SmartOClock: Workload- and Risk-Aware Overclocking in the Cloud

ISCA 2024

Jovan Stojkovic*, Pulkit Misra, Íñigo Goiri, Sam Whitlock, Esha Choukse, Mayukh Das, Chetan Bansal, Jason Lee, Zoey Sun, Haoran Qiu*, Reed Zimmermann†, Savyasachi Samal, Brijesh Warrier, Ashish Raniwala, Ricardo Bianchini

Microsoft, *University of Illinois at Urbana-Champaign, †University of Texas Austin

Cloud Services are Heavily Overprovisioned

- User-facing workload with diurnal patterns and occasional spikes

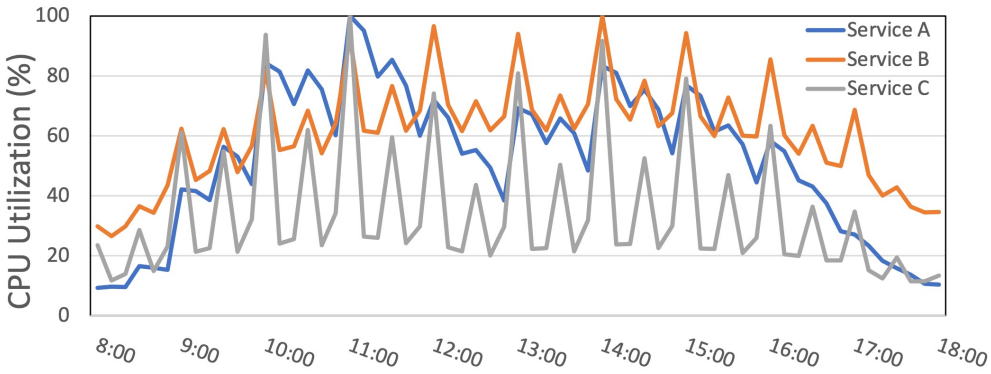

3 large Microsoft services

~1M virtual cores

Cloud Services are Heavily Overprovisioned

- User-facing workload with diurnal patterns and occasional spikes
  - Stringent SLO requirements → provisioned for the peak

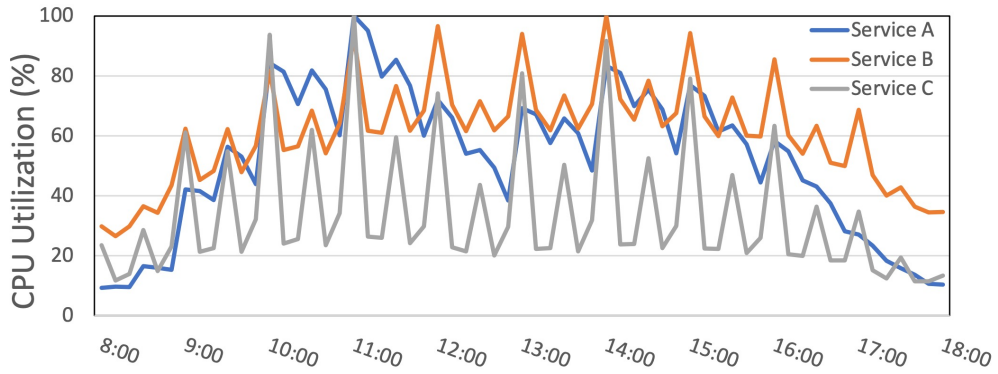

3 large Microsoft services

~1M virtual cores

Cloud Services are Heavily Overprovisioned

- User-facing workload with diurnal patterns and occasional spikes
  - Stringent SLO requirements → provisioned for the peak
- Auto-scaling and on-demand provisioning not a remedy

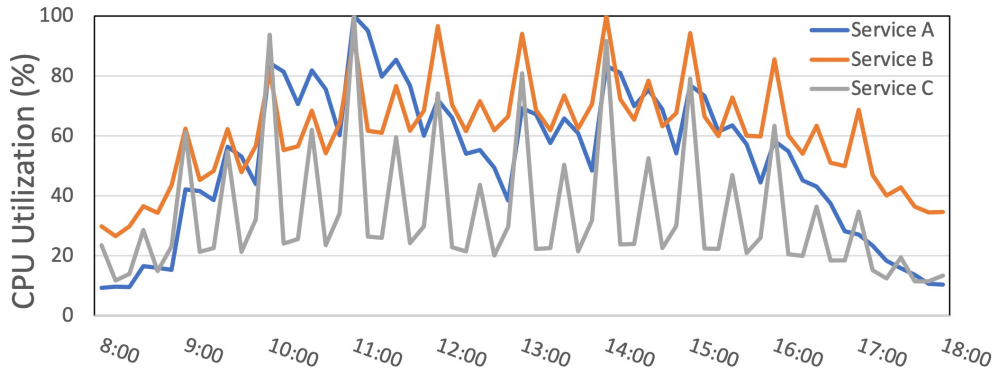

3 large Microsoft services

~1M virtual cores

Overclocking to the Rescue?

- Overclock at load spikes
  - Improves performance + save cost

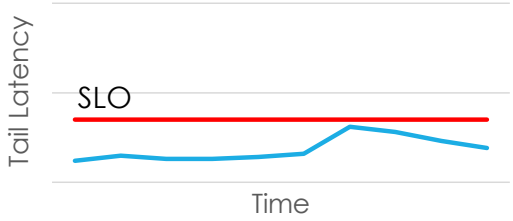

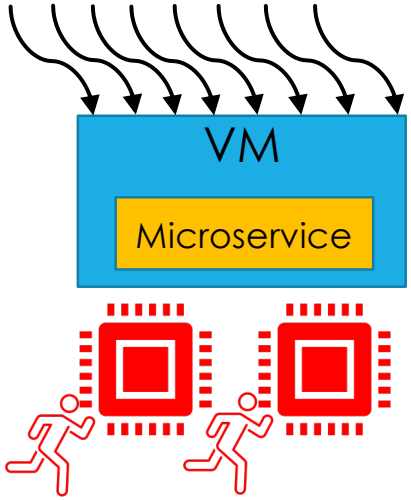

Overclocking to the Rescue?

- Overclock at load spikes
  - Improves performance + save cost
- Overclocking is not free
  - Increases power draw
  - Increases component wear-out
  - Overclocking is opportunistic

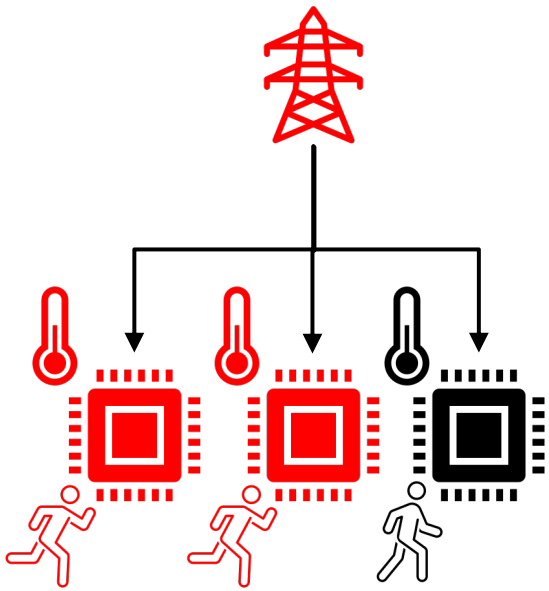

Overclocking to the Rescue?

- Overclock at load spikes
  - Impr
- Overcl
  - Incre
  - Incre
  - Over

How to get the benefits of overclocking while managing the risks?

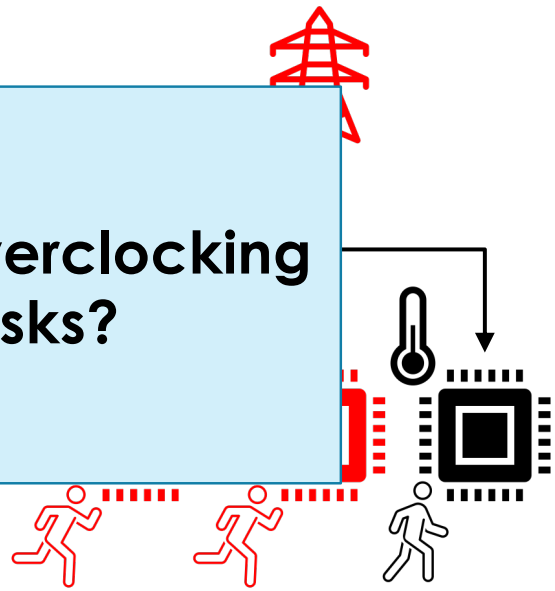

Contributions

- Thorough characterization of the environment
- Propose **SmartOClock**
- Evaluation
  - Real 36-server overclockable cluster
  - Production workloads

Outline of this talk

- **Challenges and opportunities for overclocking**
- SmartOClock
- Evaluation results

How to Manage Overclocking in the Cloud?

1. **When is it beneficial for workloads to be overclocked?**

How to Manage Overclocking in the Cloud?

1. When is it beneficial for workloads to be overclocked?
2. **Are there enough resources for overclocking in the cloud?**

How to Manage Overclocking in the Cloud?

1. When is it beneficial for workloads to be overclocked?
2. Are there enough resources for overclocking in the cloud?
3. **How to overclock without exceeding the resource limits?**

How to Manage Overclocking in the Cloud?

1. When is it beneficial for workloads to be overclocked?
2. Are there enough resources for overclocking in the cloud?
3. How to overclock without exceeding the resource limits?
4. **How to act when the resource limits are exceeded?**

Overclocking opaque-box VMs is inefficient

- Microservices – target metric is tail latency

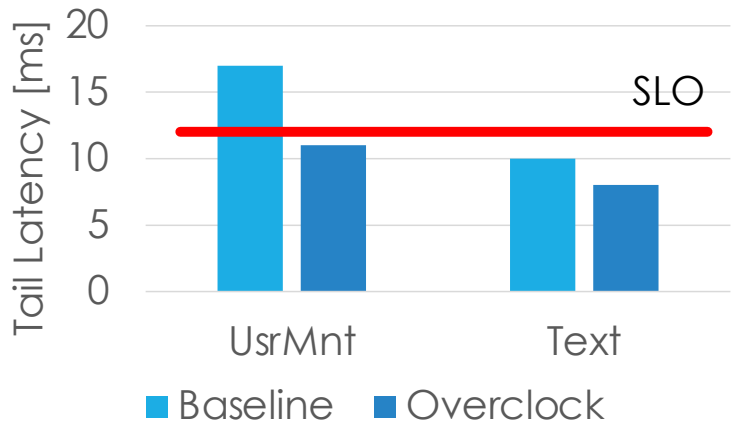

Overclocking opaque-box VMs is inefficient

- Microservices – target metric is tail latency
- Workload-agnostic signals (e.g., CPU utilization) not enough

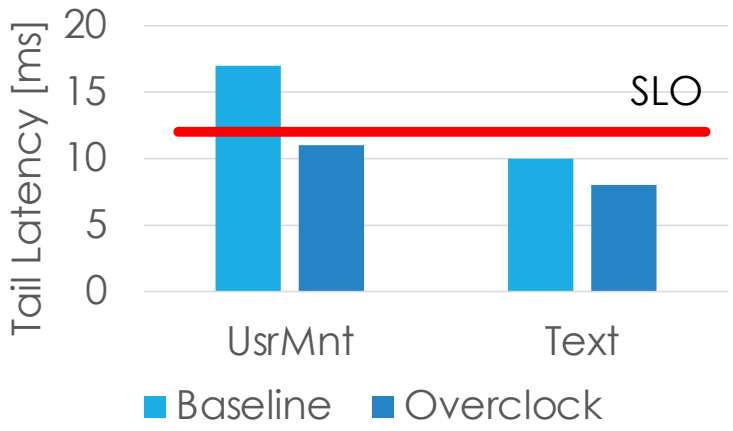

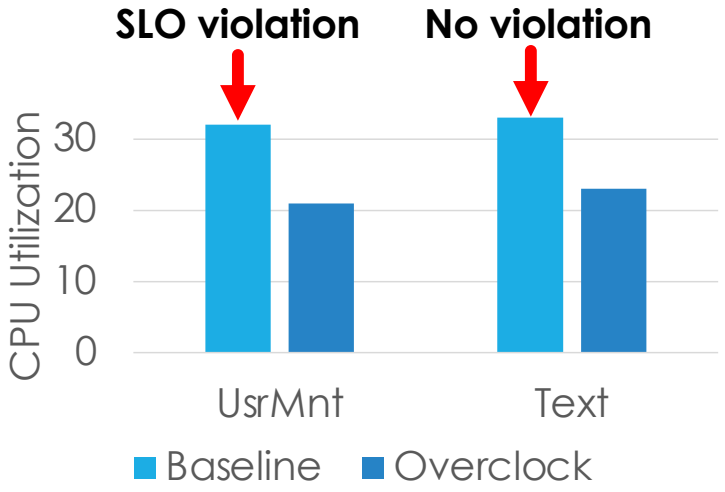

Overclocking opaque-box VMs is inefficient

- Microservices – target metric is tail latency
- Worklo… …gh …ion

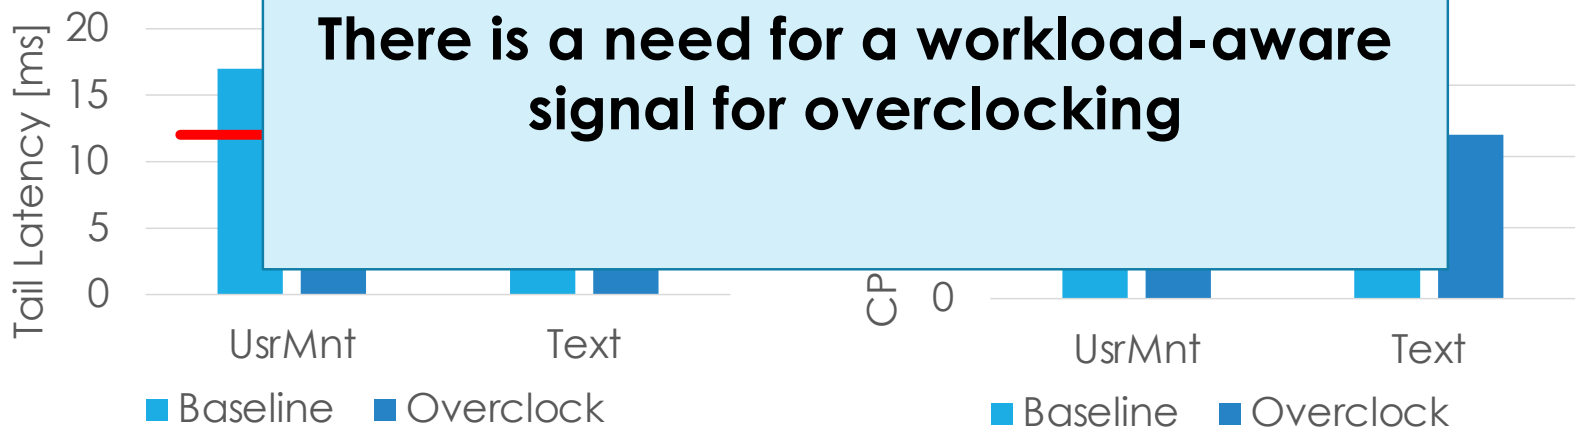

There is a need for a workload-aware signal for overclocking

There is power headroom across Azure

- Analyze power consumption of Azure Fleet
- Power usage typically low (~60%), can spike (up to 99%)

There is power headroom across Azure

- Analyze power consumption of Azure Fleet
- Power usage typically low (~60%), can spike (up to 99%)

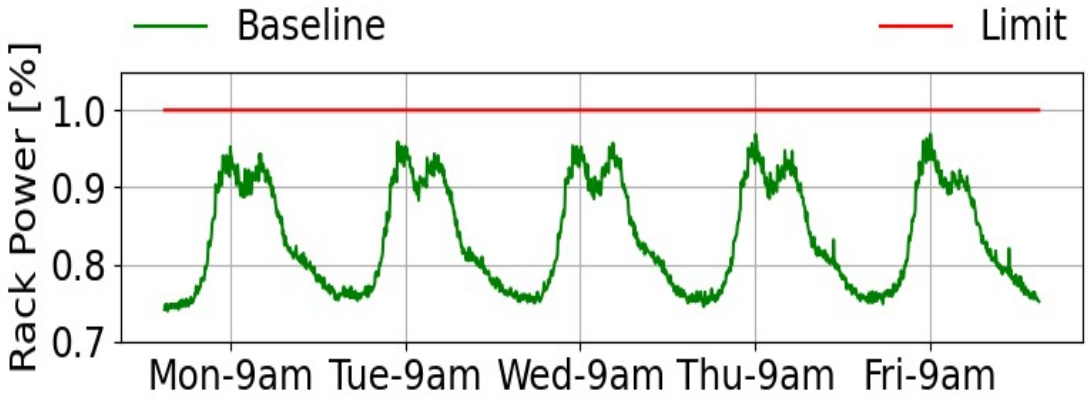

There is power headroom across Azure

- Analyze power consumption of Azure Fleet
- Power usage typically low (~60%), can spike (up to 99%)

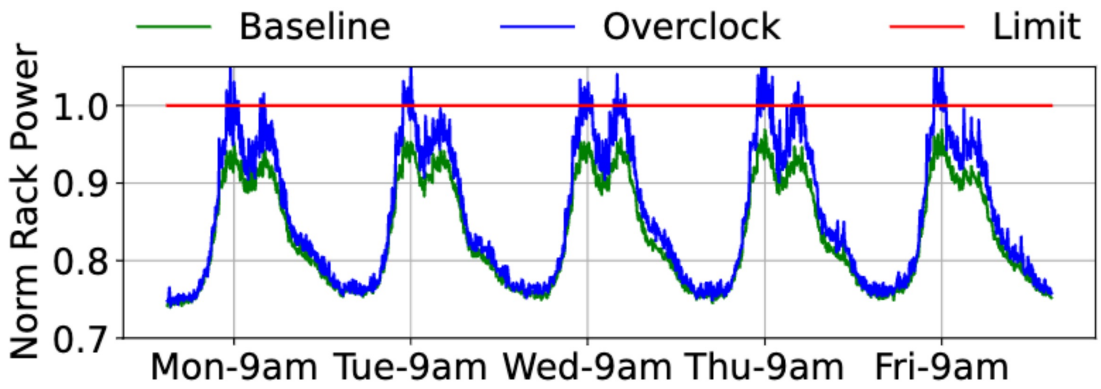

There is power headroom across Azure

- Analyze power consumption of Azure Fleet
- Power usage typically low (~60%), can spike (up to 99%)

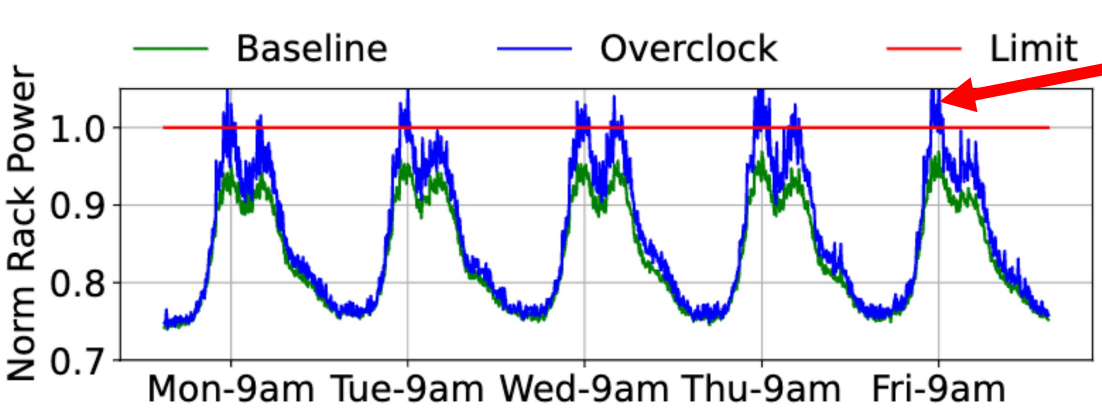

Power Capping
(degrades perf by reducing frequency)

There is power headroom across Azure

- Analyze power consumption of Azure Fleet
- Power u

Naively overclocking can cause power capping events

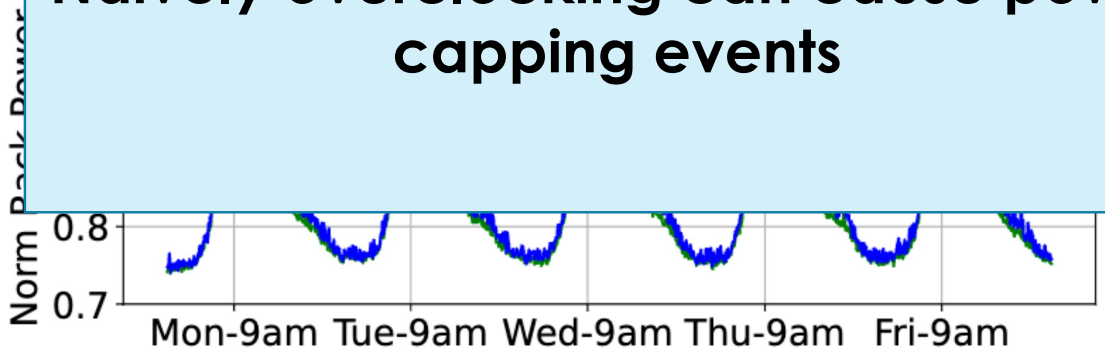

ower Capping
(degrades perf by
educing frequency)

There is lifetime headroom across Azure

- Analyze CPU utilization of Azure Fleet
- TSMC reliability model

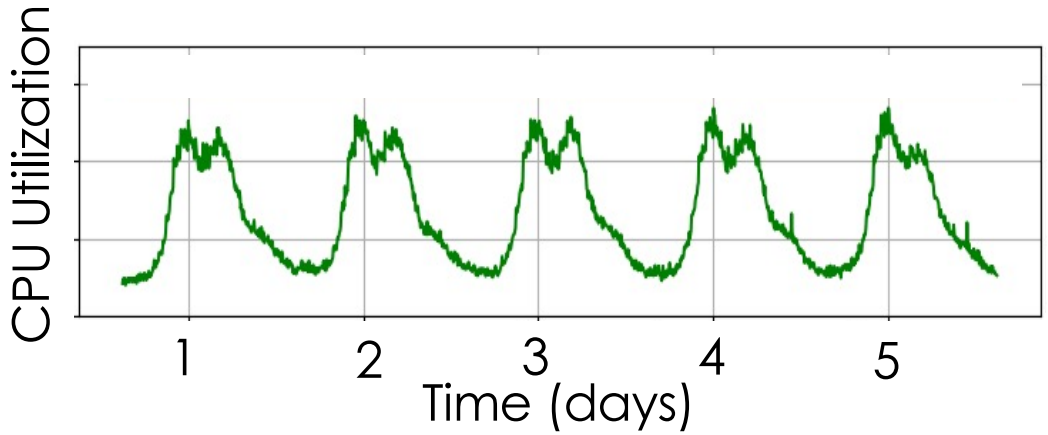

There is lifetime headroom across Azure

- Analyze CPU utilization of Azure Fleet
- TSMC reliability model

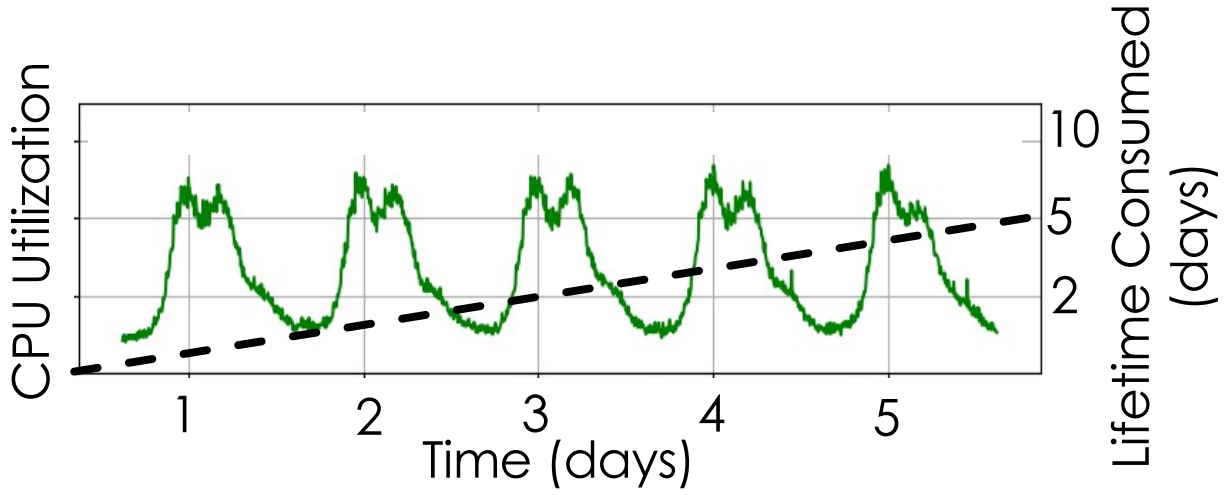

There is lifetime headroom across Azure

- Analyze CPU utilization of Azure Fleet
- TSMC reliability model

No-OC

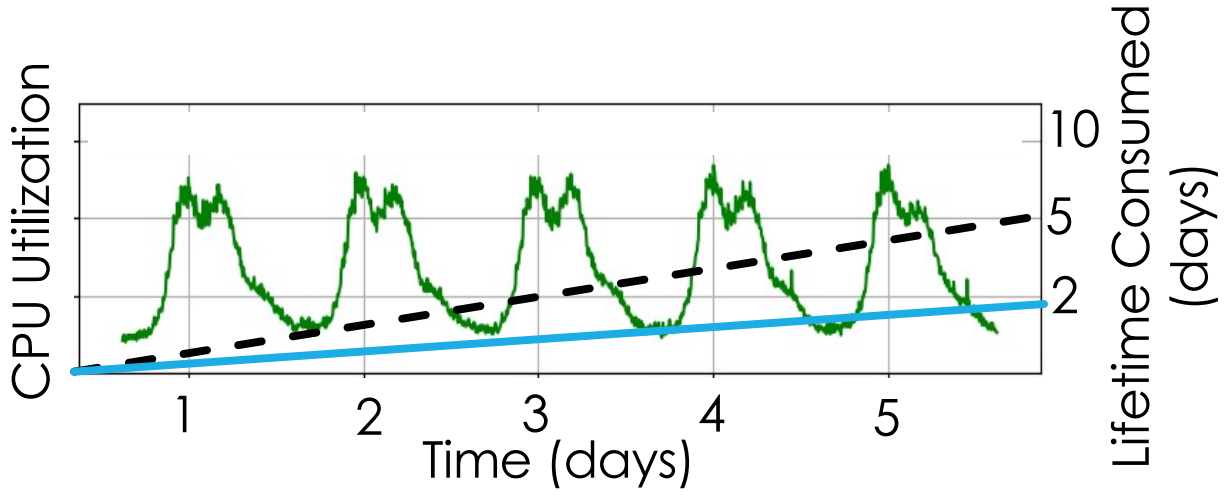

There is lifetime headroom across Azure

- Analyze CPU utilization of Azure Fleet
- TSMC reliability model

No-OC **Always-OC**

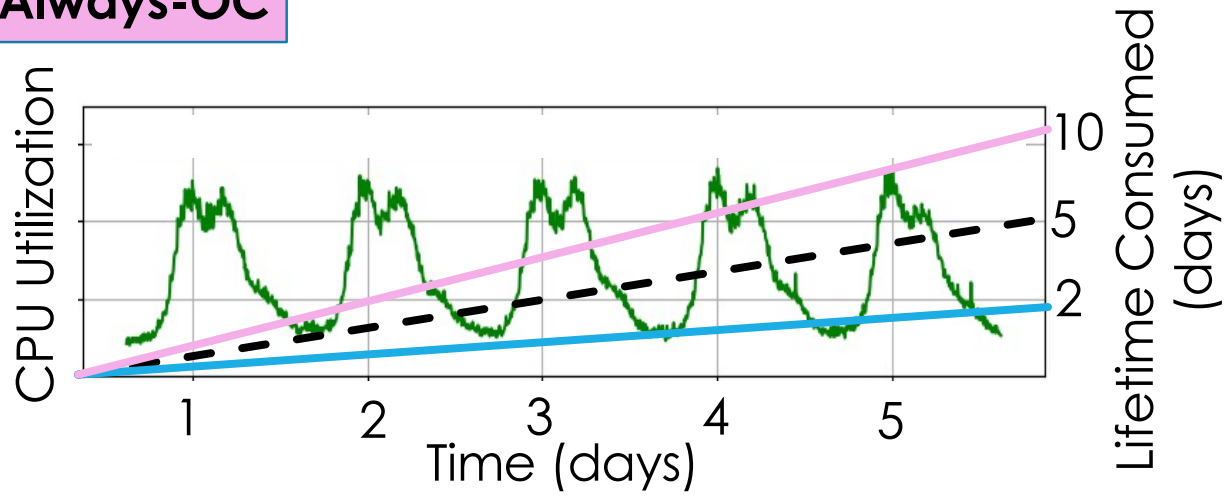

There is lifetime headroom across Azure

- Analyze CPU utilization of Azure Fleet
- TSMC reliability model

No-OC **Always-OC** **OC-Aware**

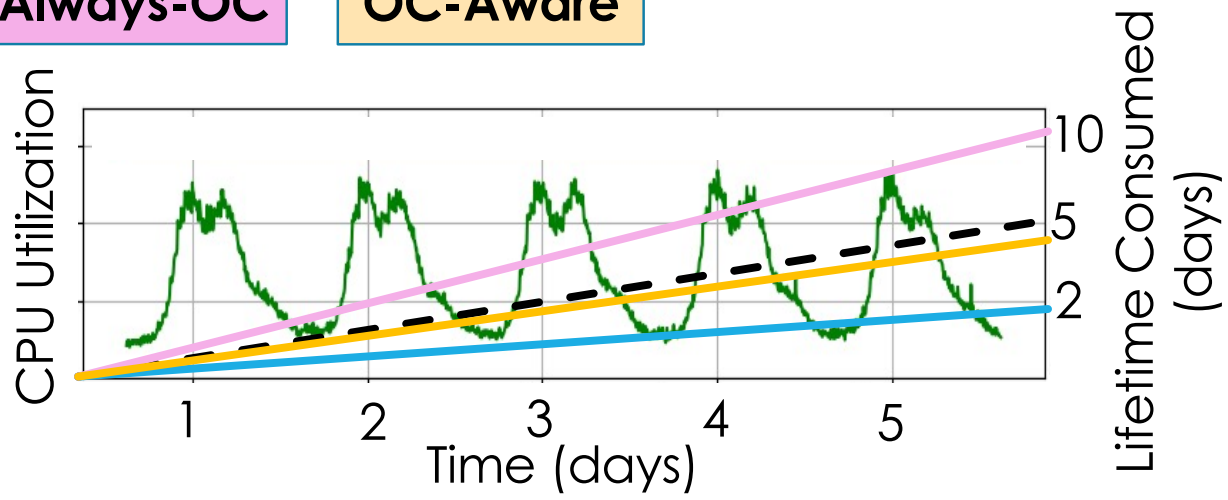

There is lifetime headroom across Azure

- Analyze CPU utilization of Azure Fleet
- TSMC re

No-OC

Naively overclocking can cause premature server wear-out

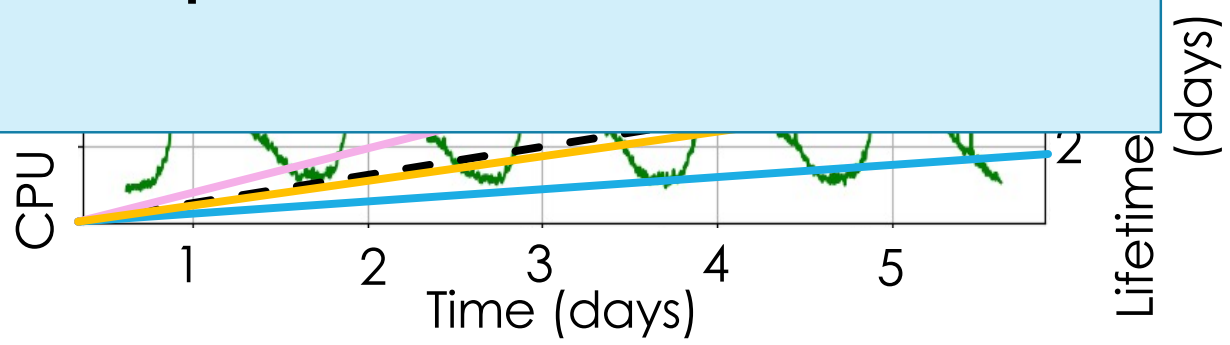

Rack power is predictable in Azure

- Highly ***periodic*** power consumption behavior across days

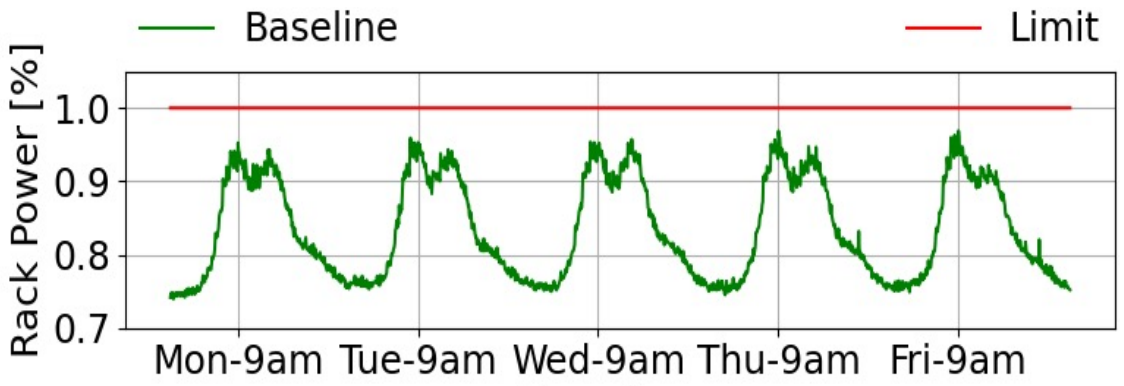

Rack power is predictable in Azure

- Highly ***periodic*** power consumption behavior across days

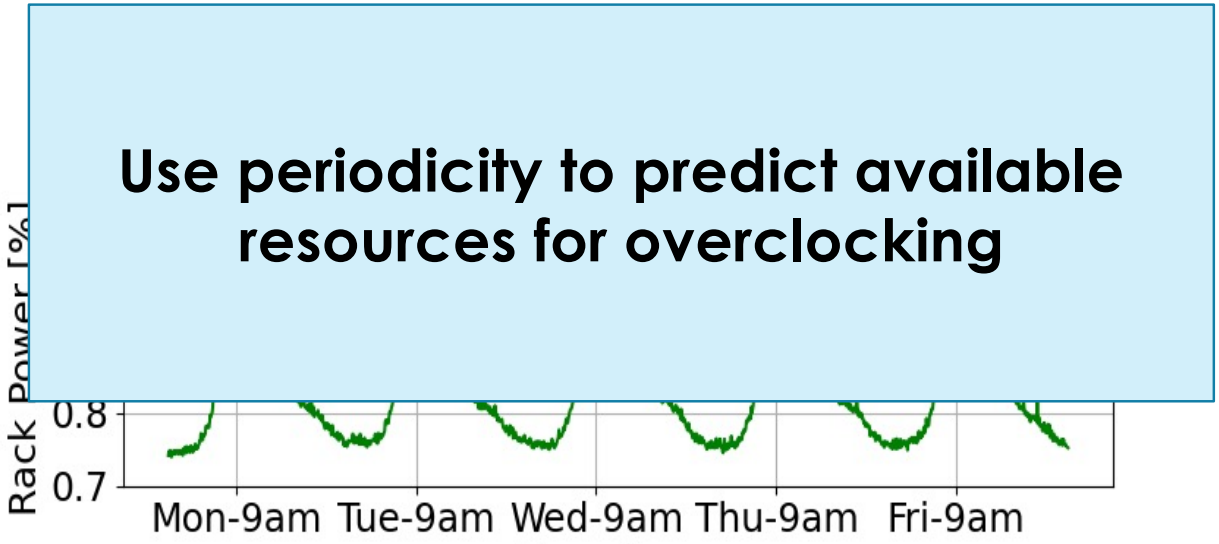

Servers have heterogeneous power profiles

- Servers' needs are heterogeneous and dynamic

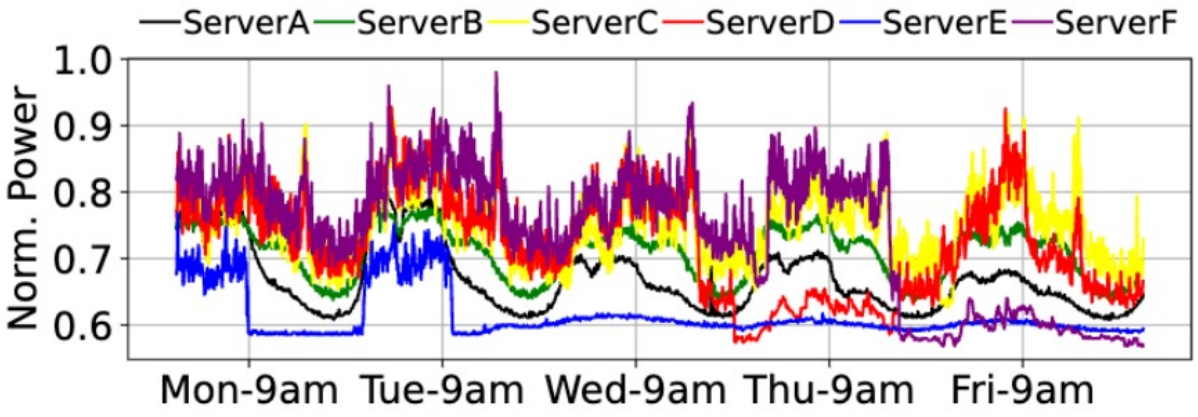

Servers have heterogeneous power profiles

- Servers' needs are heterogeneous and dynamic

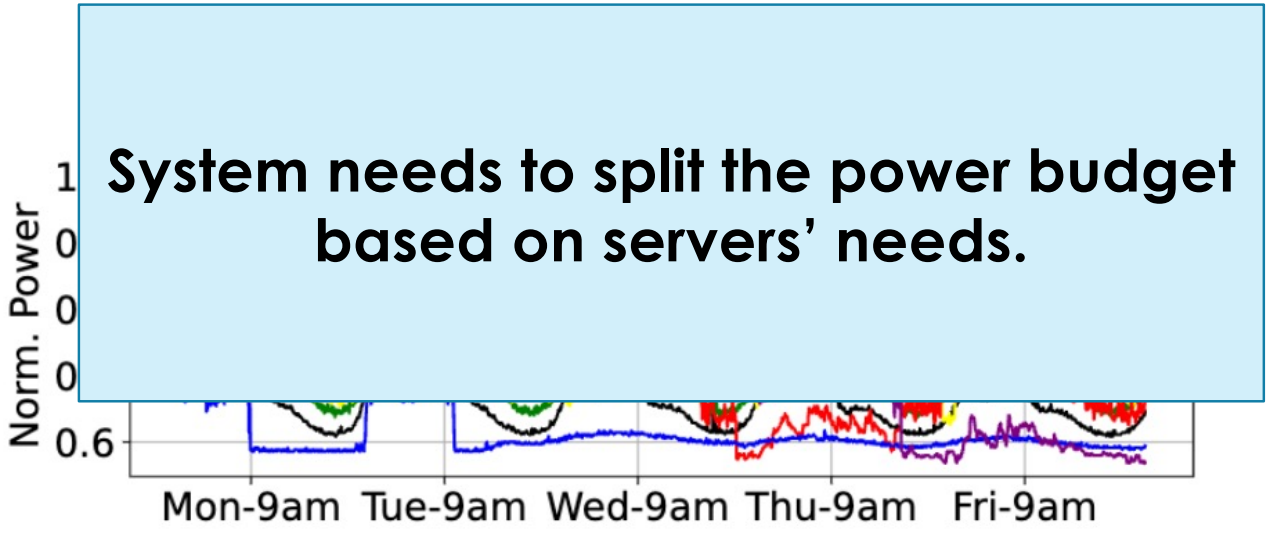

Outline

- Challenges and opportunities for overclocking
- **SmartOClock**
- Evaluation results

SmartOClock: Overclocking Management Framework for Cloud Platforms

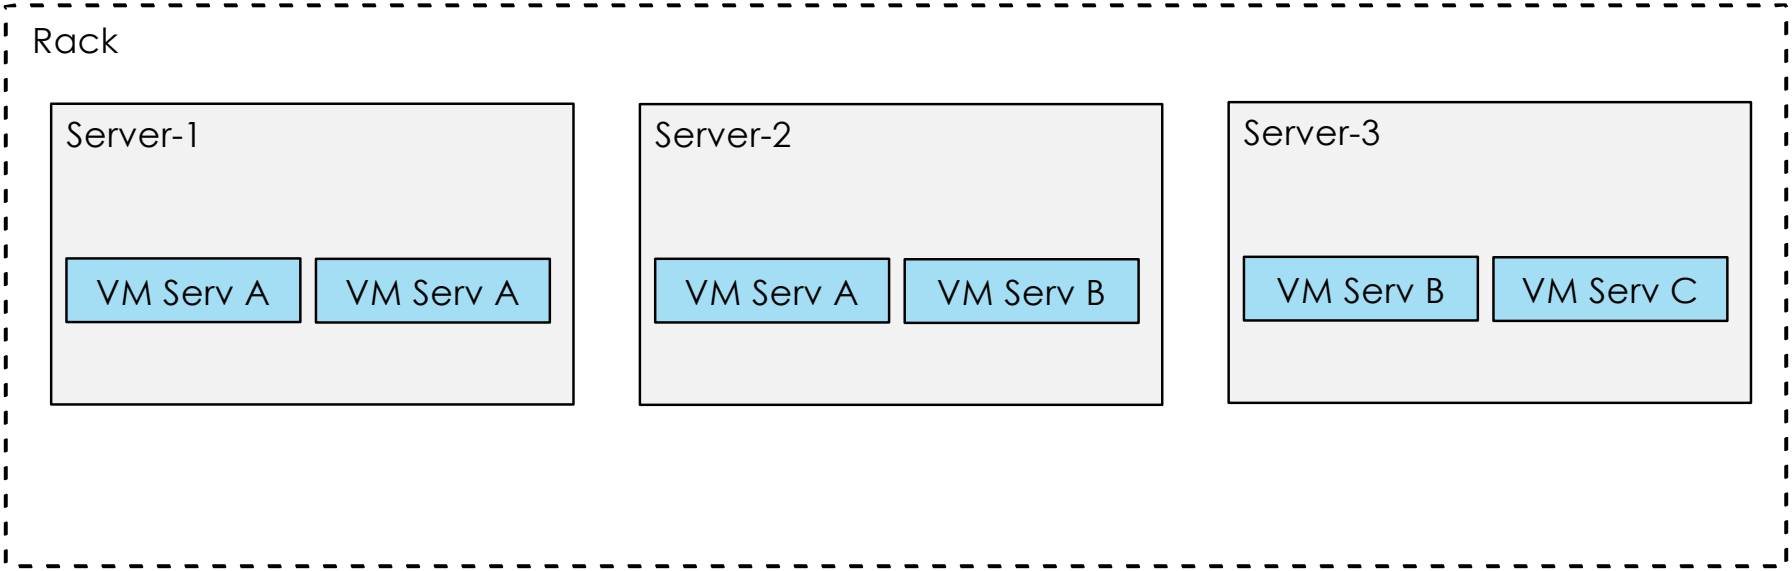

SmartOClock: Overclocking Management Framework for Cloud Platforms

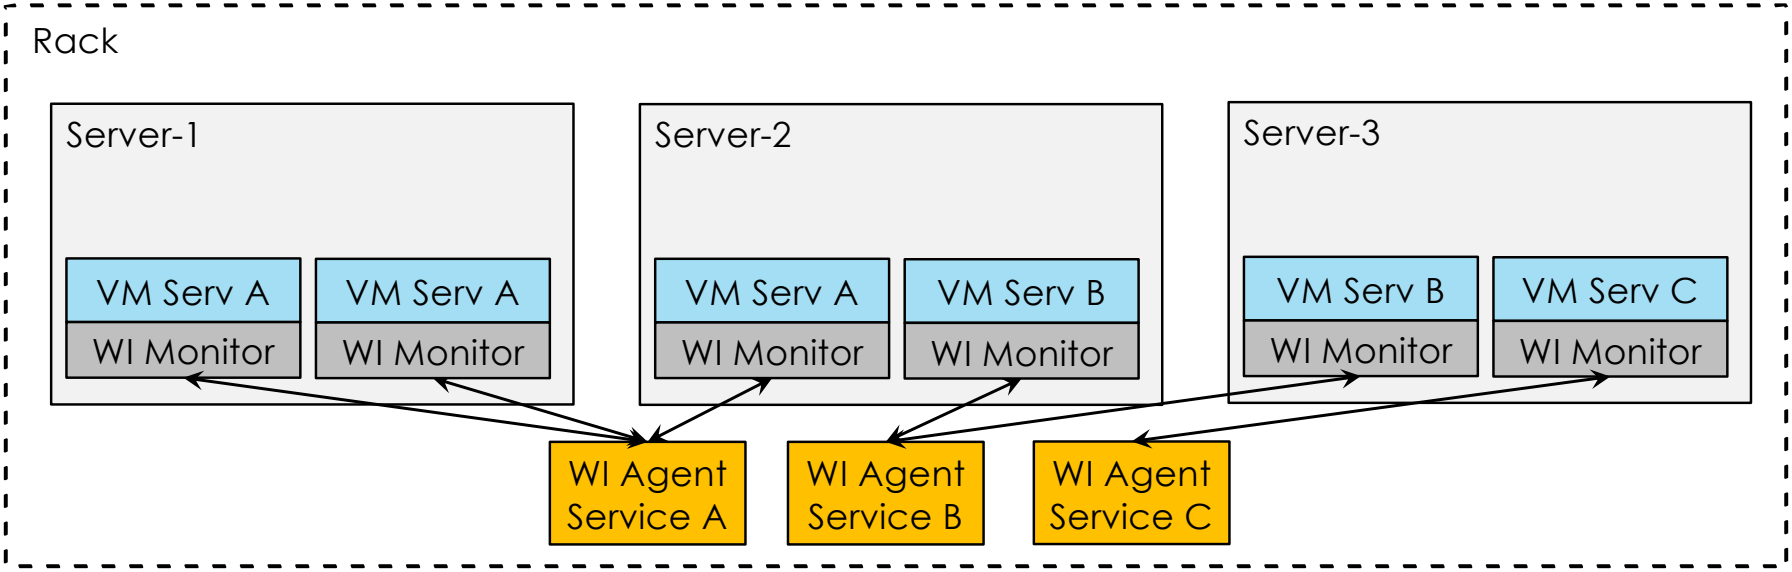

SmartOClock: Overclocking Management Framework for Cloud Platforms

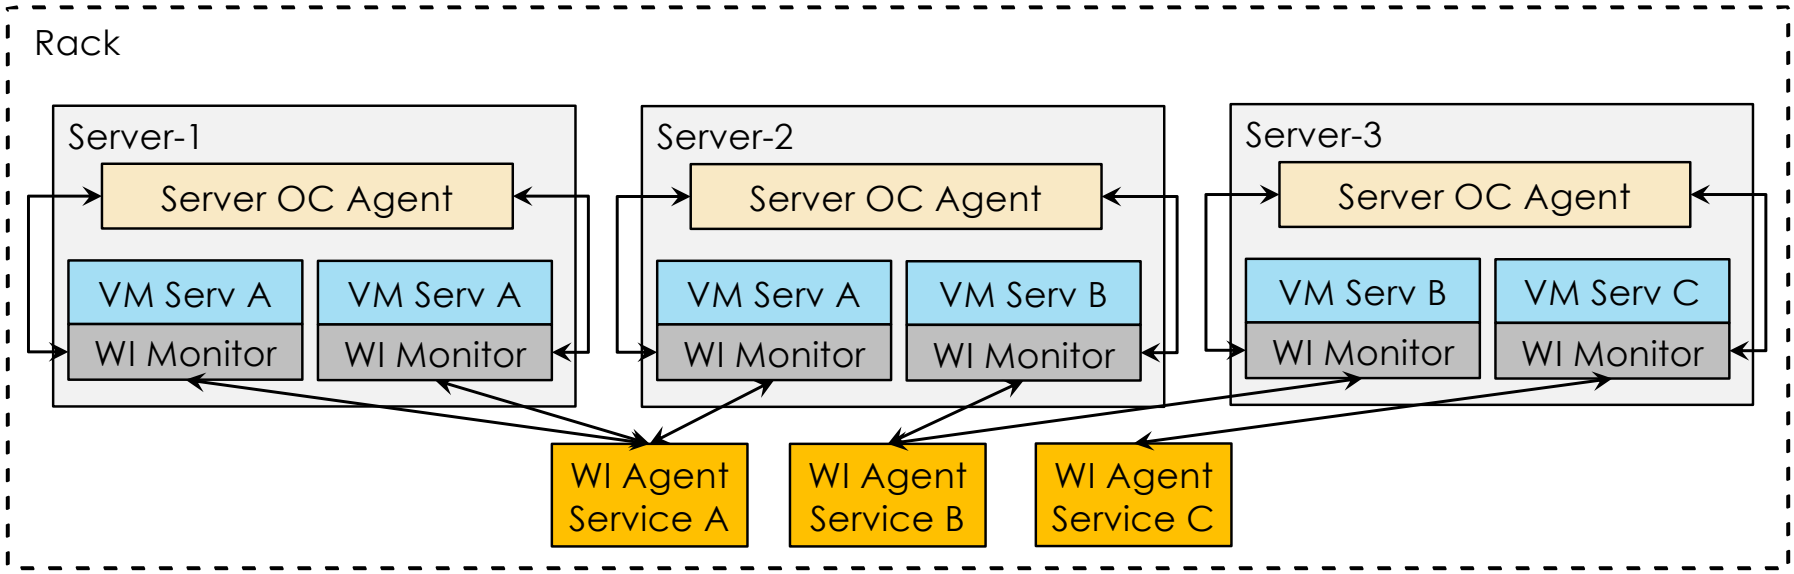

SmartOClock: Overclocking Management Framework for Cloud Platforms

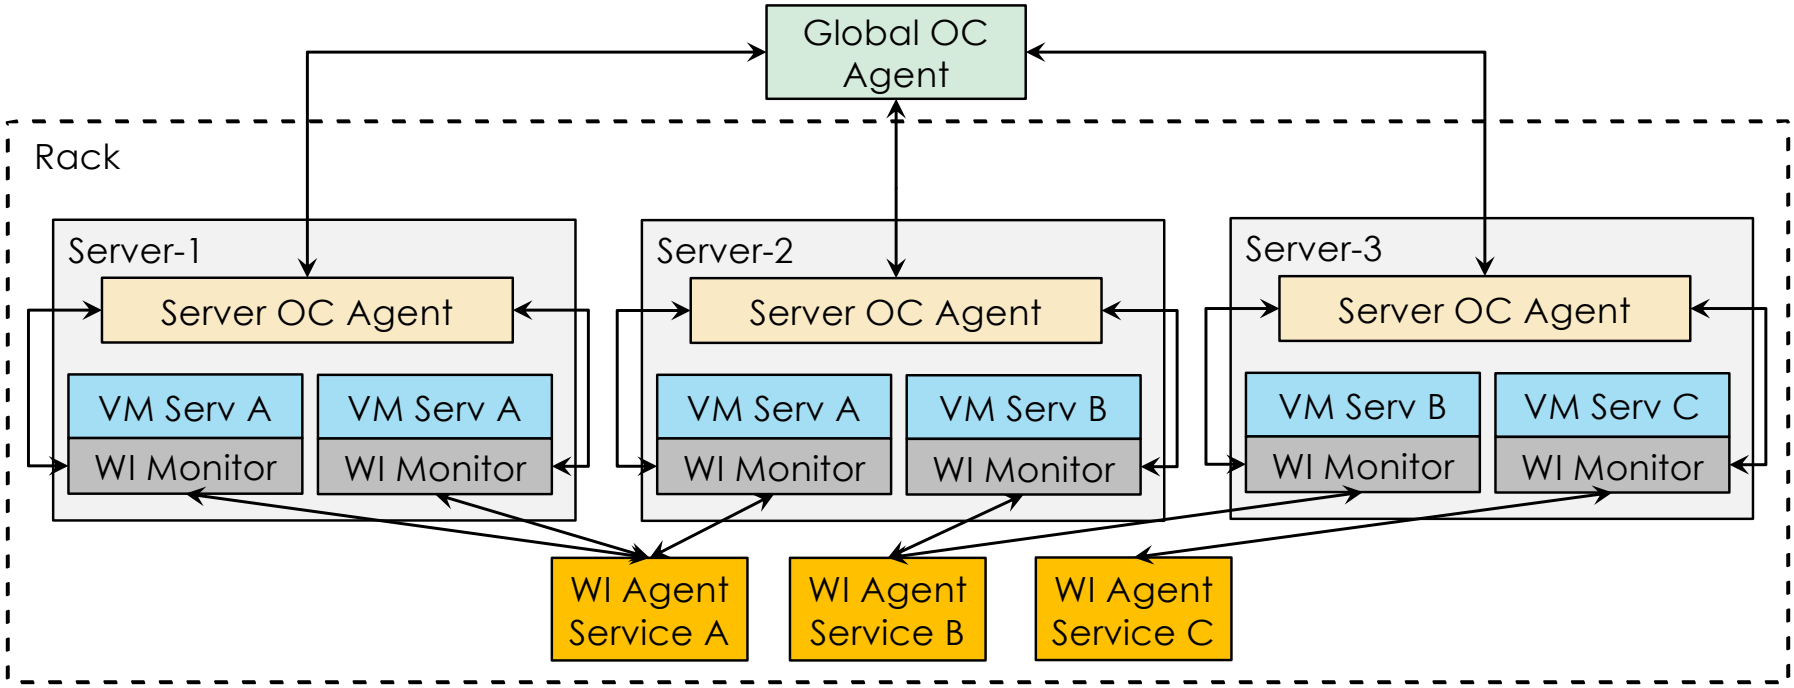

SmartOClock: Workload-Intelligent Overclocking

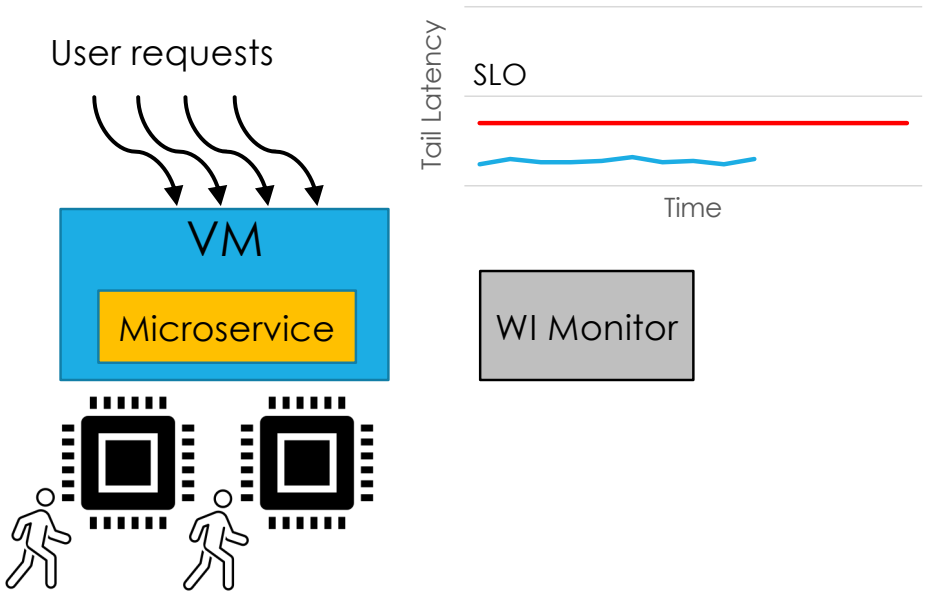

SmartOClock: Workload-Intelligent Overclocking

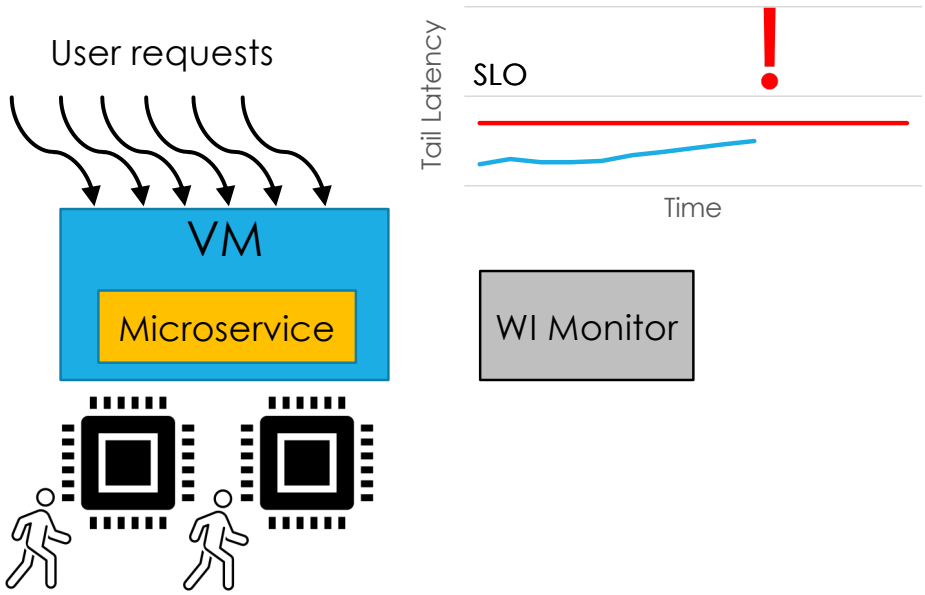

SmartOClock: Workload-Intelligent Overclocking

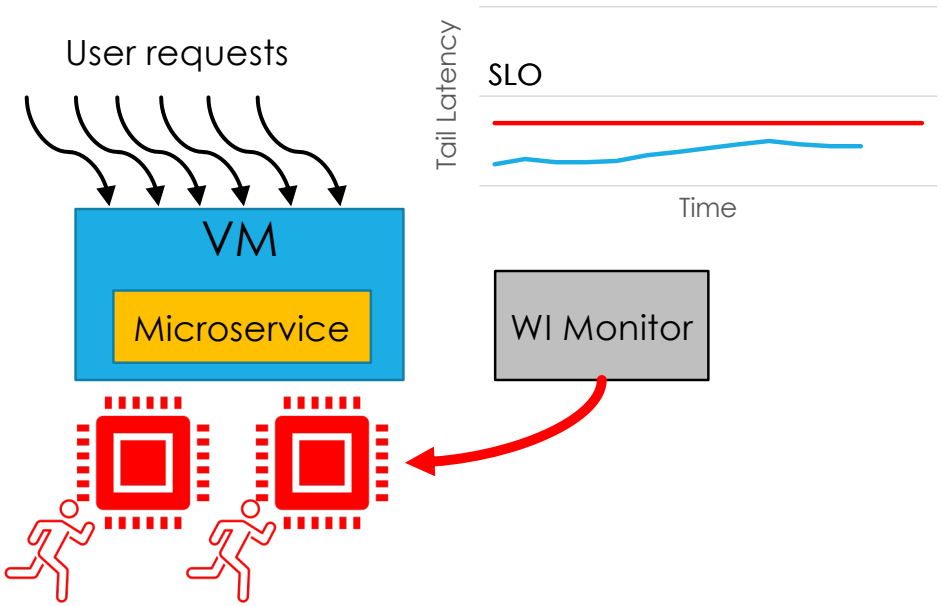

SmartOClock: Workload-Intelligent Overclocking

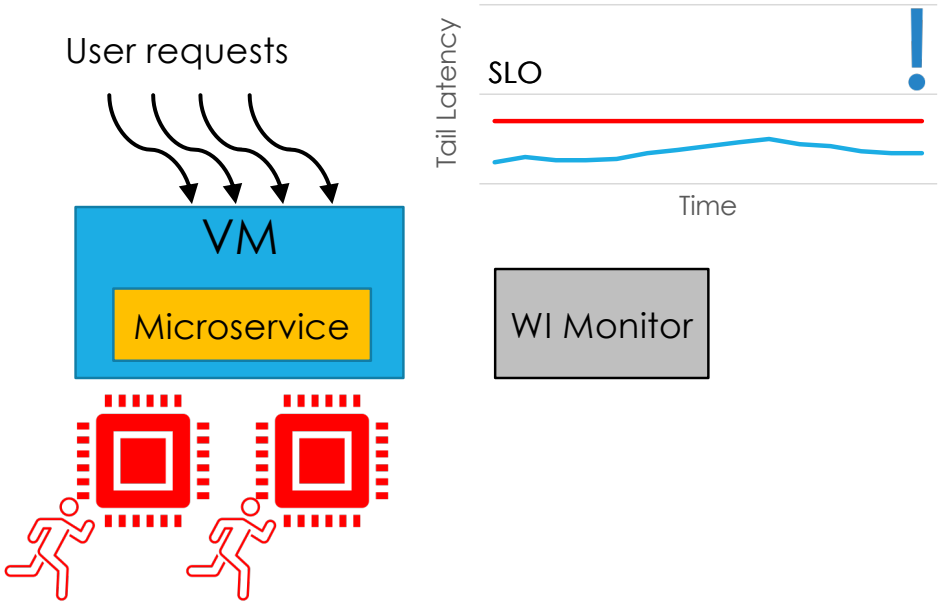

SmartOClock: Workload-Intelligent Overclocking

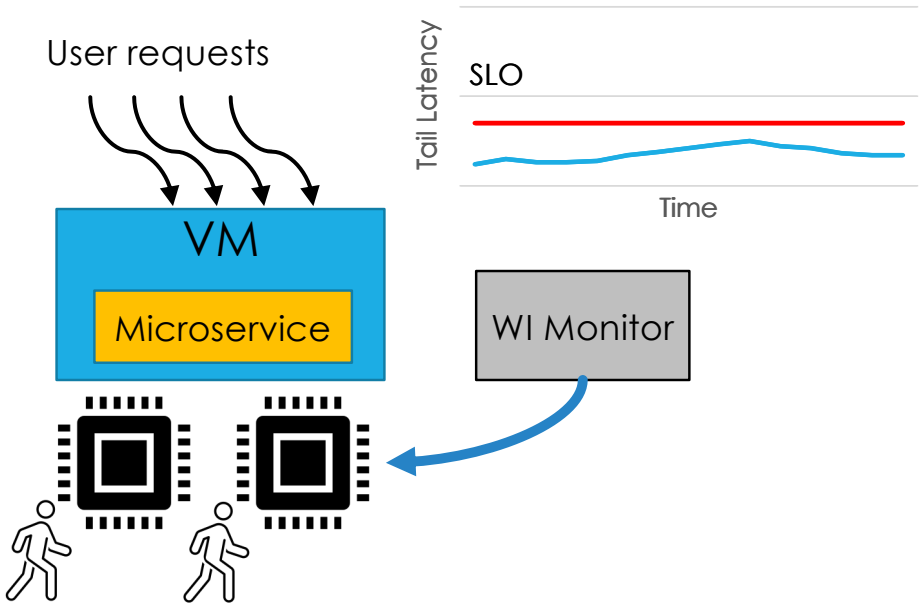

SmartOClock: Prediction-Based Overclocking Admission

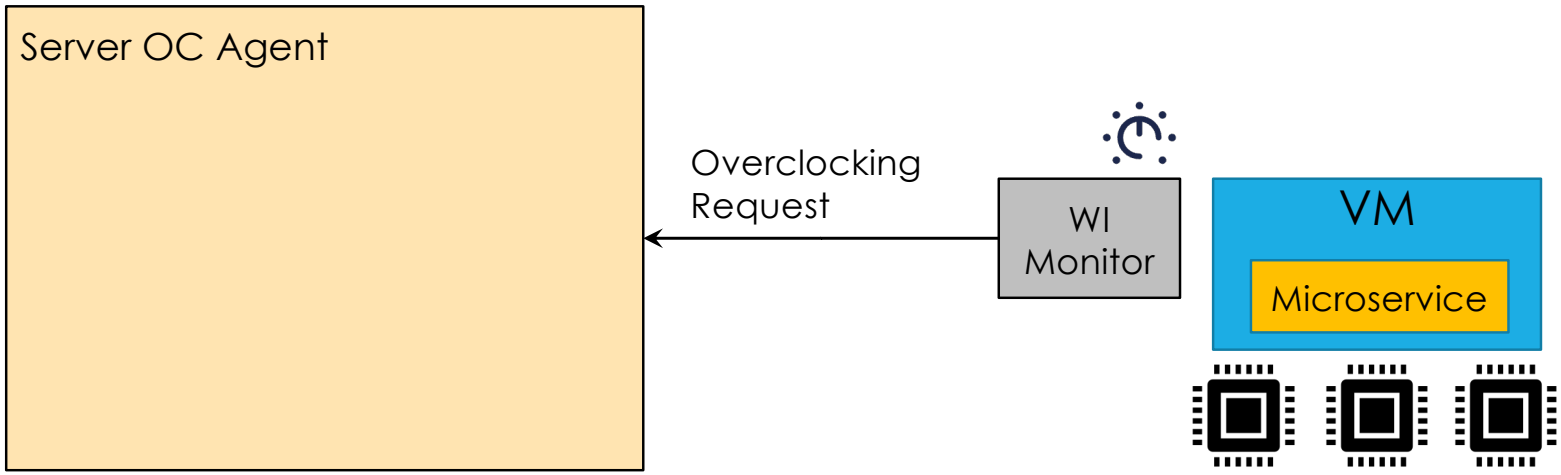

SmartOClock: Prediction-Based Overclocking Admission

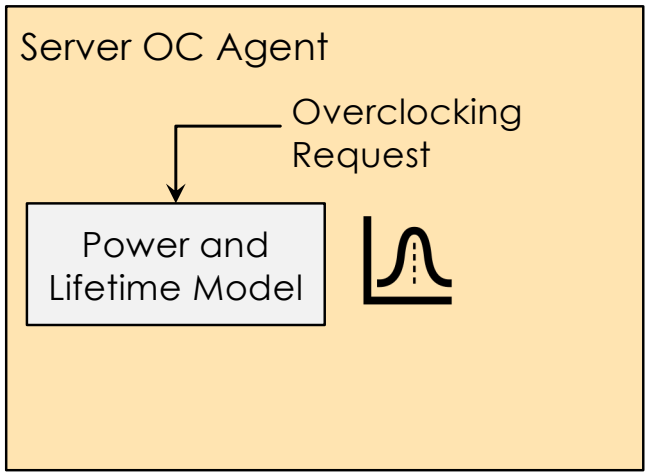

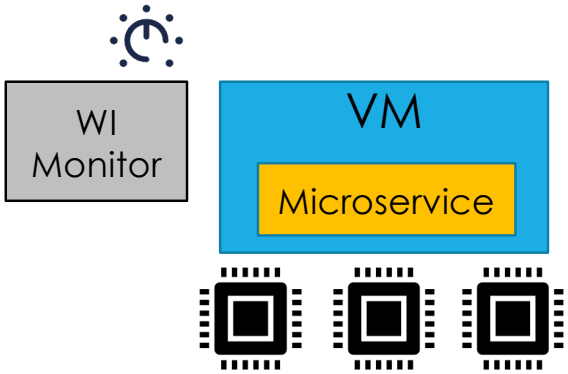

SmartOClock: Prediction-Based Overclocking Admission

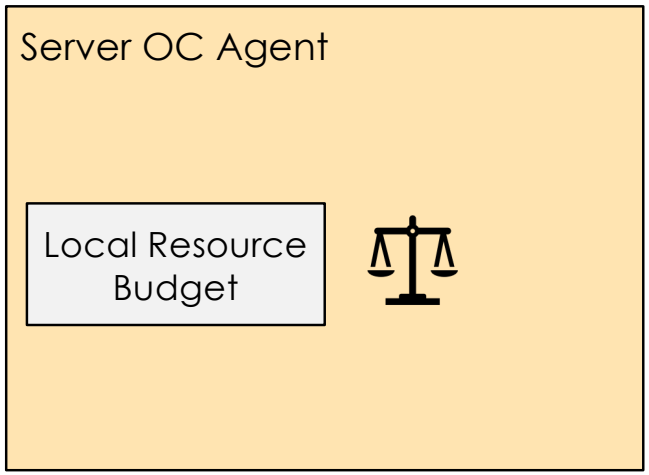

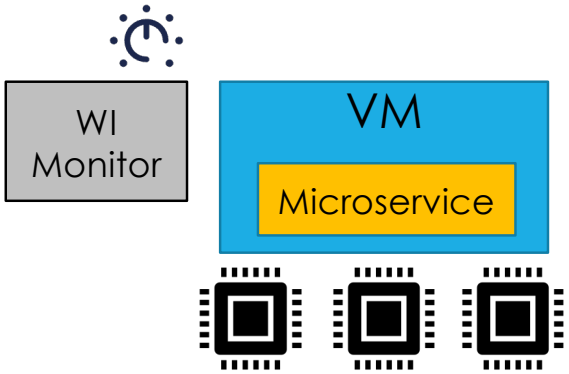

SmartOClock: Prediction-Based Overclocking Admission

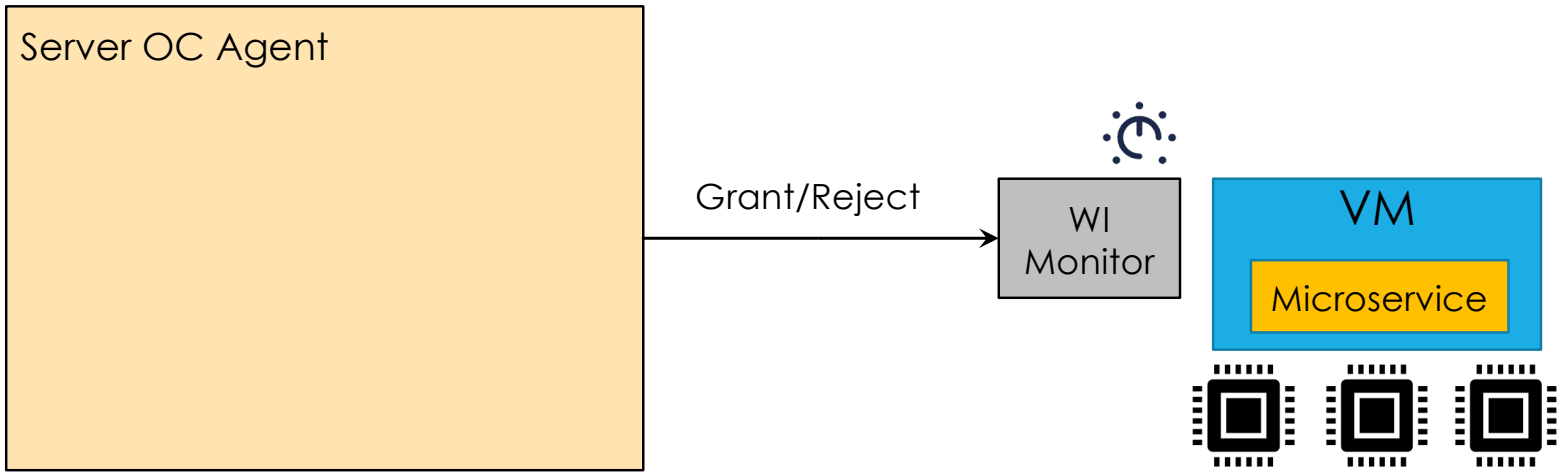

SmartOClock: Prediction-Based Overclocking Admission

Server OC Agent

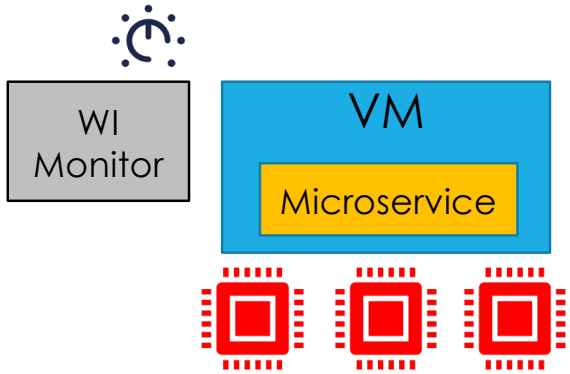

SmartOClock: Heterogeneous Overclock-Budgeting

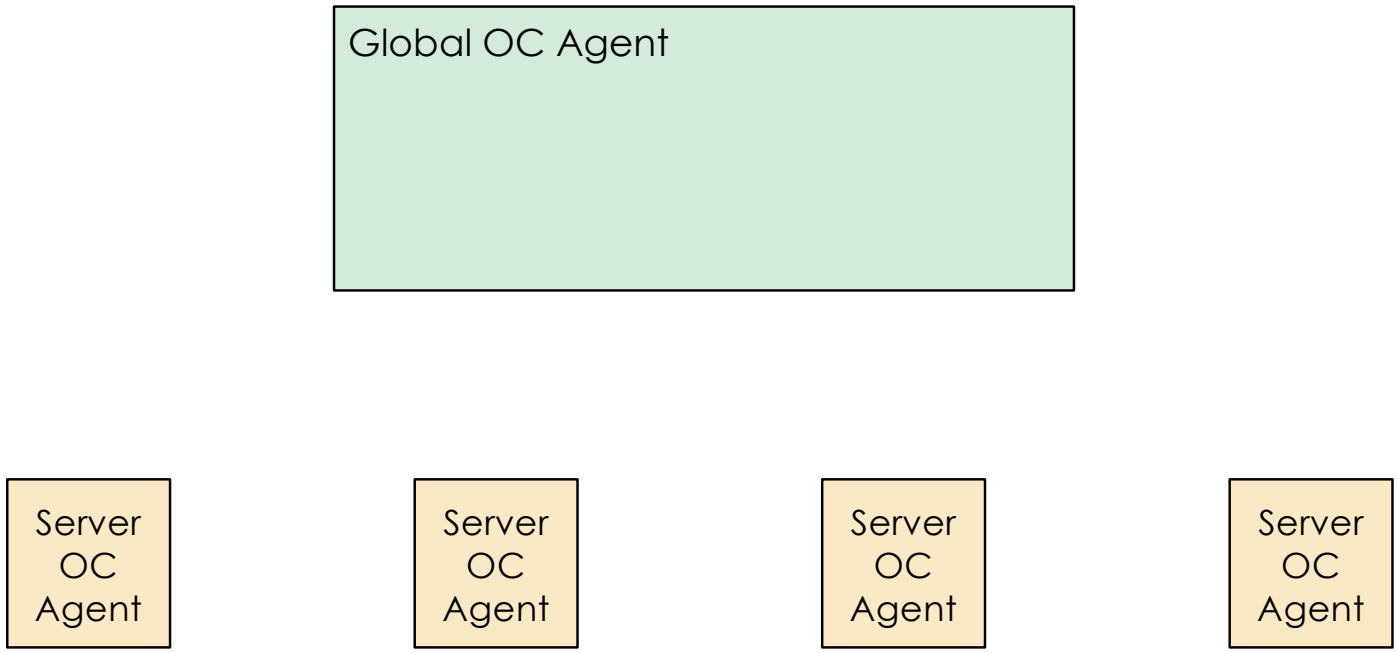

SmartOClock: Heterogeneous Overclock-Budgeting

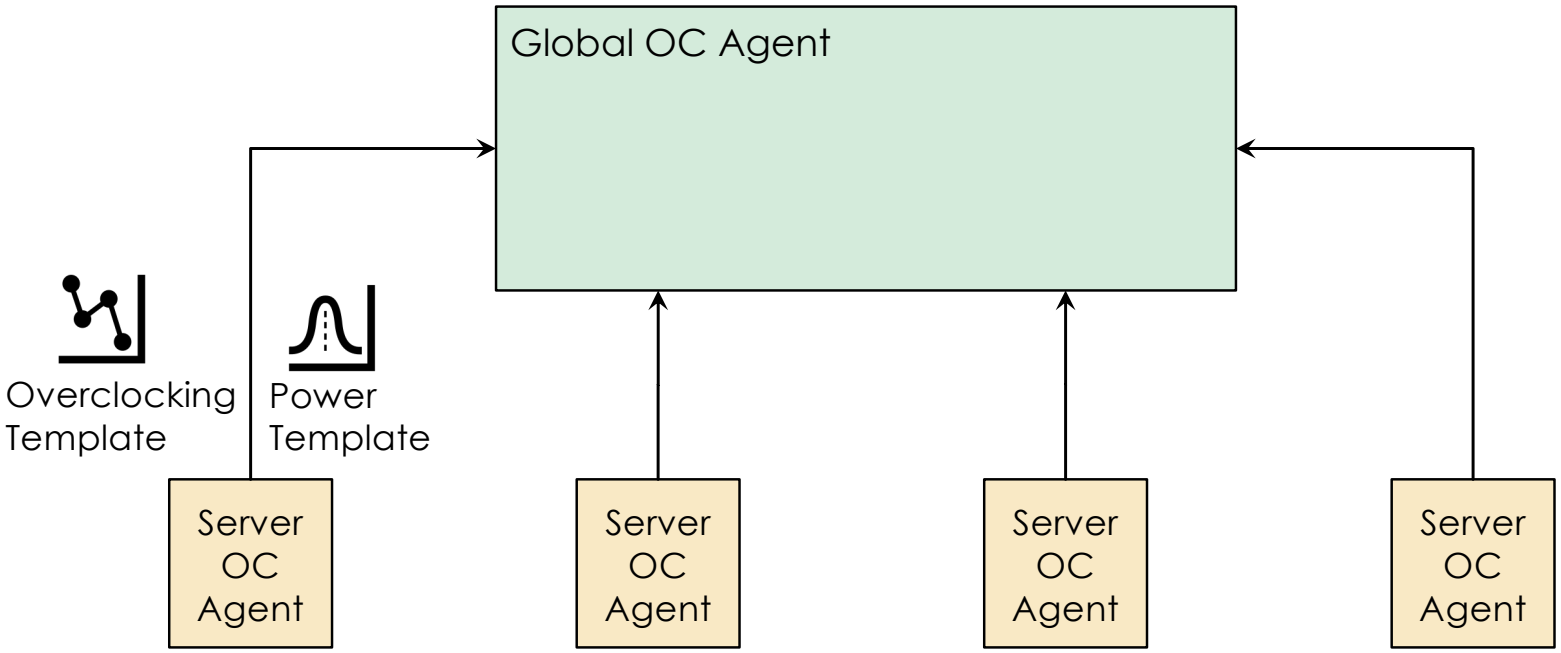

SmartOClock: Heterogeneous Overclock-Budgeting

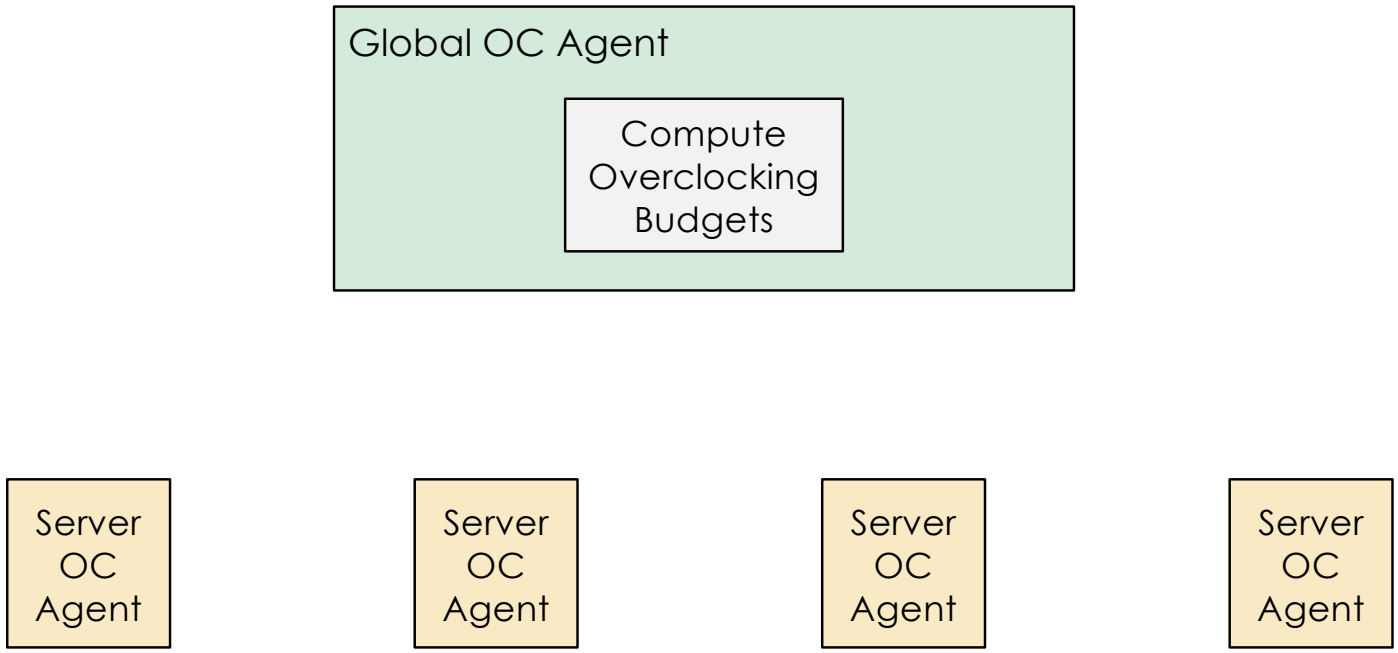

SmartOClock: Heterogeneous Overclock-Budgeting

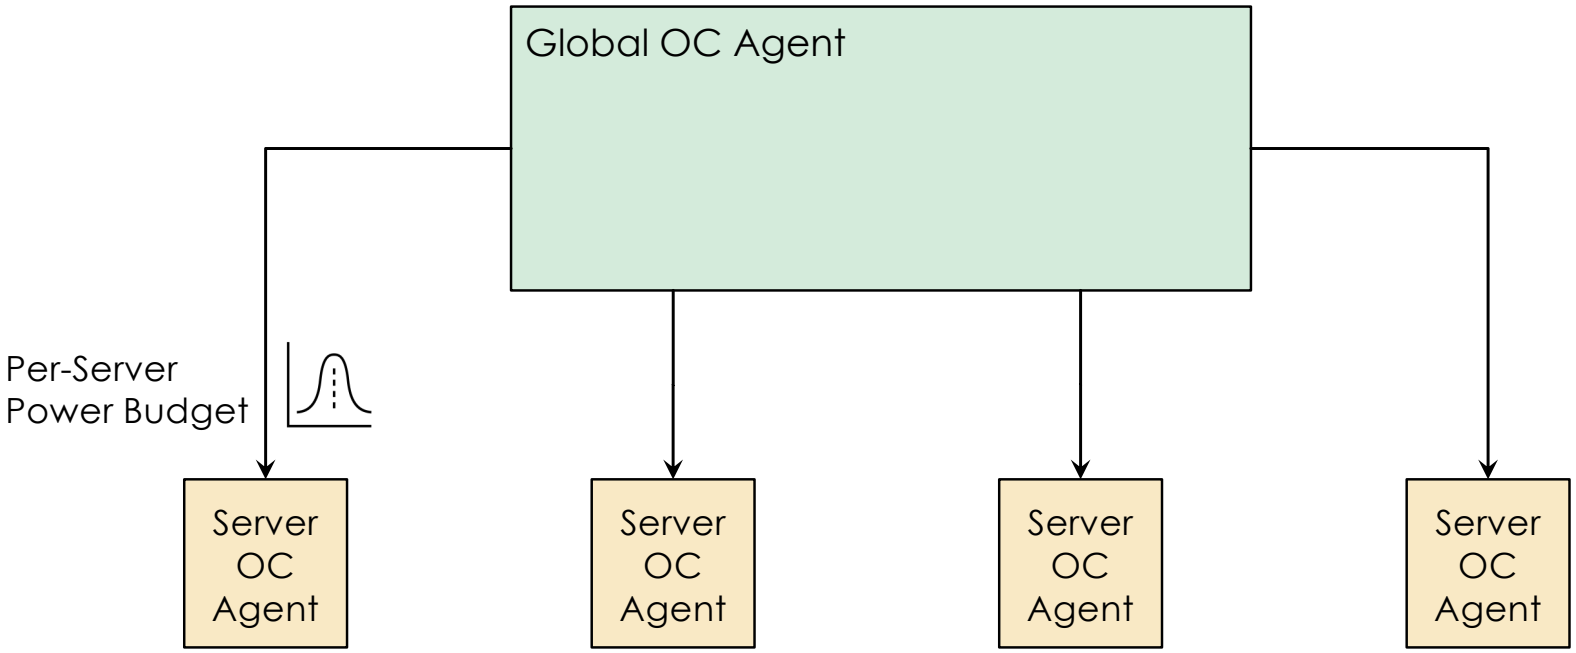

SmartOClock: Decentralized Budget Enforcement

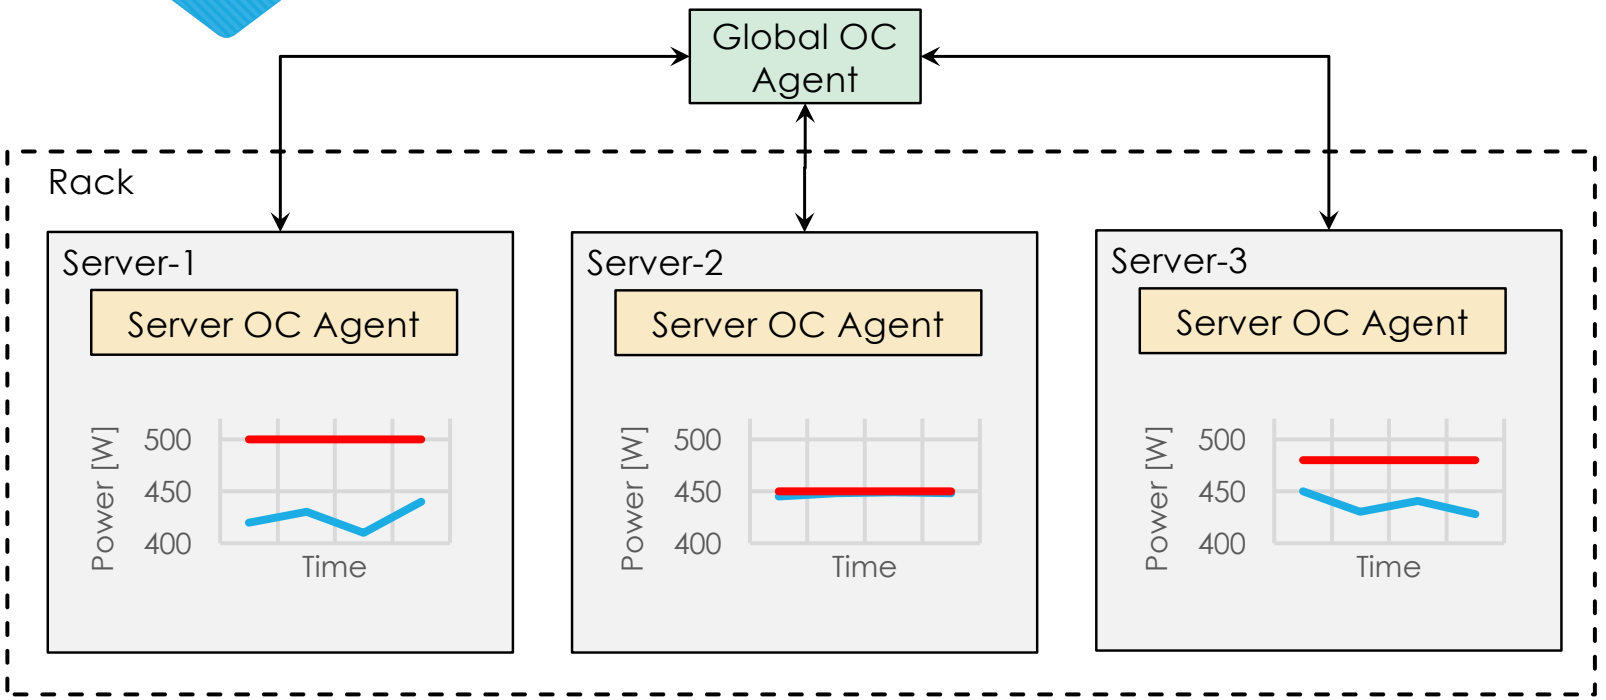

SmartOClock: Decentralized Budget Enforcement

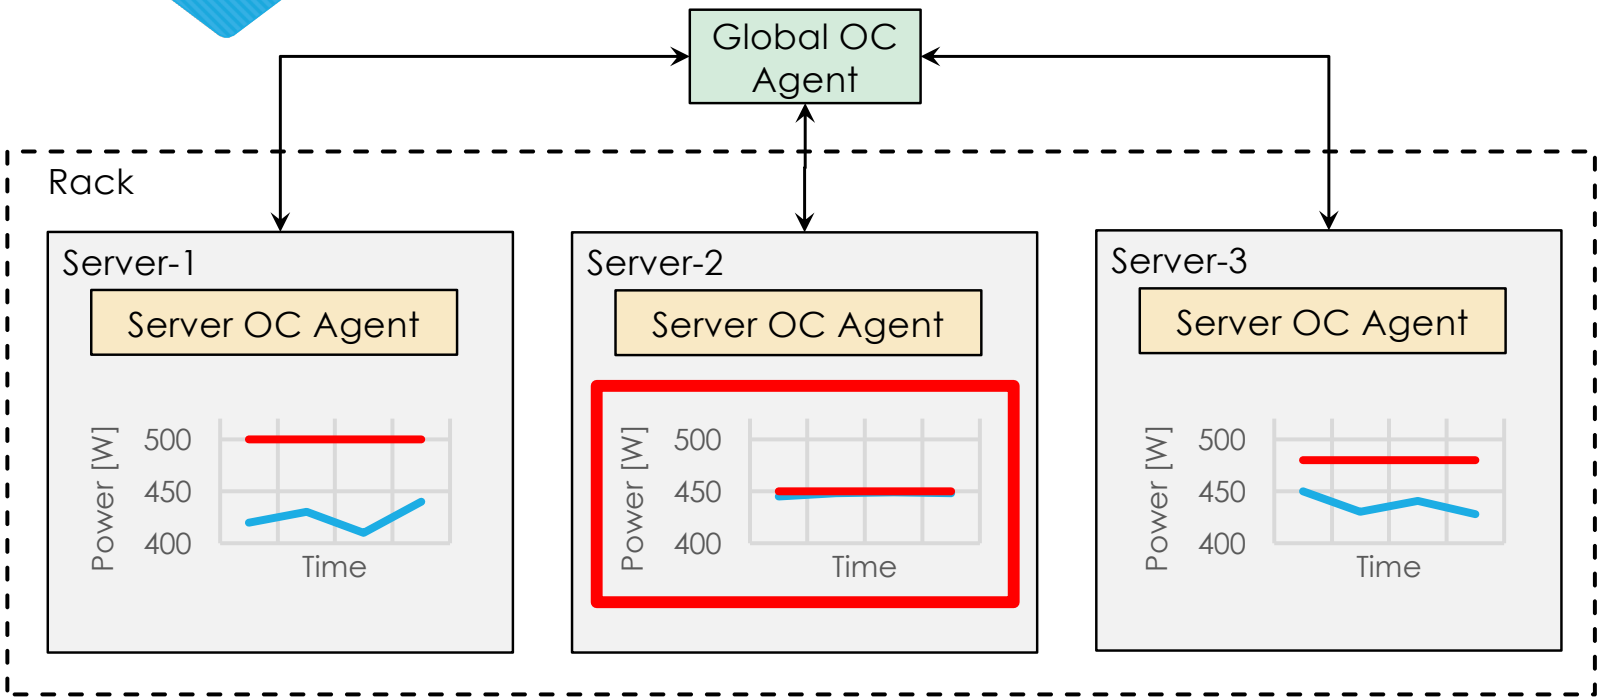

SmartOClock: Decentralized Budget Enforcement

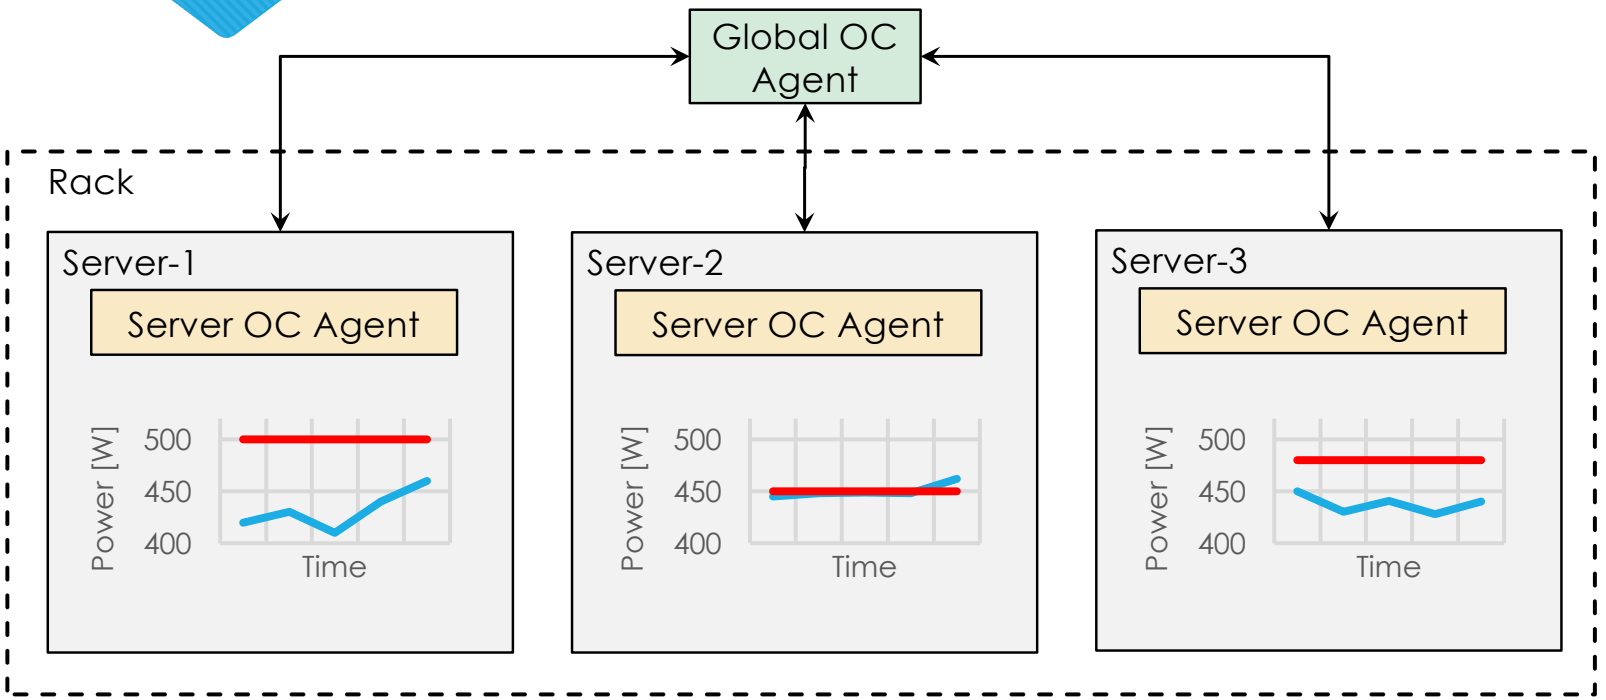

SmartOClock: Decentralized Budget Enforcement

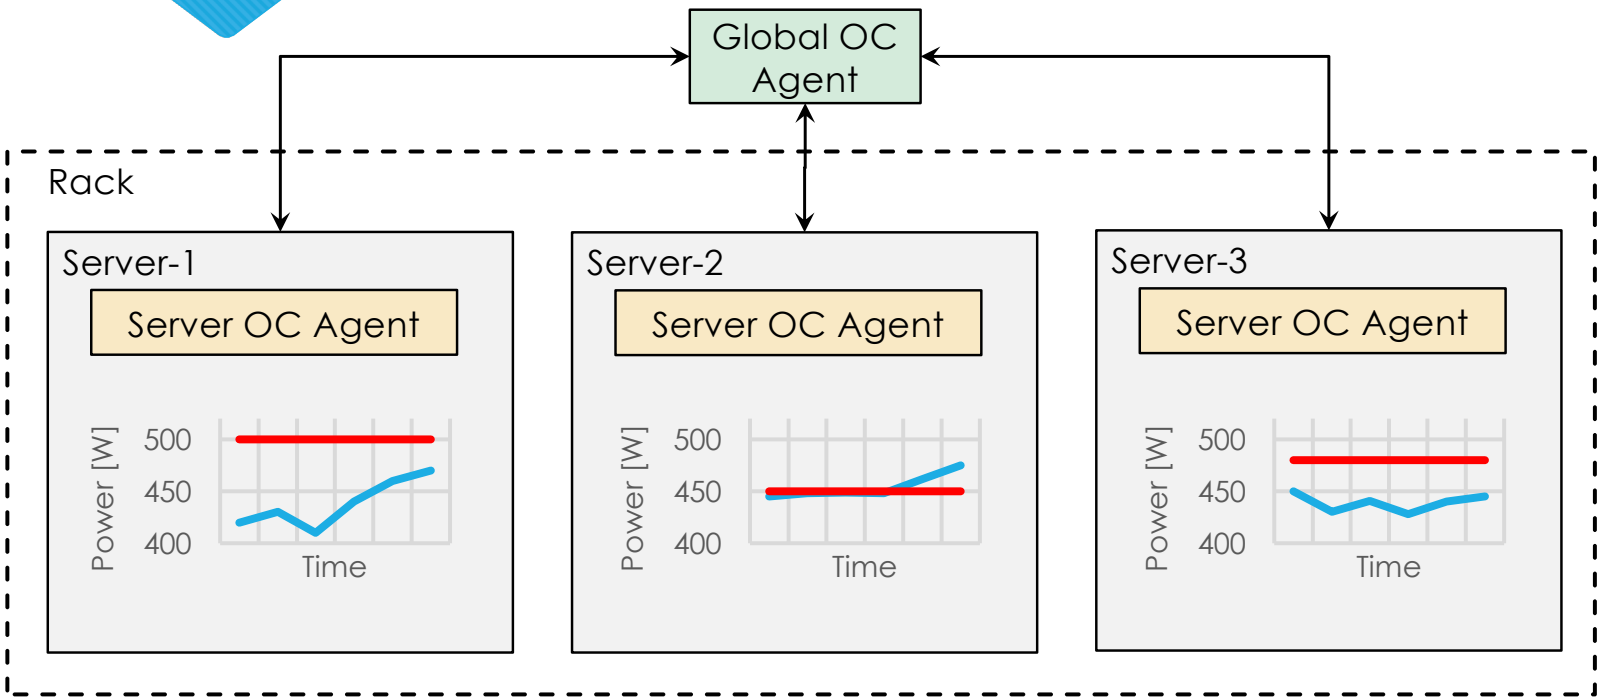

SmartOClock: Decentralized Budget Enforcement

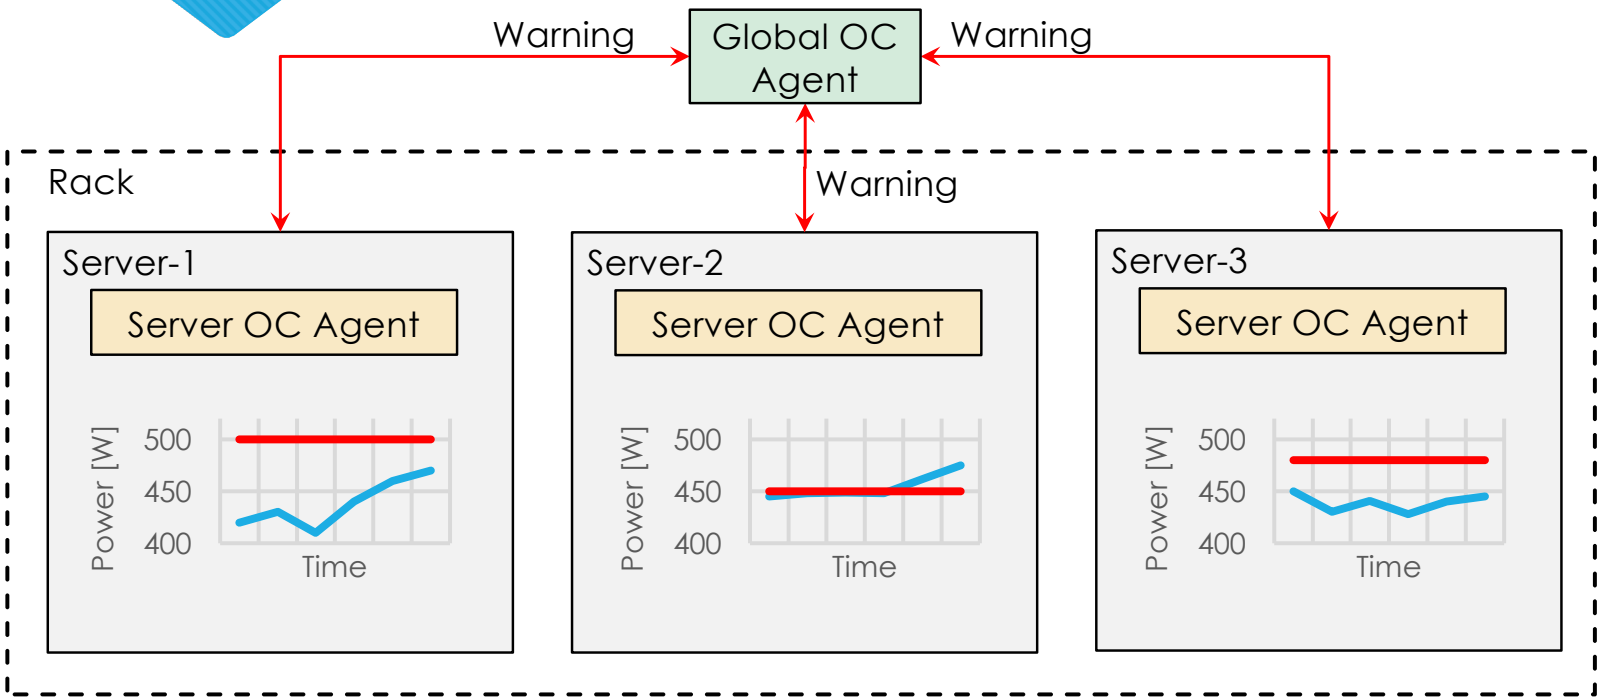

SmartOClock: Decentralized Budget Enforcement

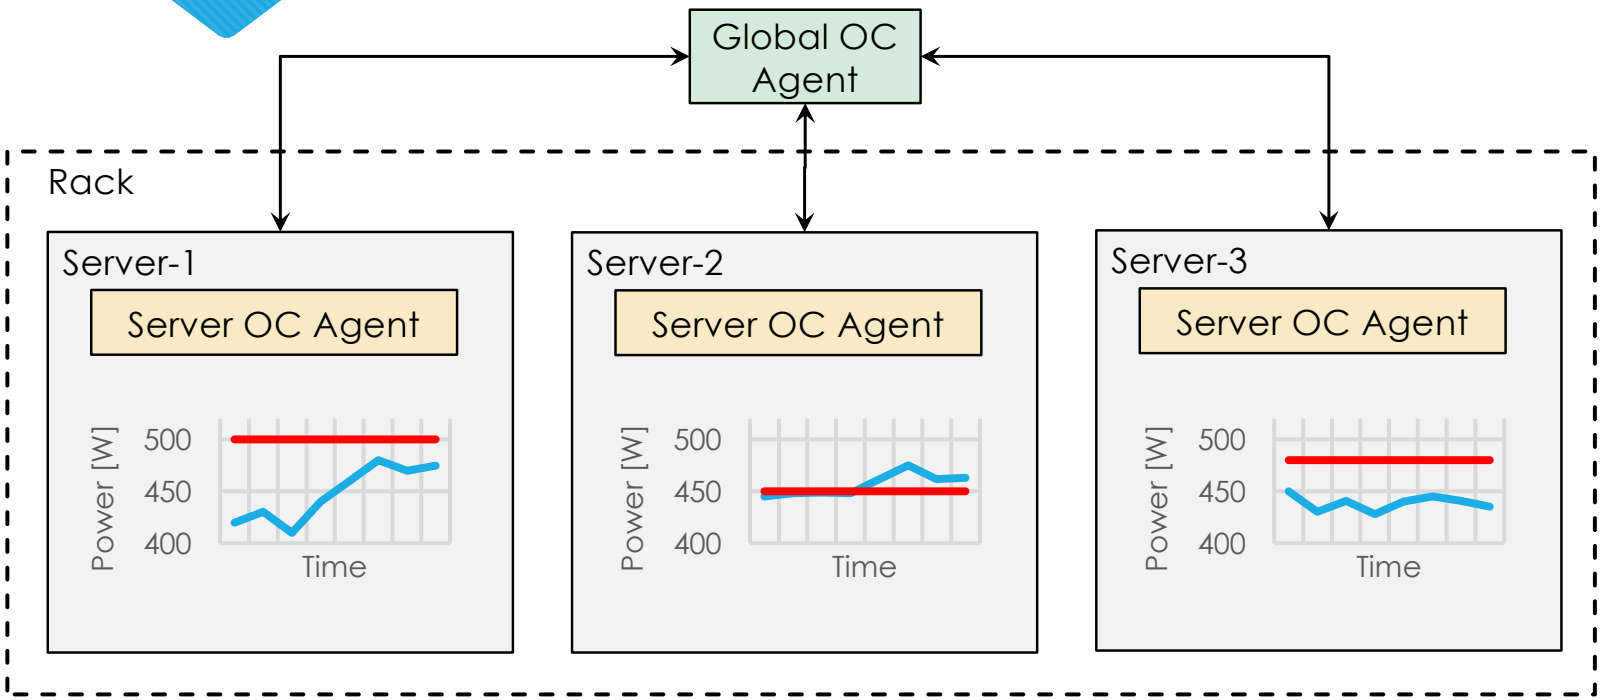

Outline of this talk

- Challenges and opportunities for overclocking
- SmartOClock
- **Evaluation results**

Evaluation Setup

- Experiments on two racks with 36 overclockable servers
  - Mimic power consumption of a production rack and its servers
  - Latency critical microservices + batch ML training

	Max Frequency	Scaling
Baseline	3.3 GHz	Fixed
ScaleOut	3.3 GHz	Auto-scale
ScaleUp	4.0 GHz	Fixed
SmartOClock	4.0 GHz	Auto-scale

Tail Latency Reduction

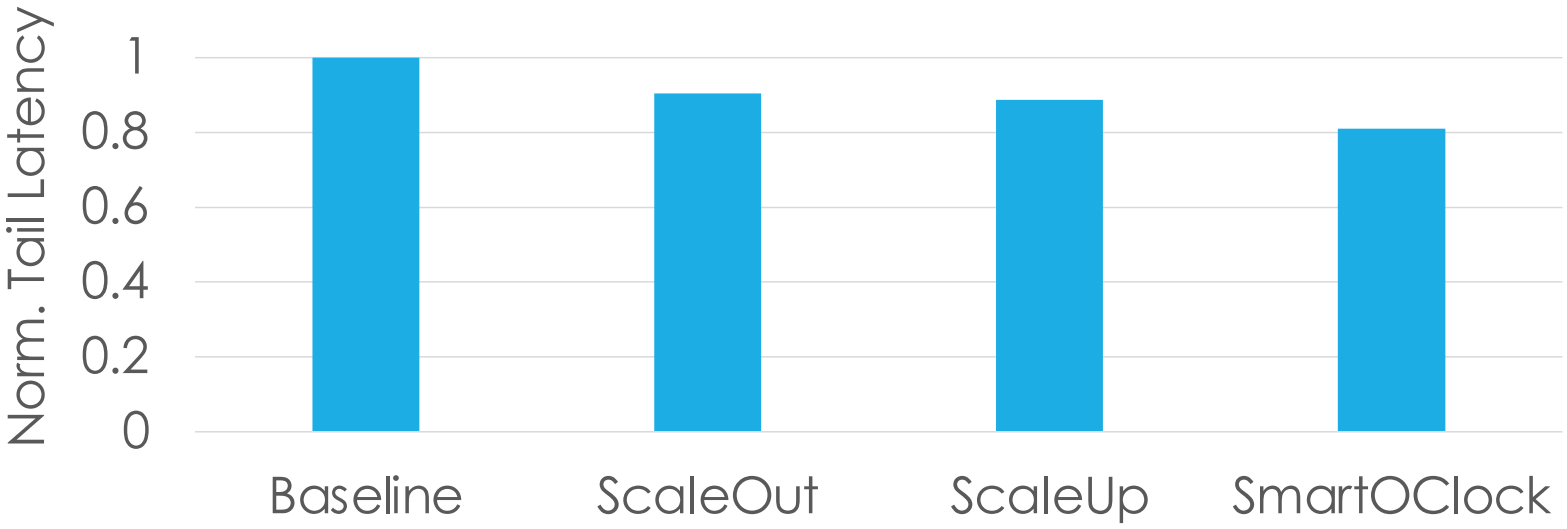

Tail Latency Reduction

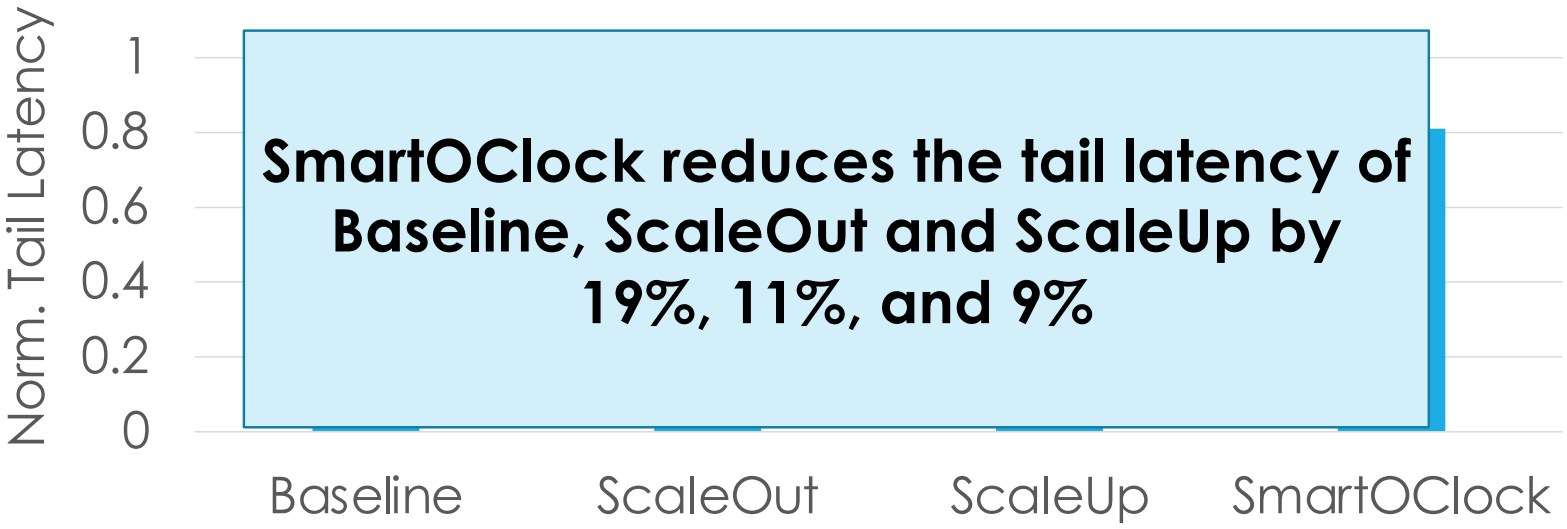

Workload Cost Reduction

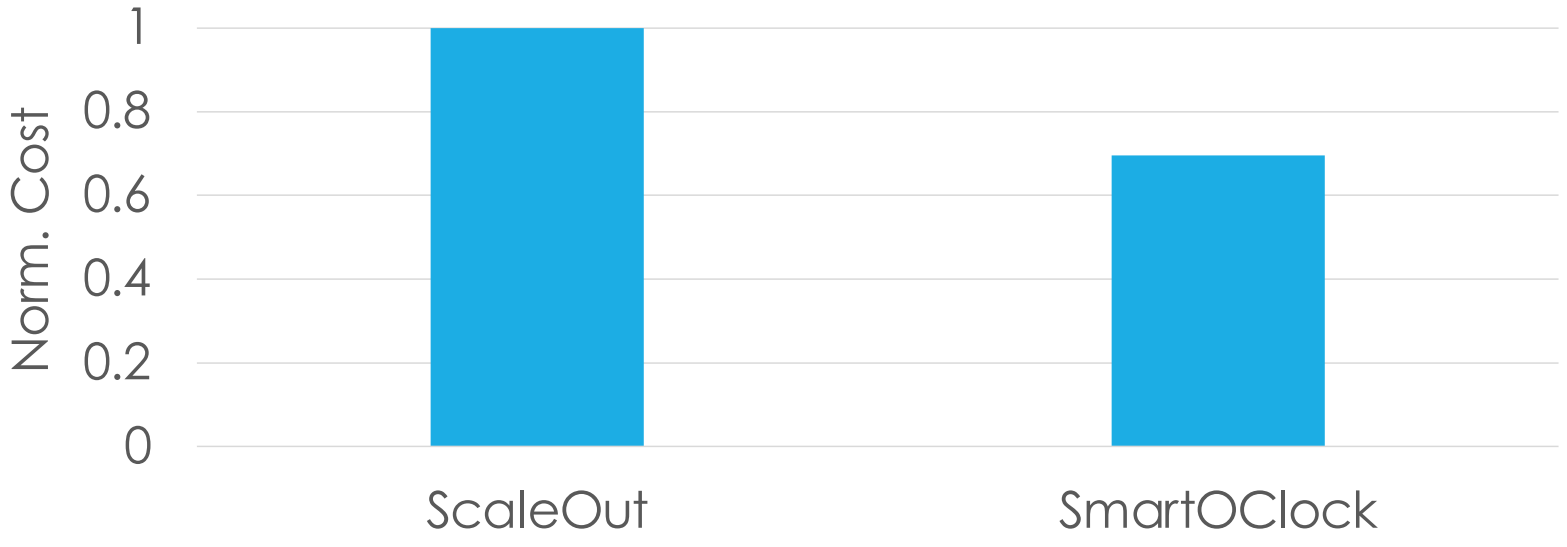

Workload Cost Reduction

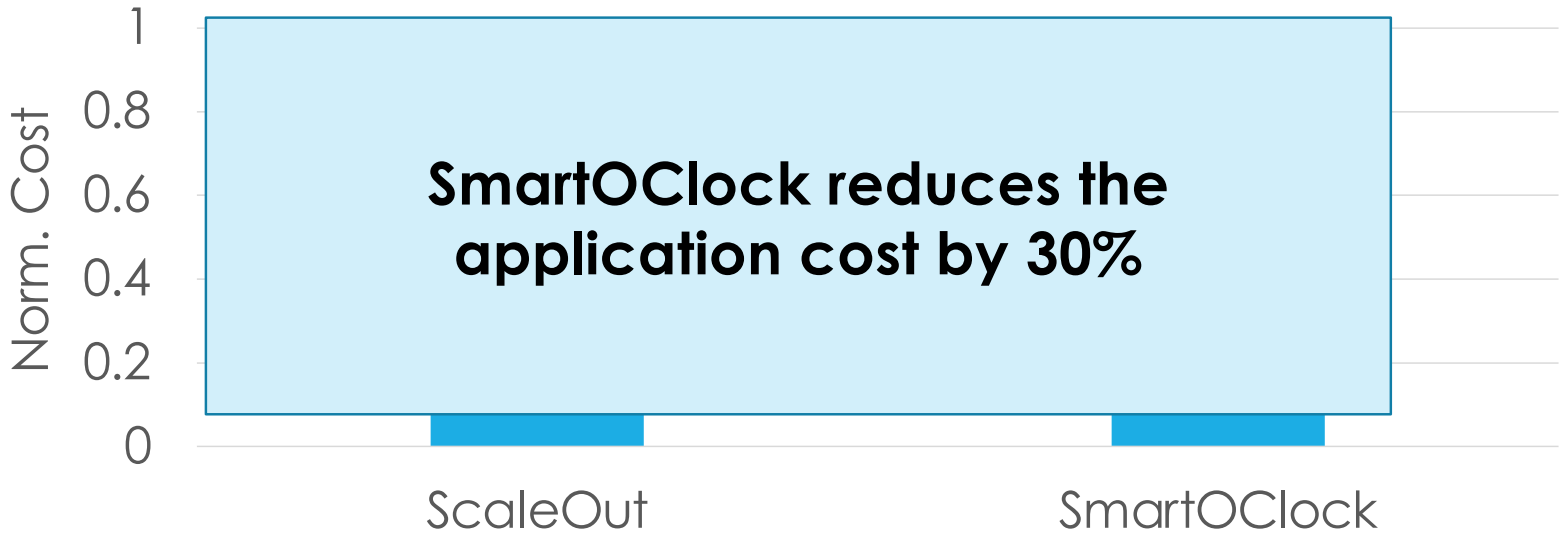

Energy Consumption Reduction

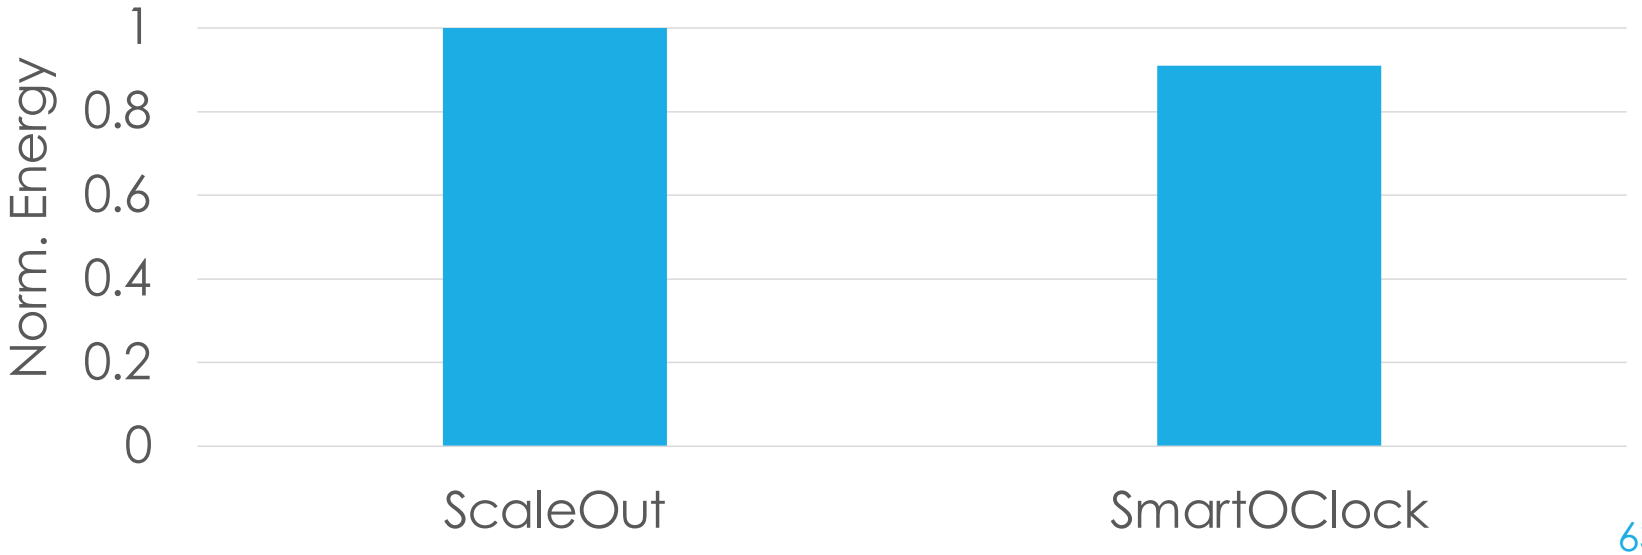

Energy Consumption Reduction

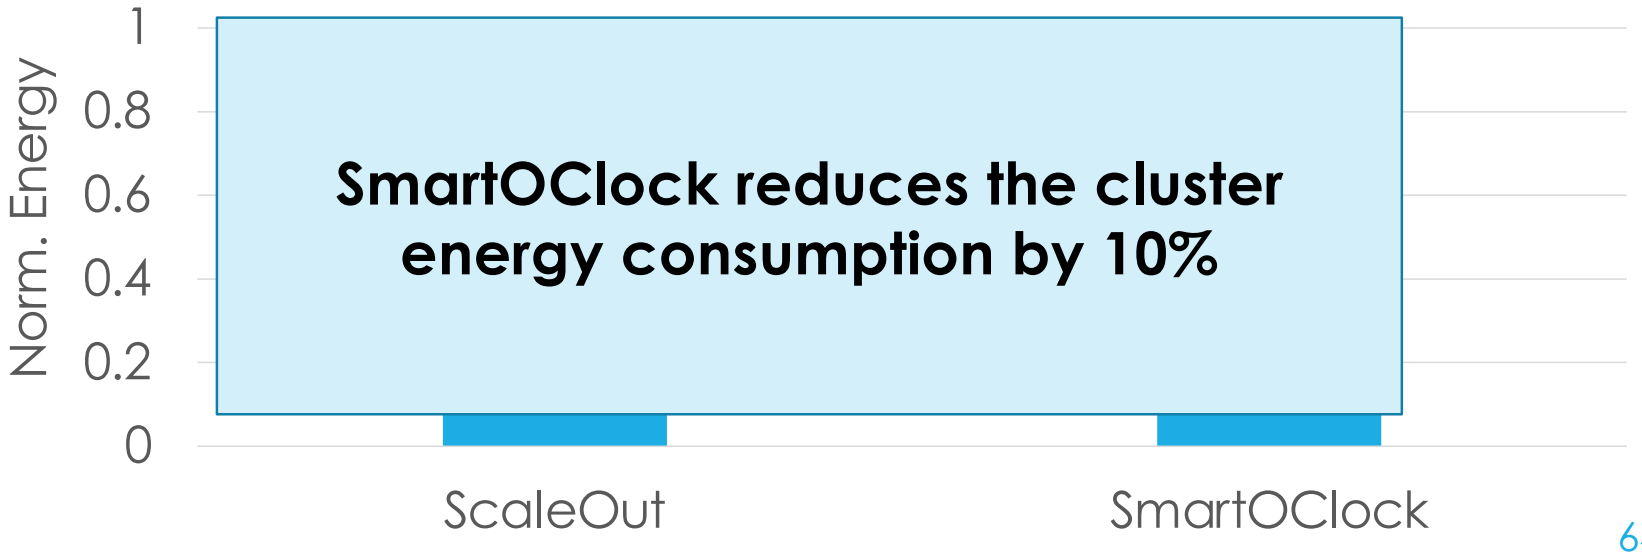

Production Workload

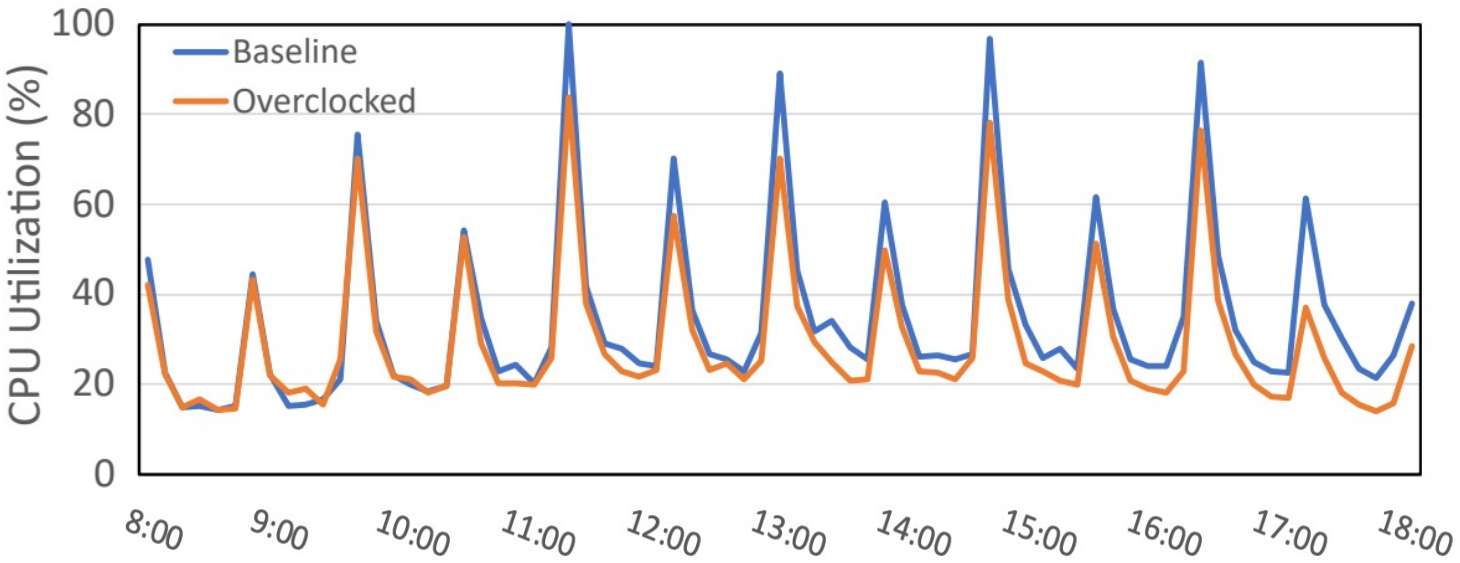

Production Workload

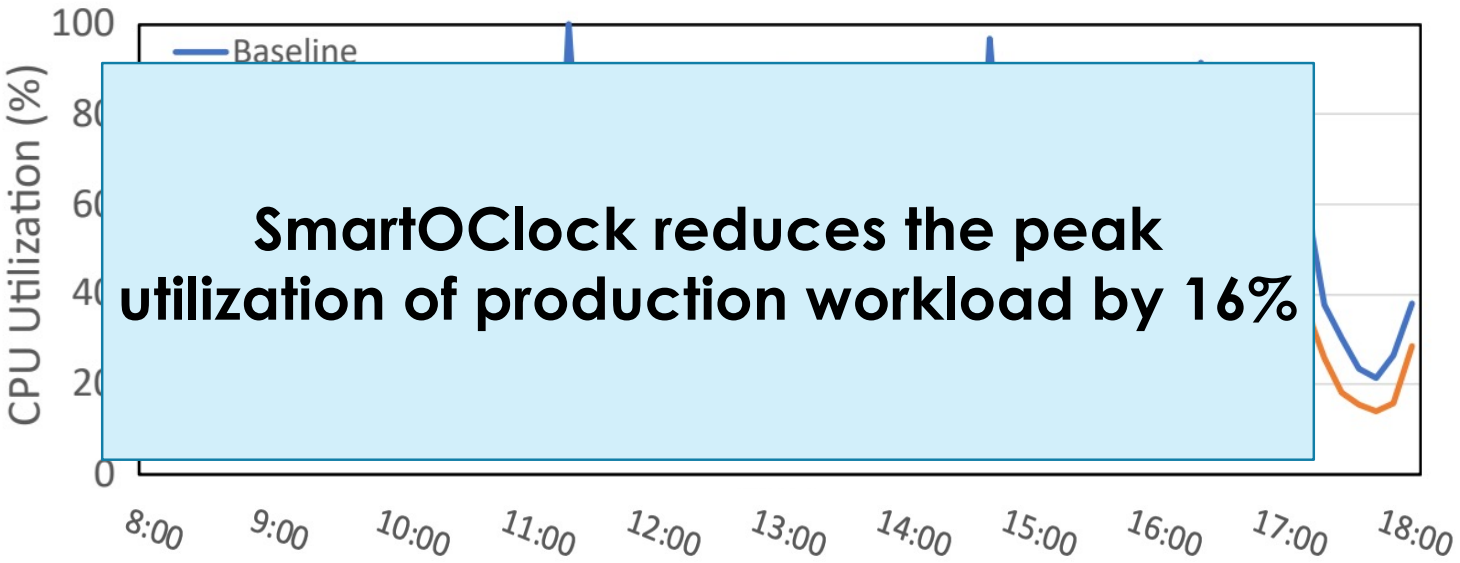

SmartOClock reduces the peak utilization of production workload by 16%

Conclusion

- Overclocking gives benefits but not for free
- **SmartOClock**: cloud overclocking management
- Evaluation shows improvements:
  - Tail latency by 9%
  - Cost by 30%
  - Energy consumption by 10%
- Lessons from production

SmartOClock: Workload- and Risk-Aware Overclocking in the Cloud

ISCA 2024

Jovan Stojkovic*, Pulkit Misra, Íñigo Goiri, Sam Whitlock, Esha Choukse, Mayukh Das, Chetan Bansal, Jason Lee, Zoey Sun, Haoran Qiu*, Reed Zimmermann†, Savyasachi Samal, Brijesh Warrier, Ashish Raniwala, Ricardo Bianchini

Microsoft, *University of Illinois at Urbana-Champaign, †University of Texas Austin