

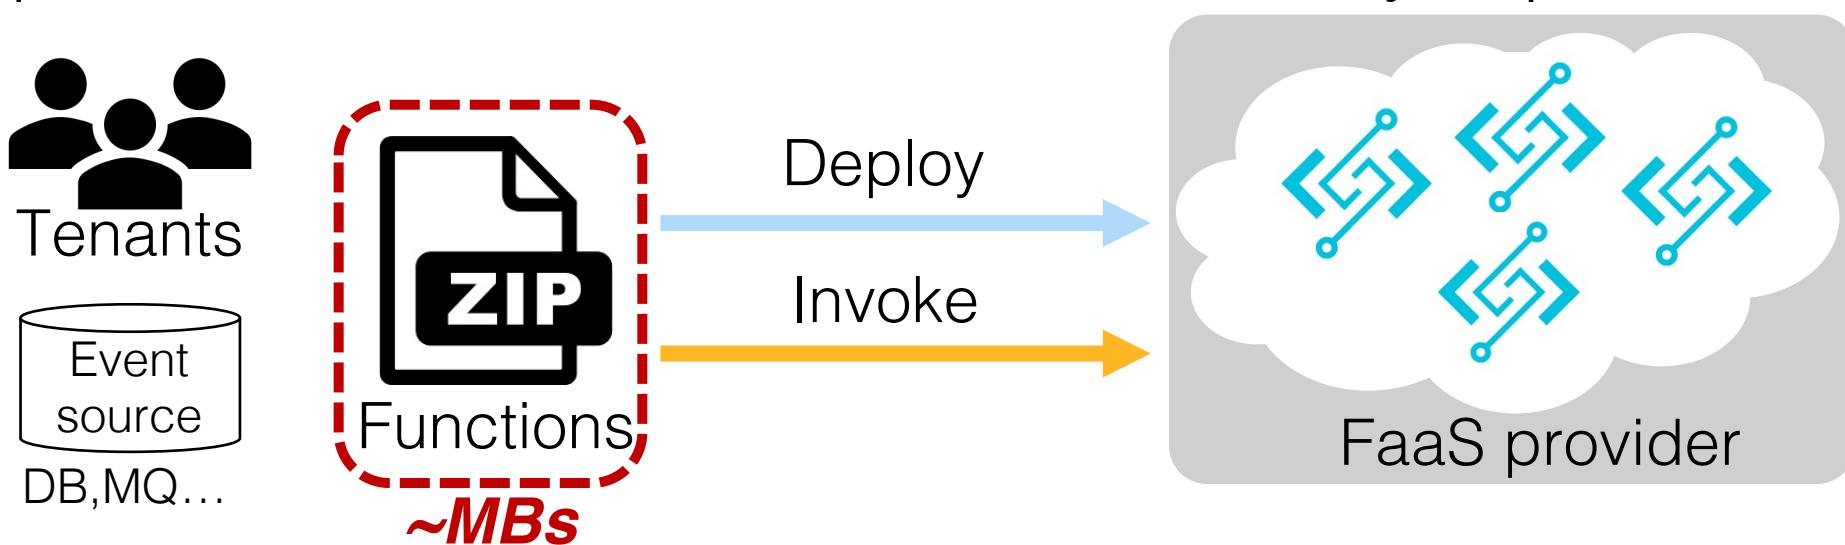
FaaSNet: Scalable and Fast Provisioning of Custom Serverless Container Runtimes at Alibaba Cloud Function Compute

Ao Wang, Shuai Chang, Huangshi Tian,
Hongqi Wang, Haoran Yang, Huiba Li, Rui Du, Yue Cheng



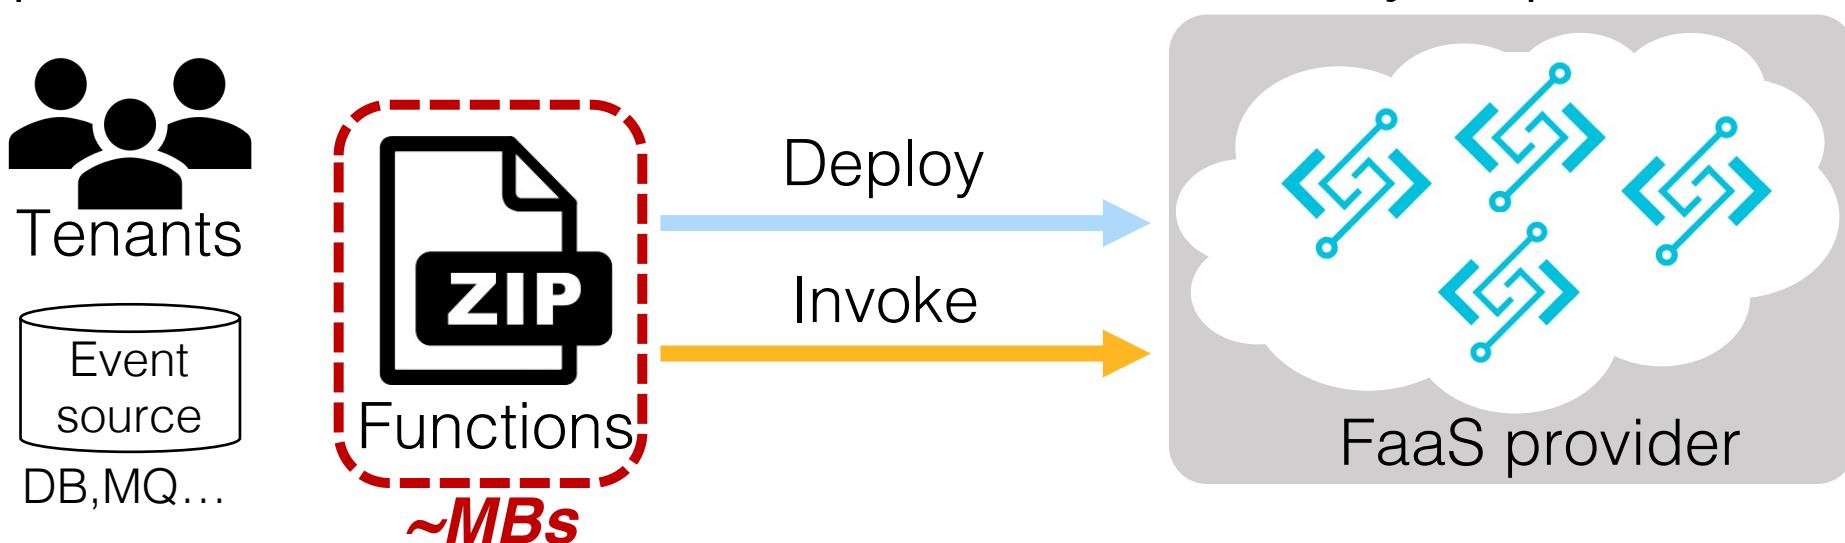
Function-as-a-Service

- FaaS enables cloud tenants to launch short-lived tasks (i.e., Lambda functions) with **high elasticity** and **fine-grained resource billing (1ms)**
- Function: basic unit of deployment. Application consists of multiple serverless functions
- Popular use cases: Backend APIs, event/async processing...

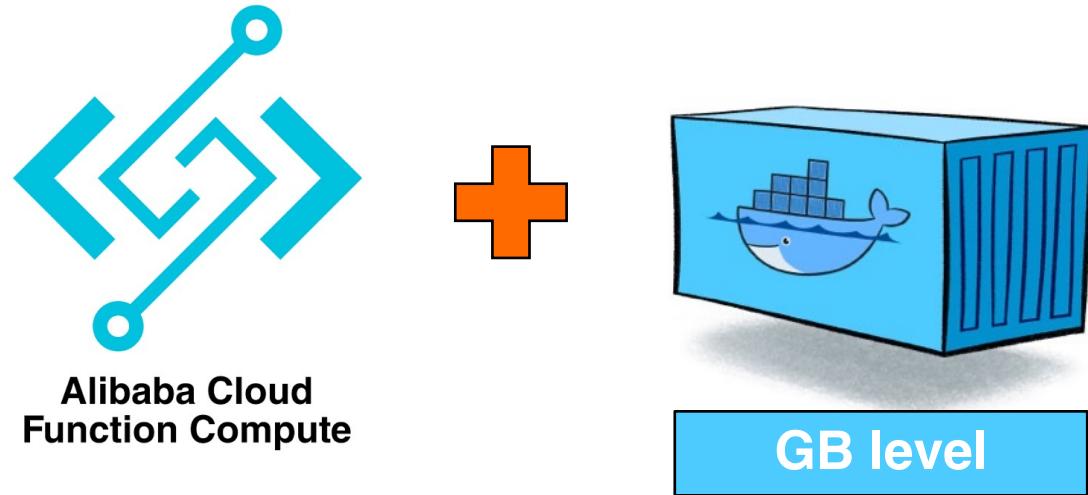


➤ **FaaS providers** normally limit tenants
code package in tens of **MB level**

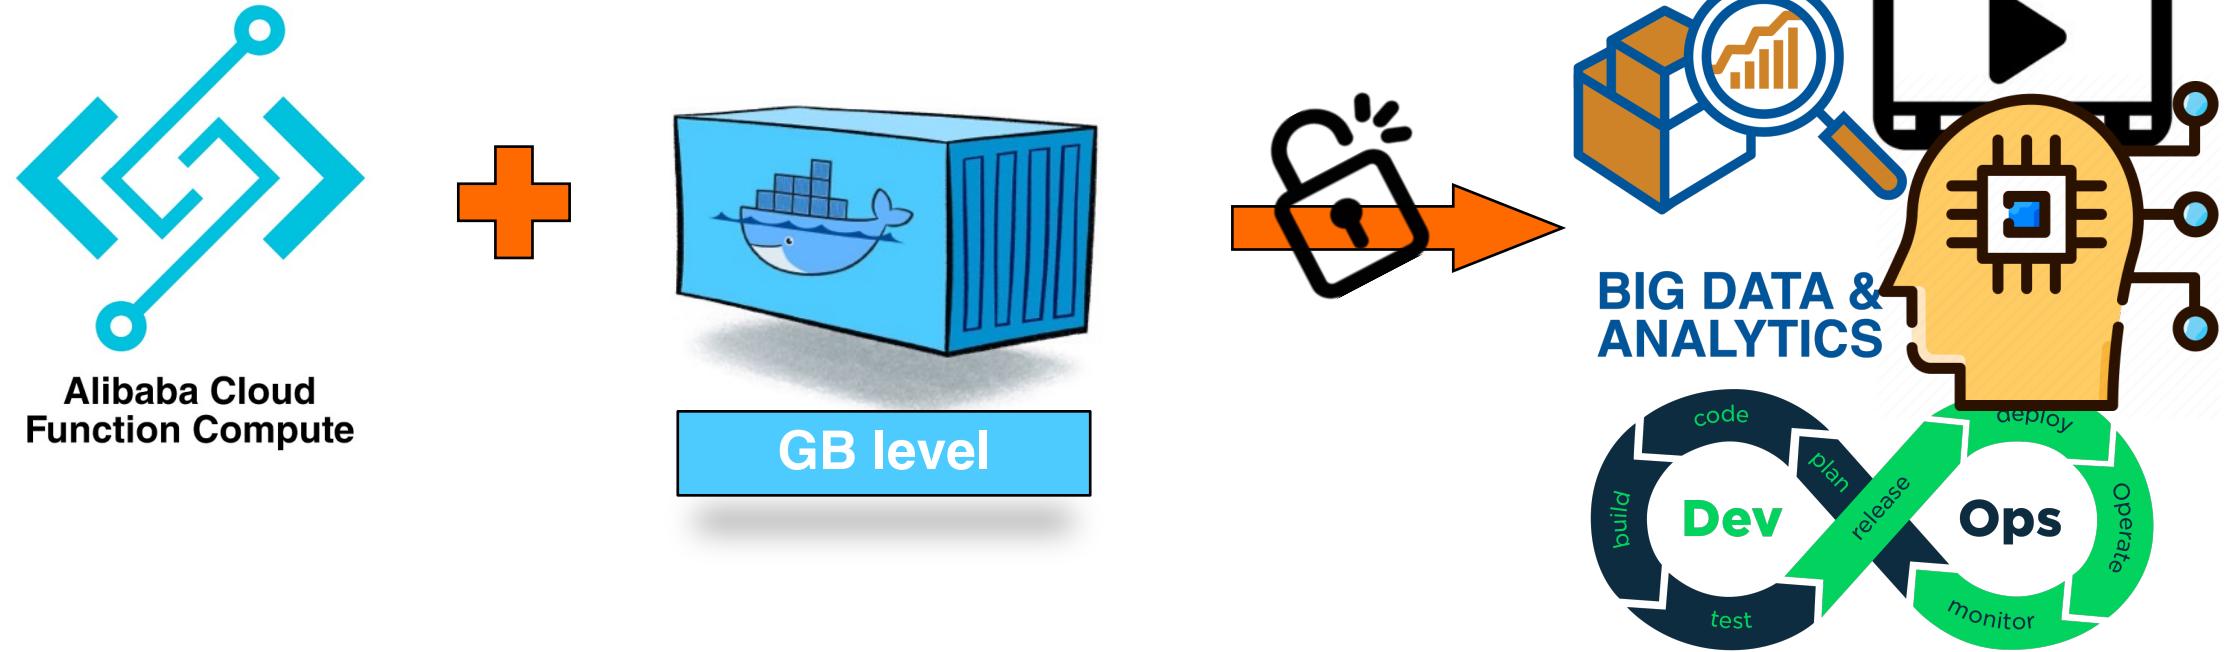
- FaaS provider enables cloud tenants to launch short-lived tasks (i.e., Lambda functions) with **high elasticity** and **fine-grained** resource billing (1ms)
- Function: basic unit of deployment. Application consists of multiple serverless functions
- Popular use cases: Backend APIs, event/async processing...



FaaS and custom-container runtimes



FaaS and custom-container runtimes



- ***FaaS + container*** unlocks new workload possibilities and makes ***serverless*** accessible to ***a broader audience***

FaaS and custom-container runtimes

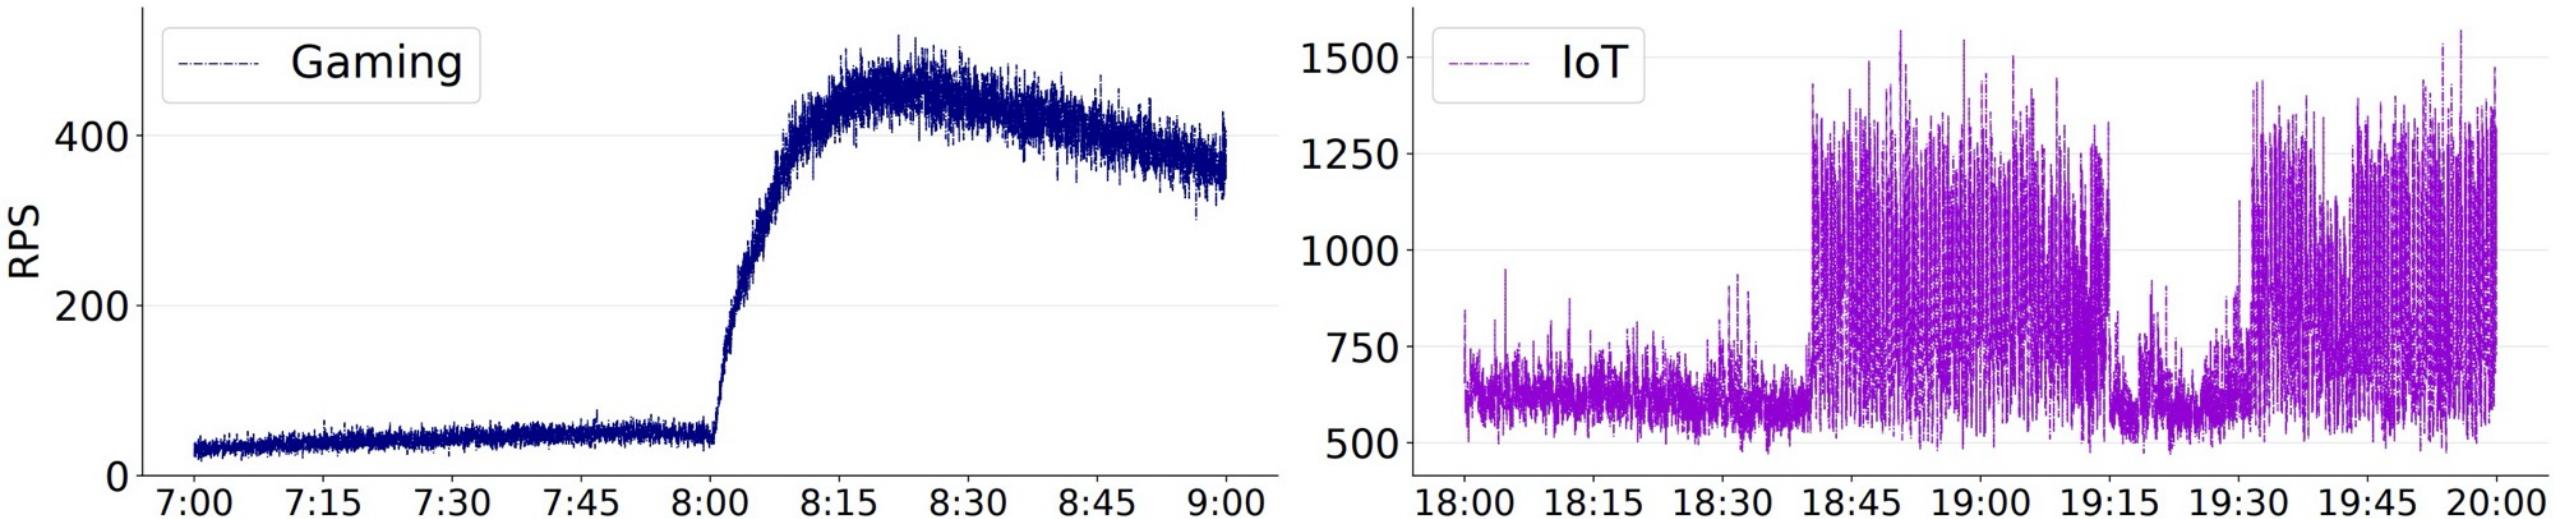
	Container	FaaS
Application transplant	Builds once, runs anywhere	Cloud vendor lock-in
Dev tools	open source ecosystems	Cloud vendor lock-in
CI/CD	open source ecosystems	Cloud vendor lock-in
Scalability	Second level	Millisecond level
Runtimes	Custom	Provided runtimes

- **FaaS + container** unlocks new workload possibilities and makes **serverless** accessible to **a broader audience**

Workload analysis

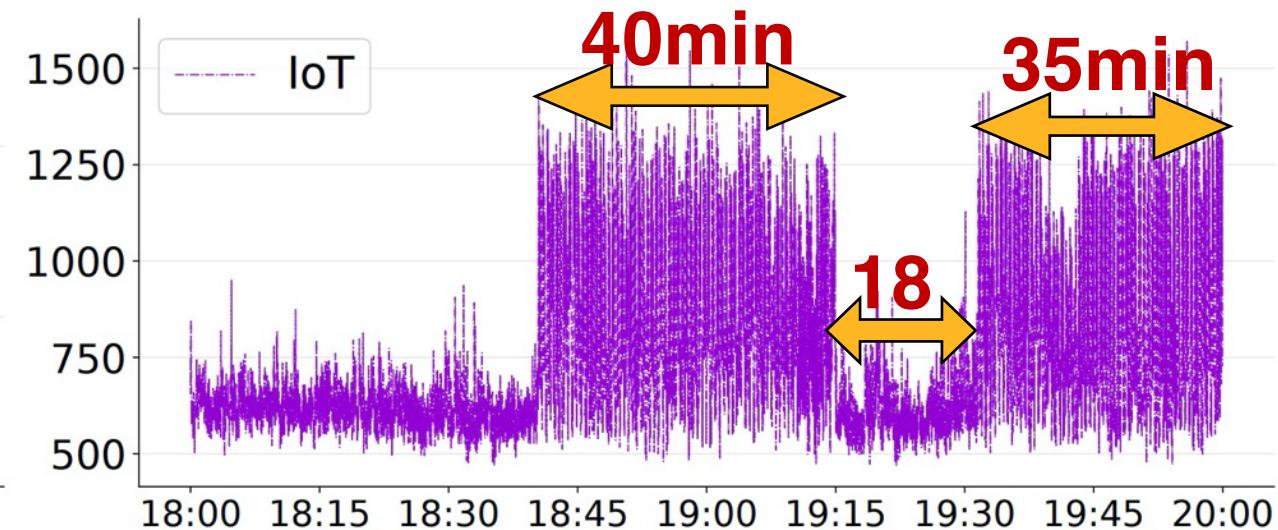
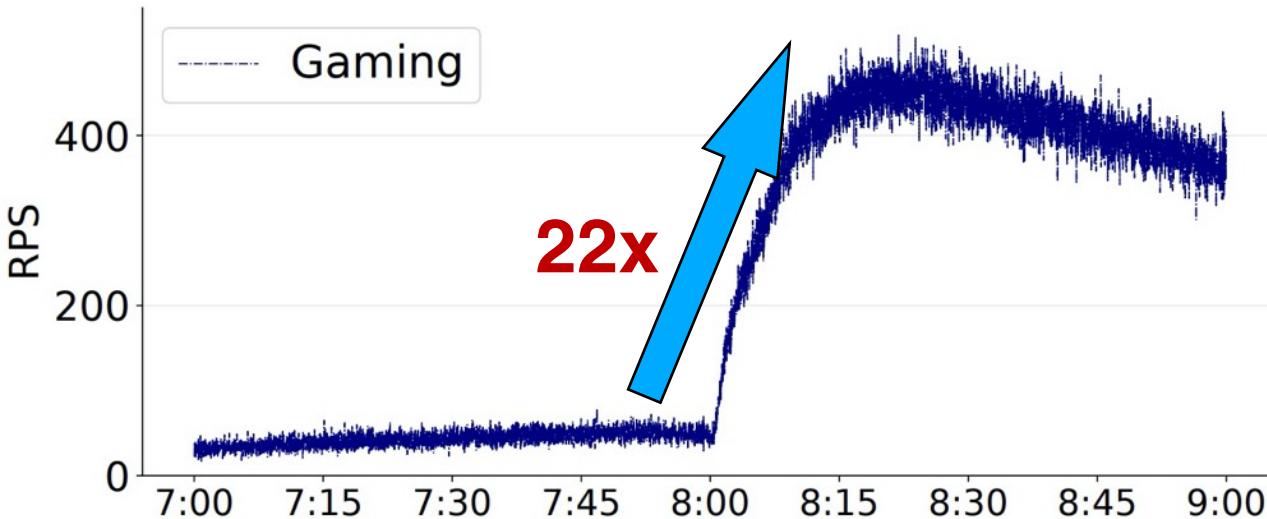
- Alibaba Cloud Function Compute 15-day-production log during May 2021
- Data centers: Beijing, Shanghai

Workload analysis



FaaS workloads are *bursty* and *dynamic*

Workload analysis



FaaS workloads are *bursty* and *dynamic*

Workload analysis

- Image pull latency distribution
- Proportion of image pull in function cold start

Workload analysis

- Image pull latency distribution
- Proportion of image pull in function cold start

Pull image in tens second level

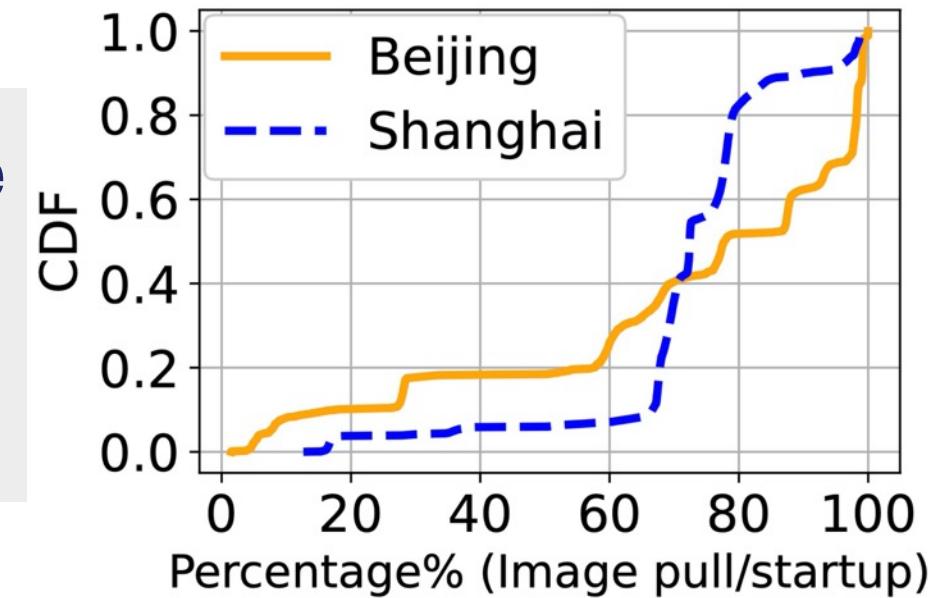
- 57% image pulls larger than 45 seconds

Workload analysis

- Image pull latency distribution
- Proportion of image pull in function cold start

Image pull dominates function startup time

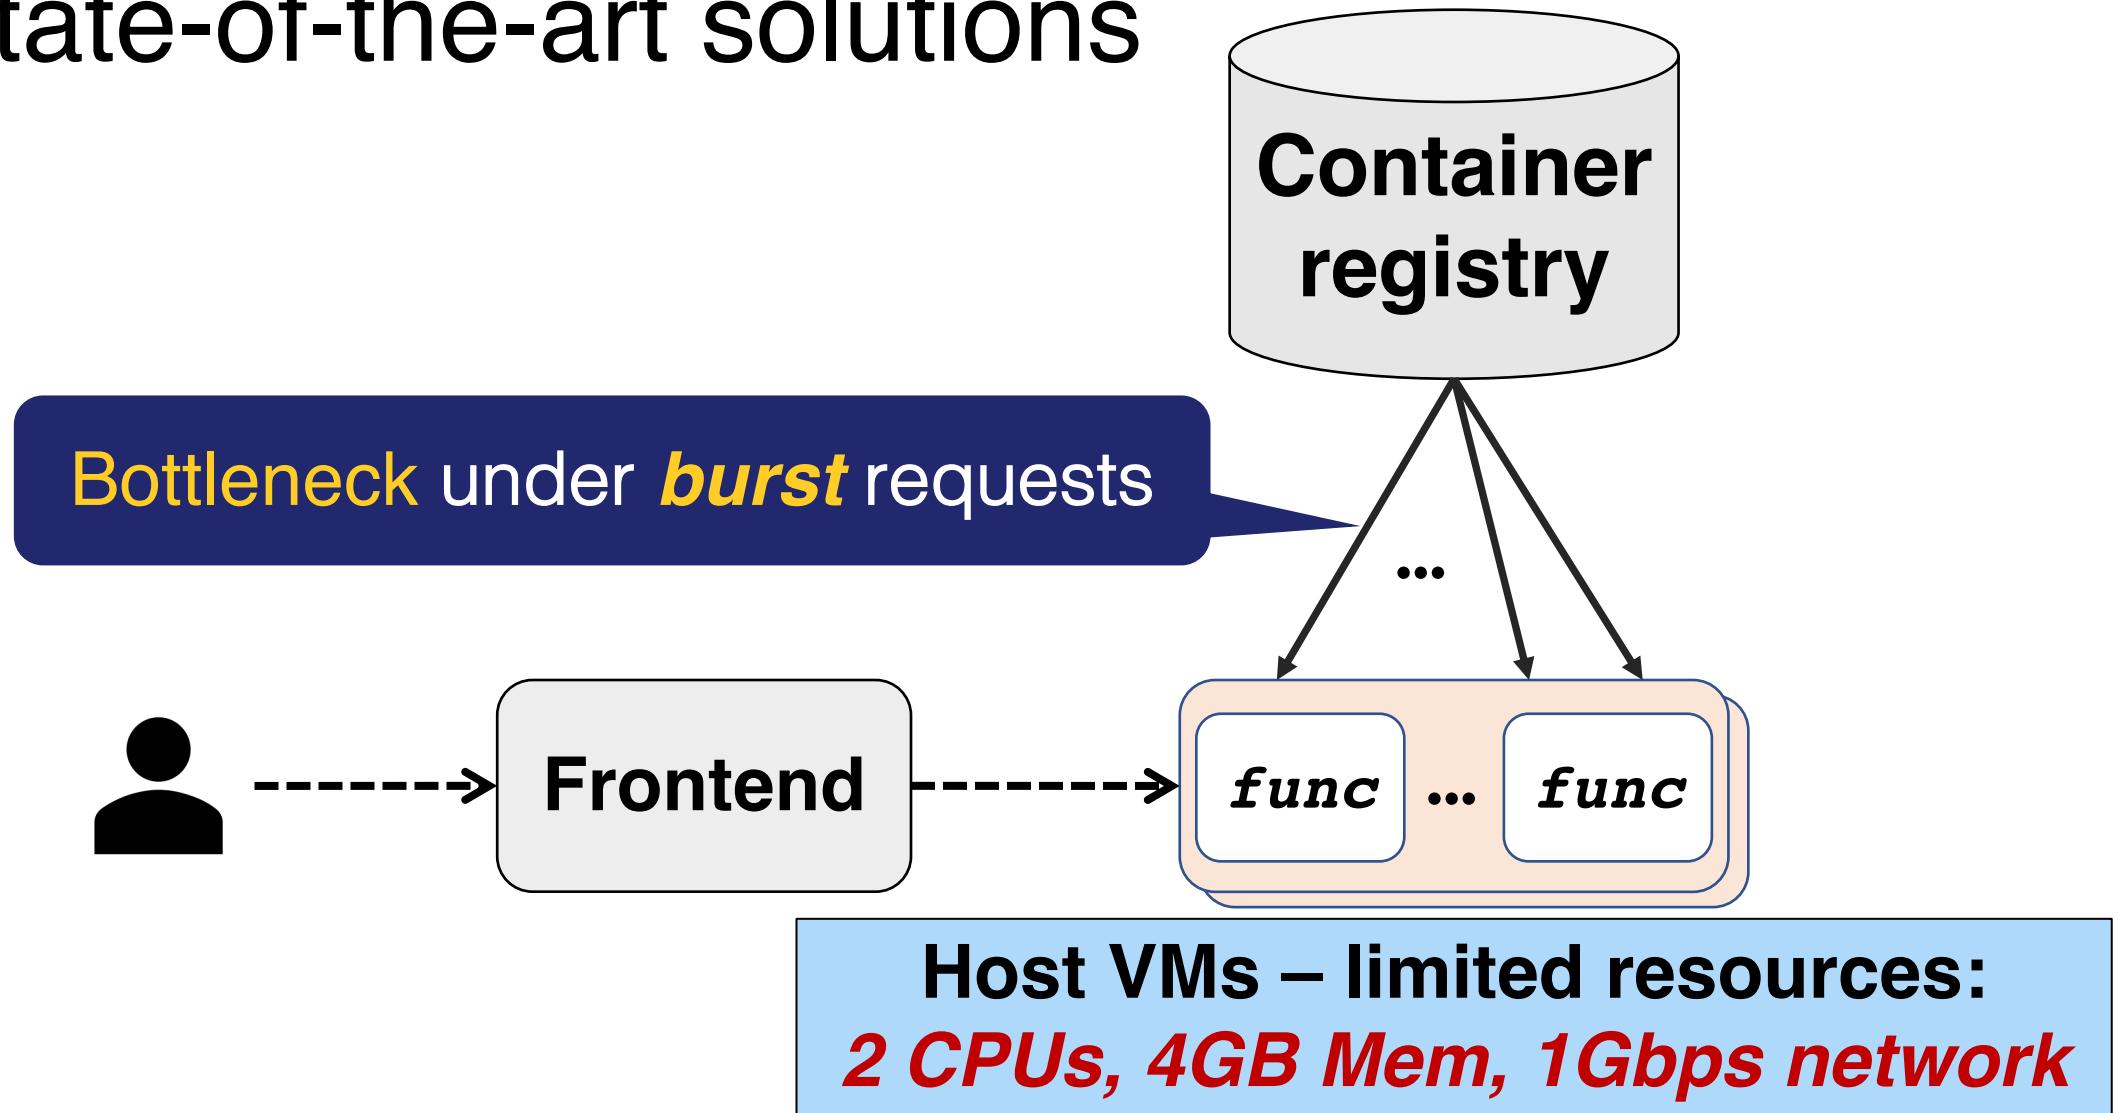
- A large fraction of startup time is spent on pulling images



Workload analysis

- Image pull latency distribution
- ***To handle workload dynamicity*** -> Scalable and resilient provisioning of large numbers of function containers
- ***To reduce cold start latency*** -> Optimize the performance of container provisioning process

State-of-the-art solutions



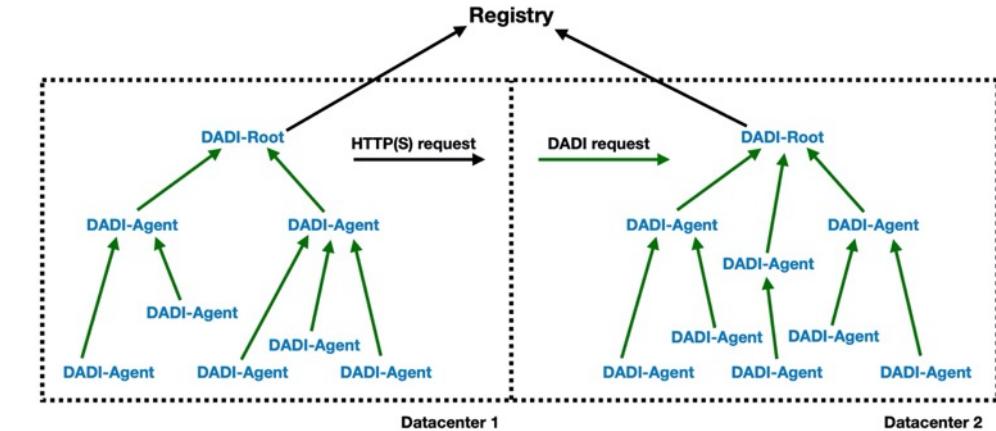
State-of-the-art solutions



Dragonfly



Kraken



Alibaba DADI [ATC '20]

Problems:

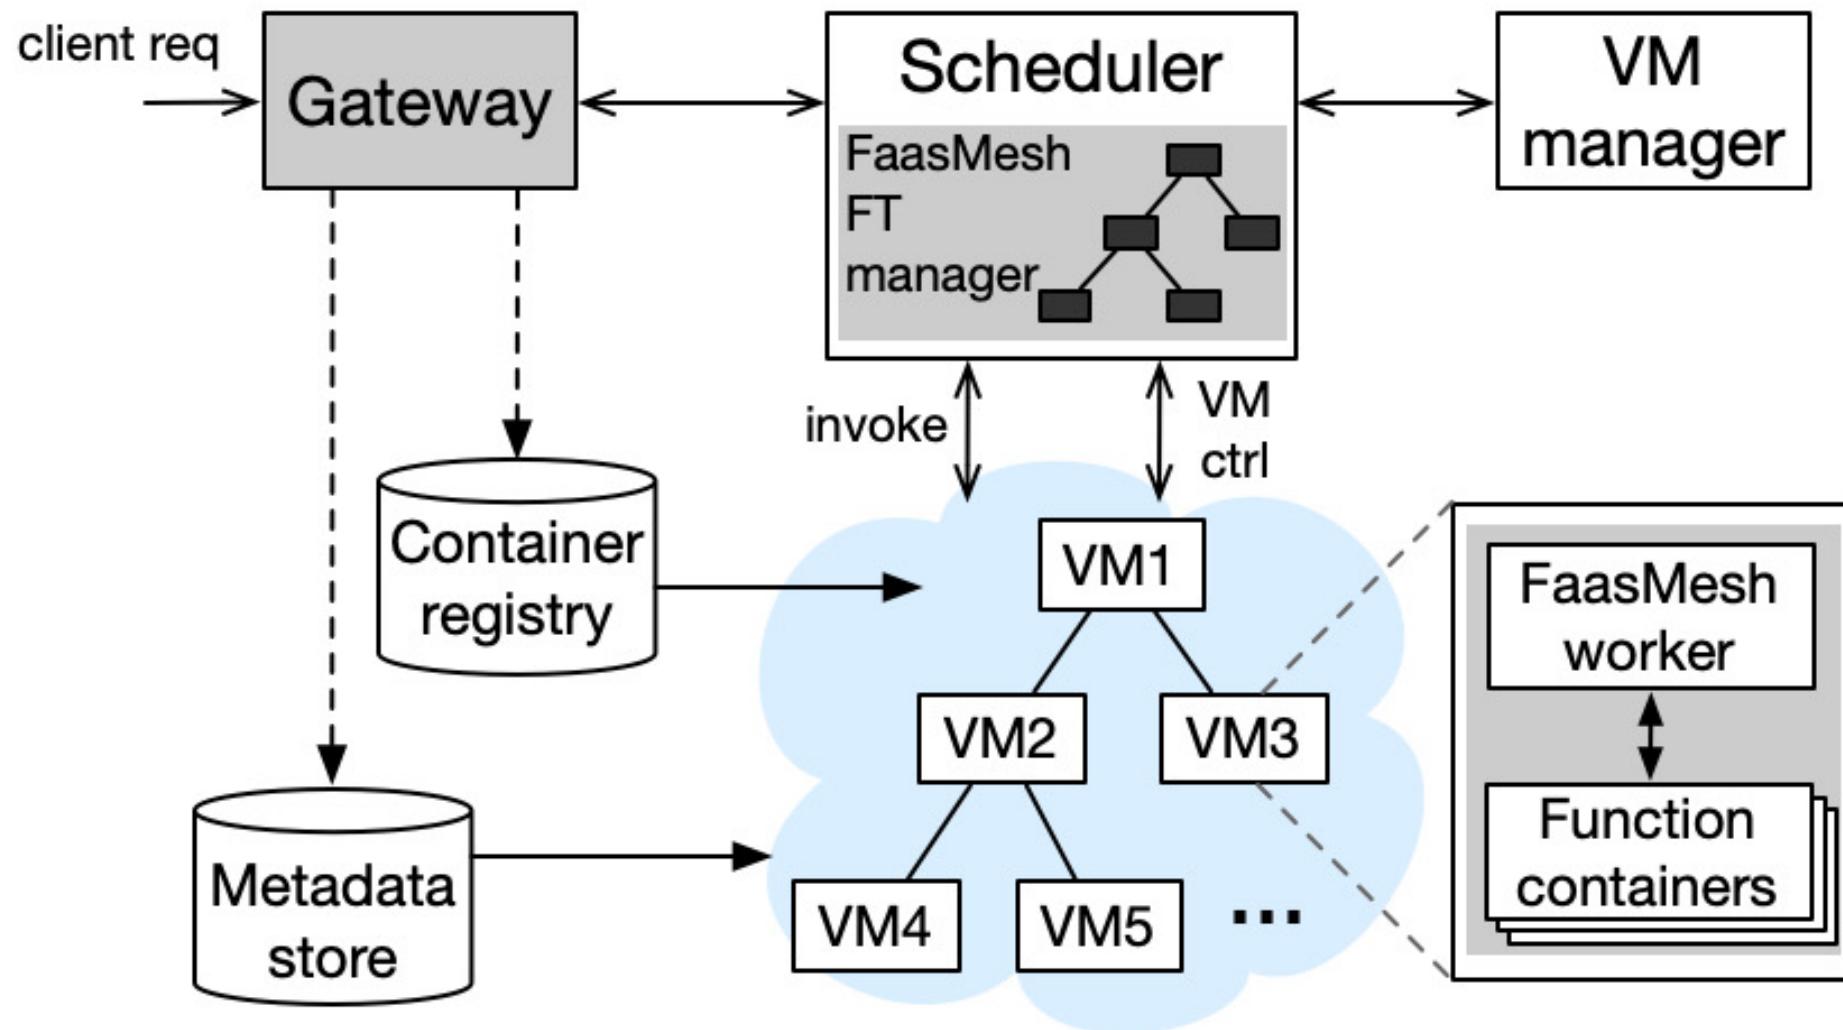
- *Extra, dedicated, centralized* components
- *Limited* VM resources
- VM's lifecycle is *unpredictable*
- *Multi-Tenancy isolation under FaaS* is not considered

FaaSNet: Scalable and Fast Provisioning of Custom Serverless Container Runtimes at Alibaba Cloud Function Compute

Agenda

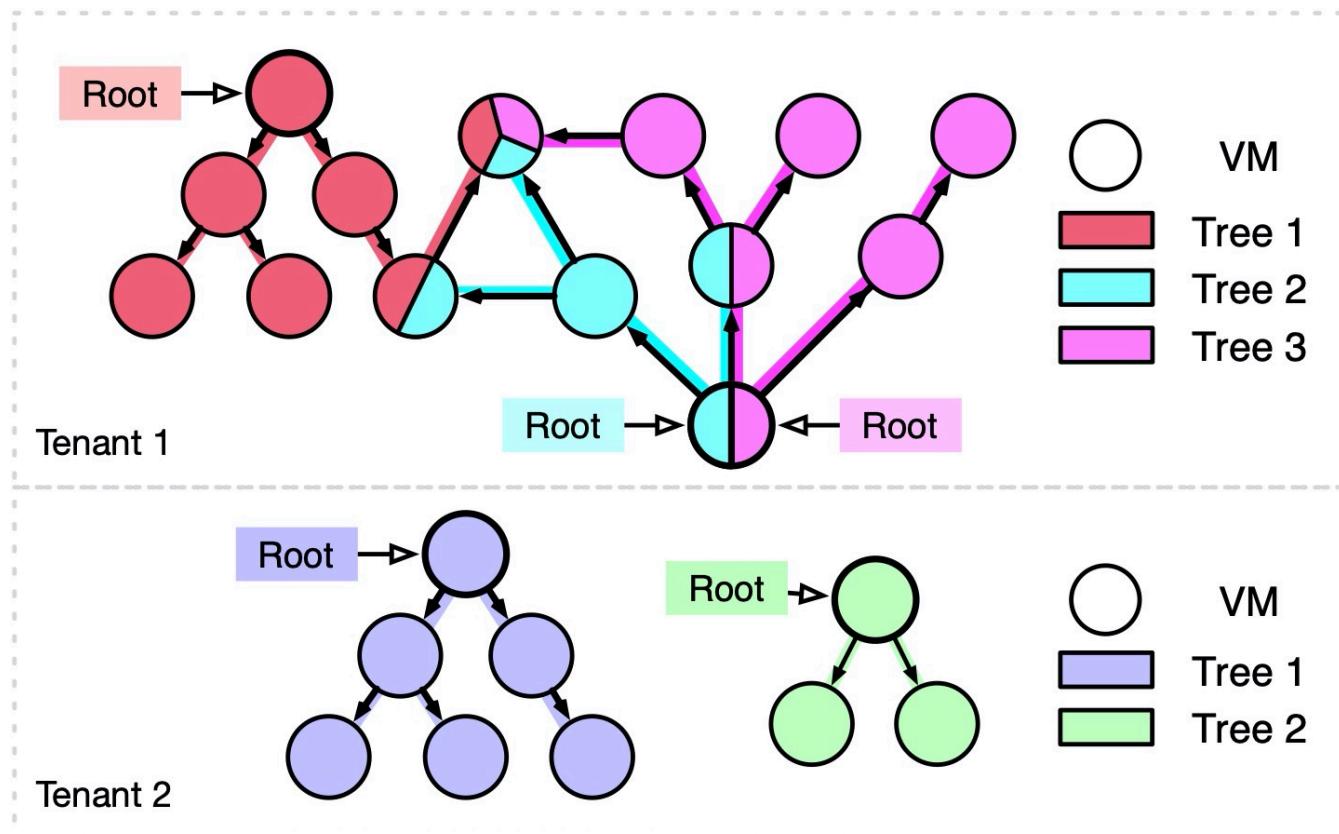
- FaaSNet design
- Evaluation
- Conclusion

FaaSNet design

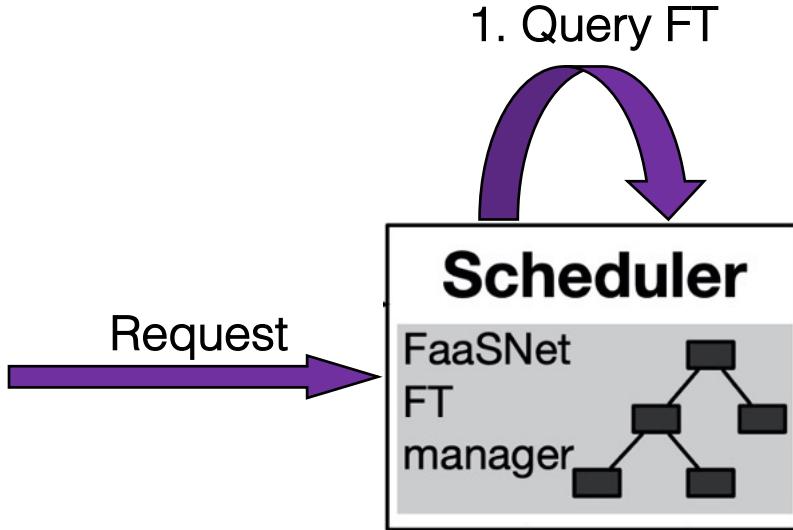


Function tree (FT)

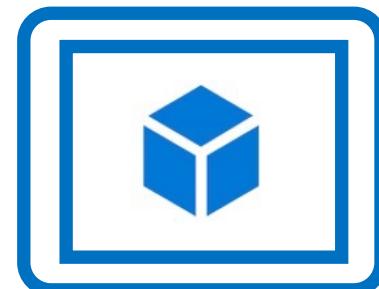
- FT is perfect self-balanced binary tree in ***Function*** level
- Exposed 2 APIs
 - insert
 - delete



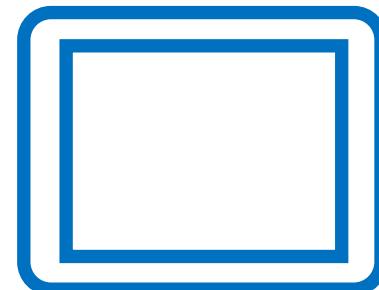
Container provisioning protocol



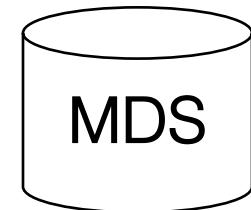
- : Image data
- : Data path
- : Control path



VM2 - Upstream

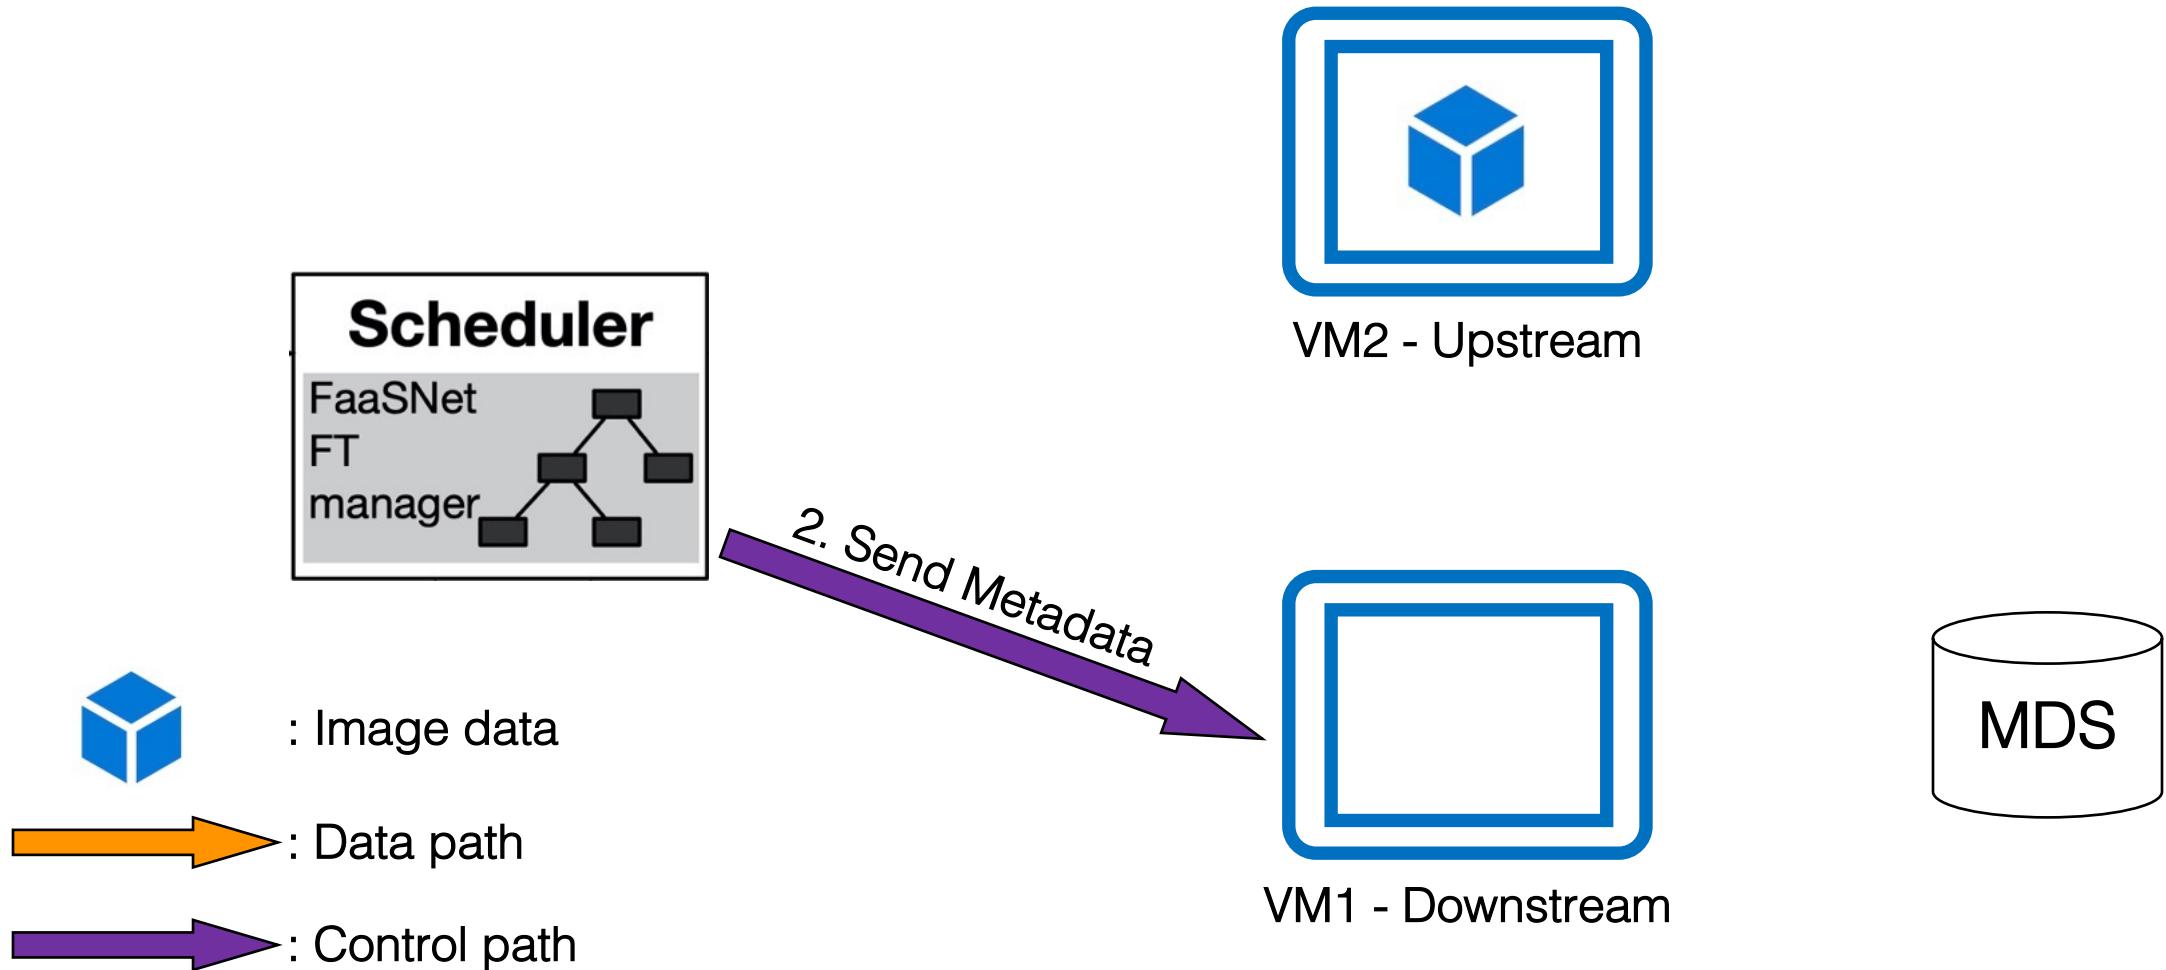


VM1 - Downstream

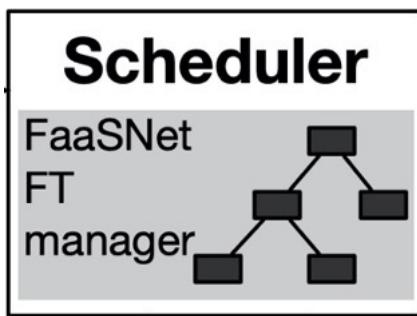


MDS

Container provisioning protocol



Container provisioning protocol



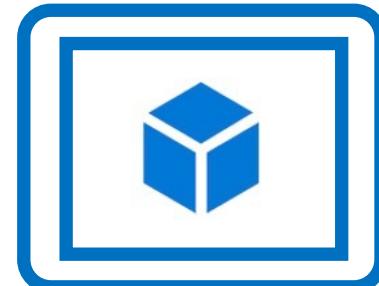
: Image data



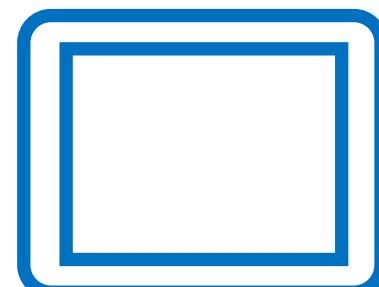
: Data path



: Control path

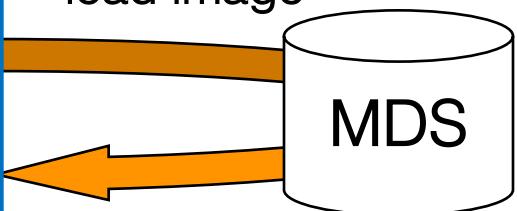


VM2 - Upstream

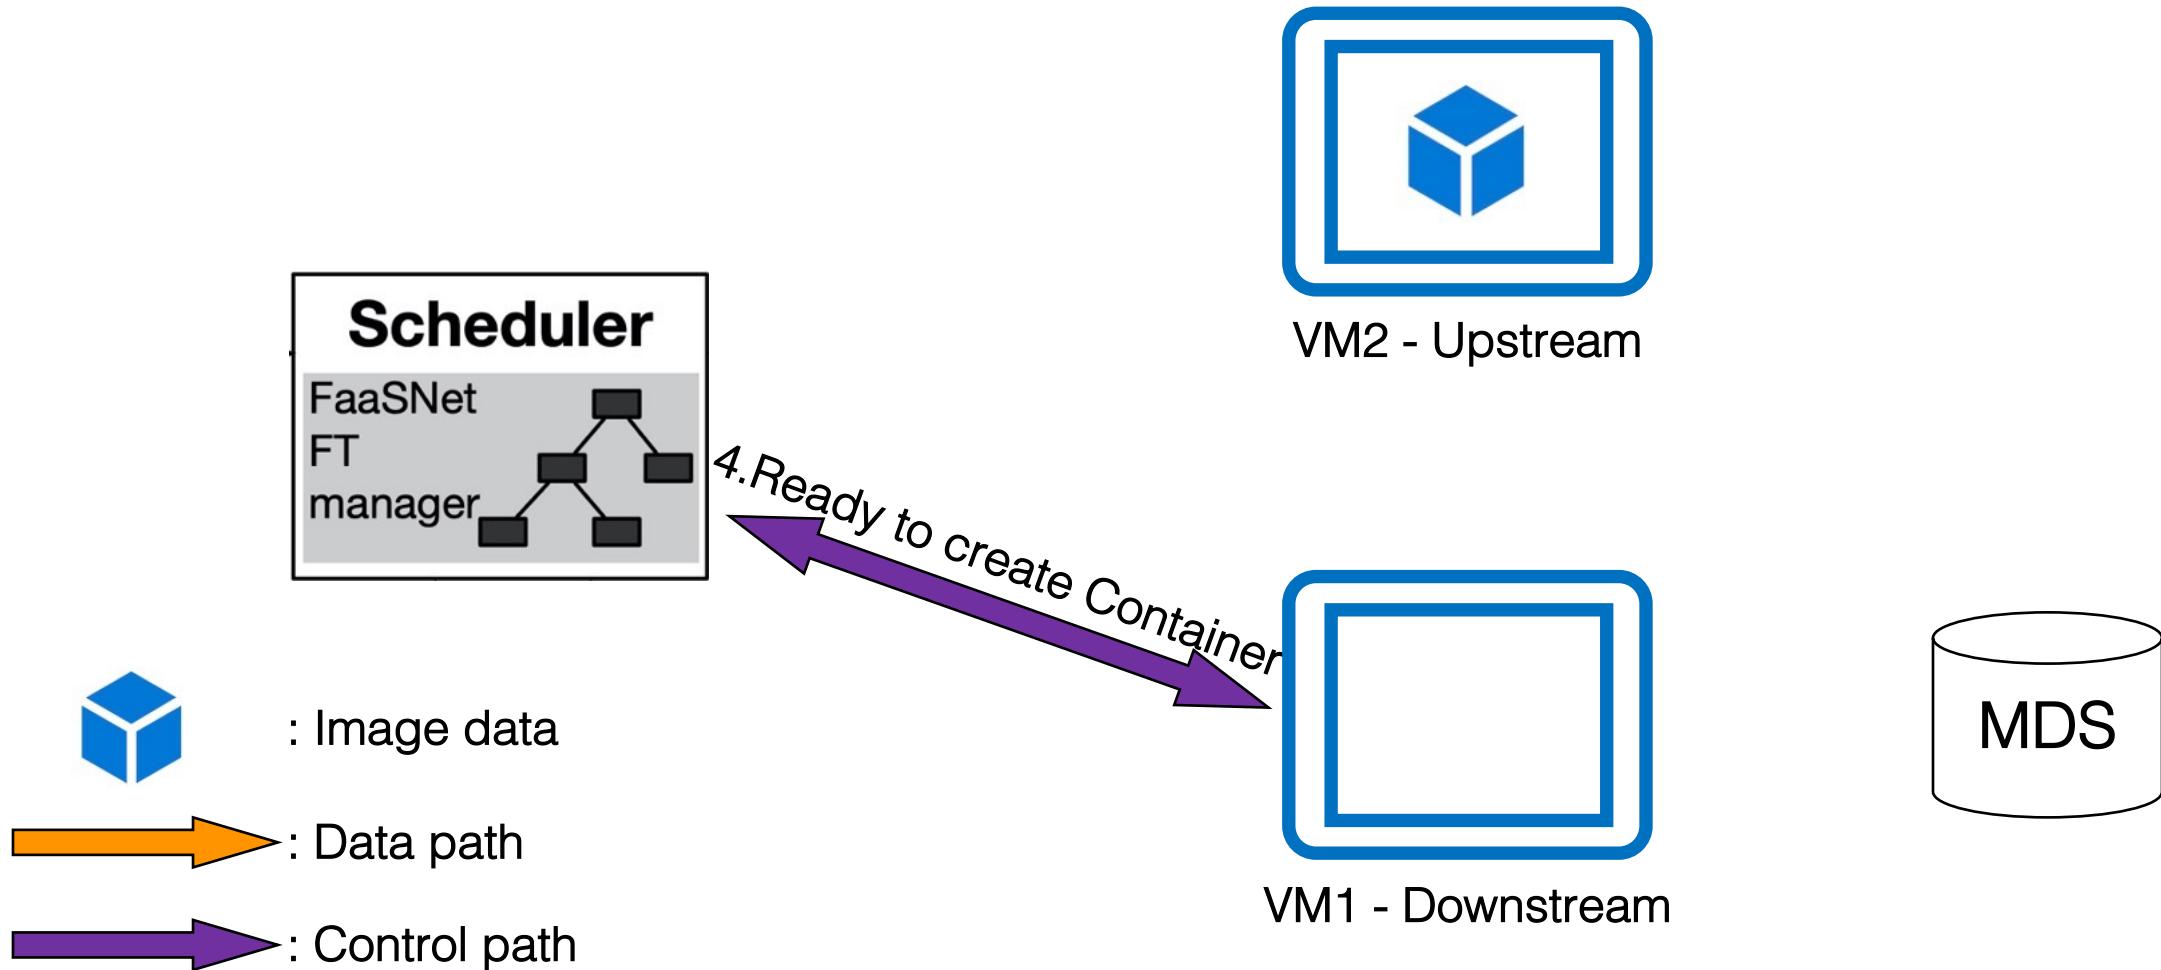


VM1 - Downstream

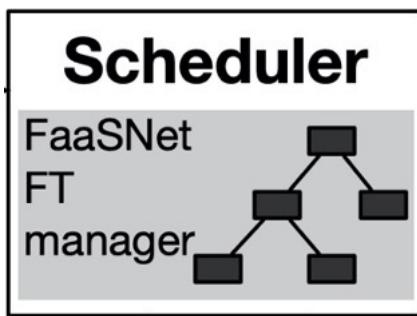
3. Download image manifest & load image



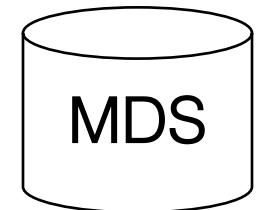
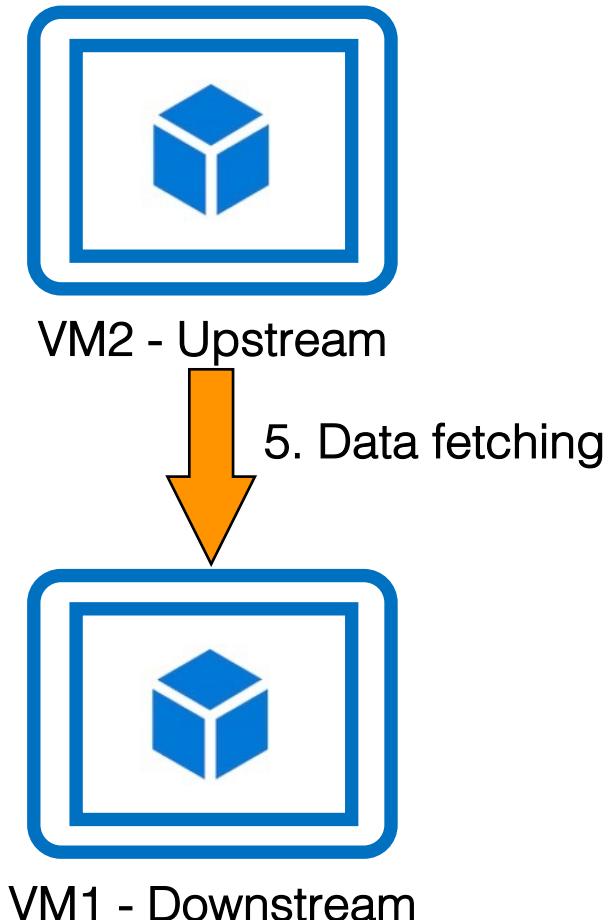
Container provisioning protocol



Container provisioning protocol



- : Image data
- : Data path
- : Control path



Agenda

- FaaSNet design
- Evaluation
- Conclusion

Experimental setup

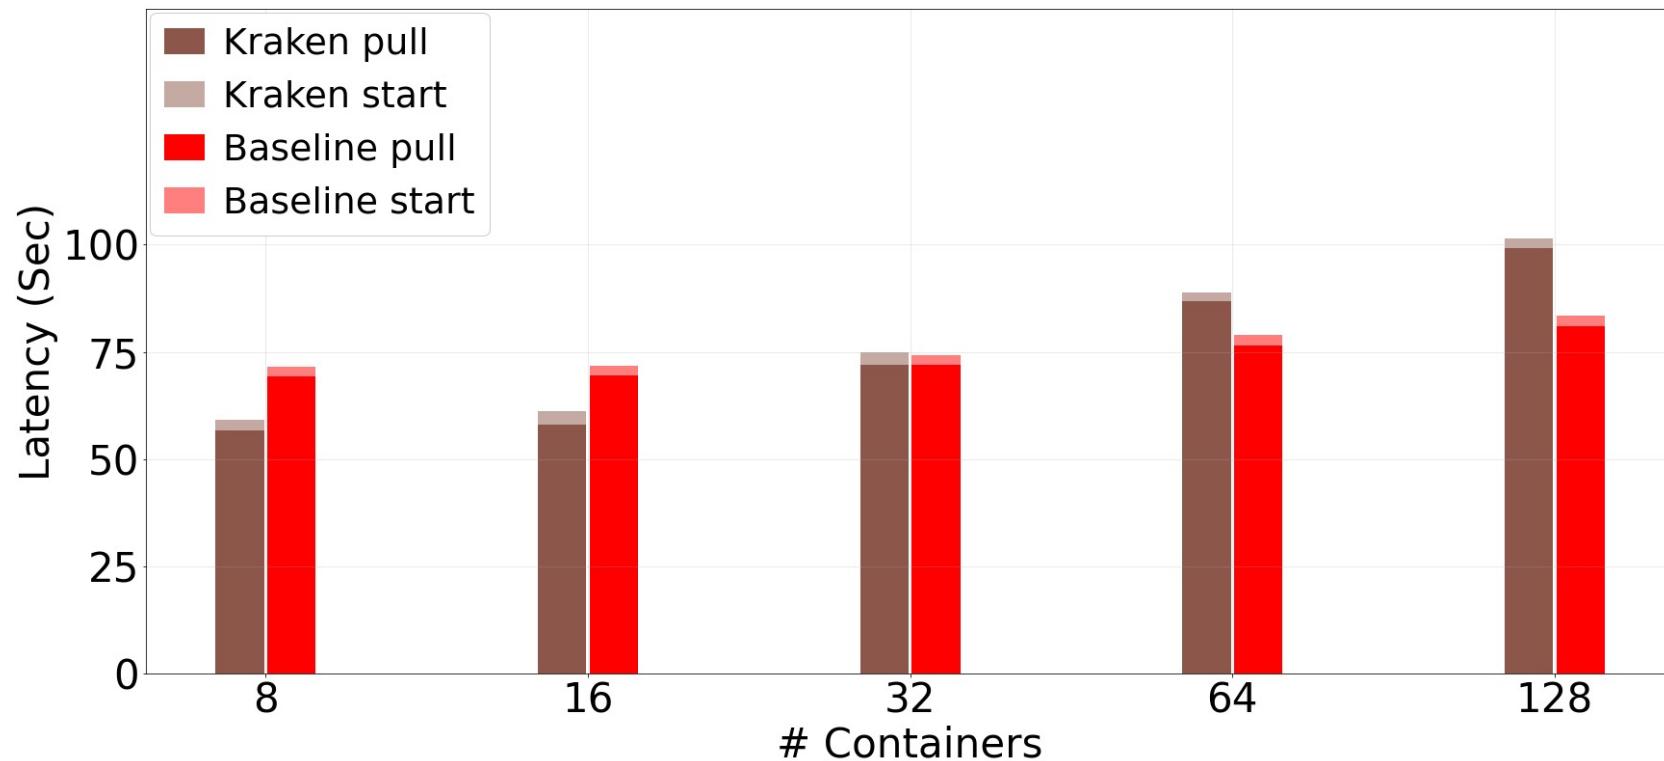
- Testbed is up to 1,000 VMs
- VM type: 2 CPUs, 4 GB memory, 1 Gbps network
 - Same as our production environment
- Example container image
 - 758 MB python-based function with ~2 sec duration

FaaSNet's performance

Kraken: Kraken dev cluster

Baseline: Original Alibaba Cloud

Function Compute



FaaSNet's performance

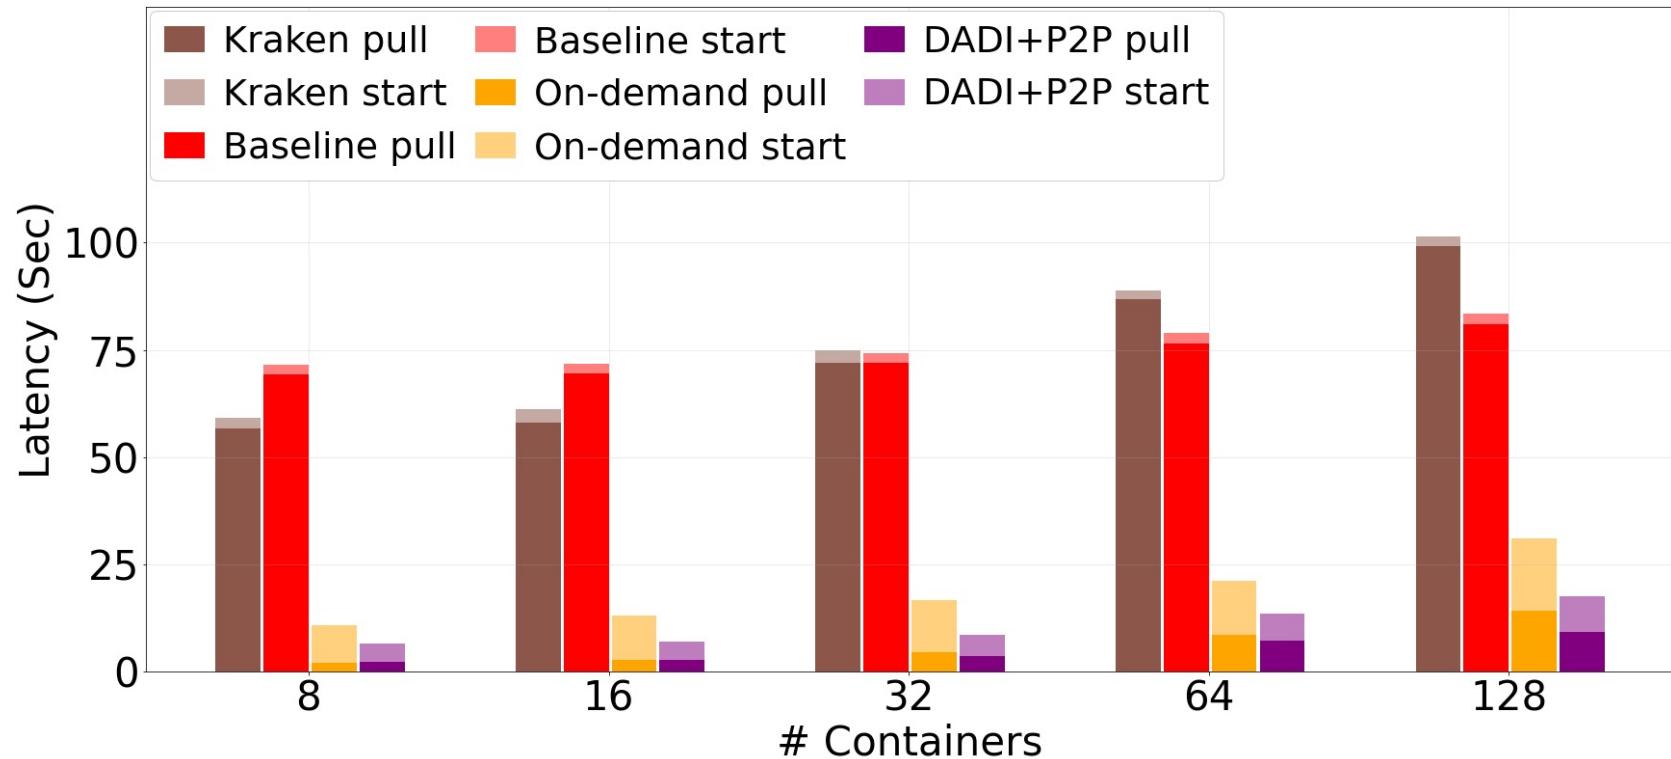
Kraken: Kraken dev cluster

Baseline: Original Alibaba Cloud

Function Compute (FC)

On-demand: FC + I/O efficient
format

DADI+P2P: FC + DADI



FaaSNet's performance

Kraken: Kraken dev cluster

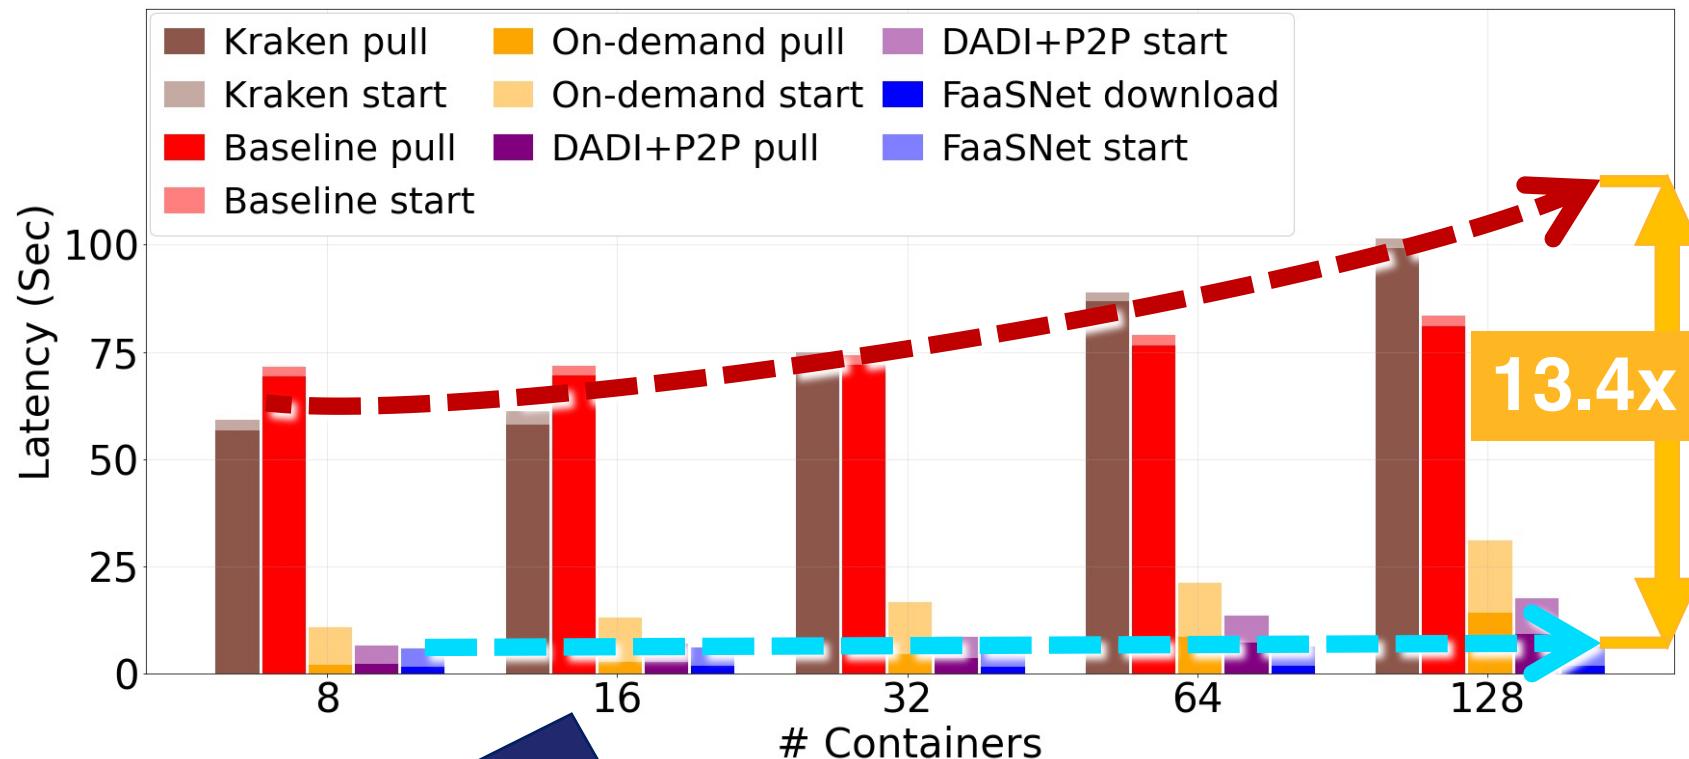
Baseline: Original Alibaba Cloud

Function Compute (FC)

On-demand: FC + I/O efficient
format

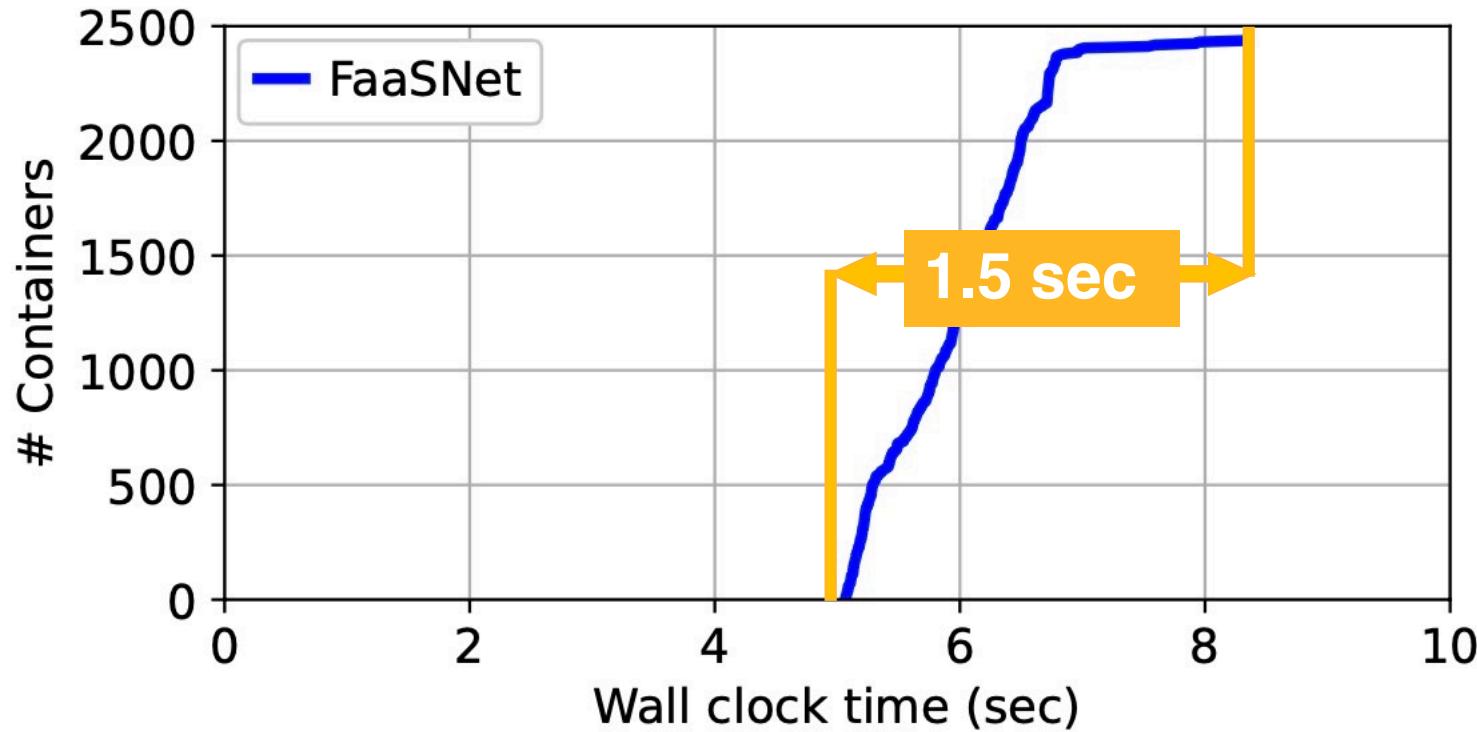
DADI+P2P: FC + DADI

FaaSNet



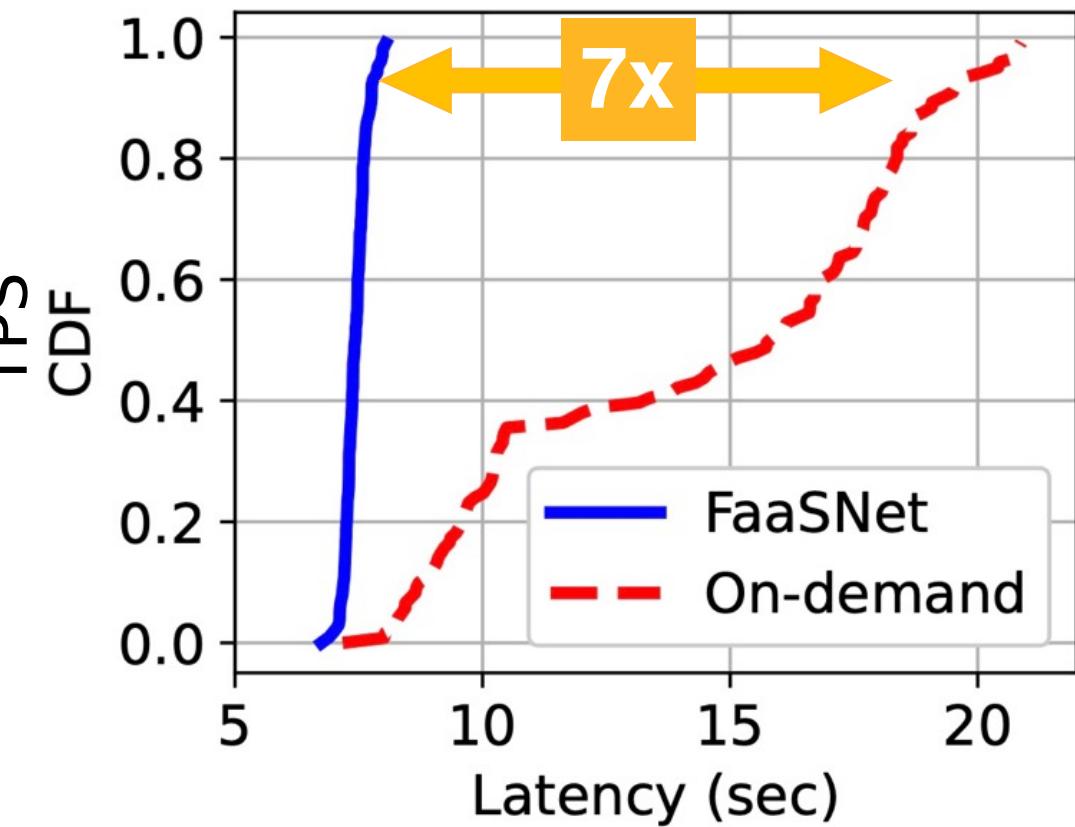
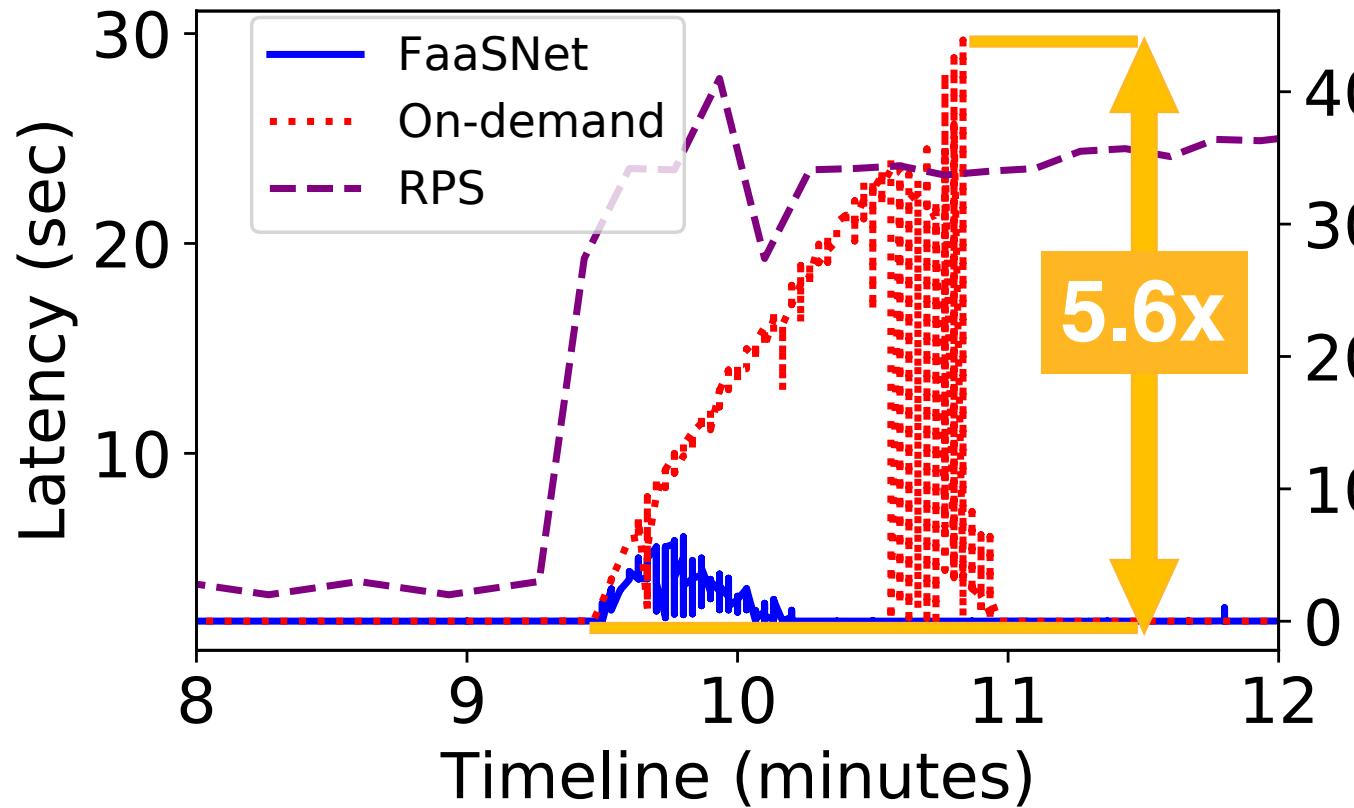
FaaSNet has ***strong scalability***

Production workload



Scales thousands of containers in seconds

Production workload



IoT trace

Conclusion

- FaaSNet is the **first** system that provides an end-to-end, integrated solution for **FaaS-optimized container runtime** provisioning (Alibaba Cloud Function Compute )
- FaaSNet scales **13.4x faster than** Alibaba Cloud's current FaaS platform



Thank you!

- Contact: Ao Wang – awang24@gmu.edu
- FT prototype & Alibaba Cloud Function Compute cold start traces
 - <https://github.com/mason-leap-lab/FaaSNet>



Q&A

- Local disk full -> cache eviction -> performance degradation
- Bandwidth issues