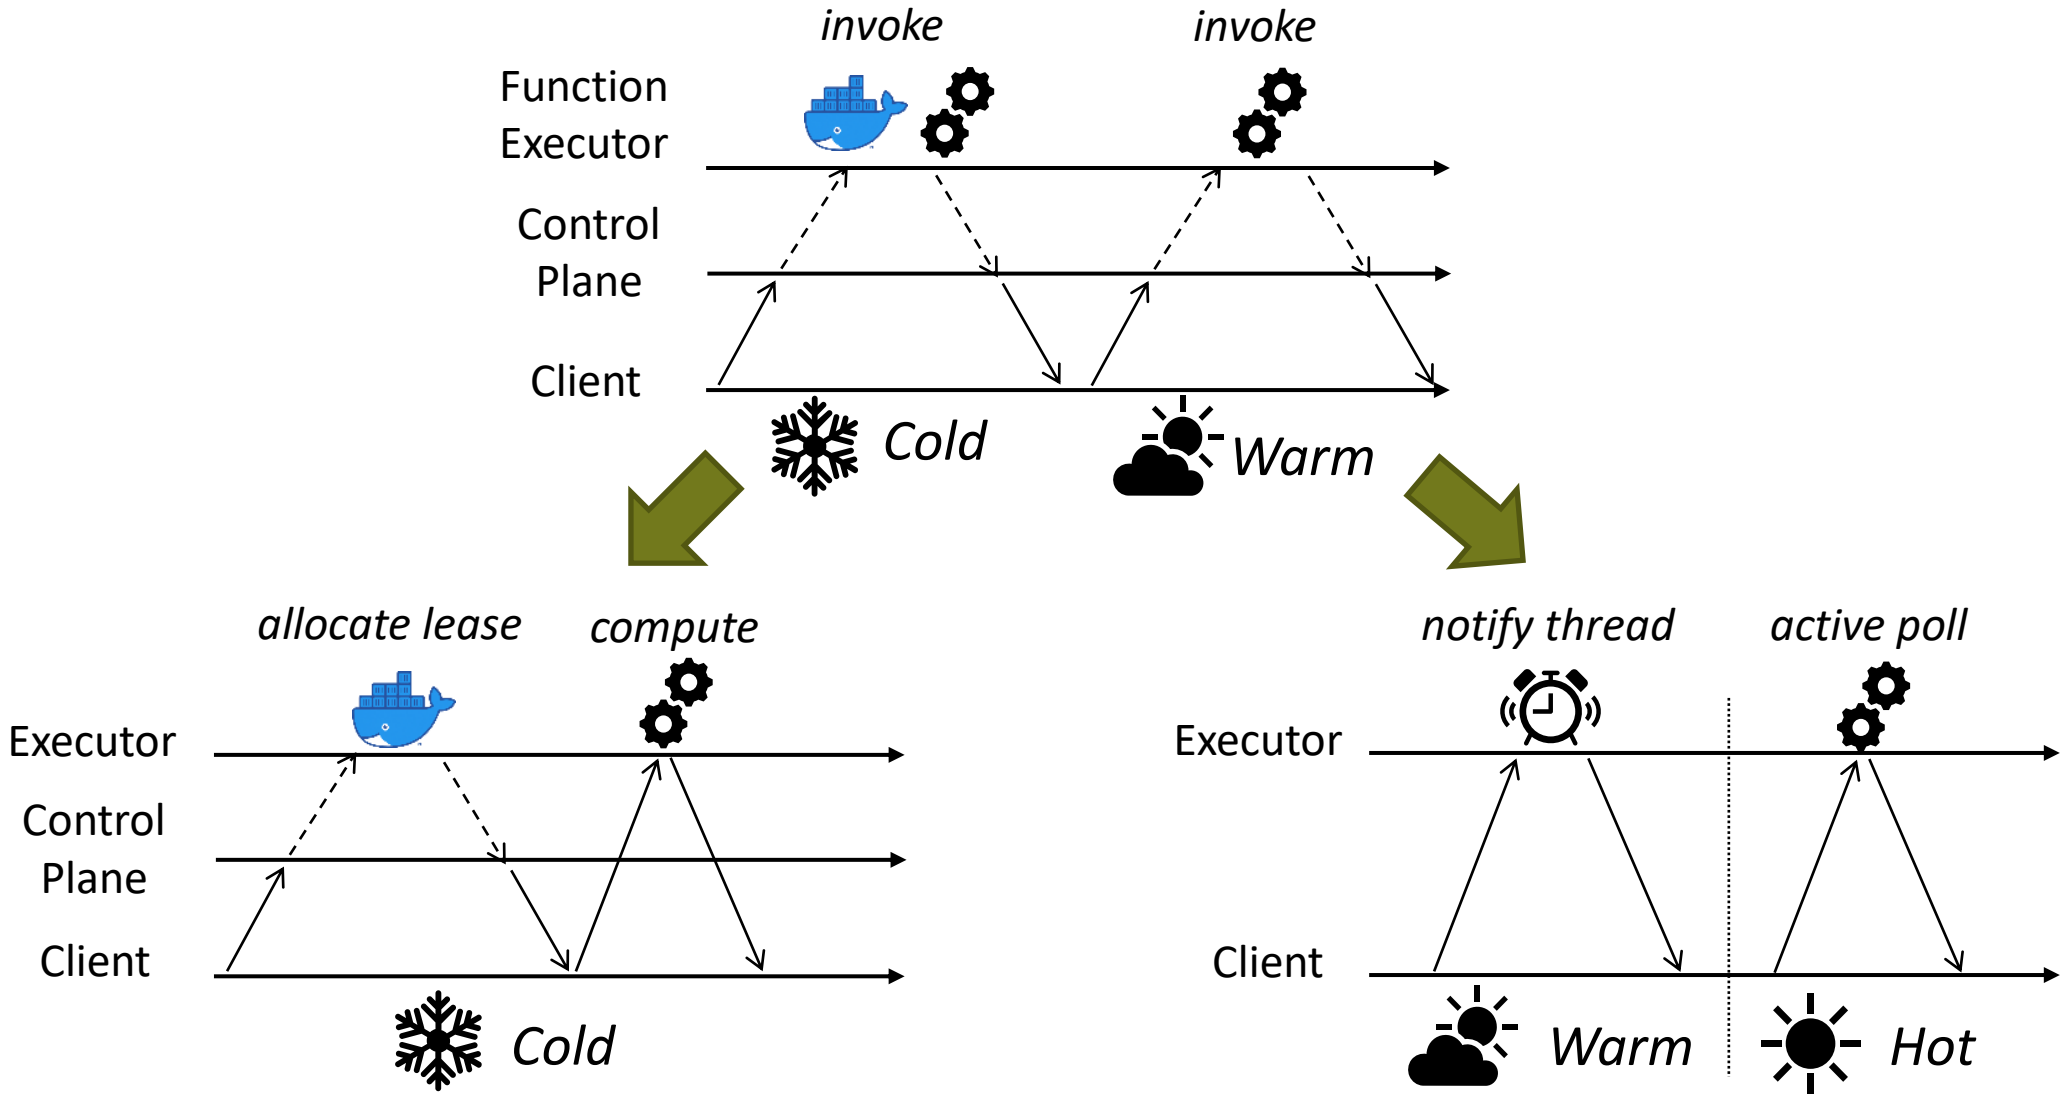

How efficient is rFaaS?

36 CPU cores, 377 GB memory.
100 Gbps Ethernet with RoCEv2 support.

How efficient is rFaaS?

36 CPU cores, 377 GB memory.
100 Gbps Ethernet with RoCEv2 support.

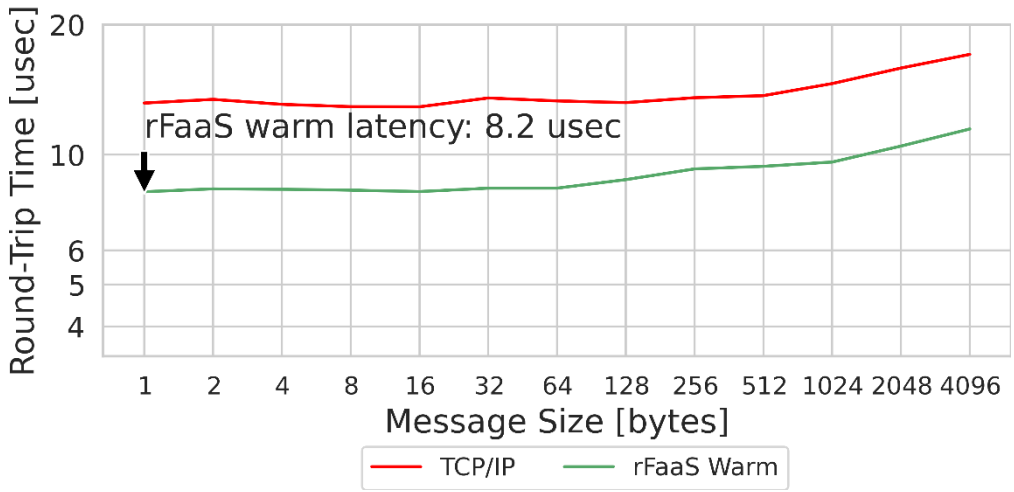

How efficient is rFaaS?

36 CPU cores, 377 GB memory.
100 Gbps Ethernet with RoCEv2 support.

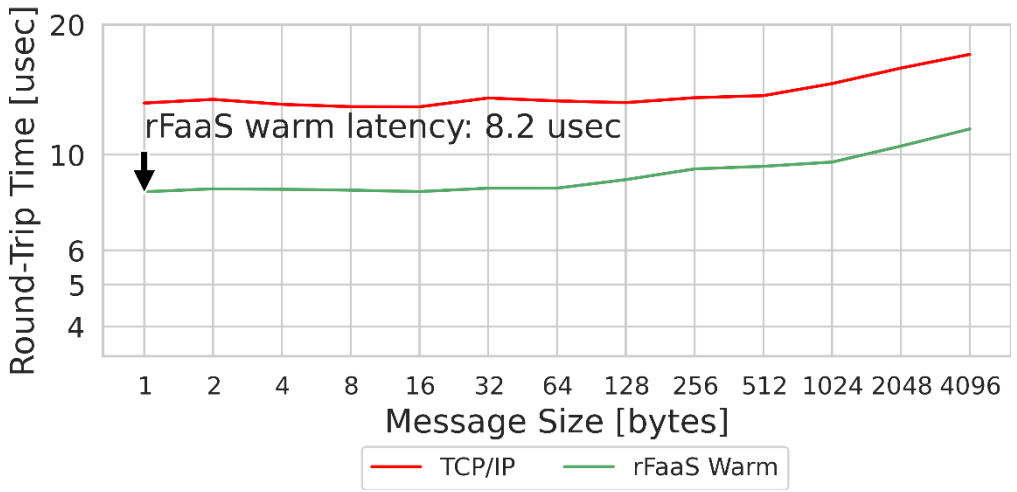

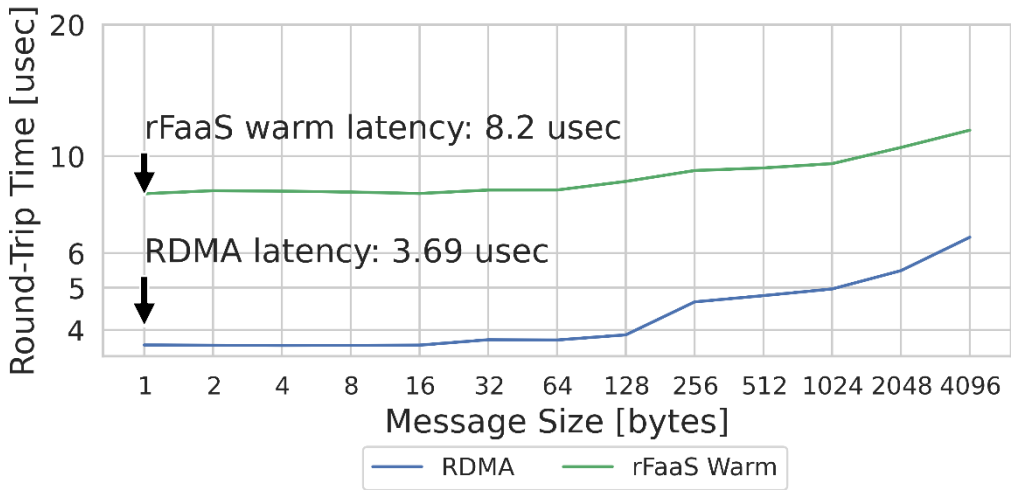

How efficient is rFaaS?

36 CPU cores, 377 GB memory.
100 Gbps Ethernet with RoCEv2 support.

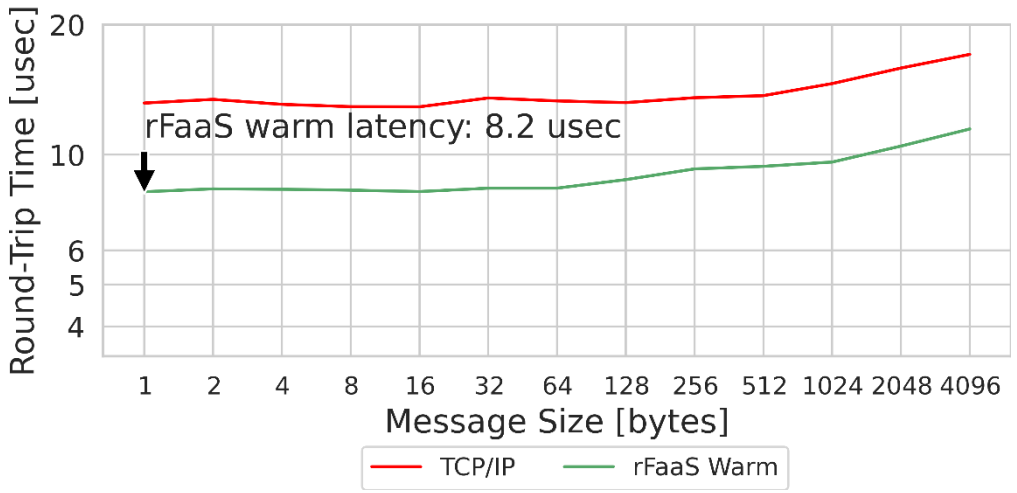

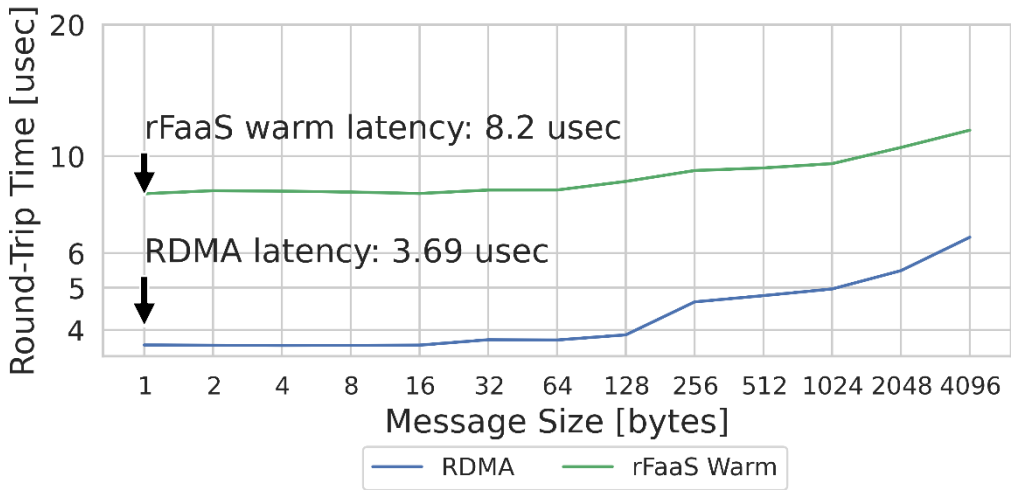

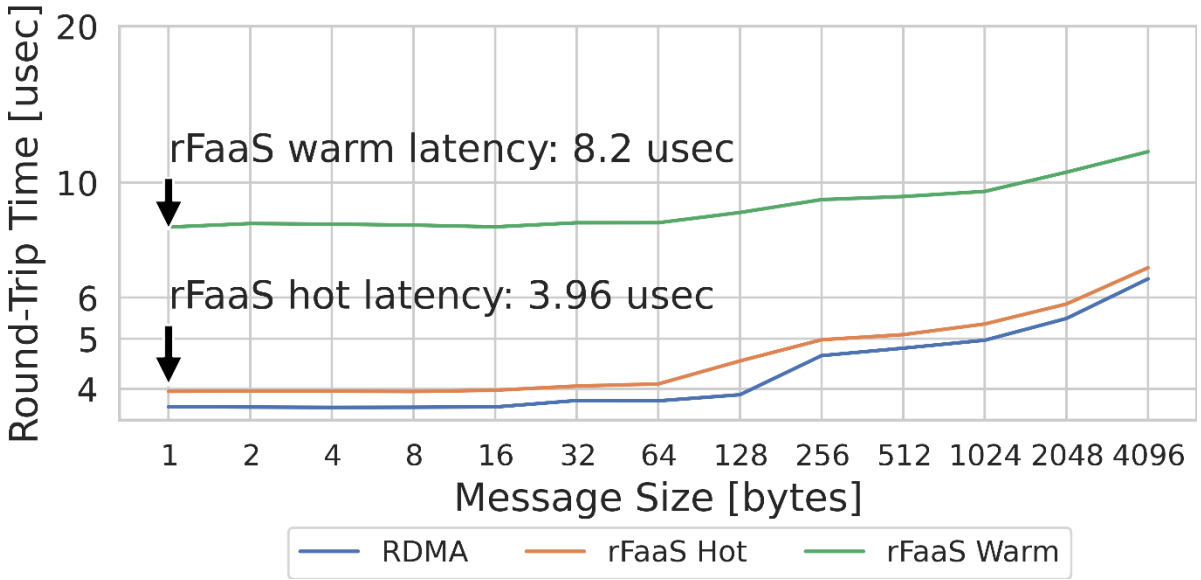

How efficient is rFaaS?

36 CPU cores, 377 GB memory.
100 Gbps Ethernet with RoCEv2 support.

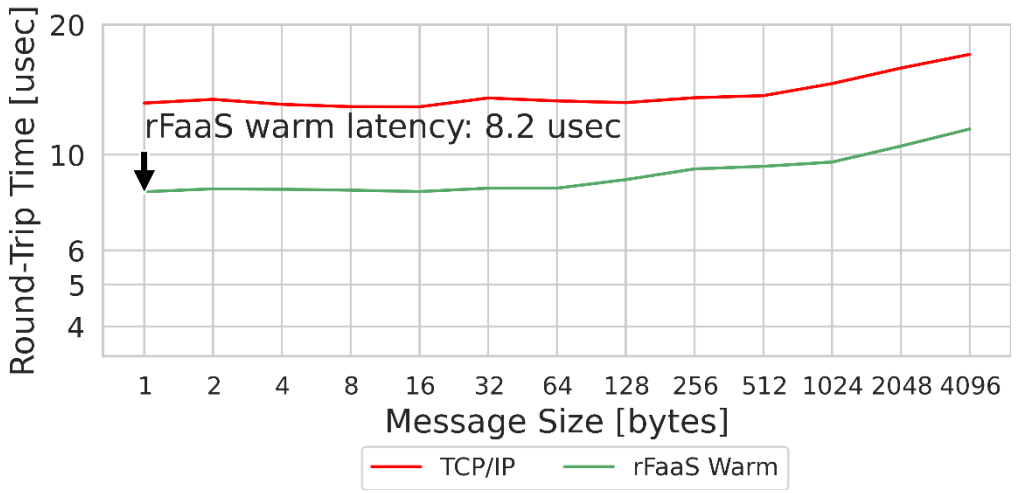

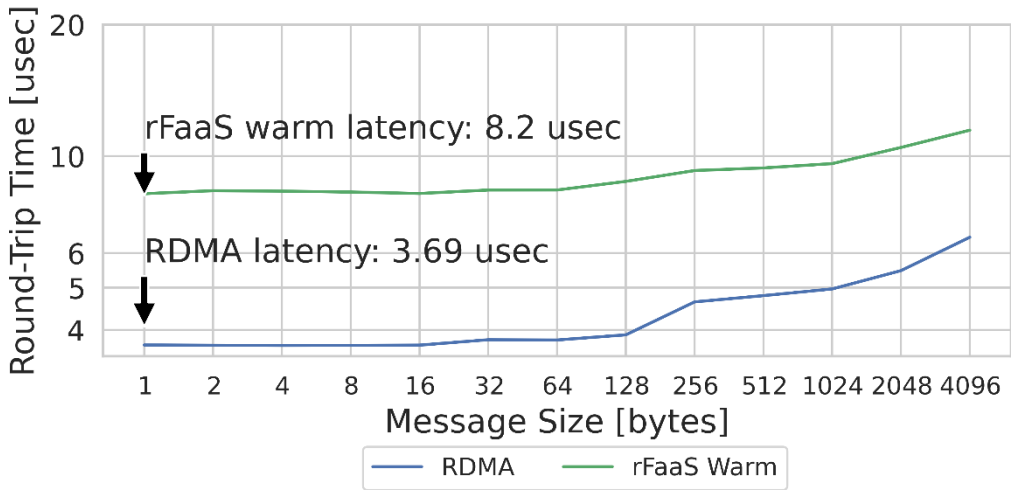

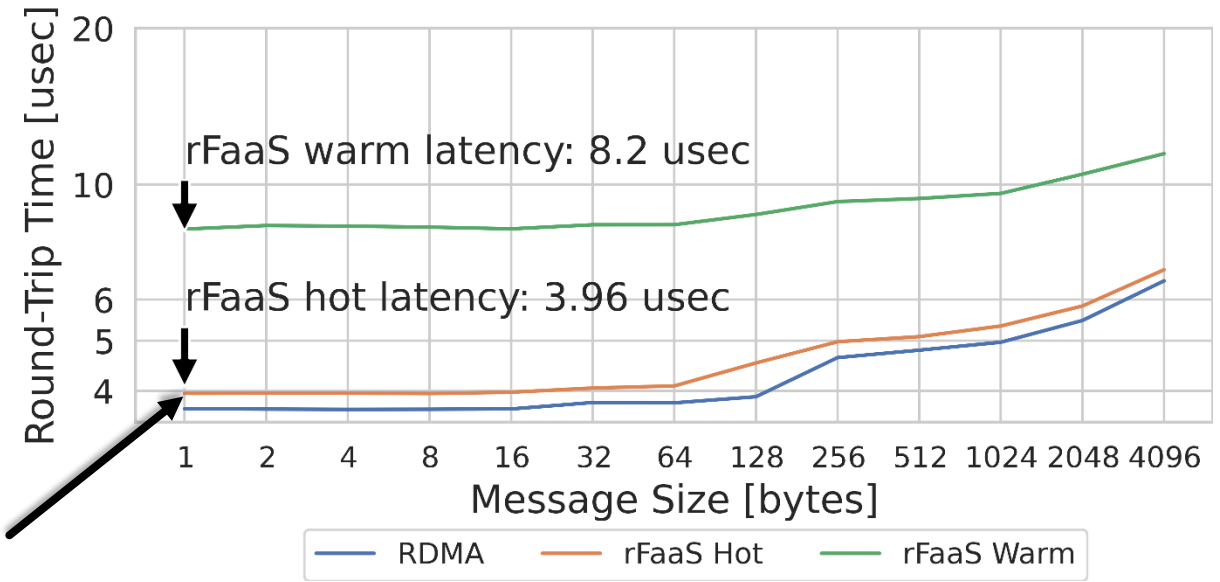

Less than **350 ns** overhead.

How fast are invocations in rFaaS?

36 CPU cores, 377 GB memory.
100 Gbps Ethernet with RoCEv2 support.

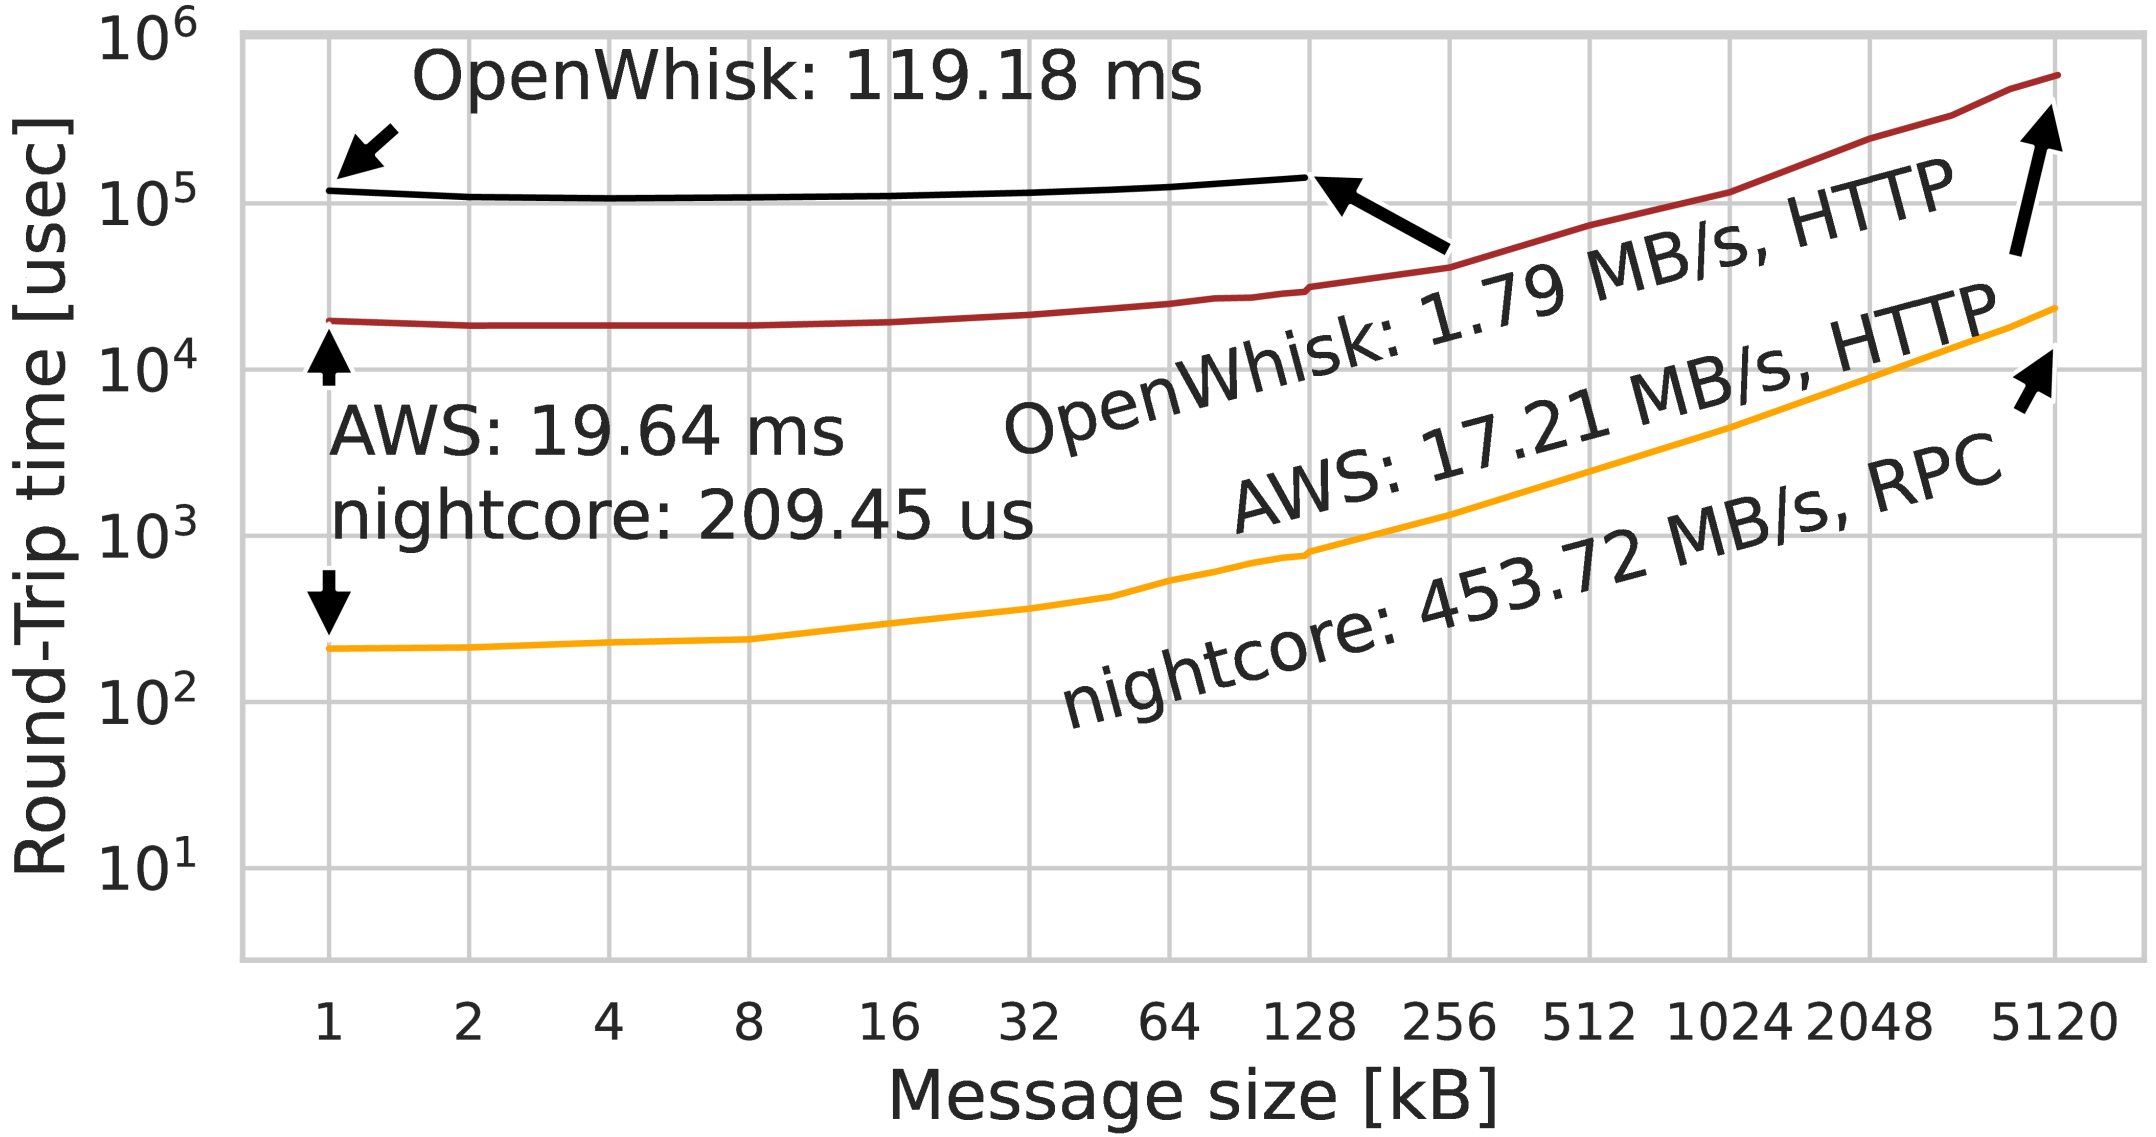

How fast are invocations in rFaaS?

36 CPU cores, 377 GB memory.
100 Gbps Ethernet with RoCEv2 support.

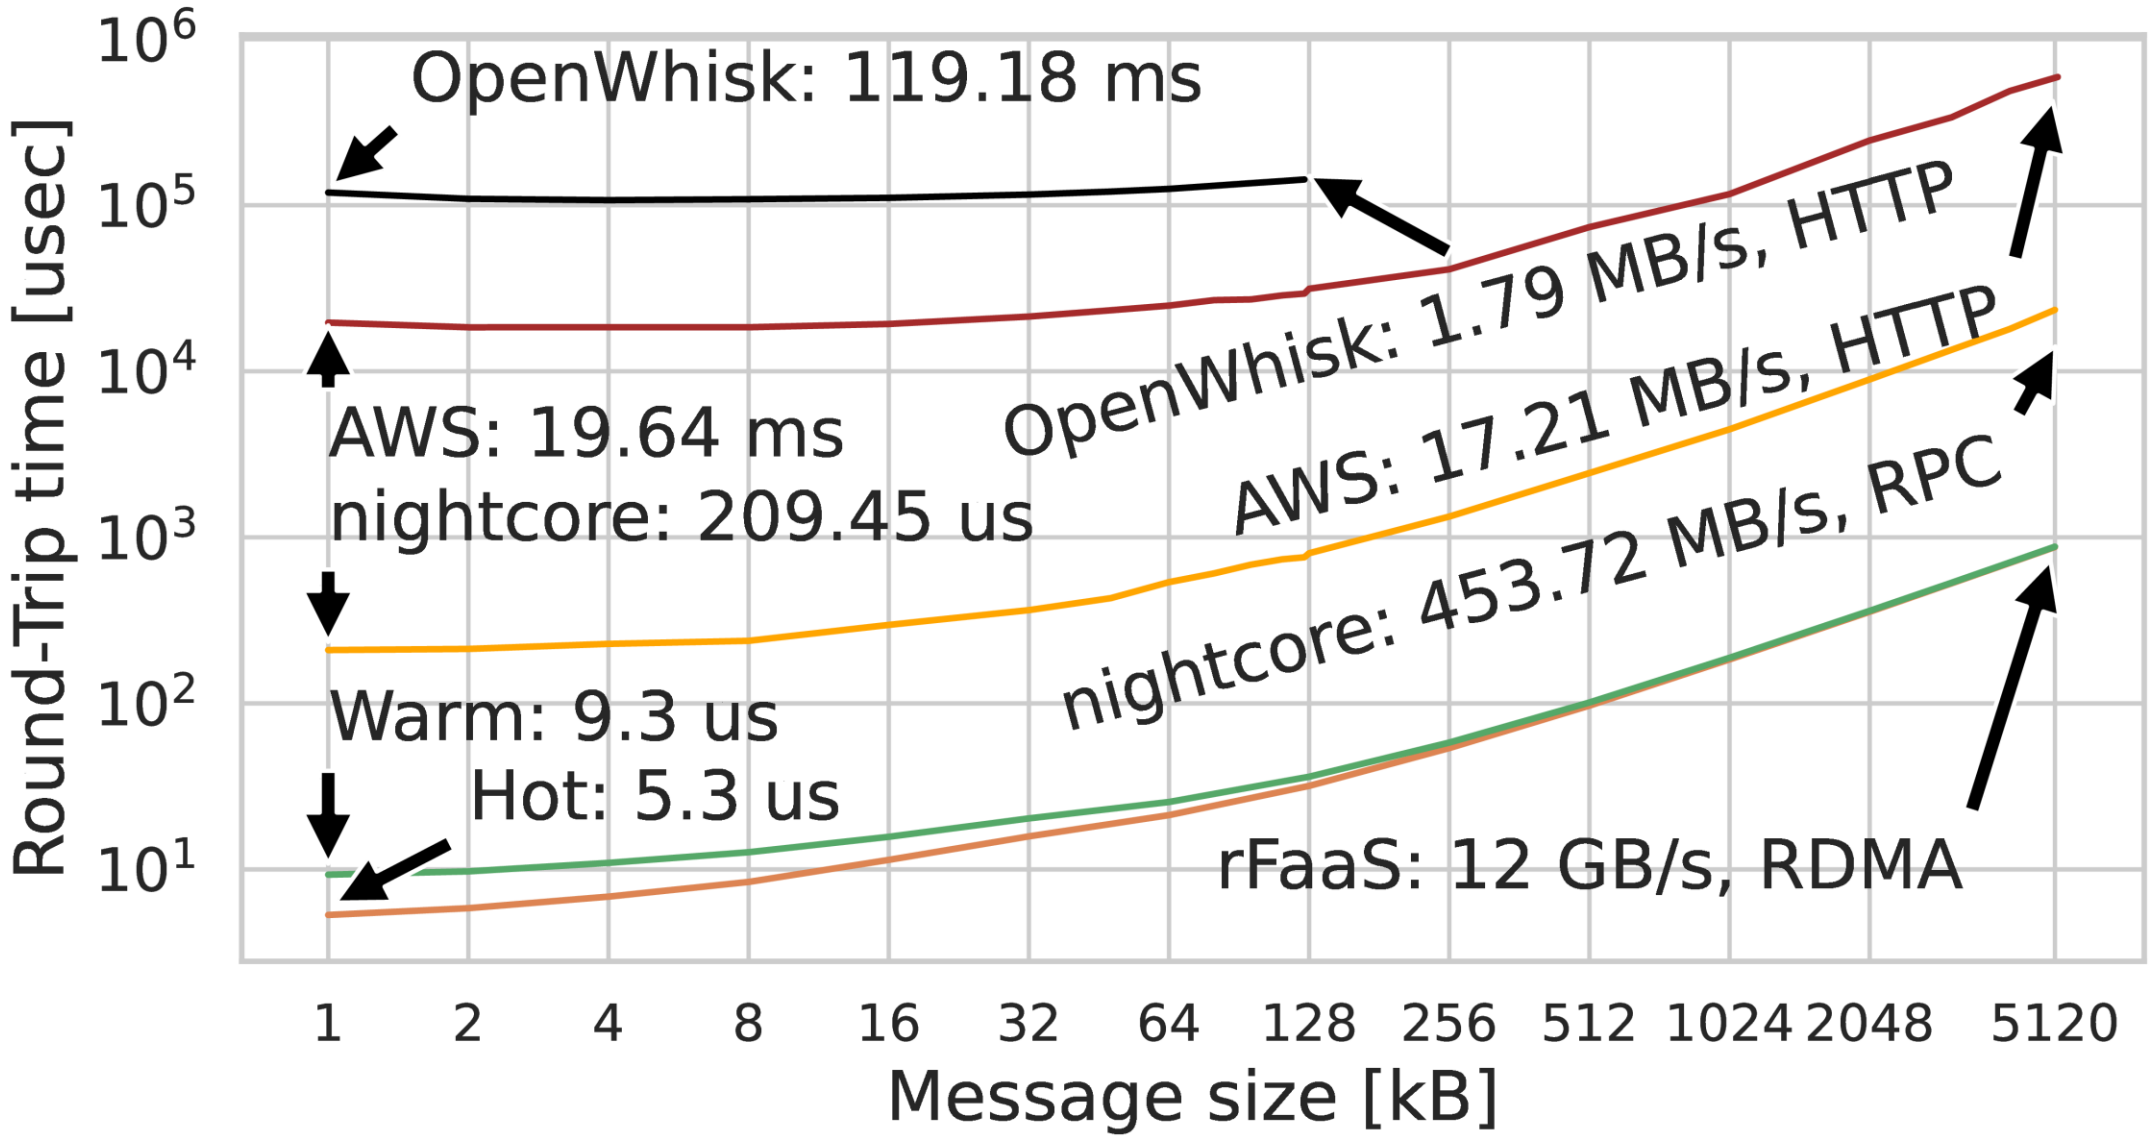

How fast are invocations in rFaaS?

36 CPU cores, 377 GB memory.
100 Gbps Ethernet with RoCEv2 support.

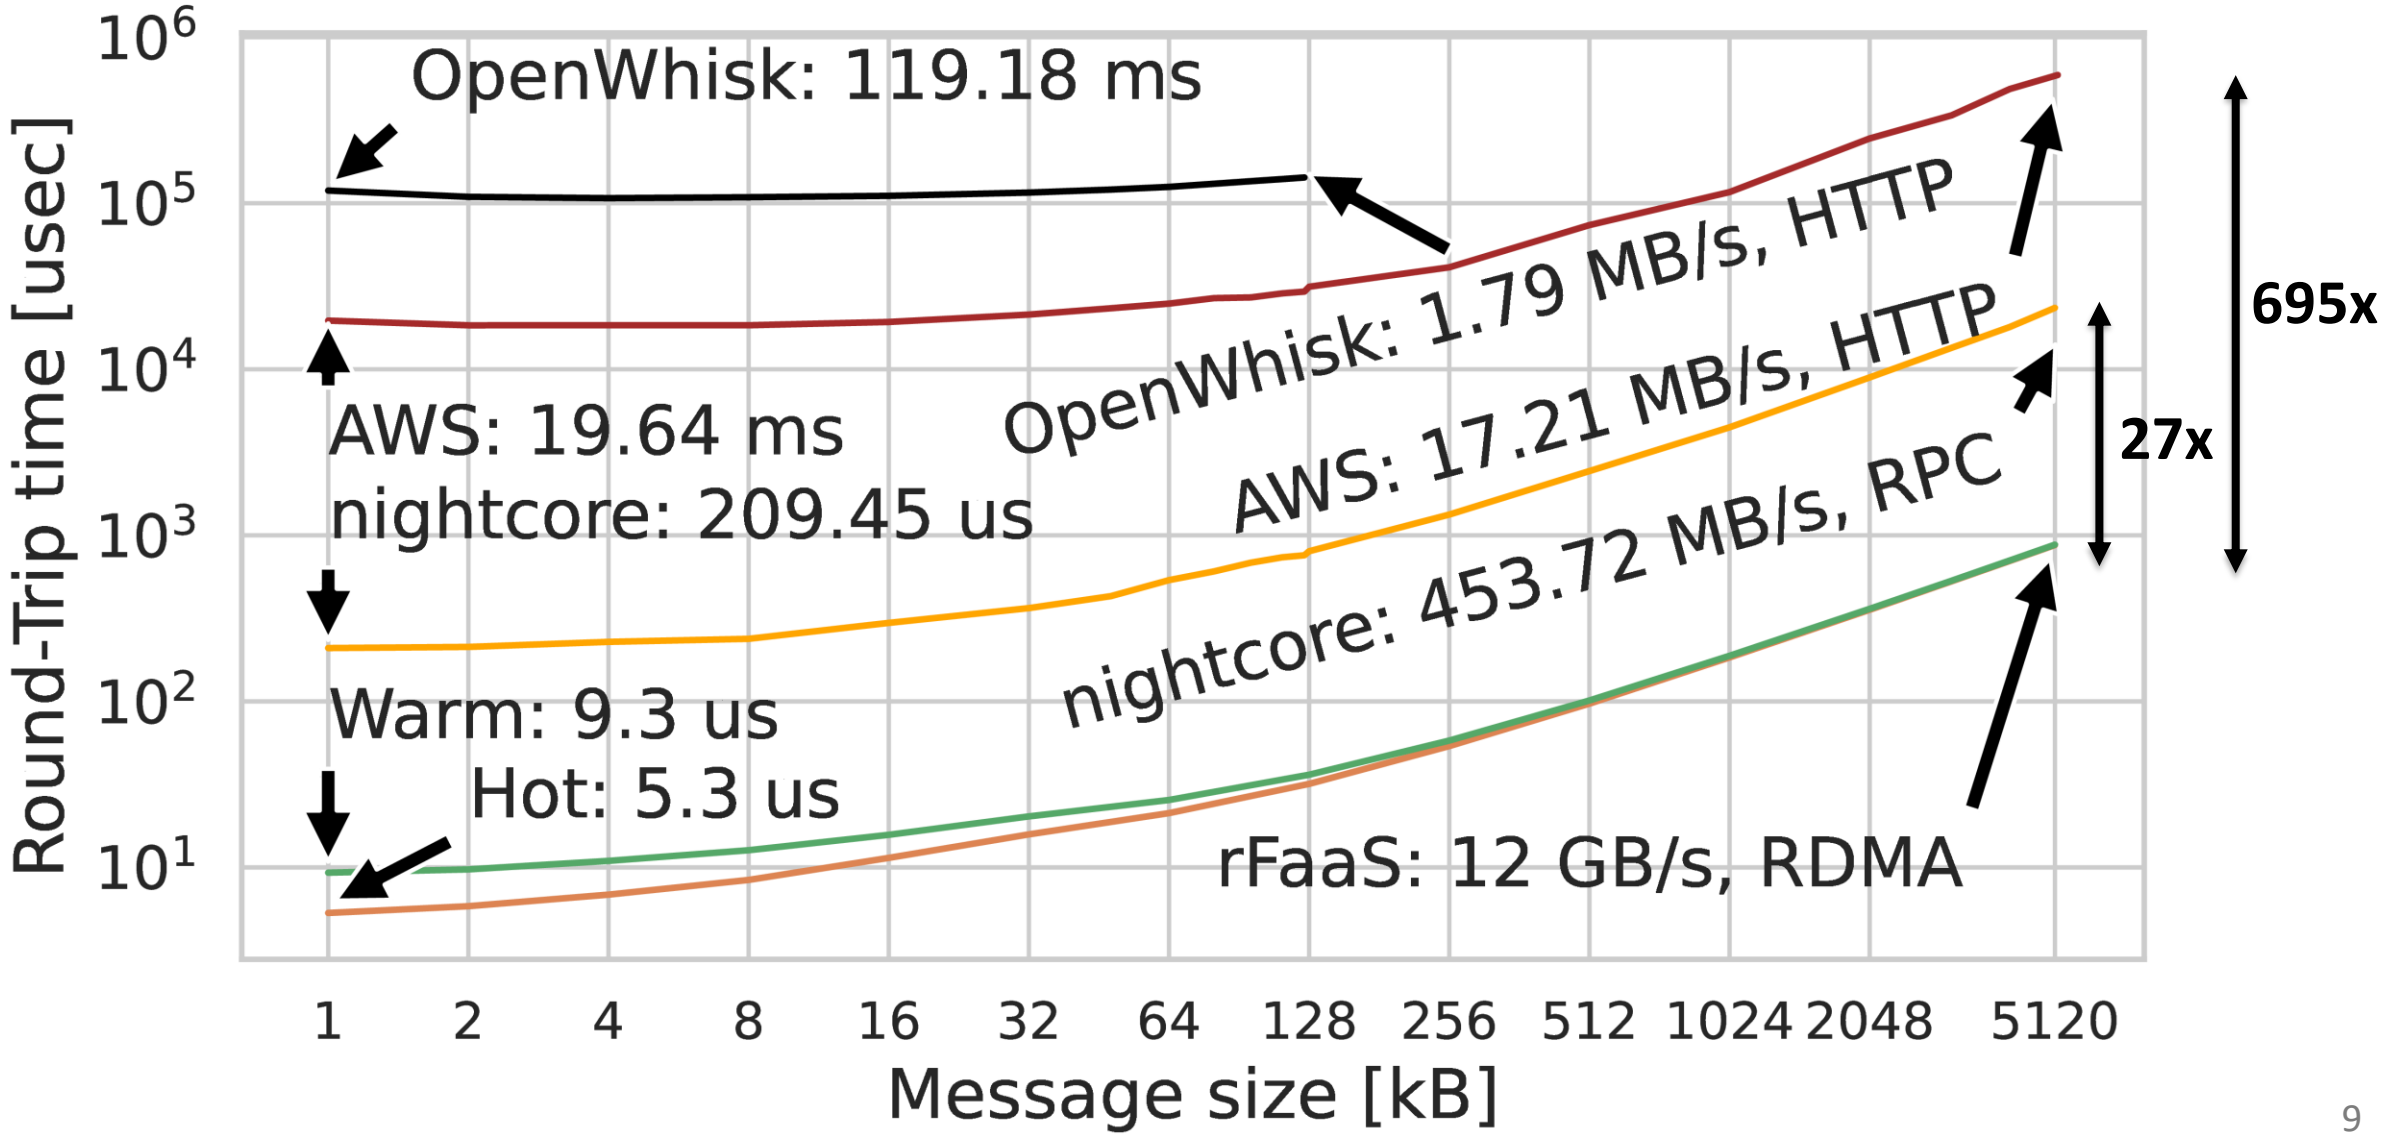

PARSEC: Black-Scholes

- 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.

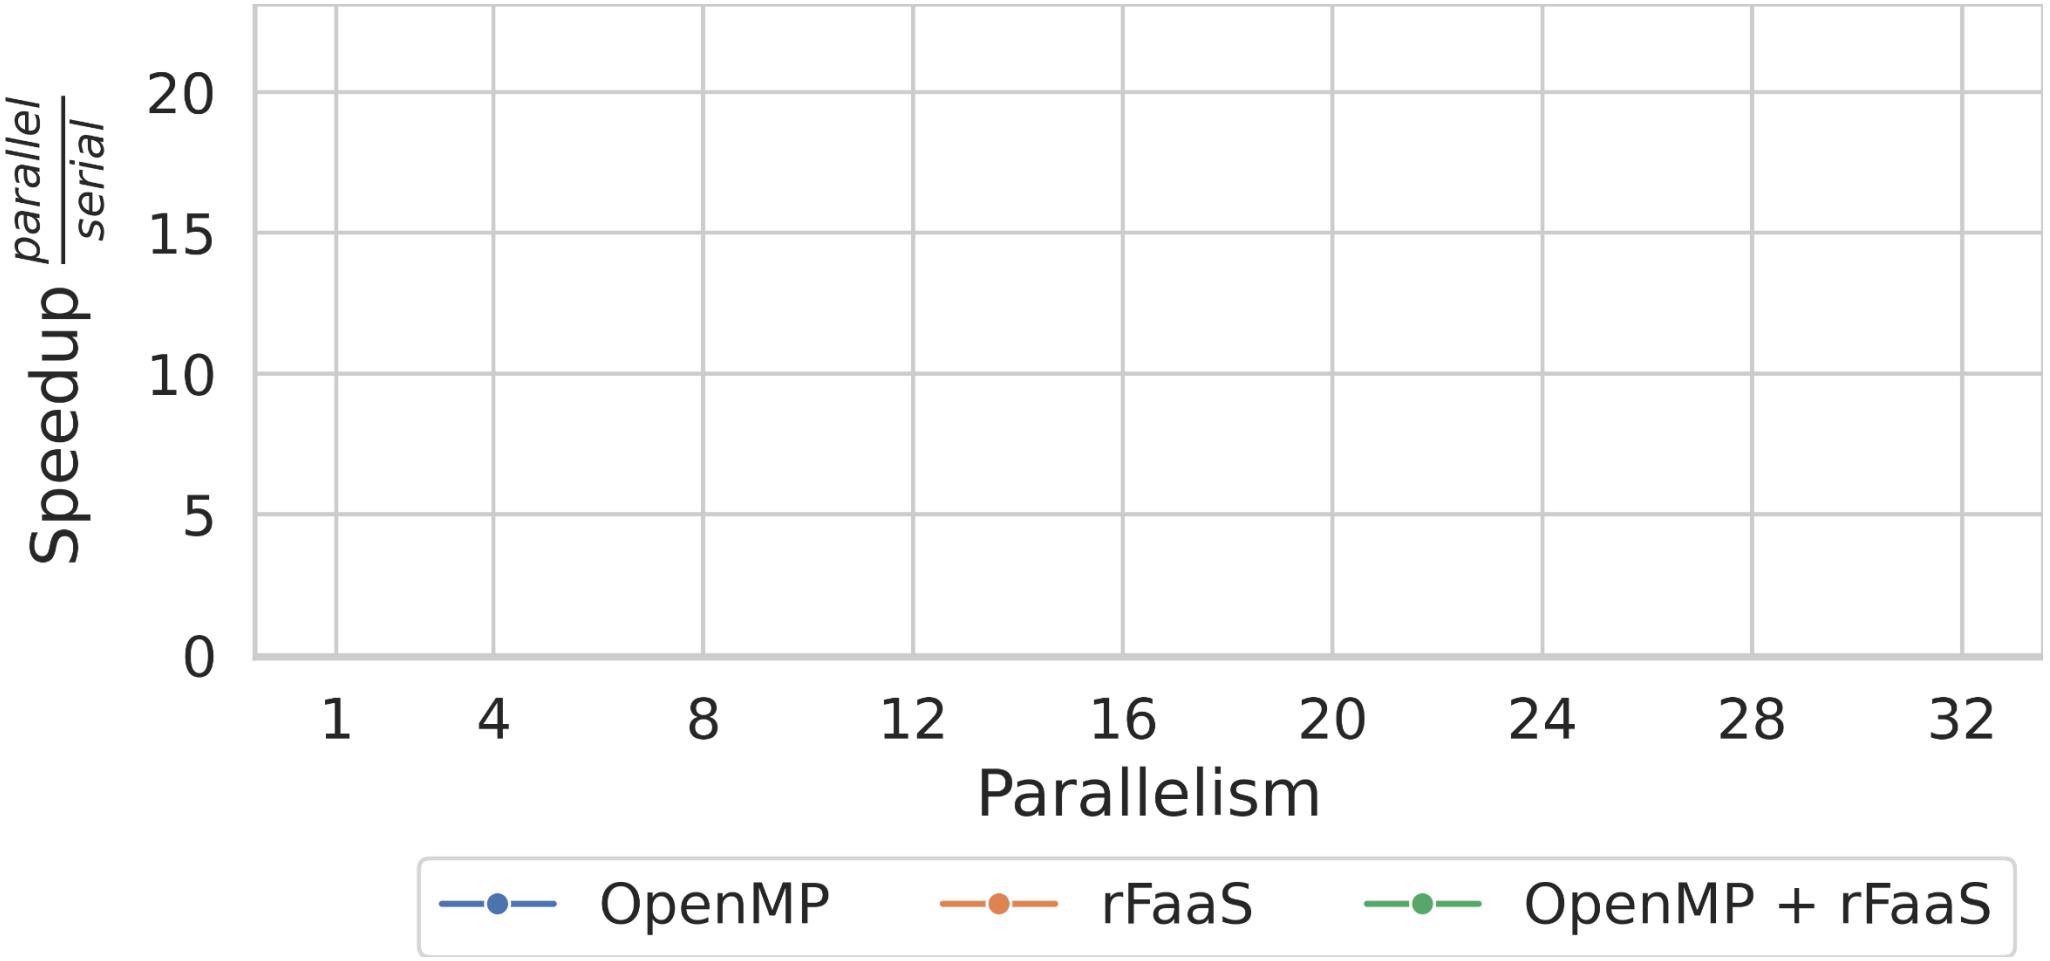

PARSEC: Black-Scholes

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

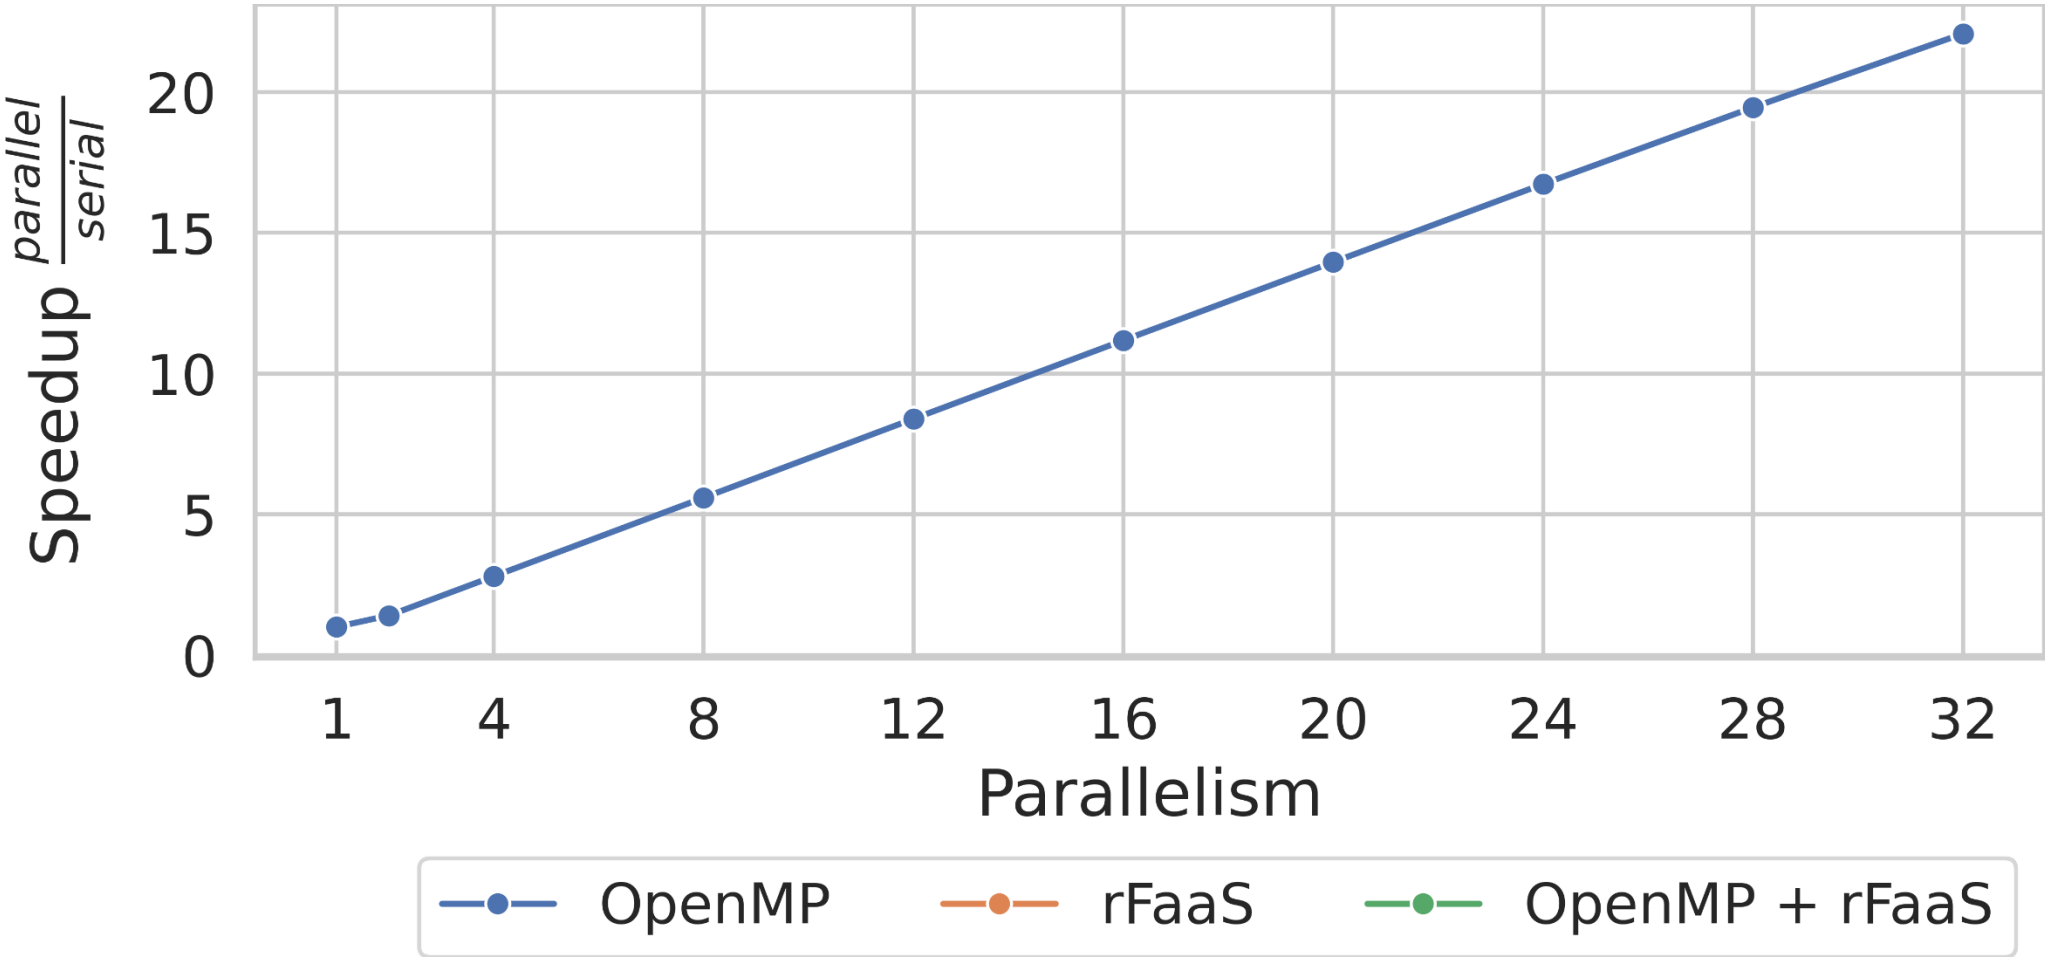

PARSEC: Black-Scholes

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

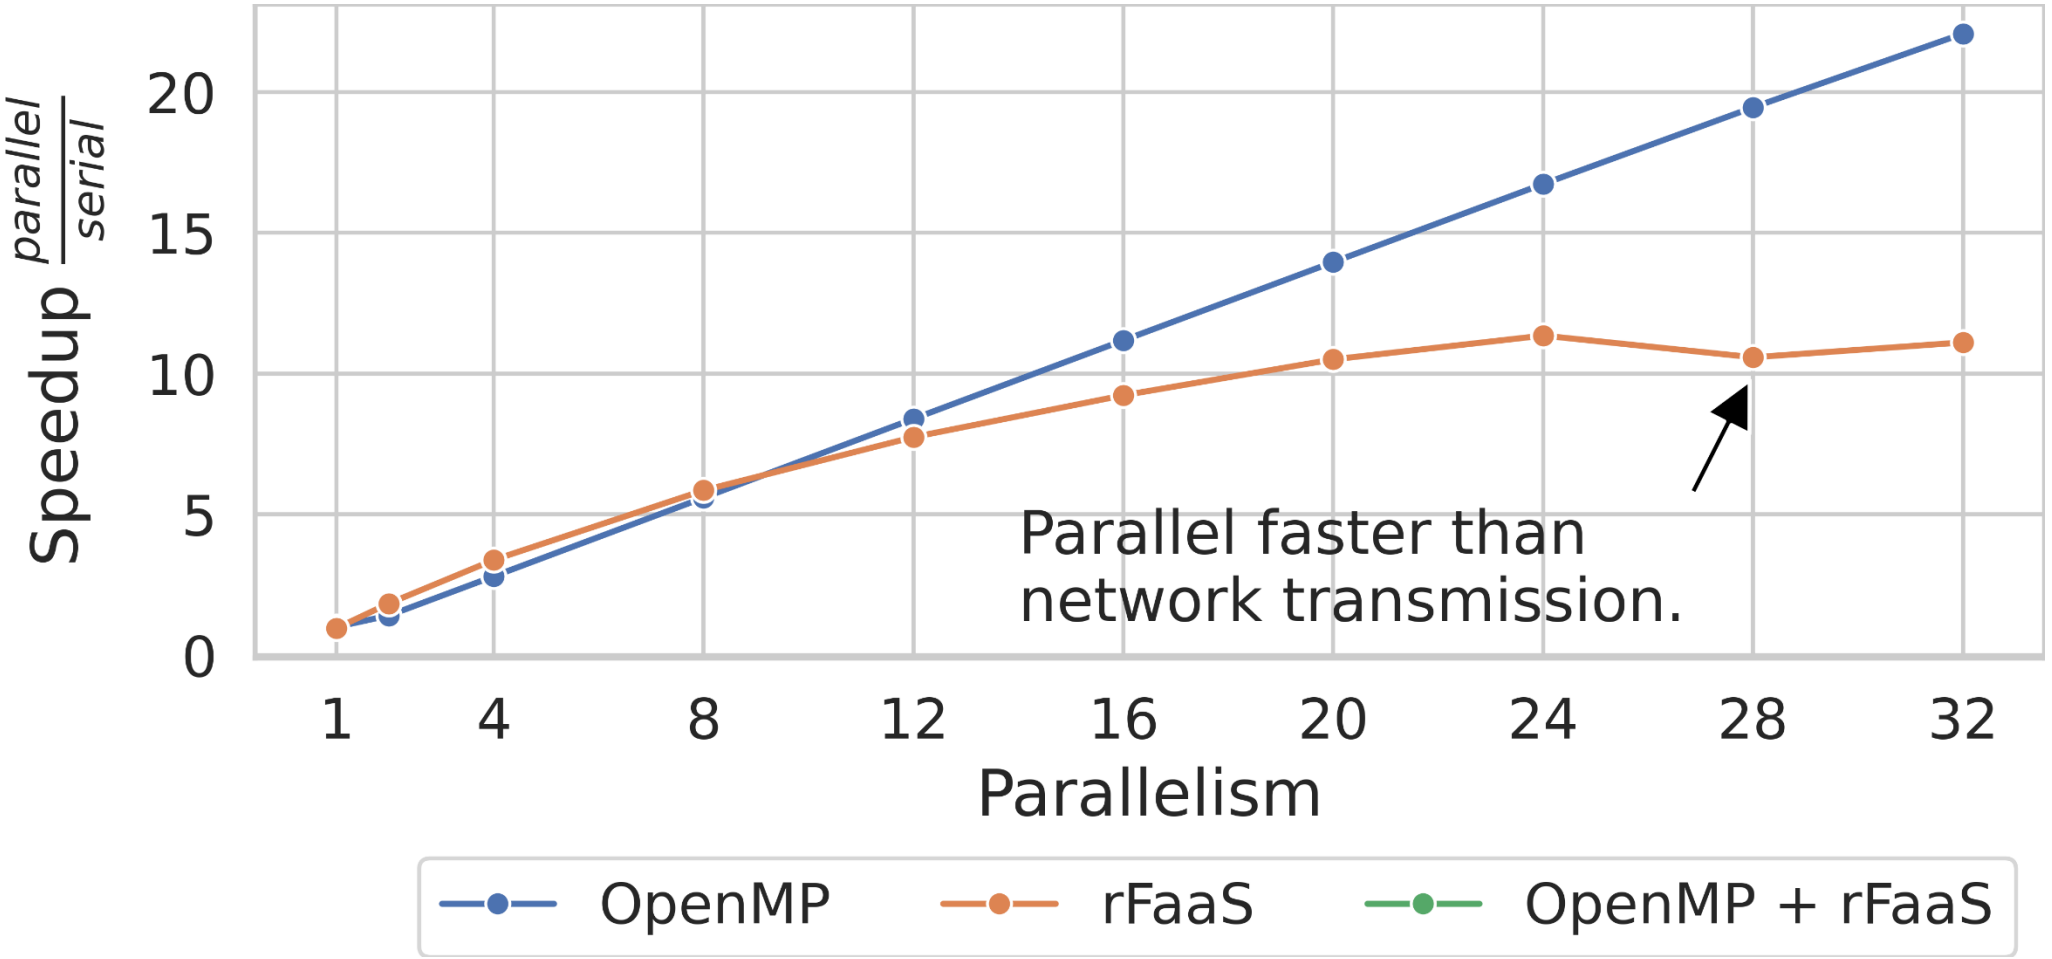

PARSEC: Black-Scholes

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

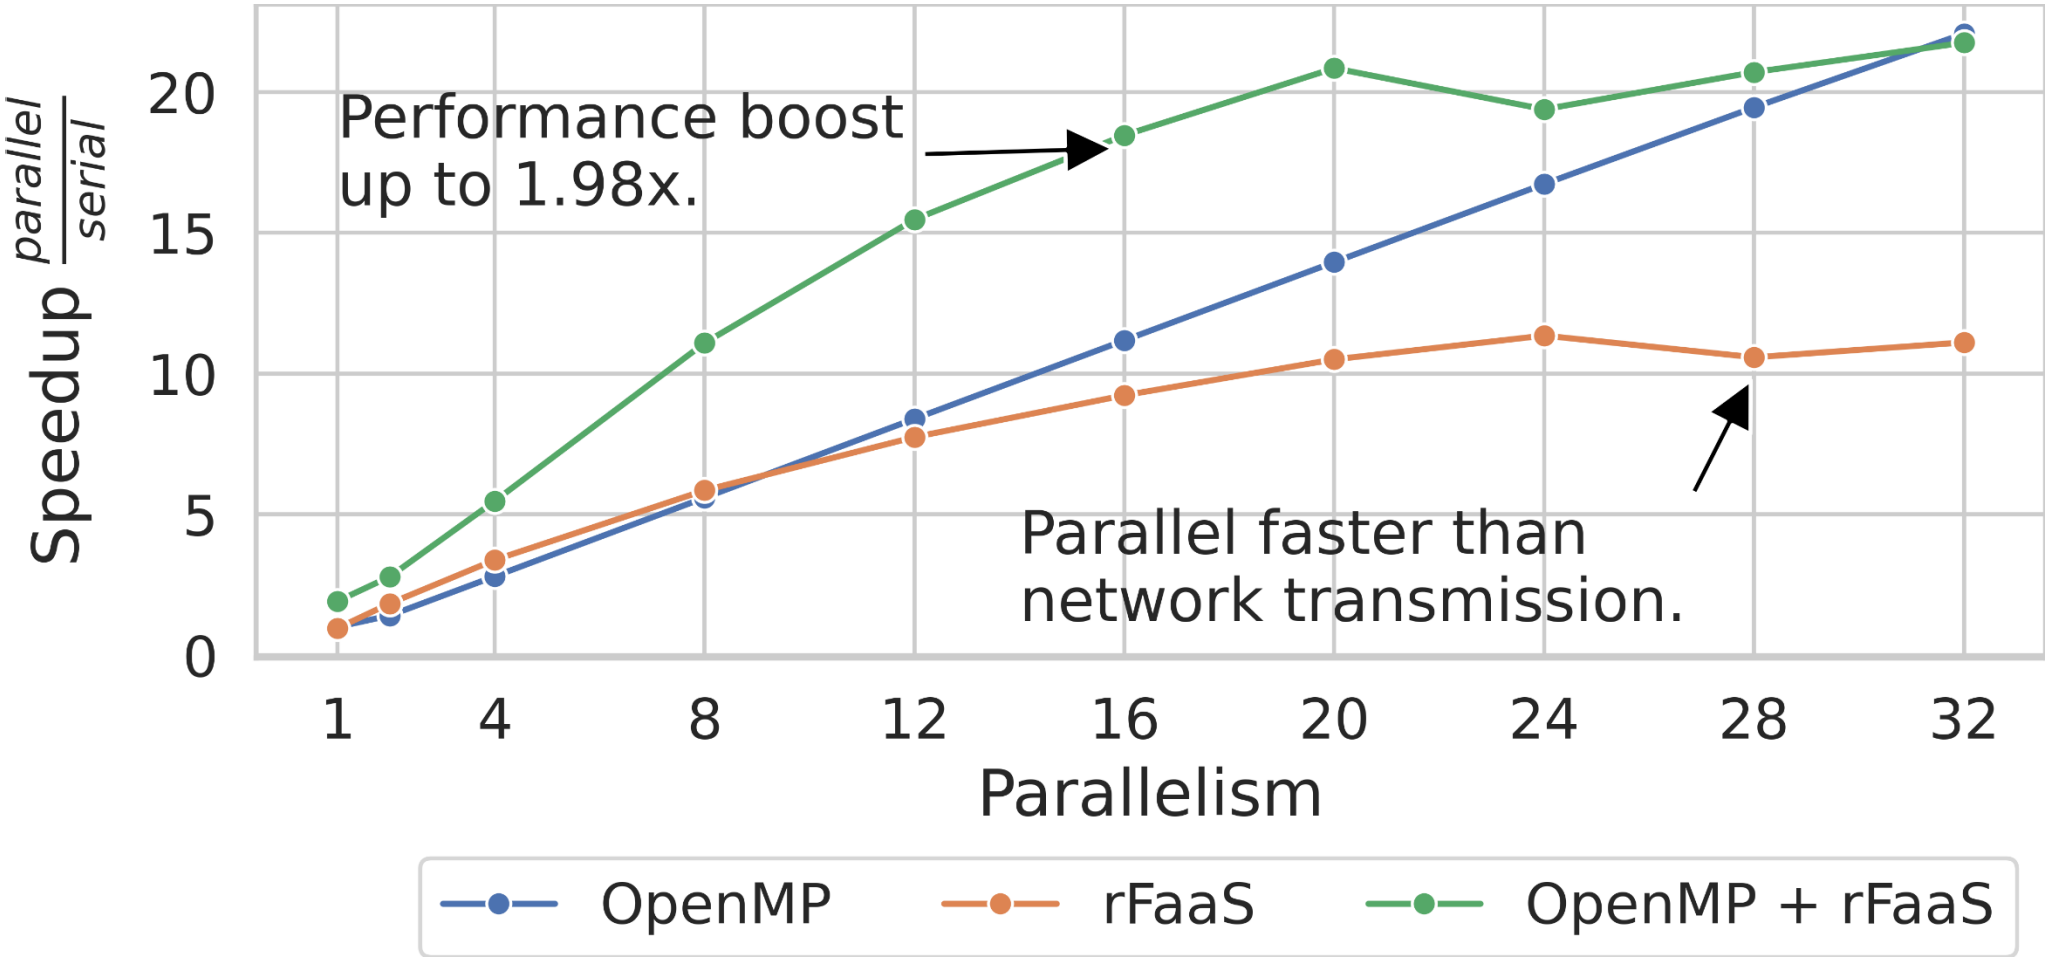

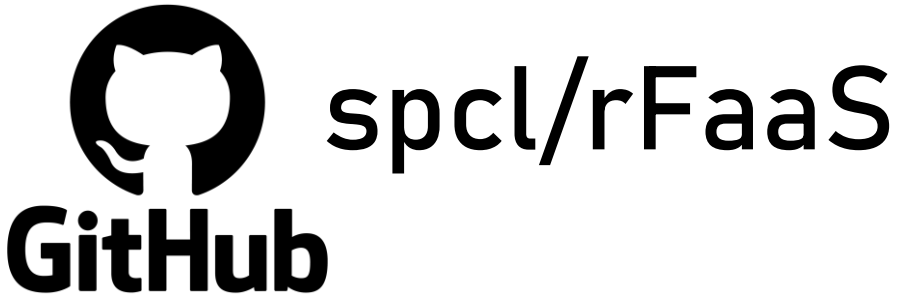

spcl/rFaaS

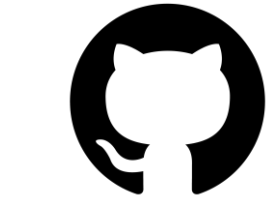

spcl/rFaaS

Containers

Docker

Sarus

Singularity

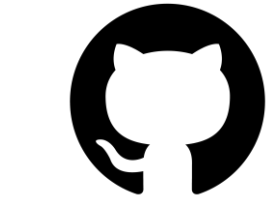

spcl/rFaaS

Containers

- Docker
- Sarus
- **Singularity**

Networks

- ibverbs
- Cray uGNI

GitHub spcl/rFaaS

Containers	Networks
Docker	ibverbs
Sarus	Cray uGNI
Singularity	aws **EFA, TCP**
	Google Summer of Code

GitHub spcl/rFaaS

Containers

- Docker
- Sarus
- **Singularity**

Networks

- ibverbs
- Cray uGNI
- **EFA, TCP**
- Google Summer of Code

Applications

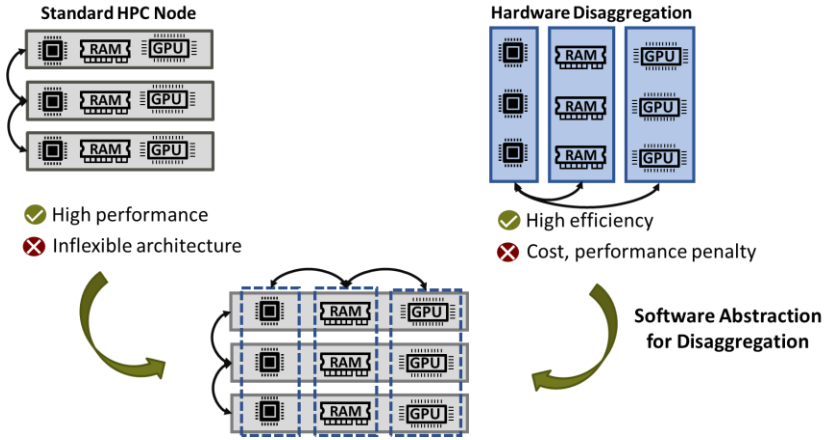

"Software Resource Disaggregation for HPC with Serverless Computing"

Conclusions

More of SPCL's research:

youtube.com/@spcl — 150+ Talks

twitter.com/spcl_eth — 1.2K+ Followers

github.com/spcl — 2K+ Stars

… or spcl.ethz.ch

Conclusions

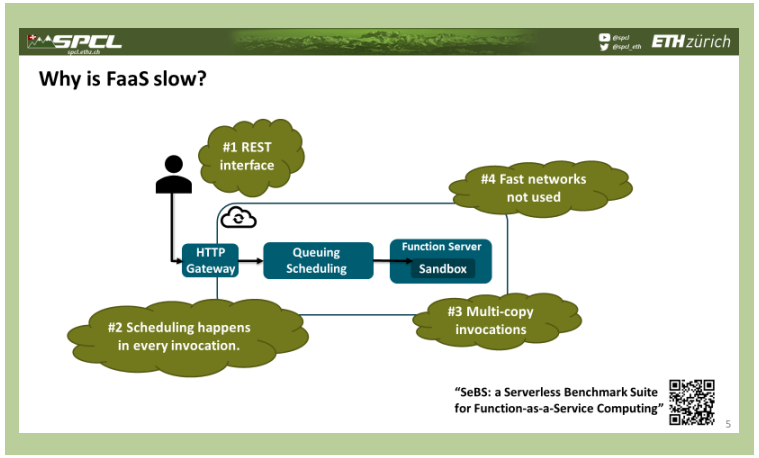

More of SPCL's research:

youtube.com/@spcl — 150+ Talks

twitter.com/spcl_eth — 1.2K+ Followers

github.com/spcl — 2K+ Stars

… or spcl.ethz.ch

Conclusions

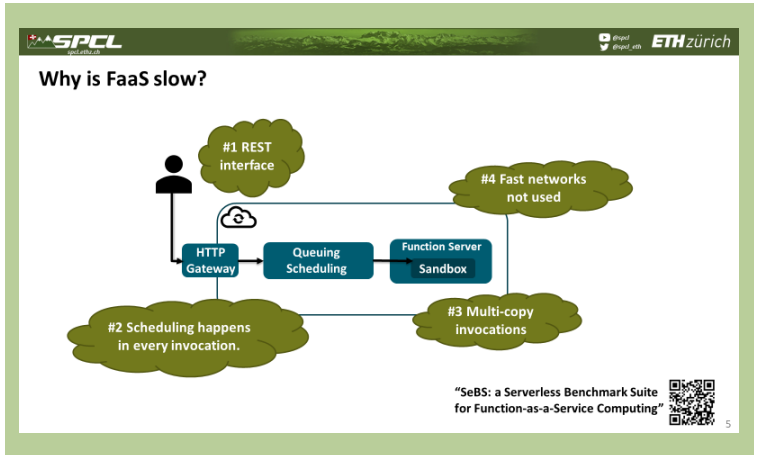

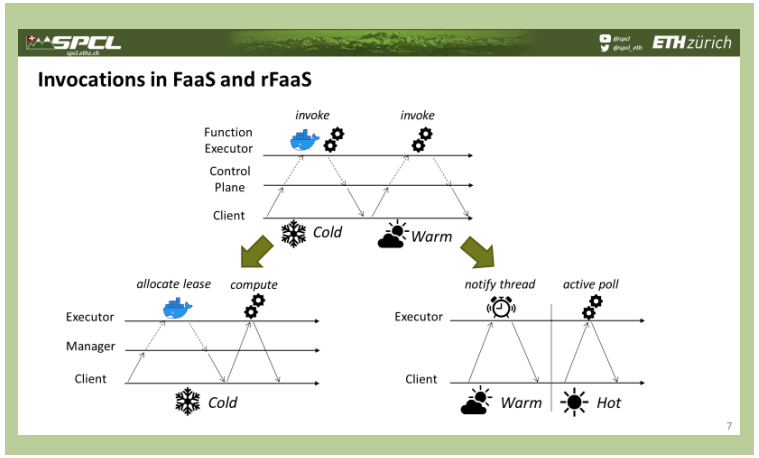

More of SPCL's research:

youtube.com/@spcl — 150+ Talks

twitter.com/spcl_eth — 1.2K+ Followers

github.com/spcl — 2K+ Stars

… or spcl.ethz.ch

Conclusions

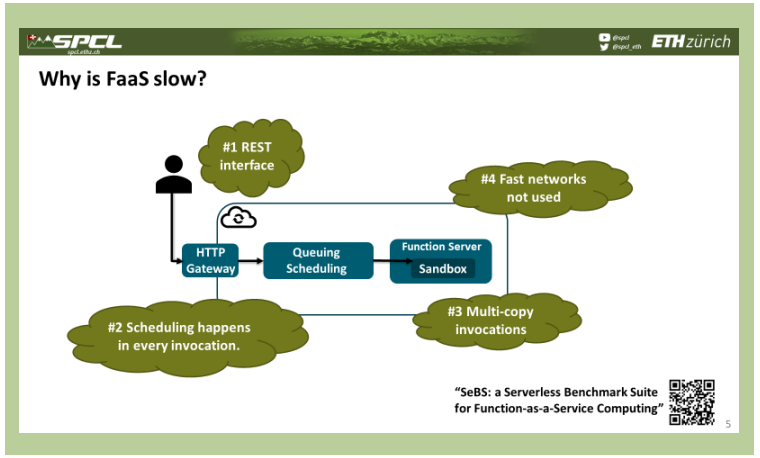

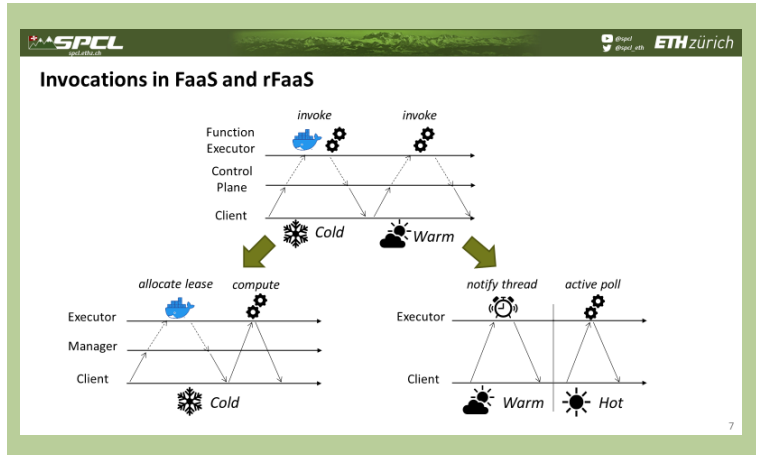

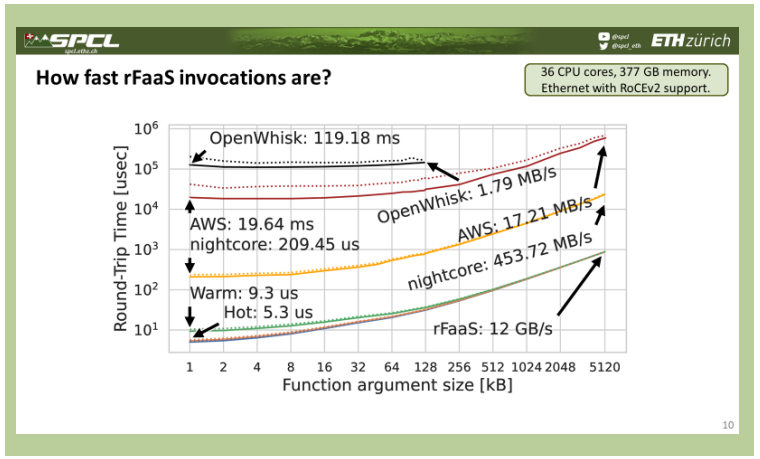

More of SPCL's research:

youtube.com/@spcl — 150+ Talks

twitter.com/spcl_eth — 1.2K+ Followers

github.com/spcl — 2K+ Stars

… or spcl.ethz.ch

Conclusions

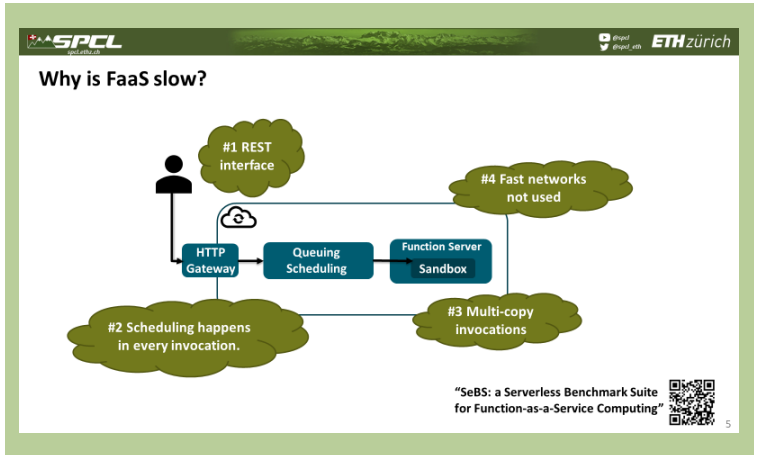

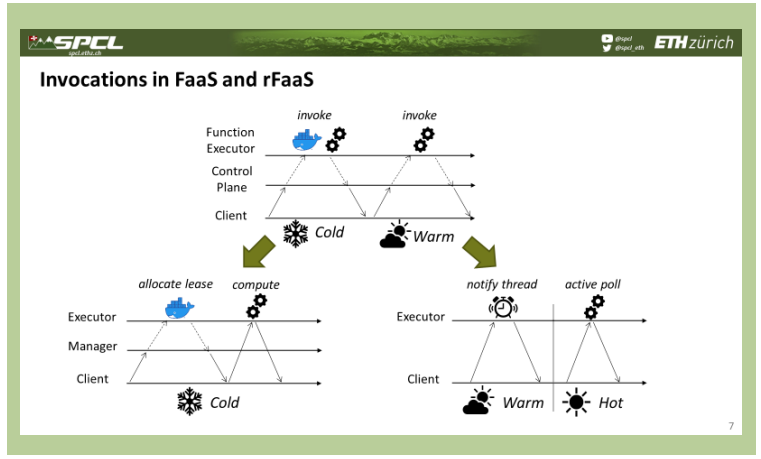

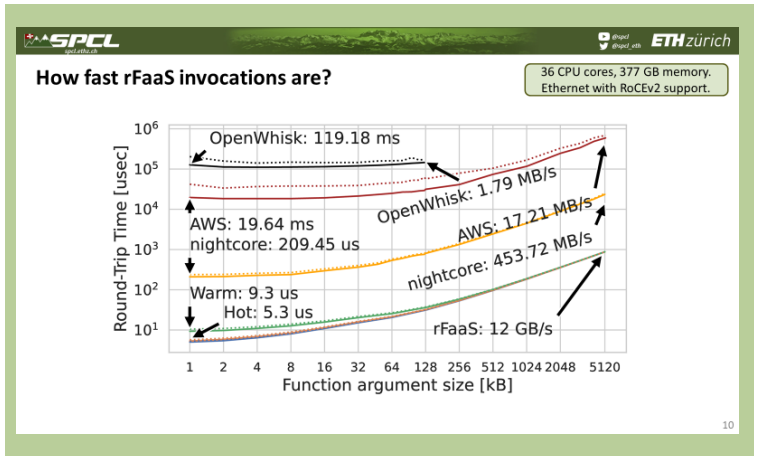

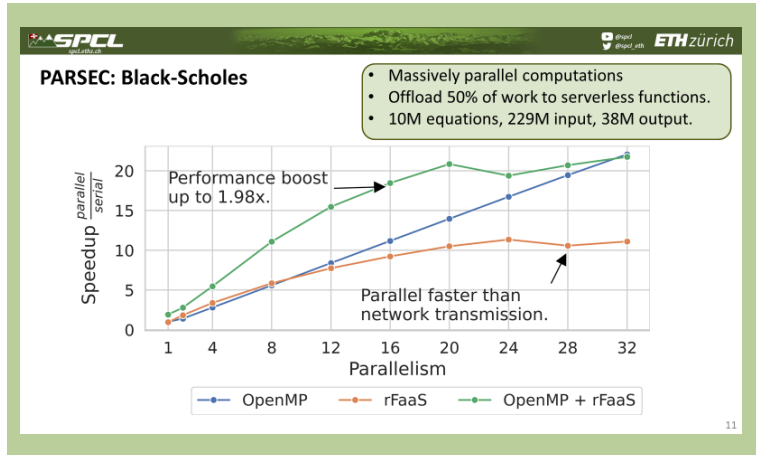

More of SPCL's research:

youtube.com/@spcl — 150+ Talks

twitter.com/spcl_eth — 1.2K+ Followers

github.com/spcl — 2K+ Stars

… or spcl.ethz.ch

Conclusions

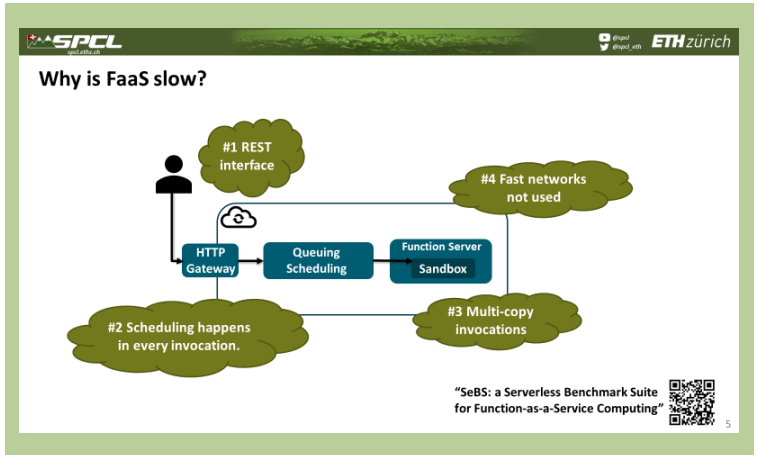

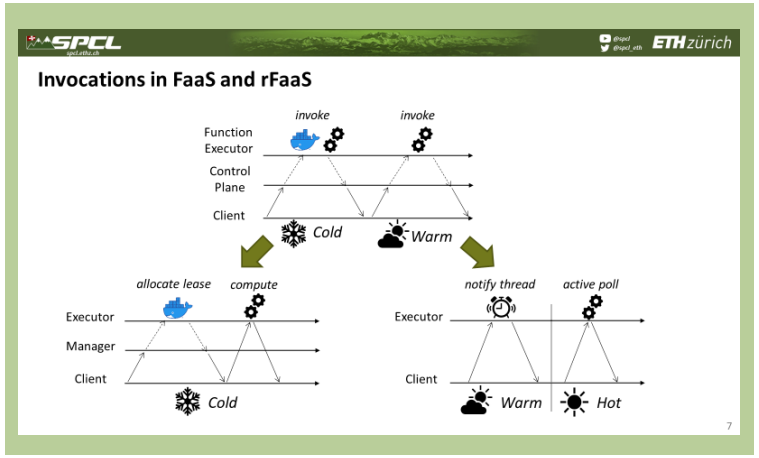

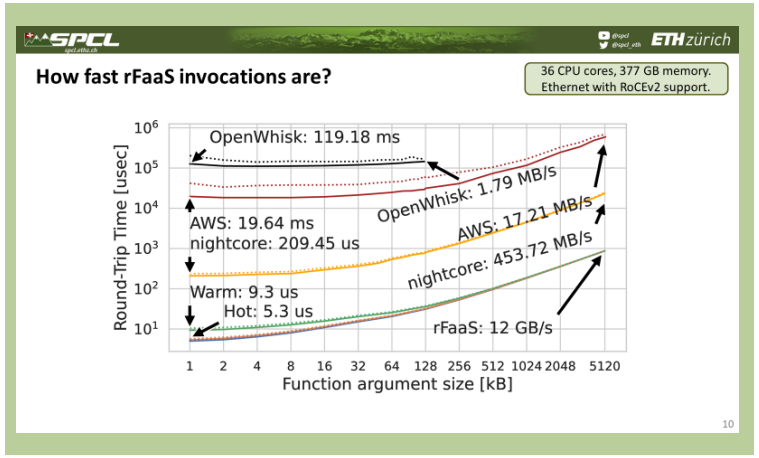

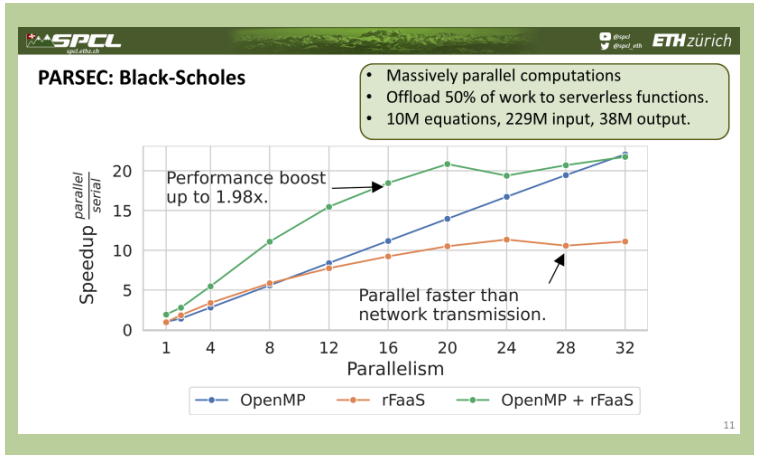

More of SPCL's research:

youtube.com/@spcl — 150+ Talks

twitter.com/spcl_eth — 1.2K+ Followers

github.com/spcl — 2K+ Stars

… or spcl.ethz.ch

Paper.

Paper artifact.