Darwin: Flexible Learning-based CDN Caching

Jiayi Chen, Nihal Sharma, Tarannum Khan, Shu Liu, Brian Chang,
Aditya Akella, Sanjay Shakkottai, Ramesh K. Sitaraman

Sigcomm 2023, New York City

Content Delivery Networks (CDNs)

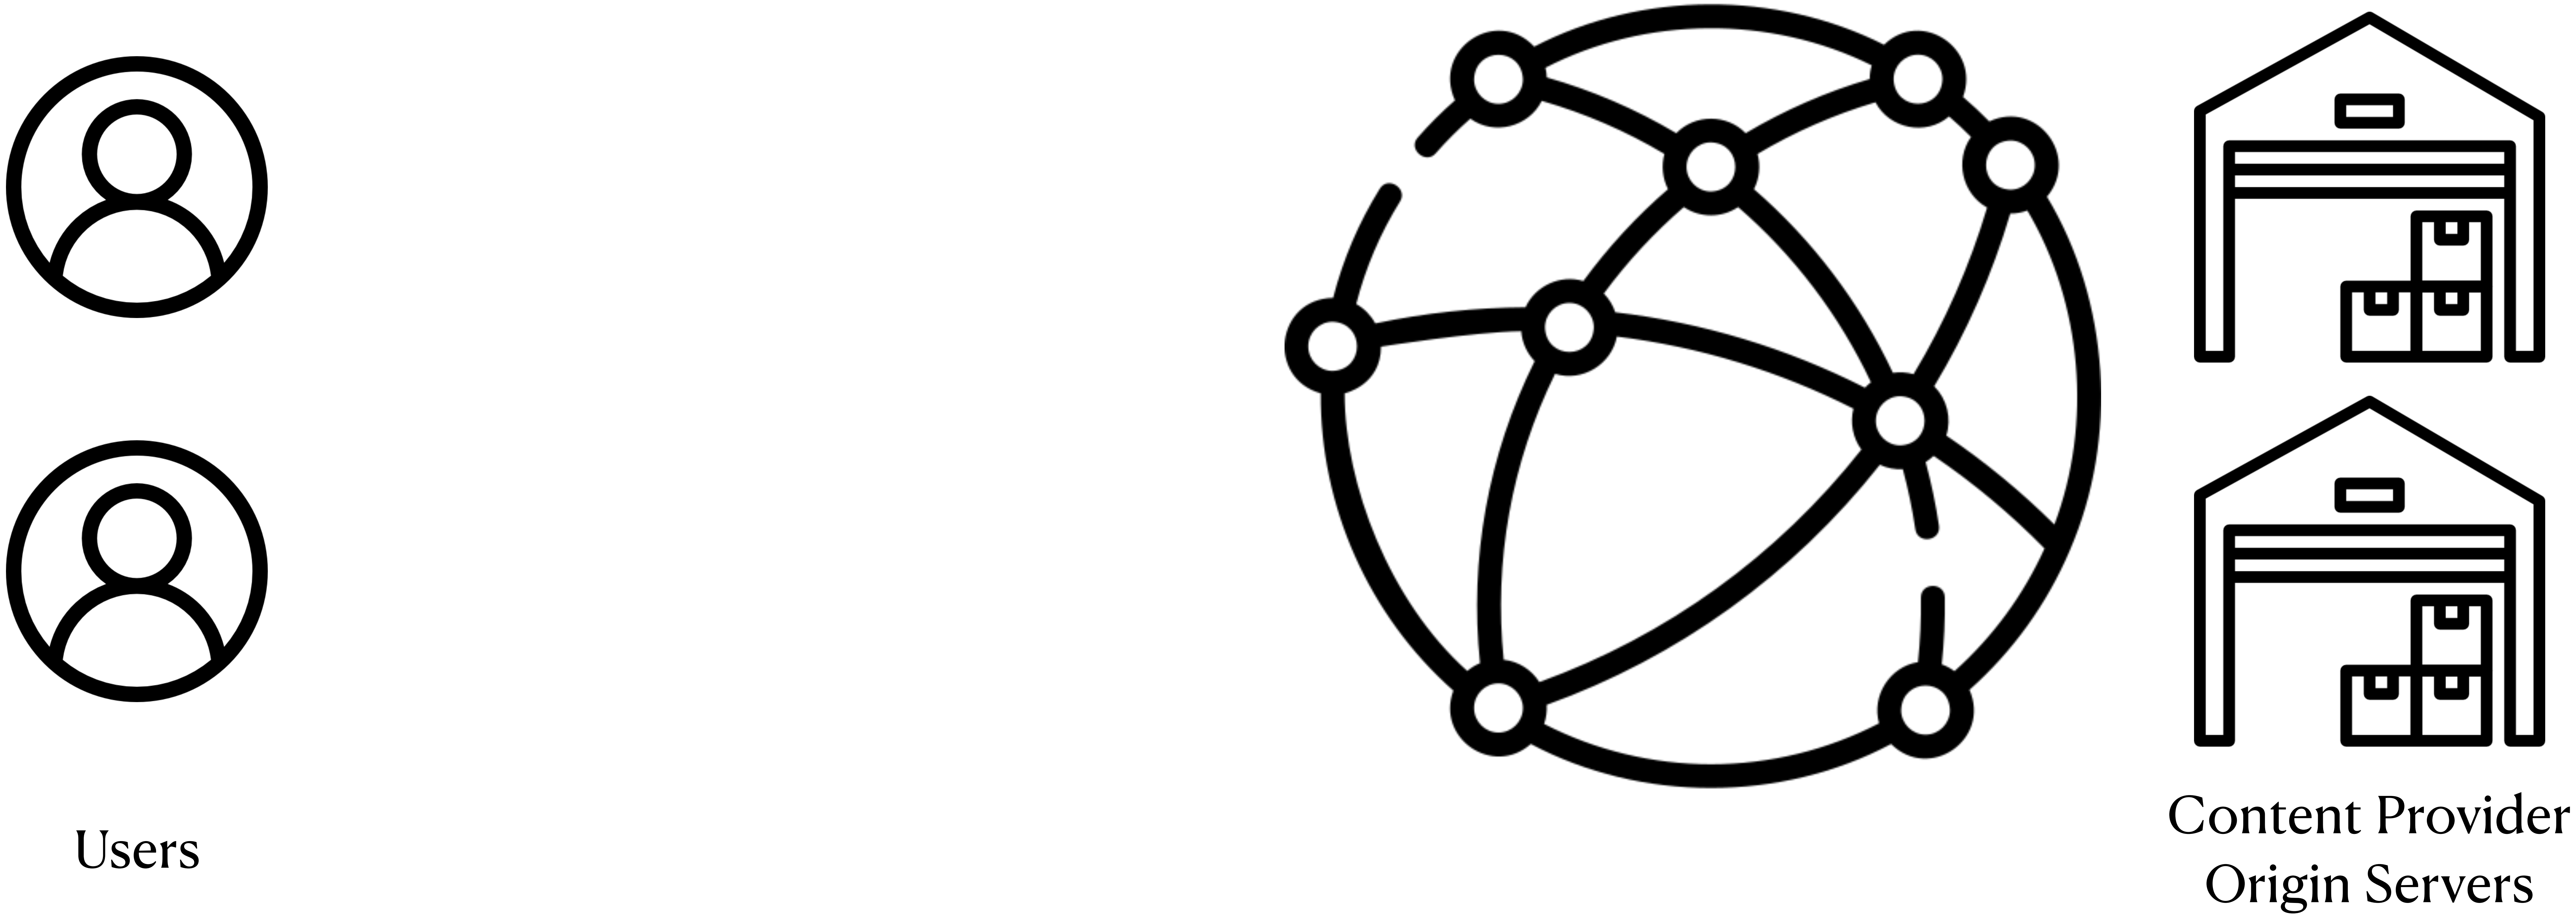

Content Delivery Networks (CDNs)

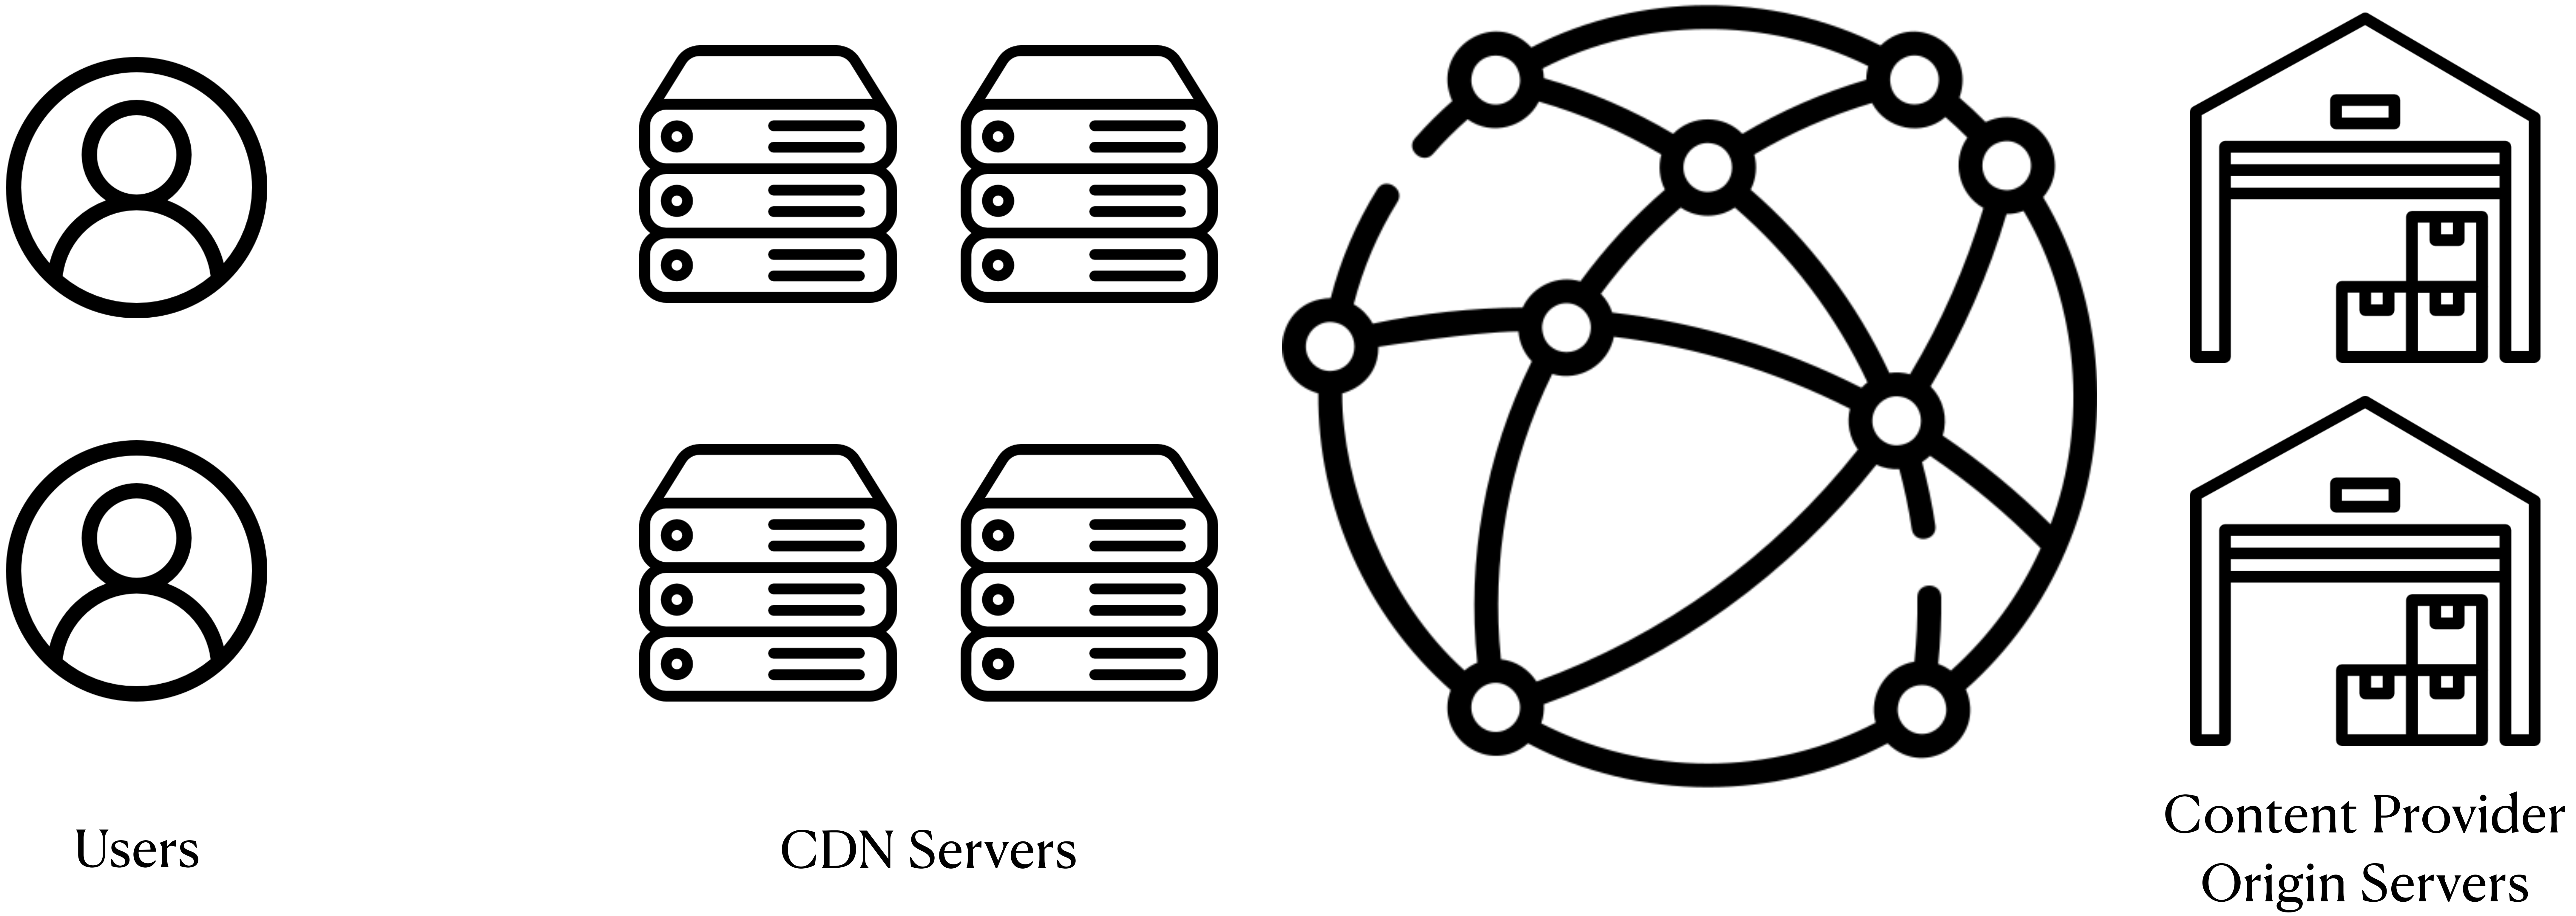

Content Delivery Networks (CDNs)

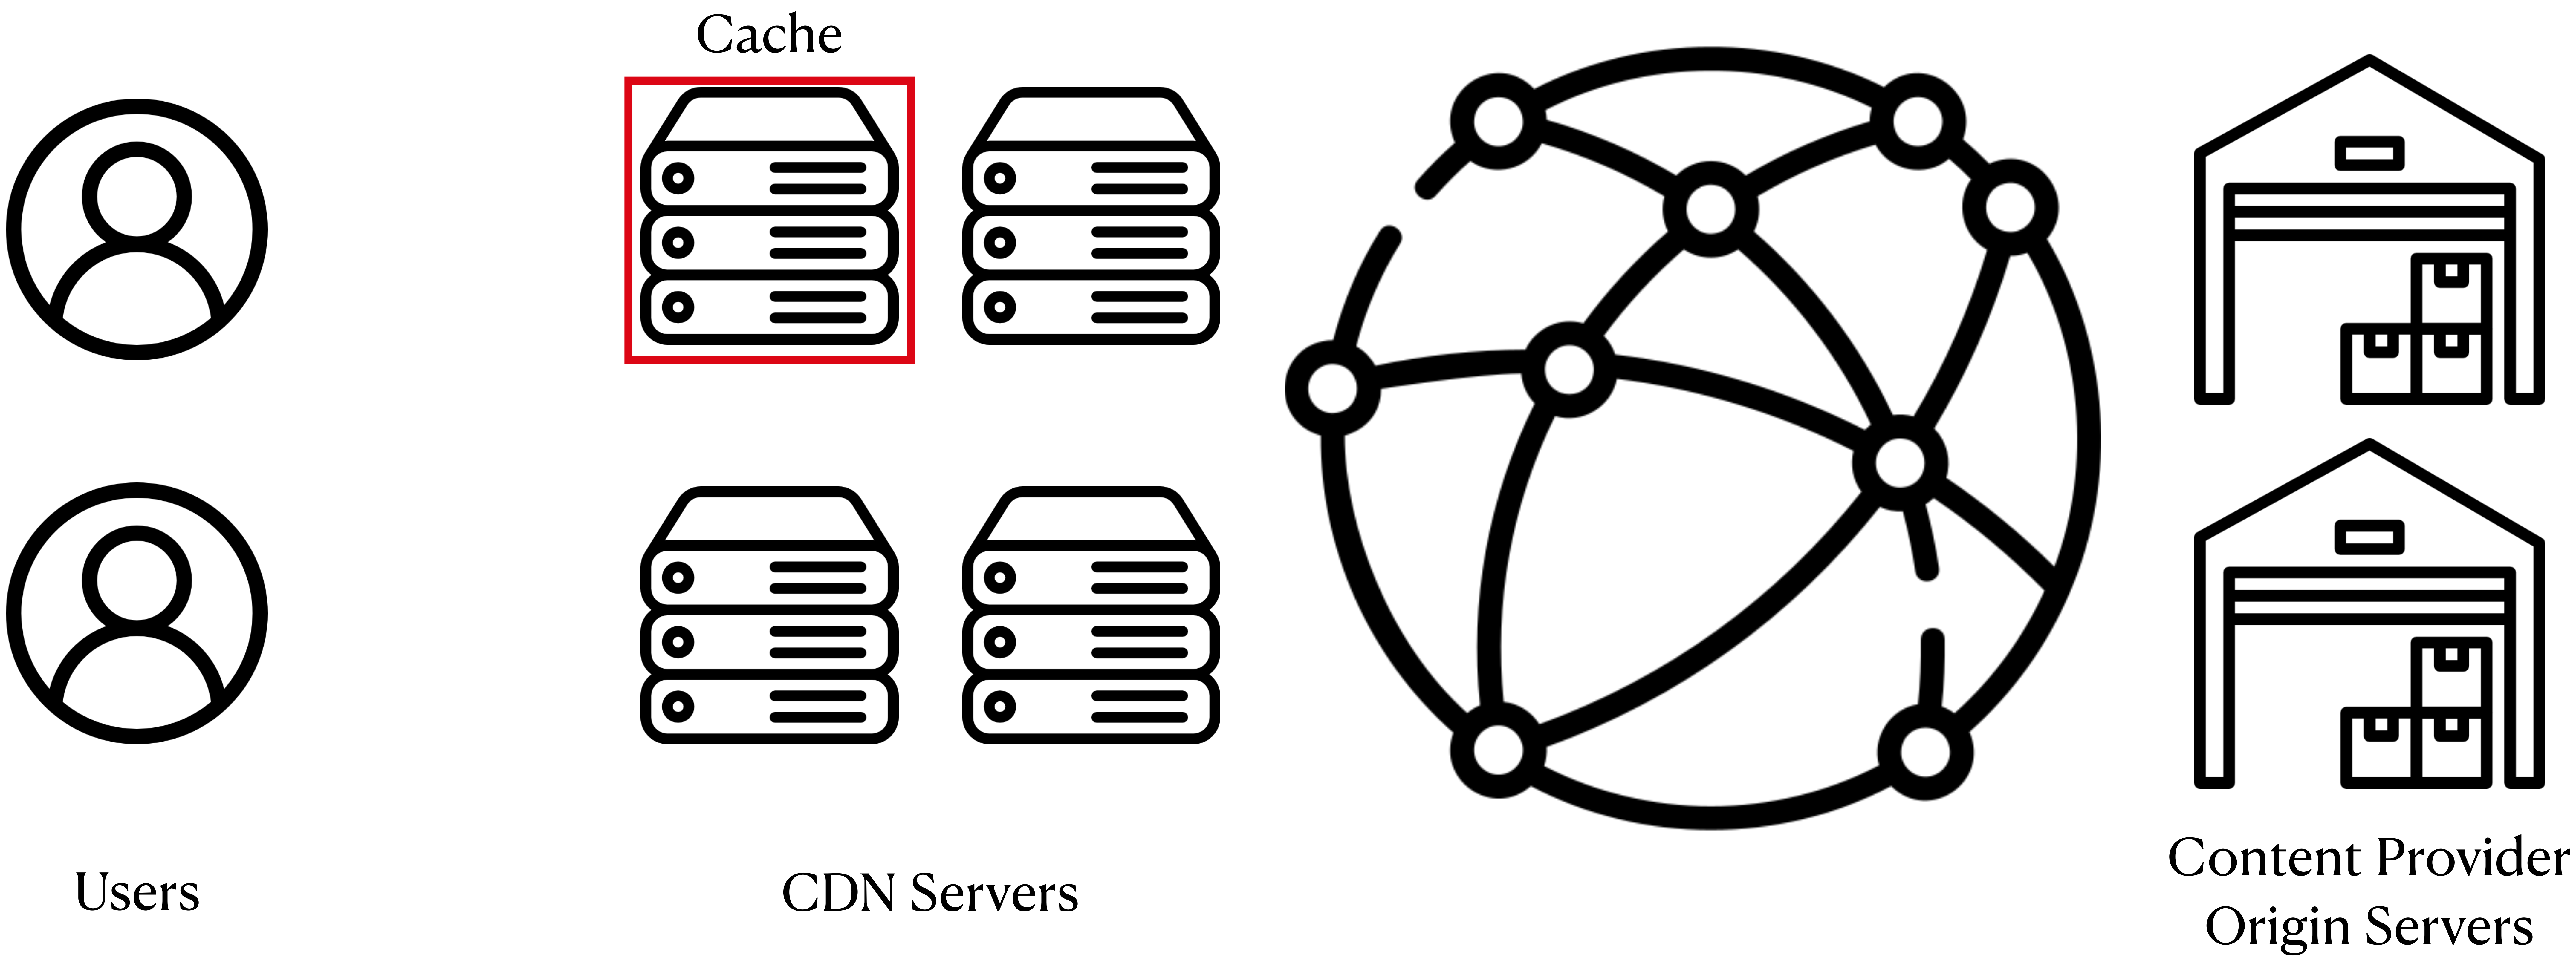

Content Delivery Networks (CDNs)

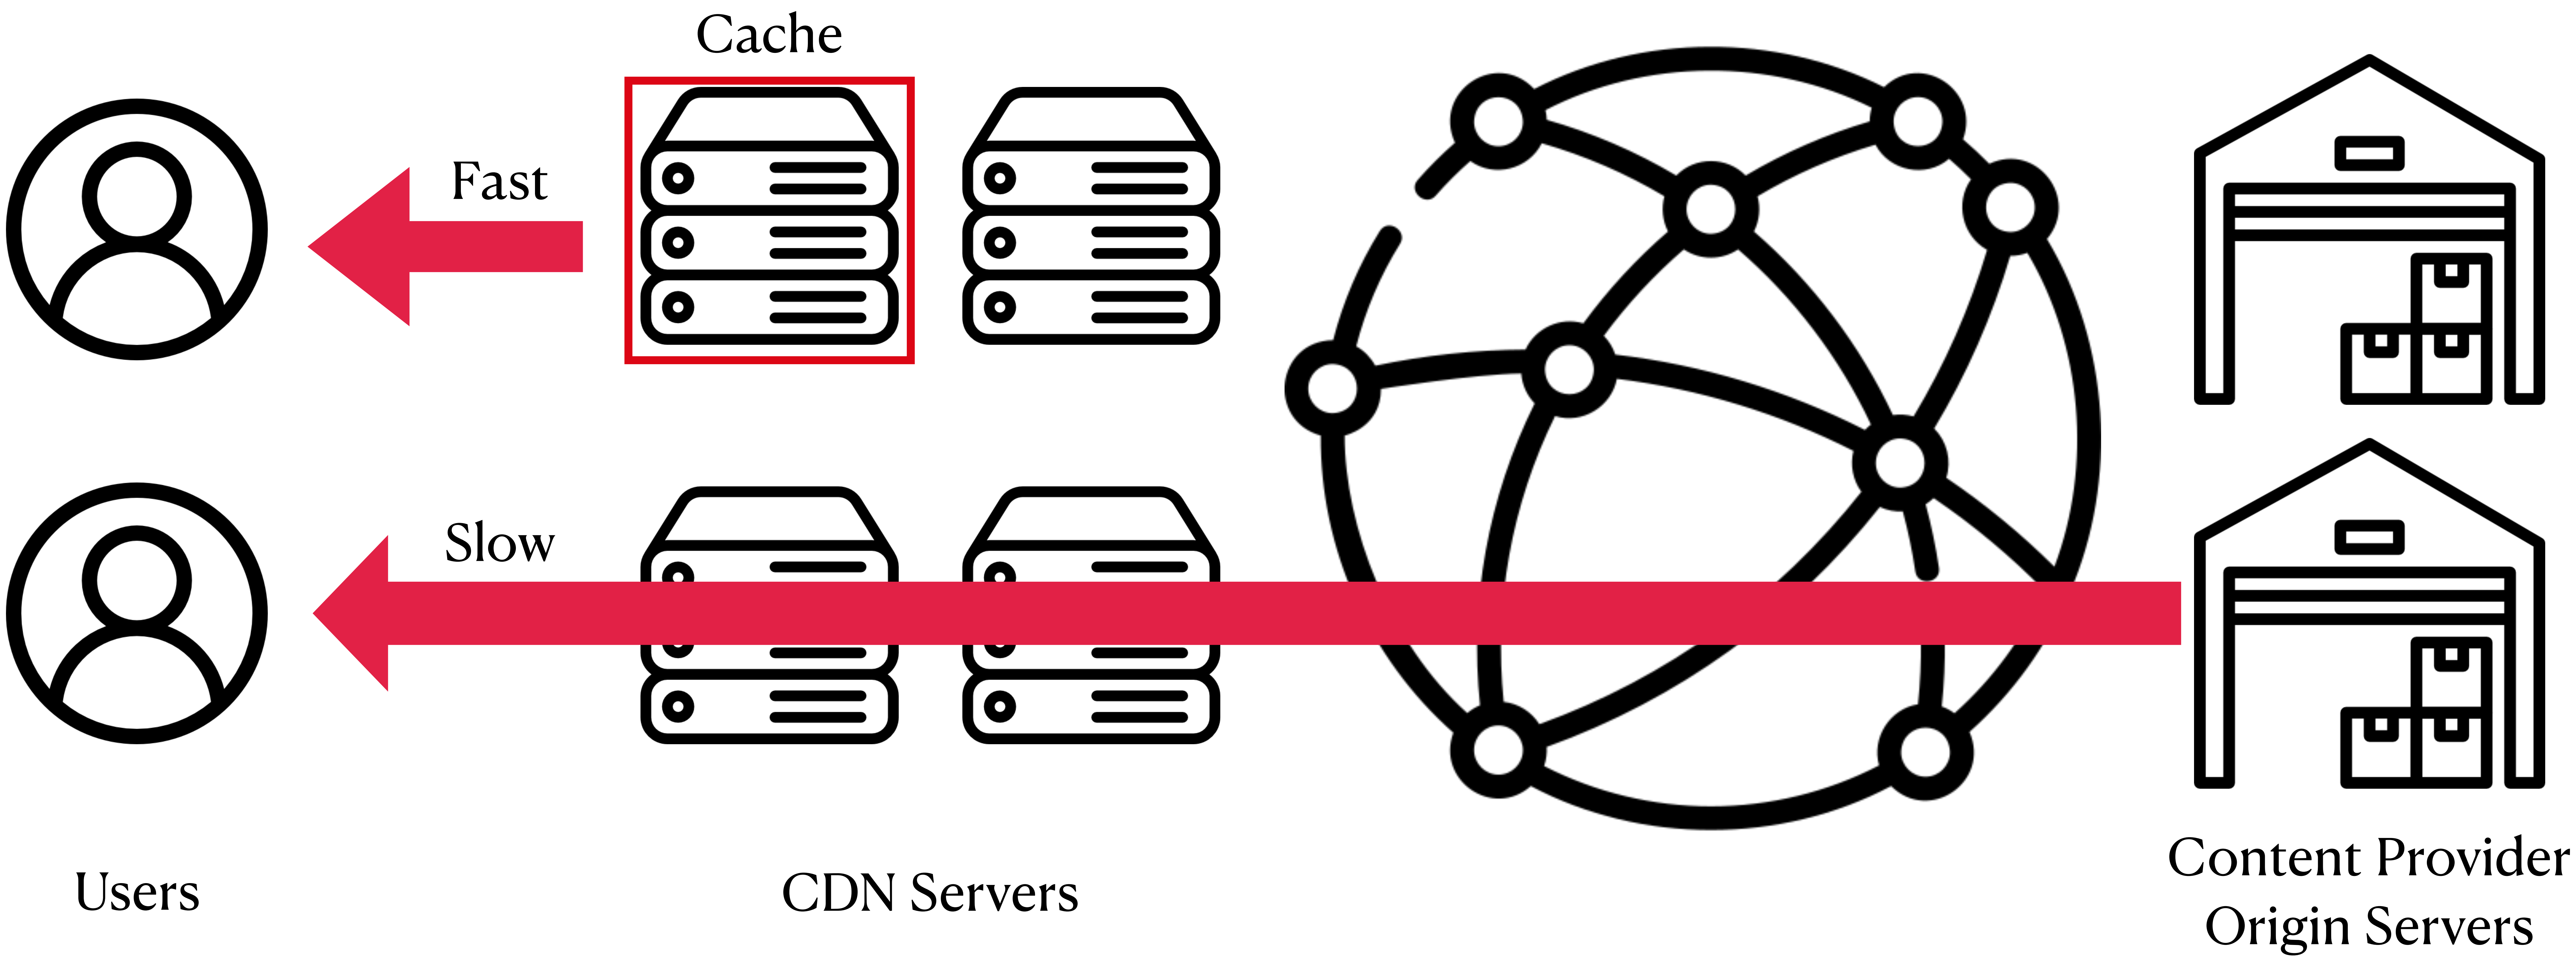

Content Delivery Networks (CDNs)

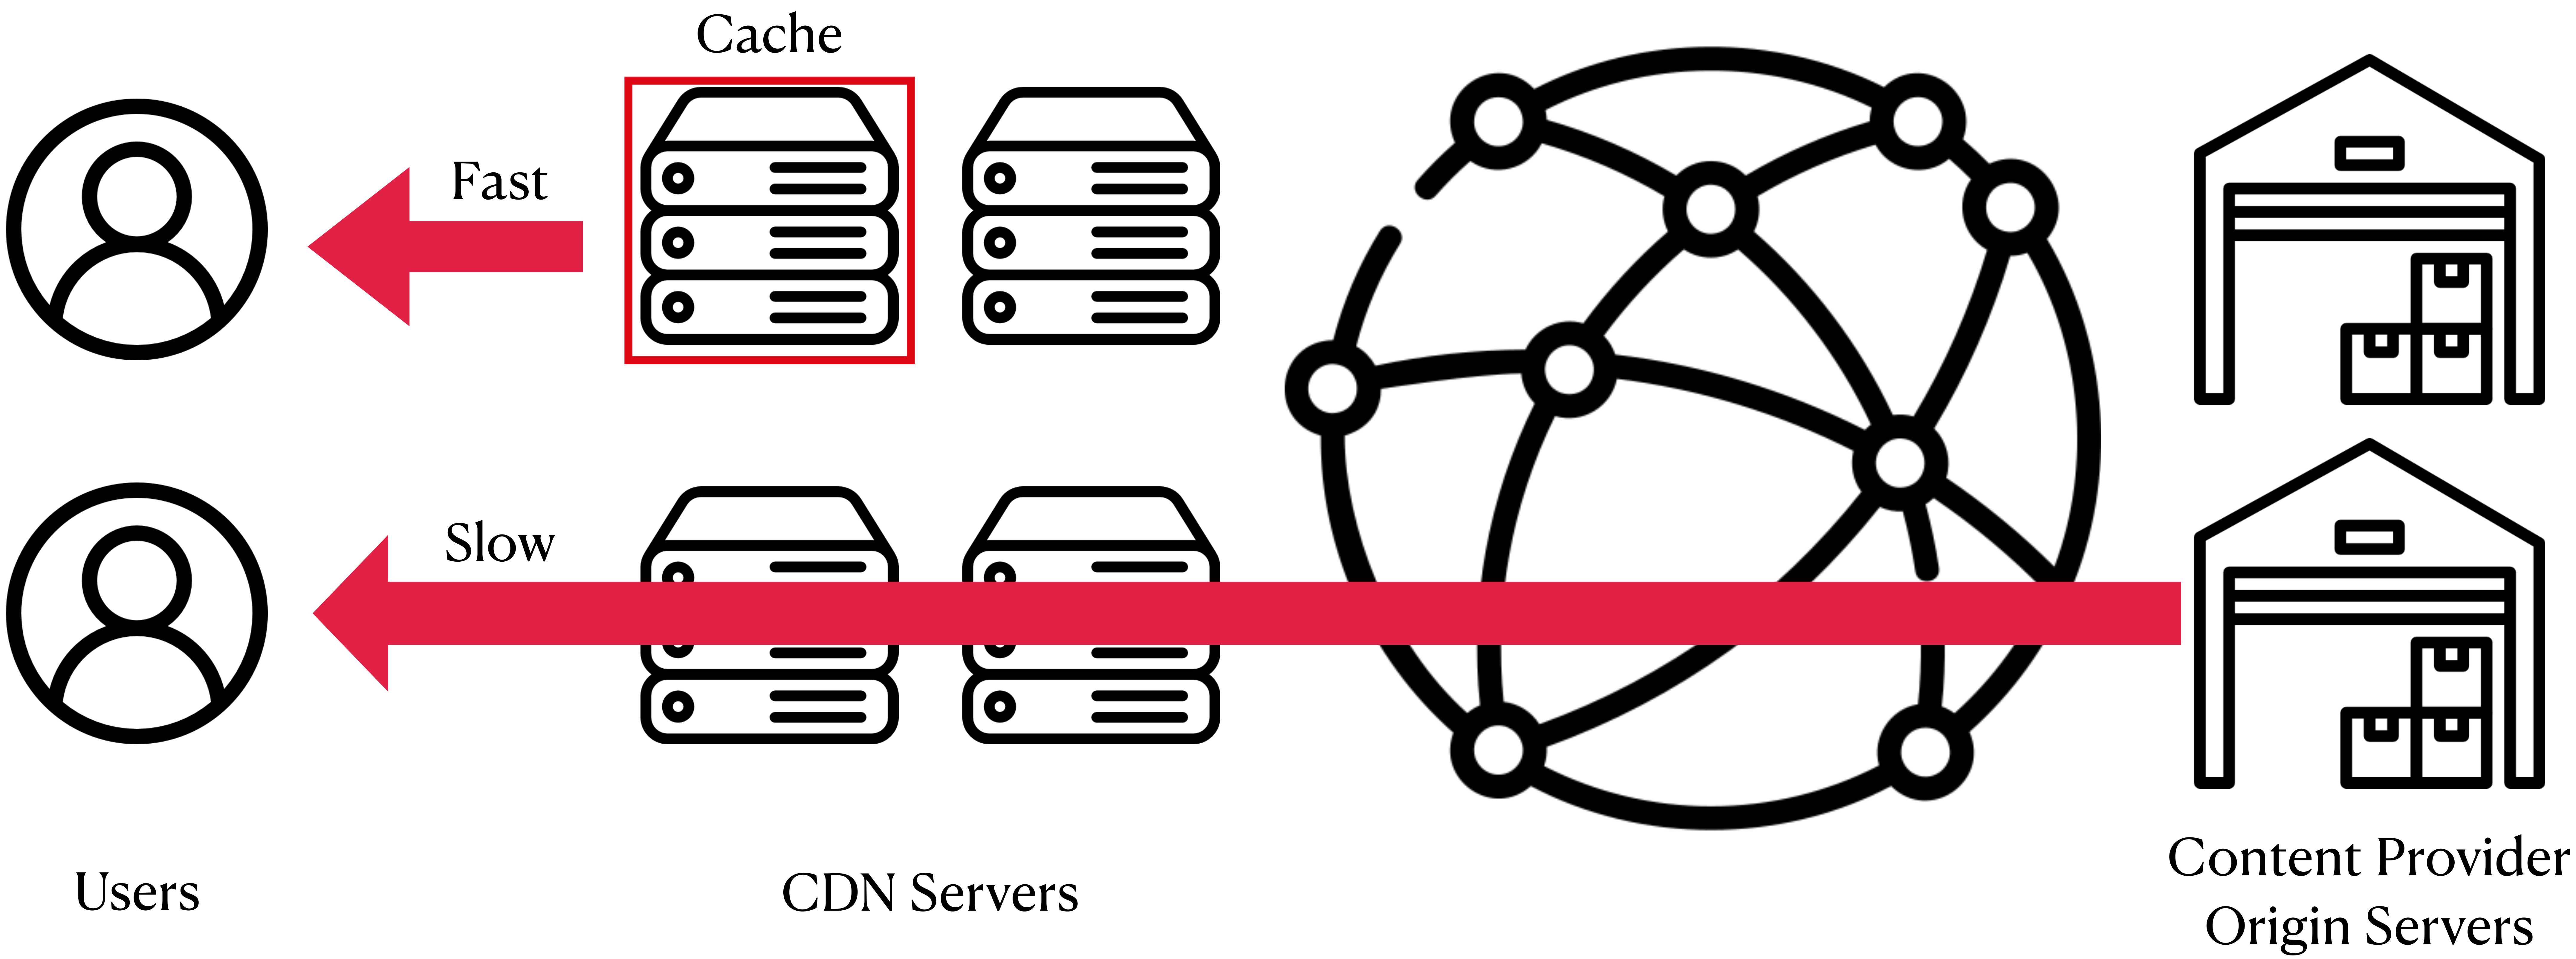

Cache management policies play an important role

Traffic Changes Make Cache Management Challenging

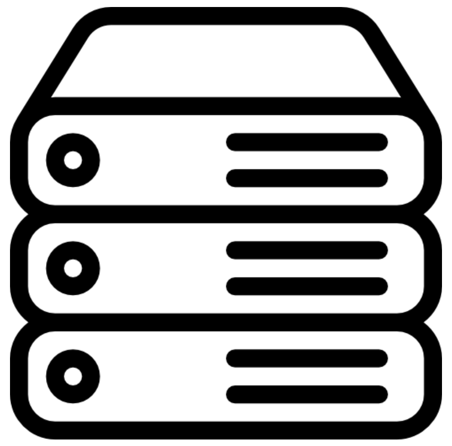

Traffic Changes Make Cache Management Challenging

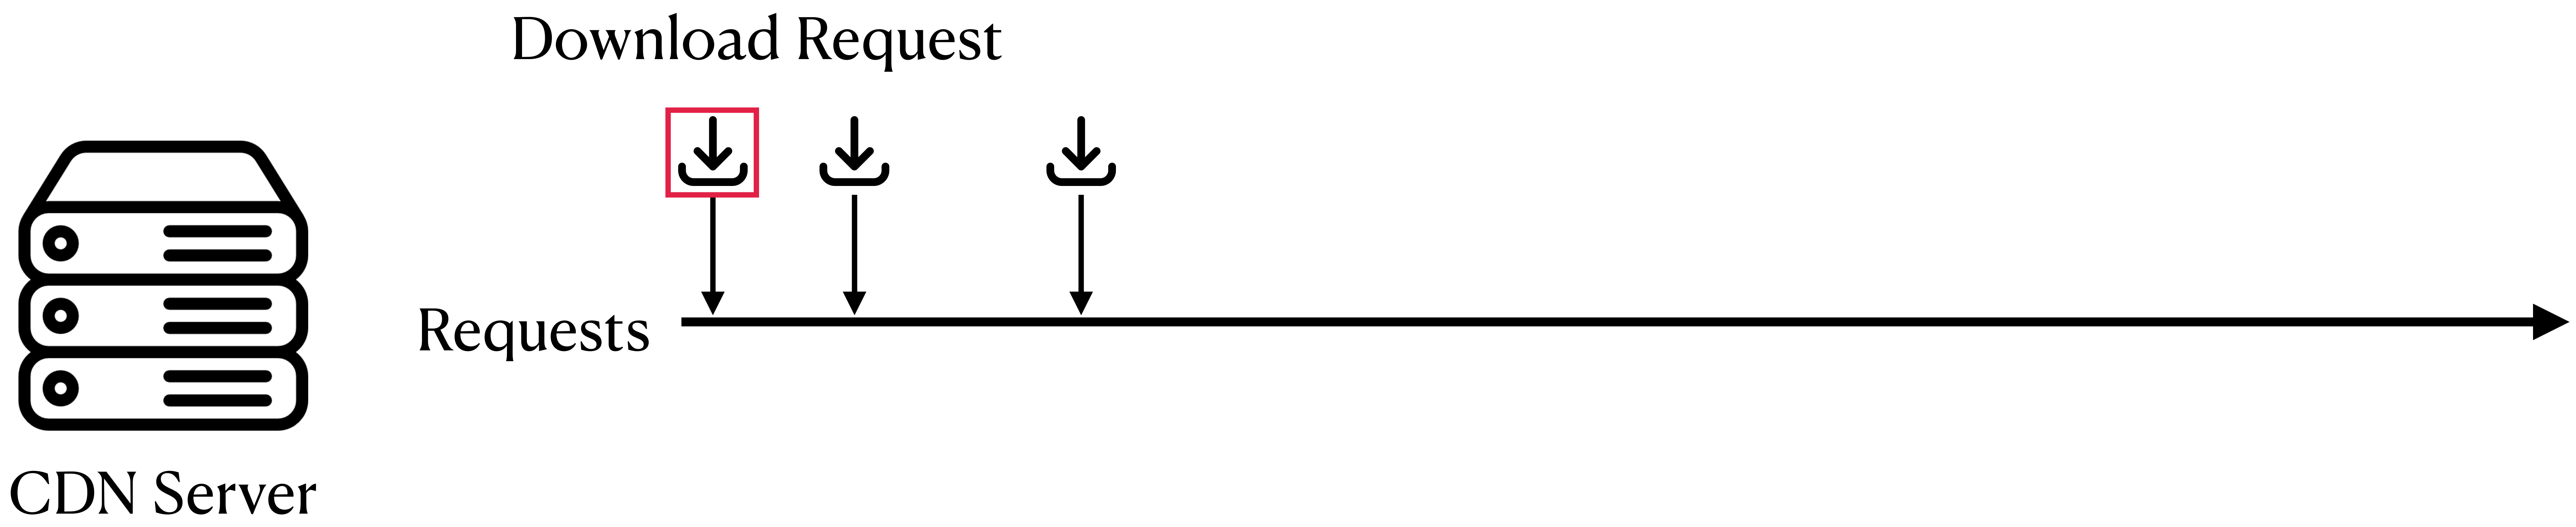

Traffic Changes Make Cache Management Challenging

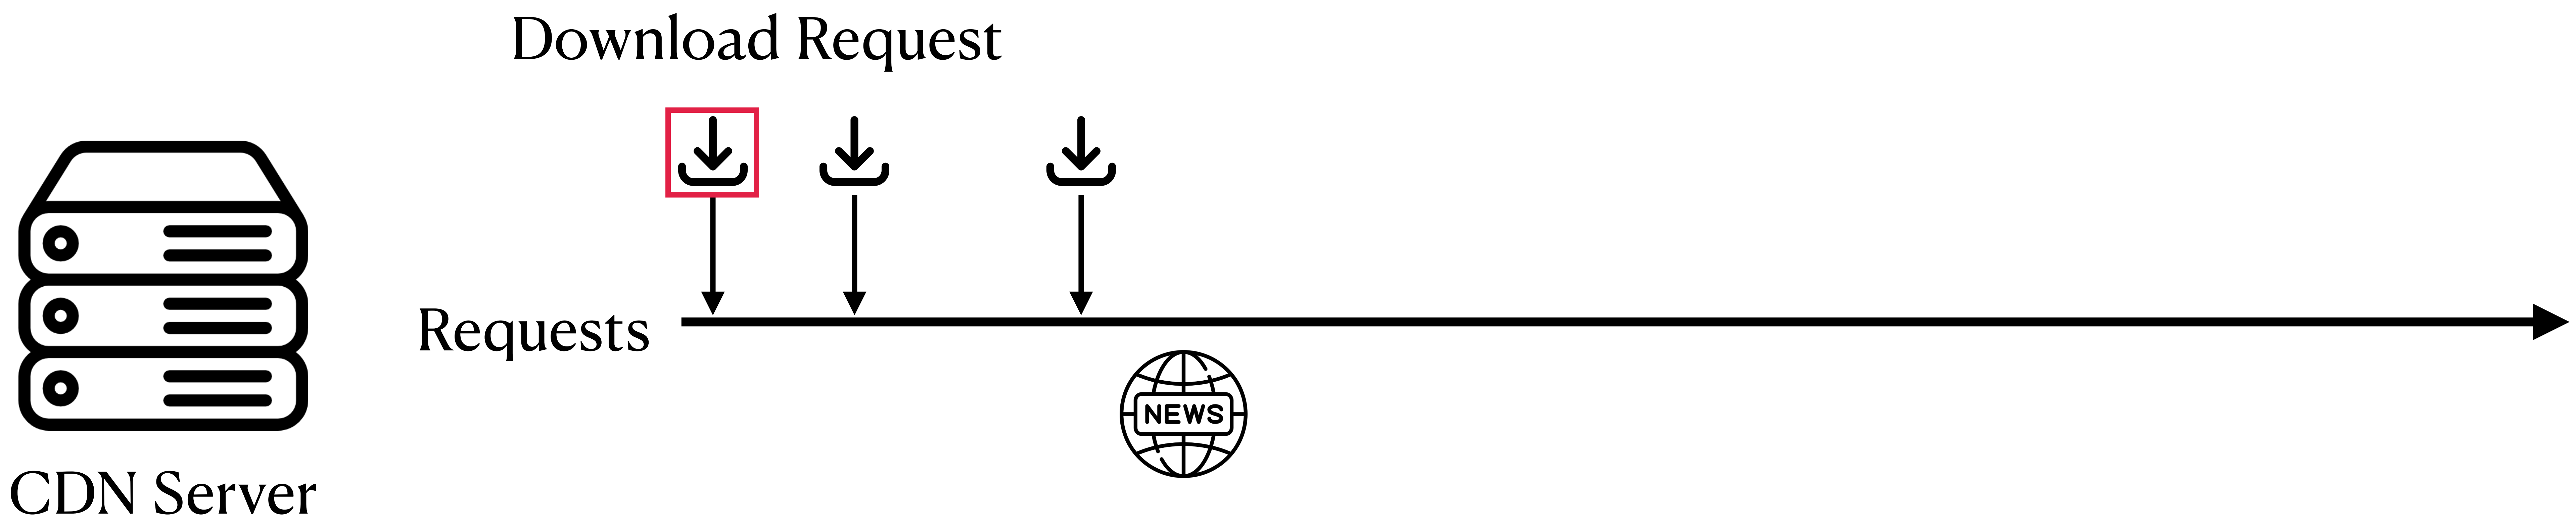

Traffic Changes Make Cache Management Challenging

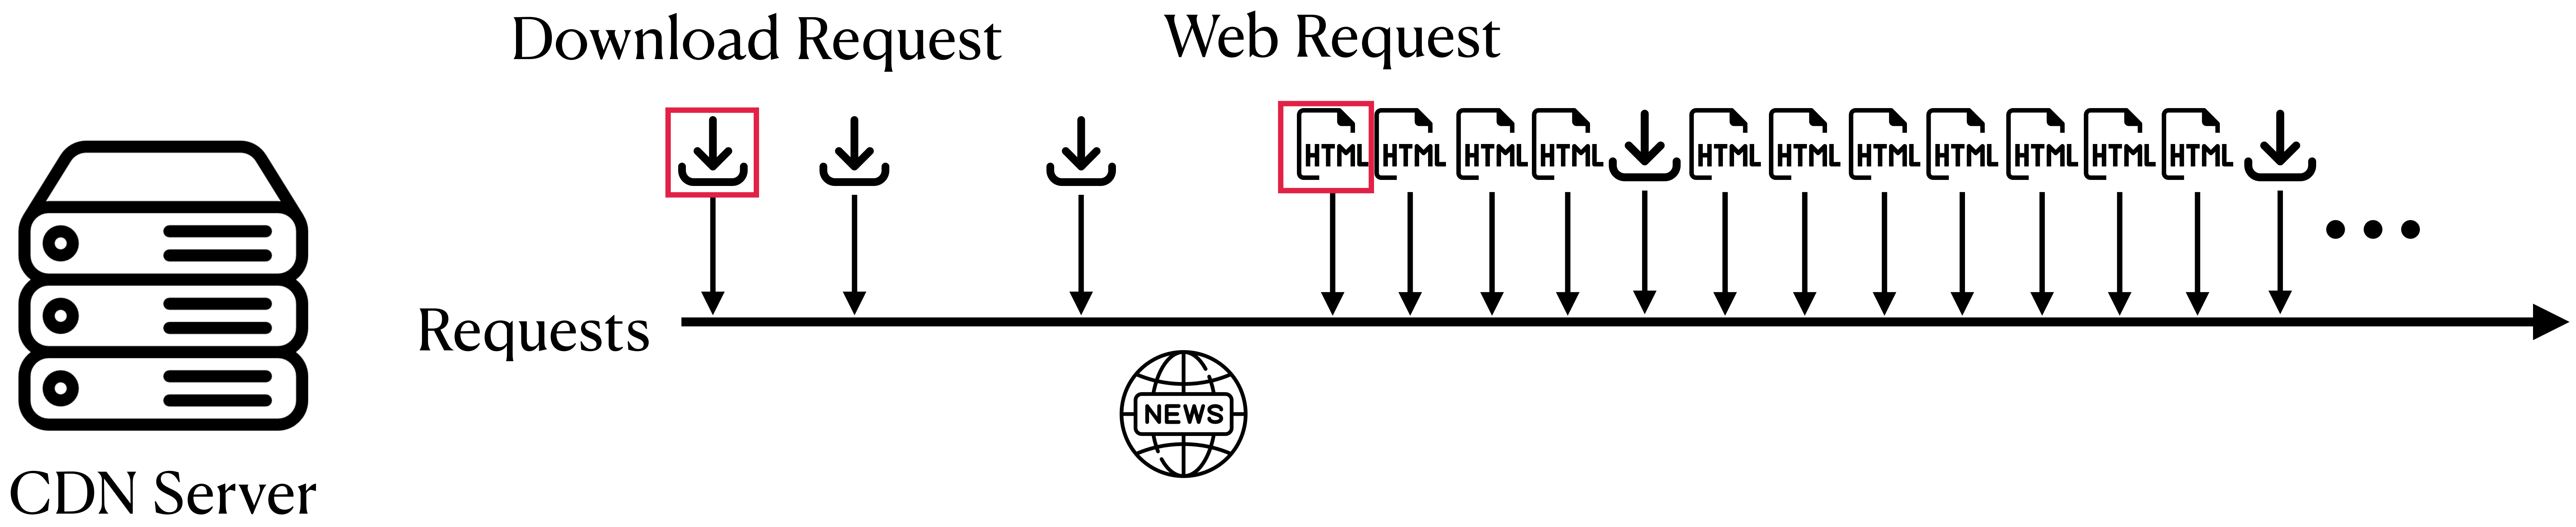

Traffic Changes Make Cache Management Challenging

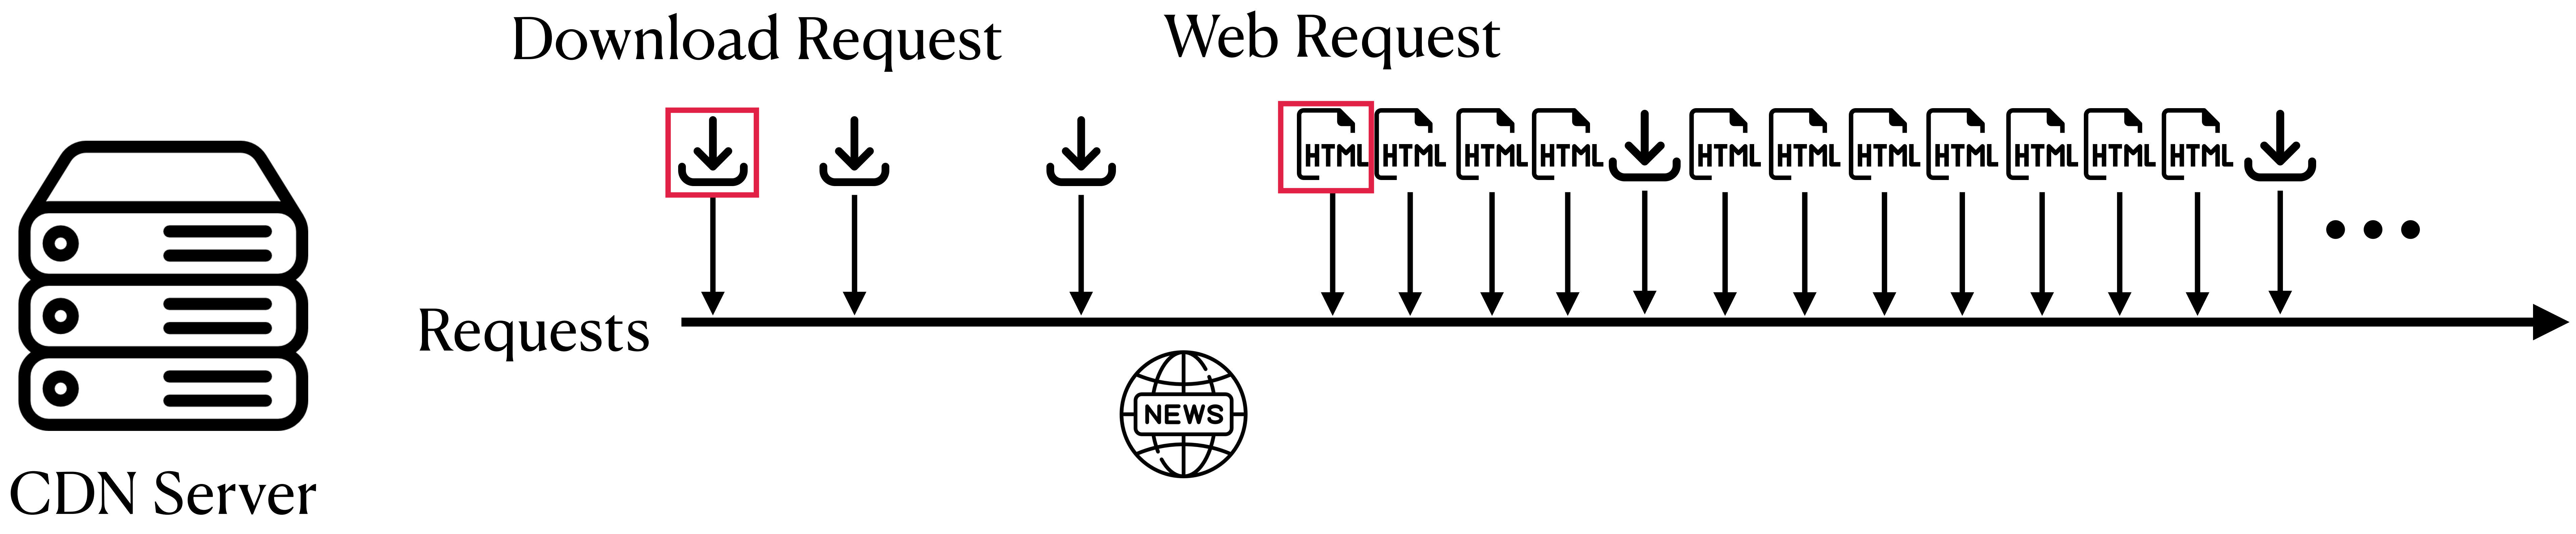

Are static cache management policies effective?

Hot Object Cache (HOC) Admission Policy

CDN Server

Hot Object Cache (HOC) Admission Policy

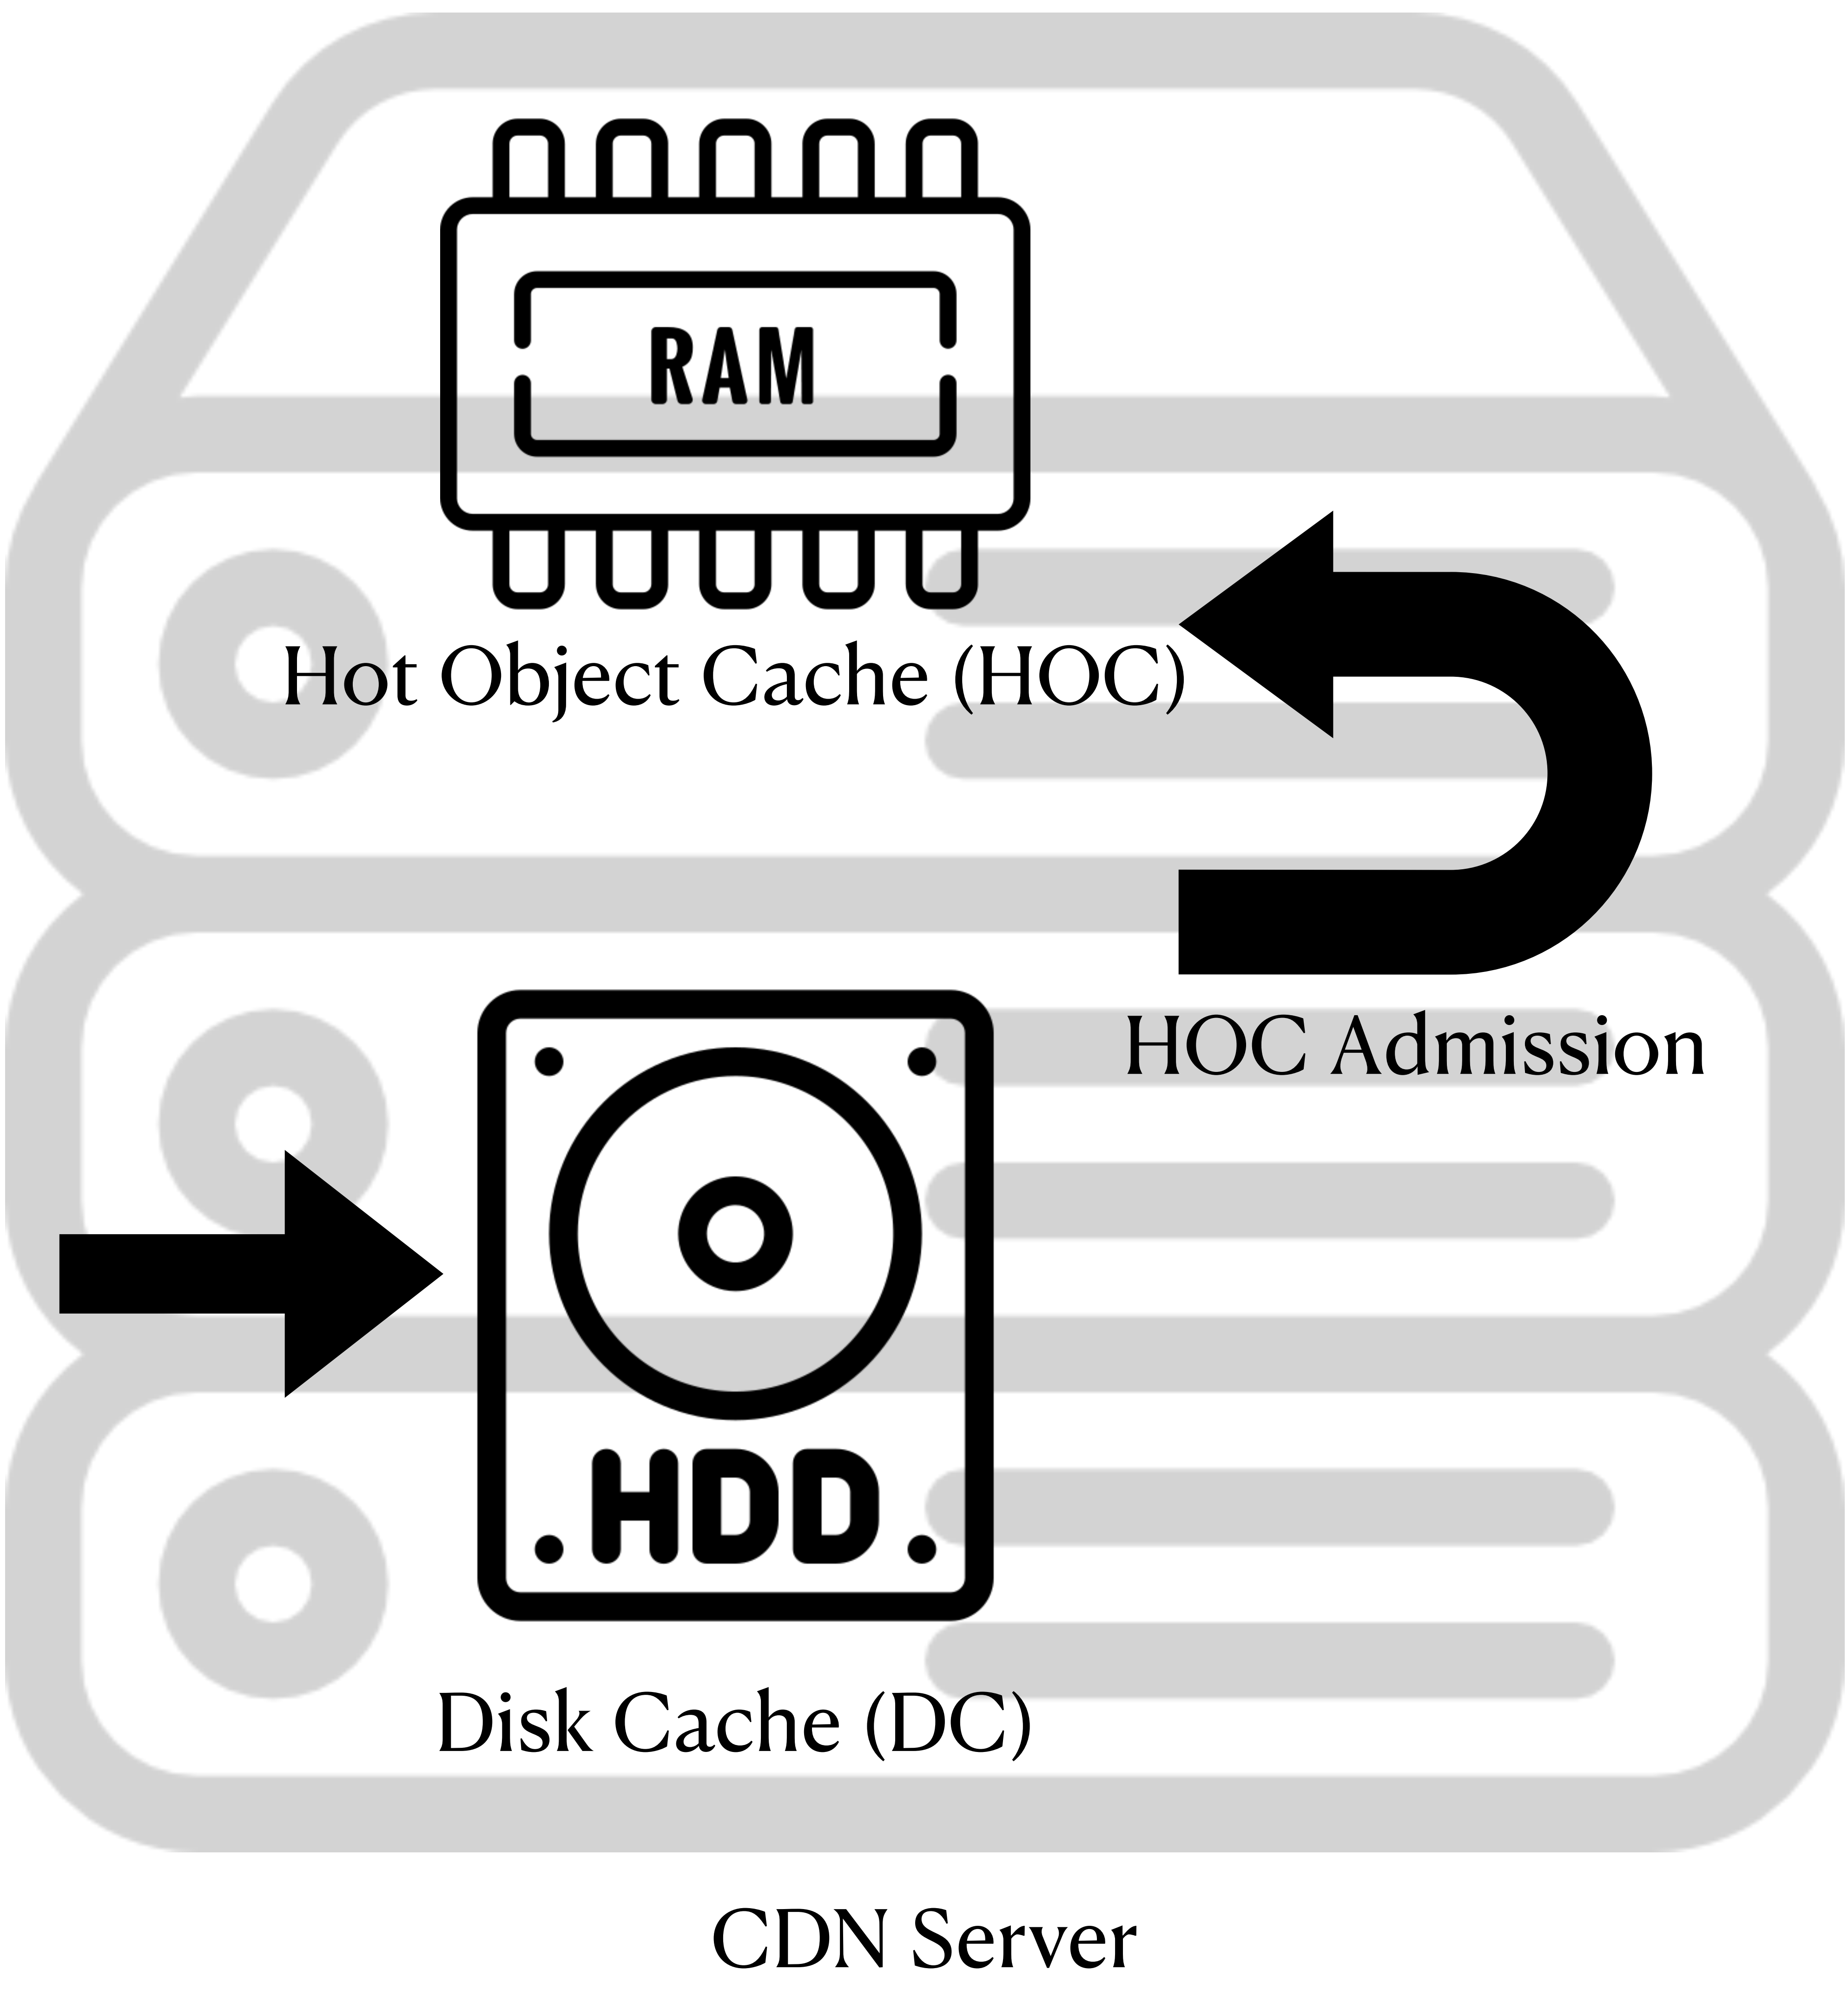

Hot Object Cache (HOC) Admission Policy

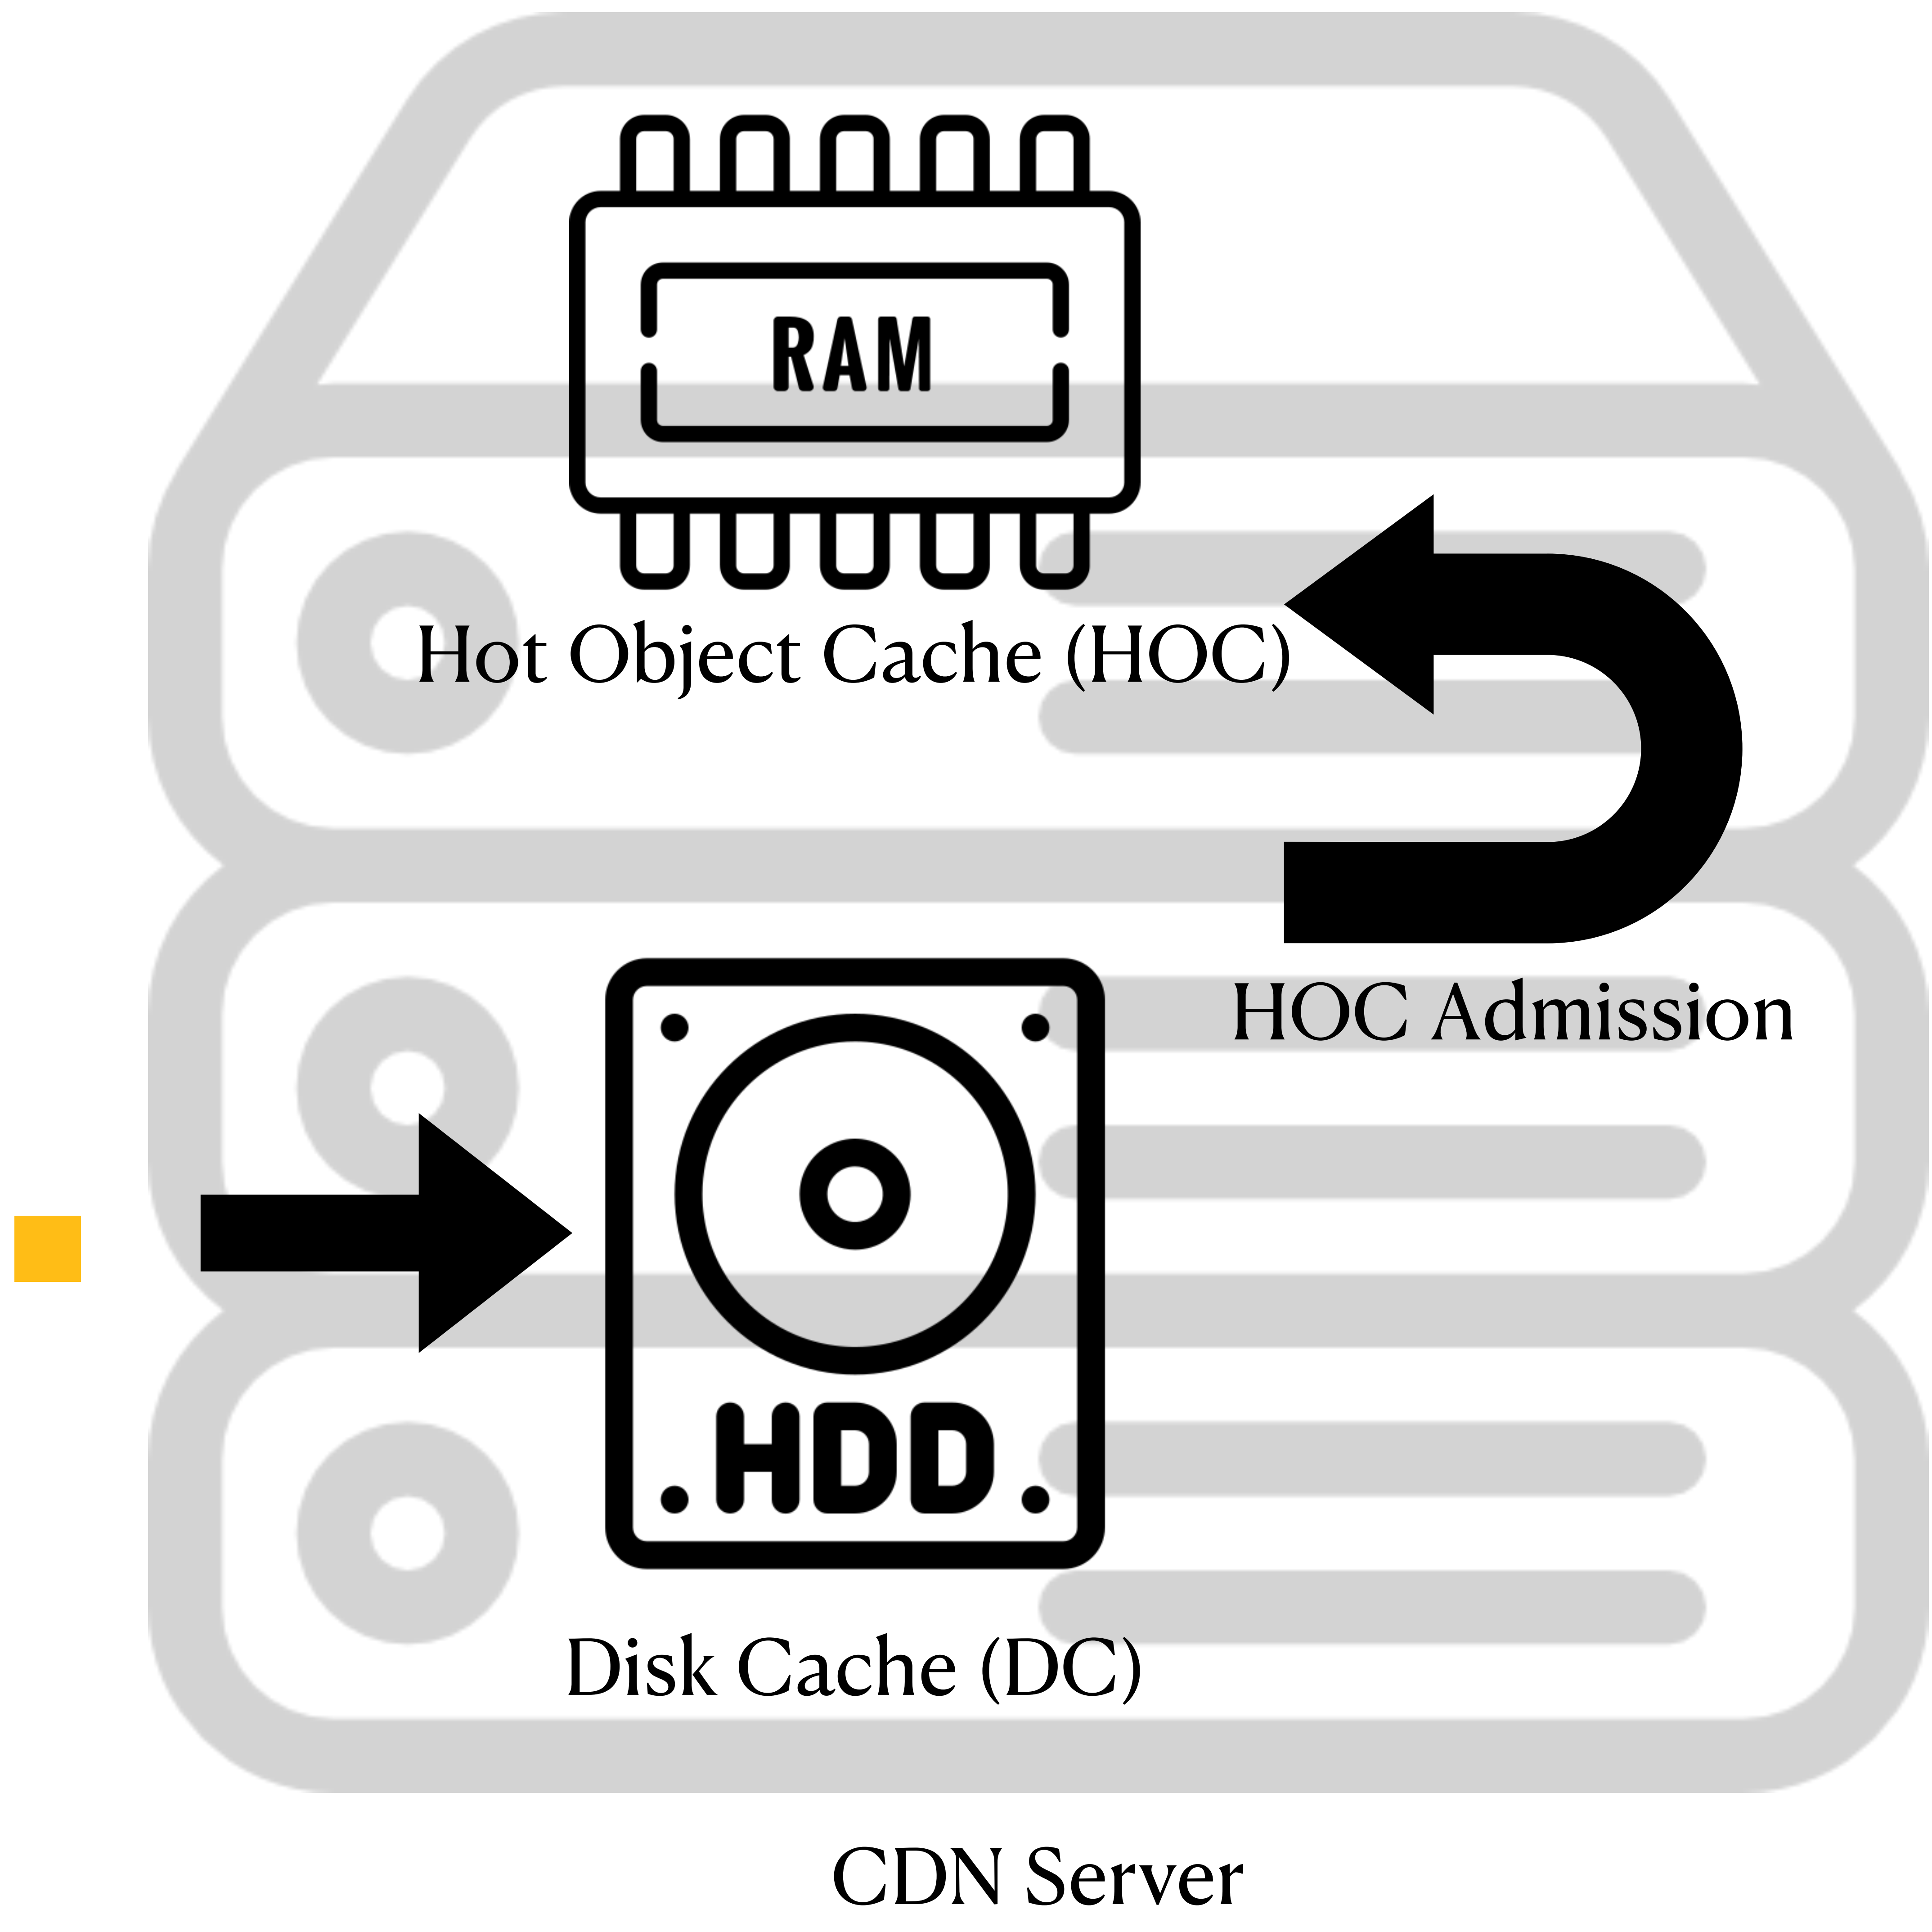

Hot Object Cache (HOC) Admission Policy

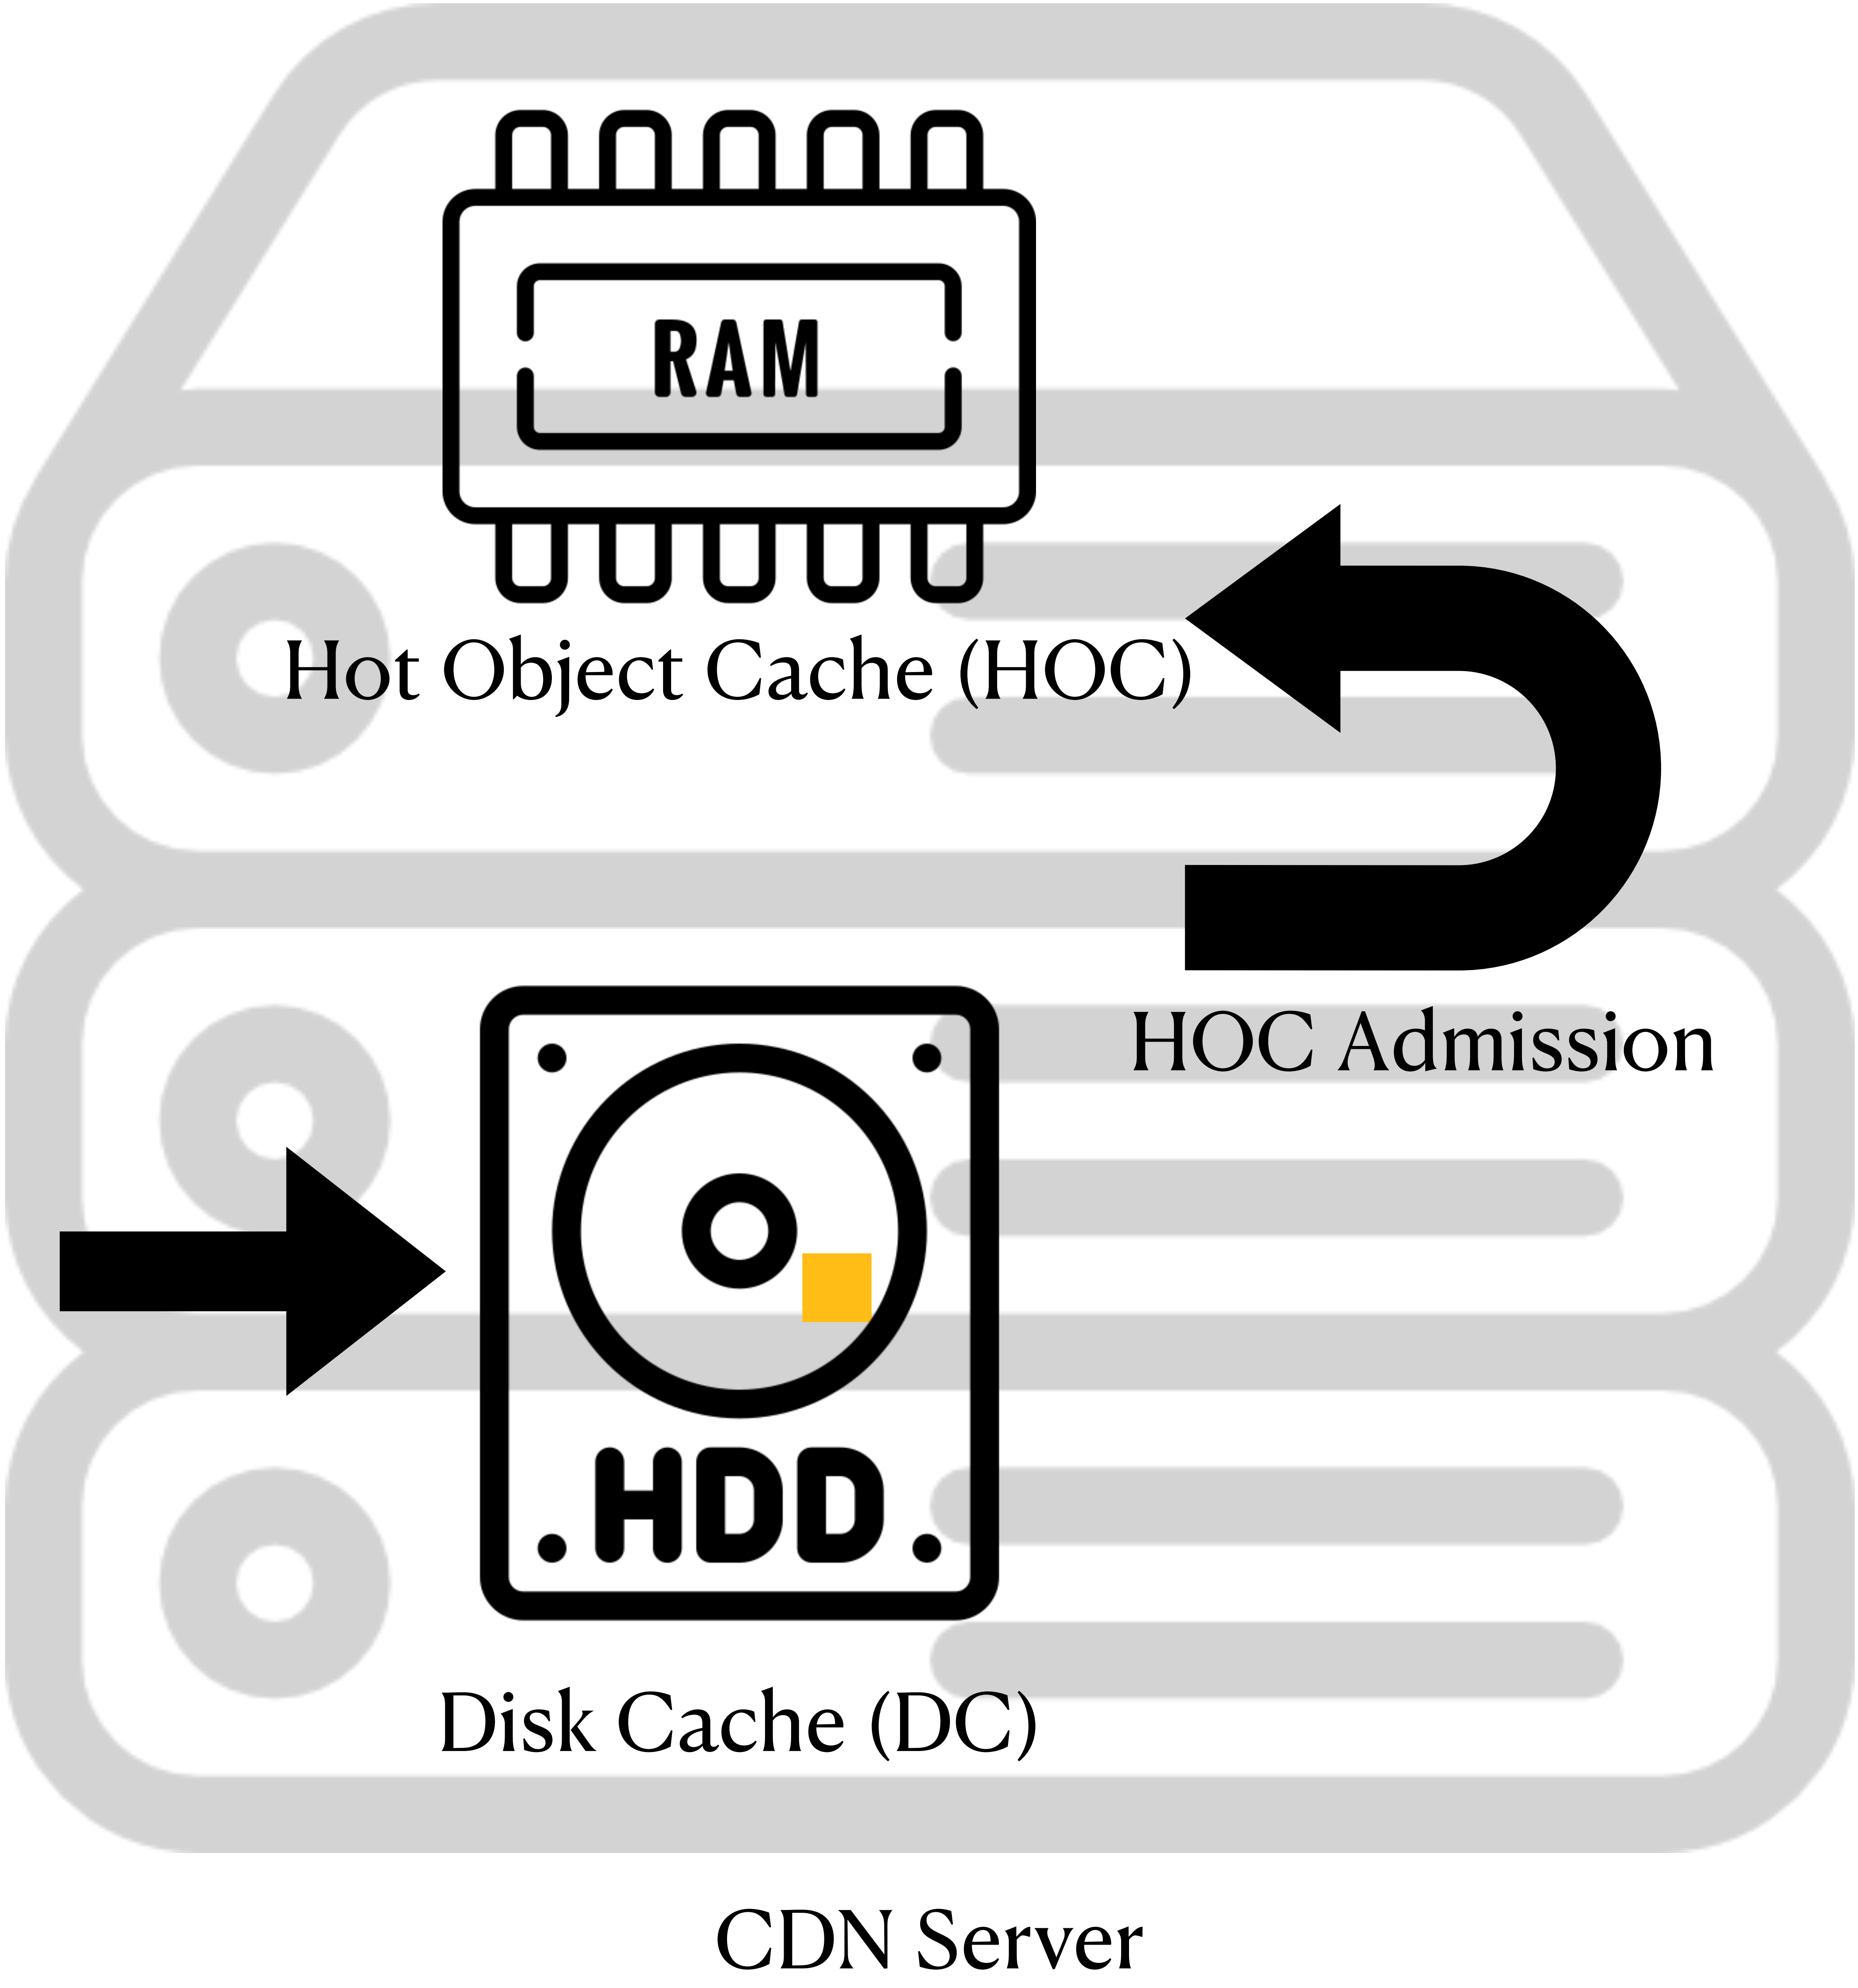

Hot Object Cache (HOC) Admission Policy

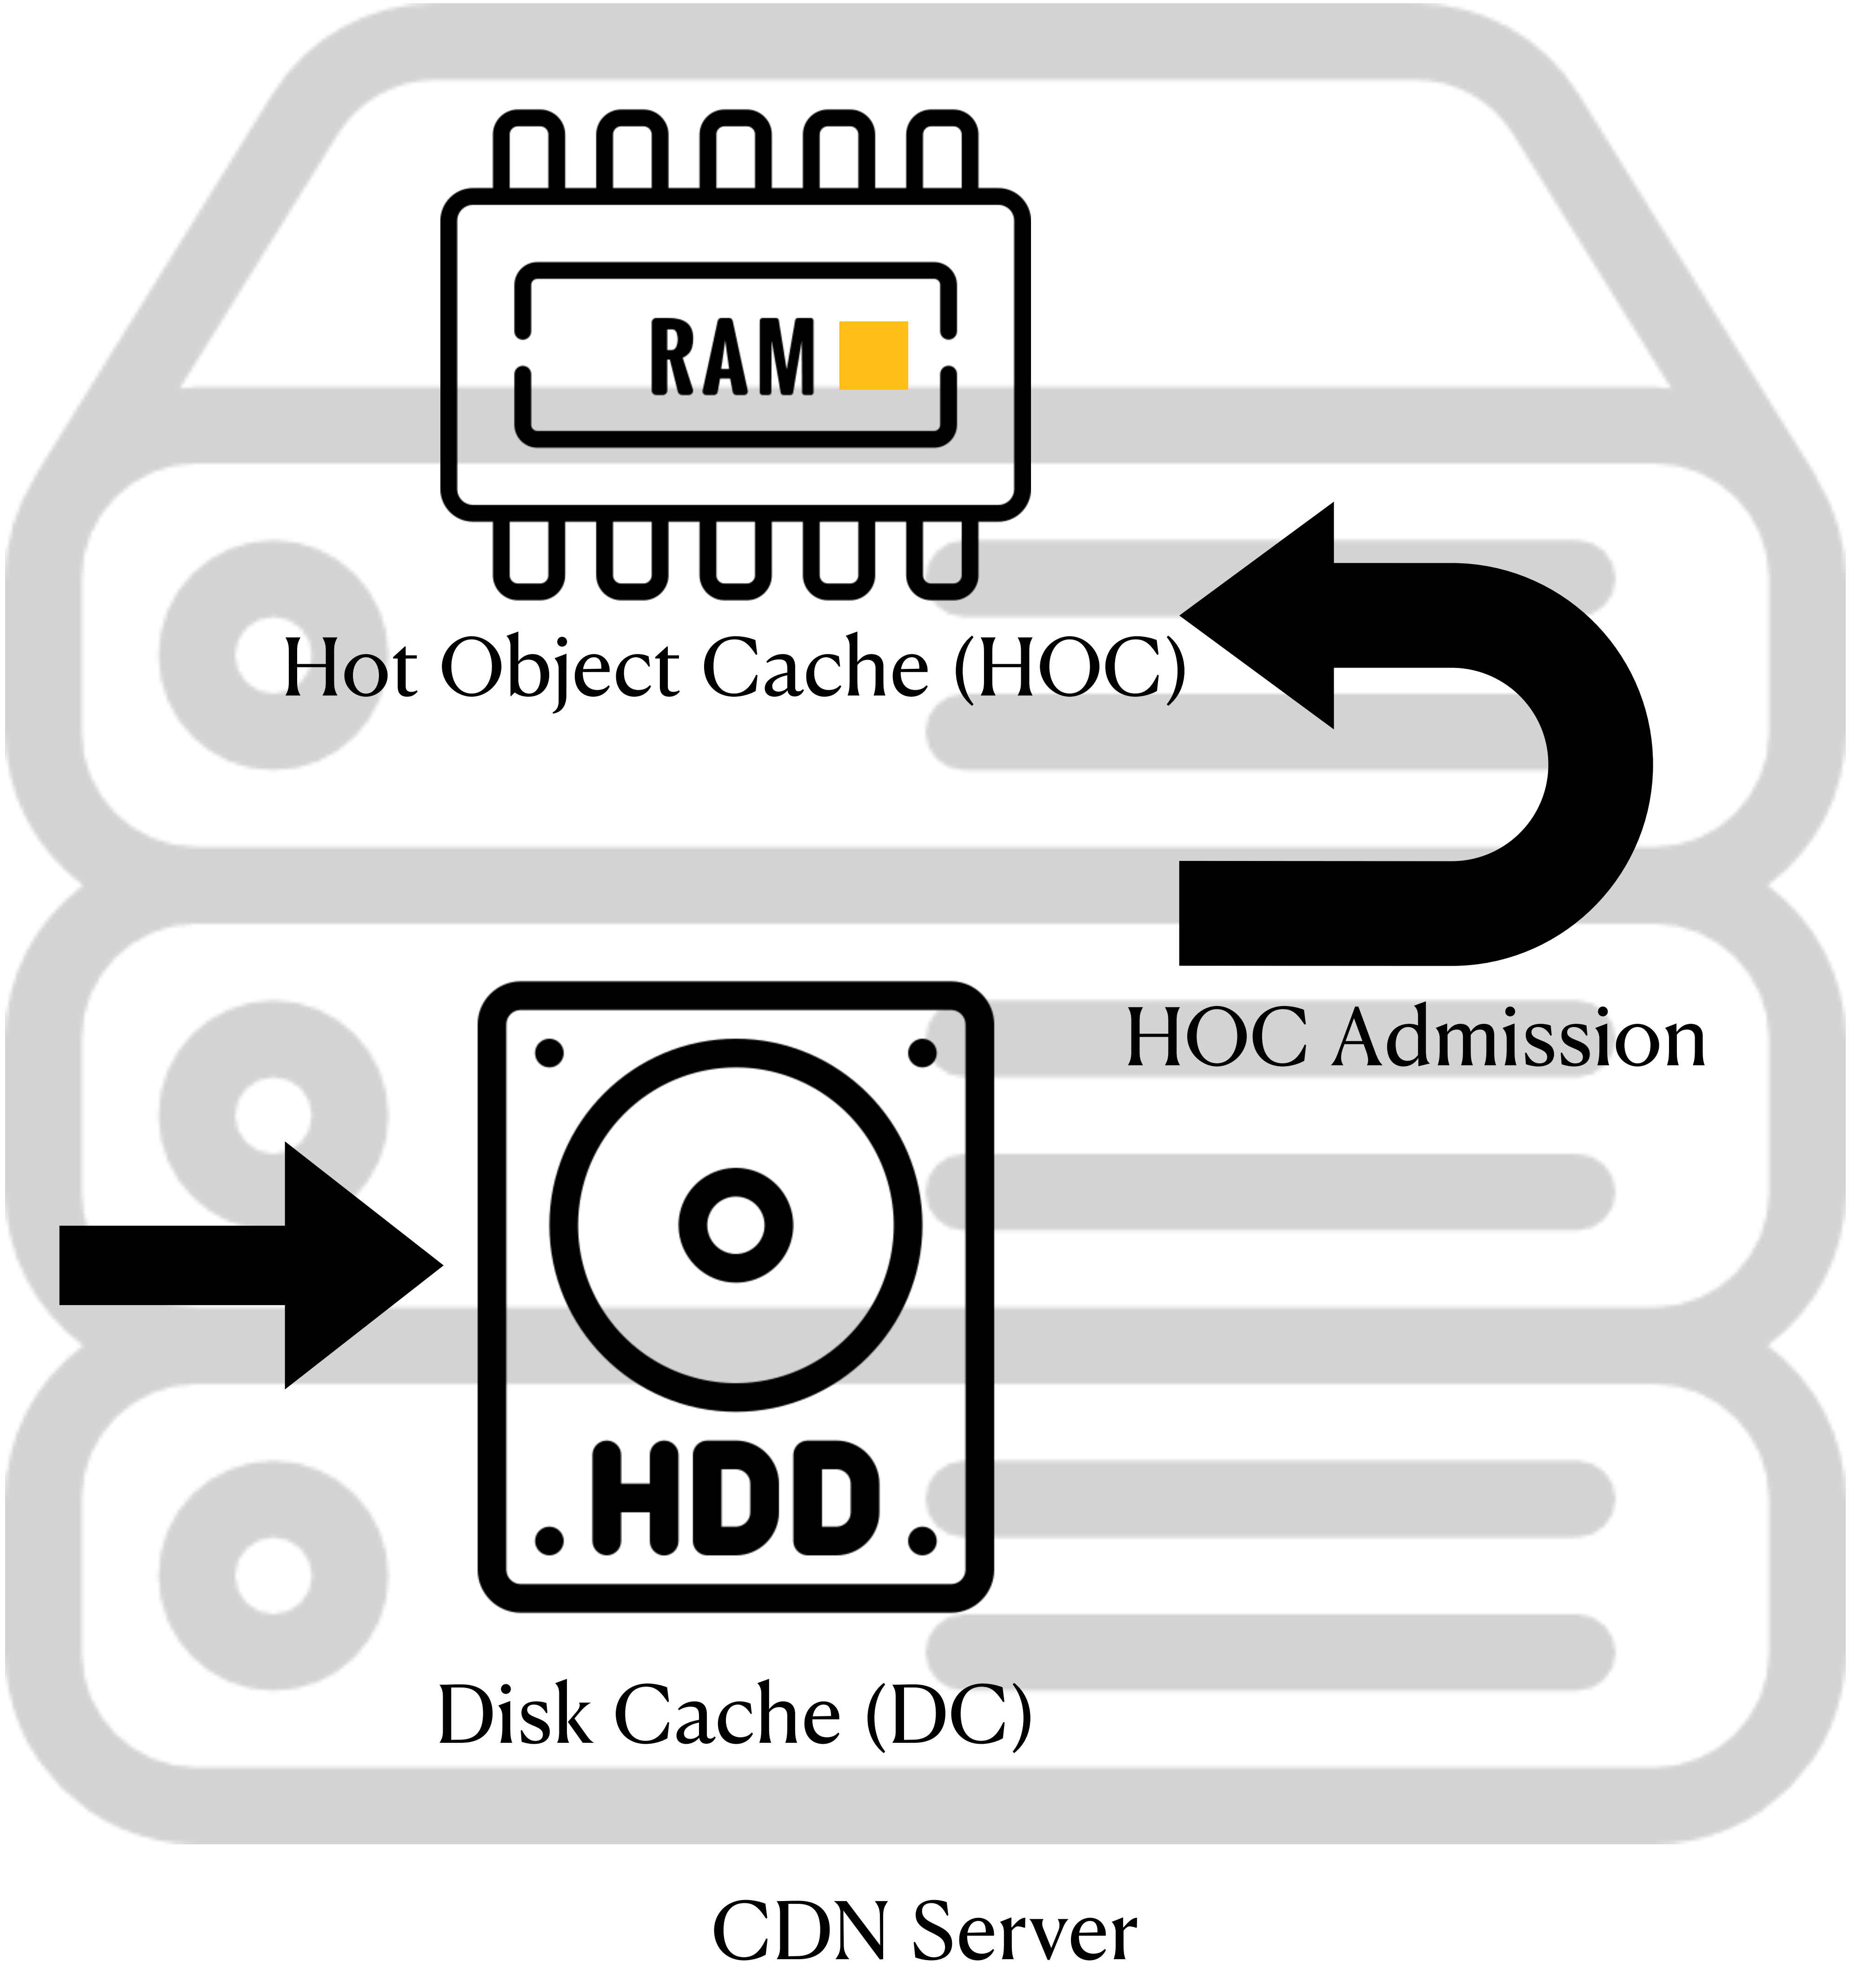

Hot Object Cache (HOC) Admission Policy

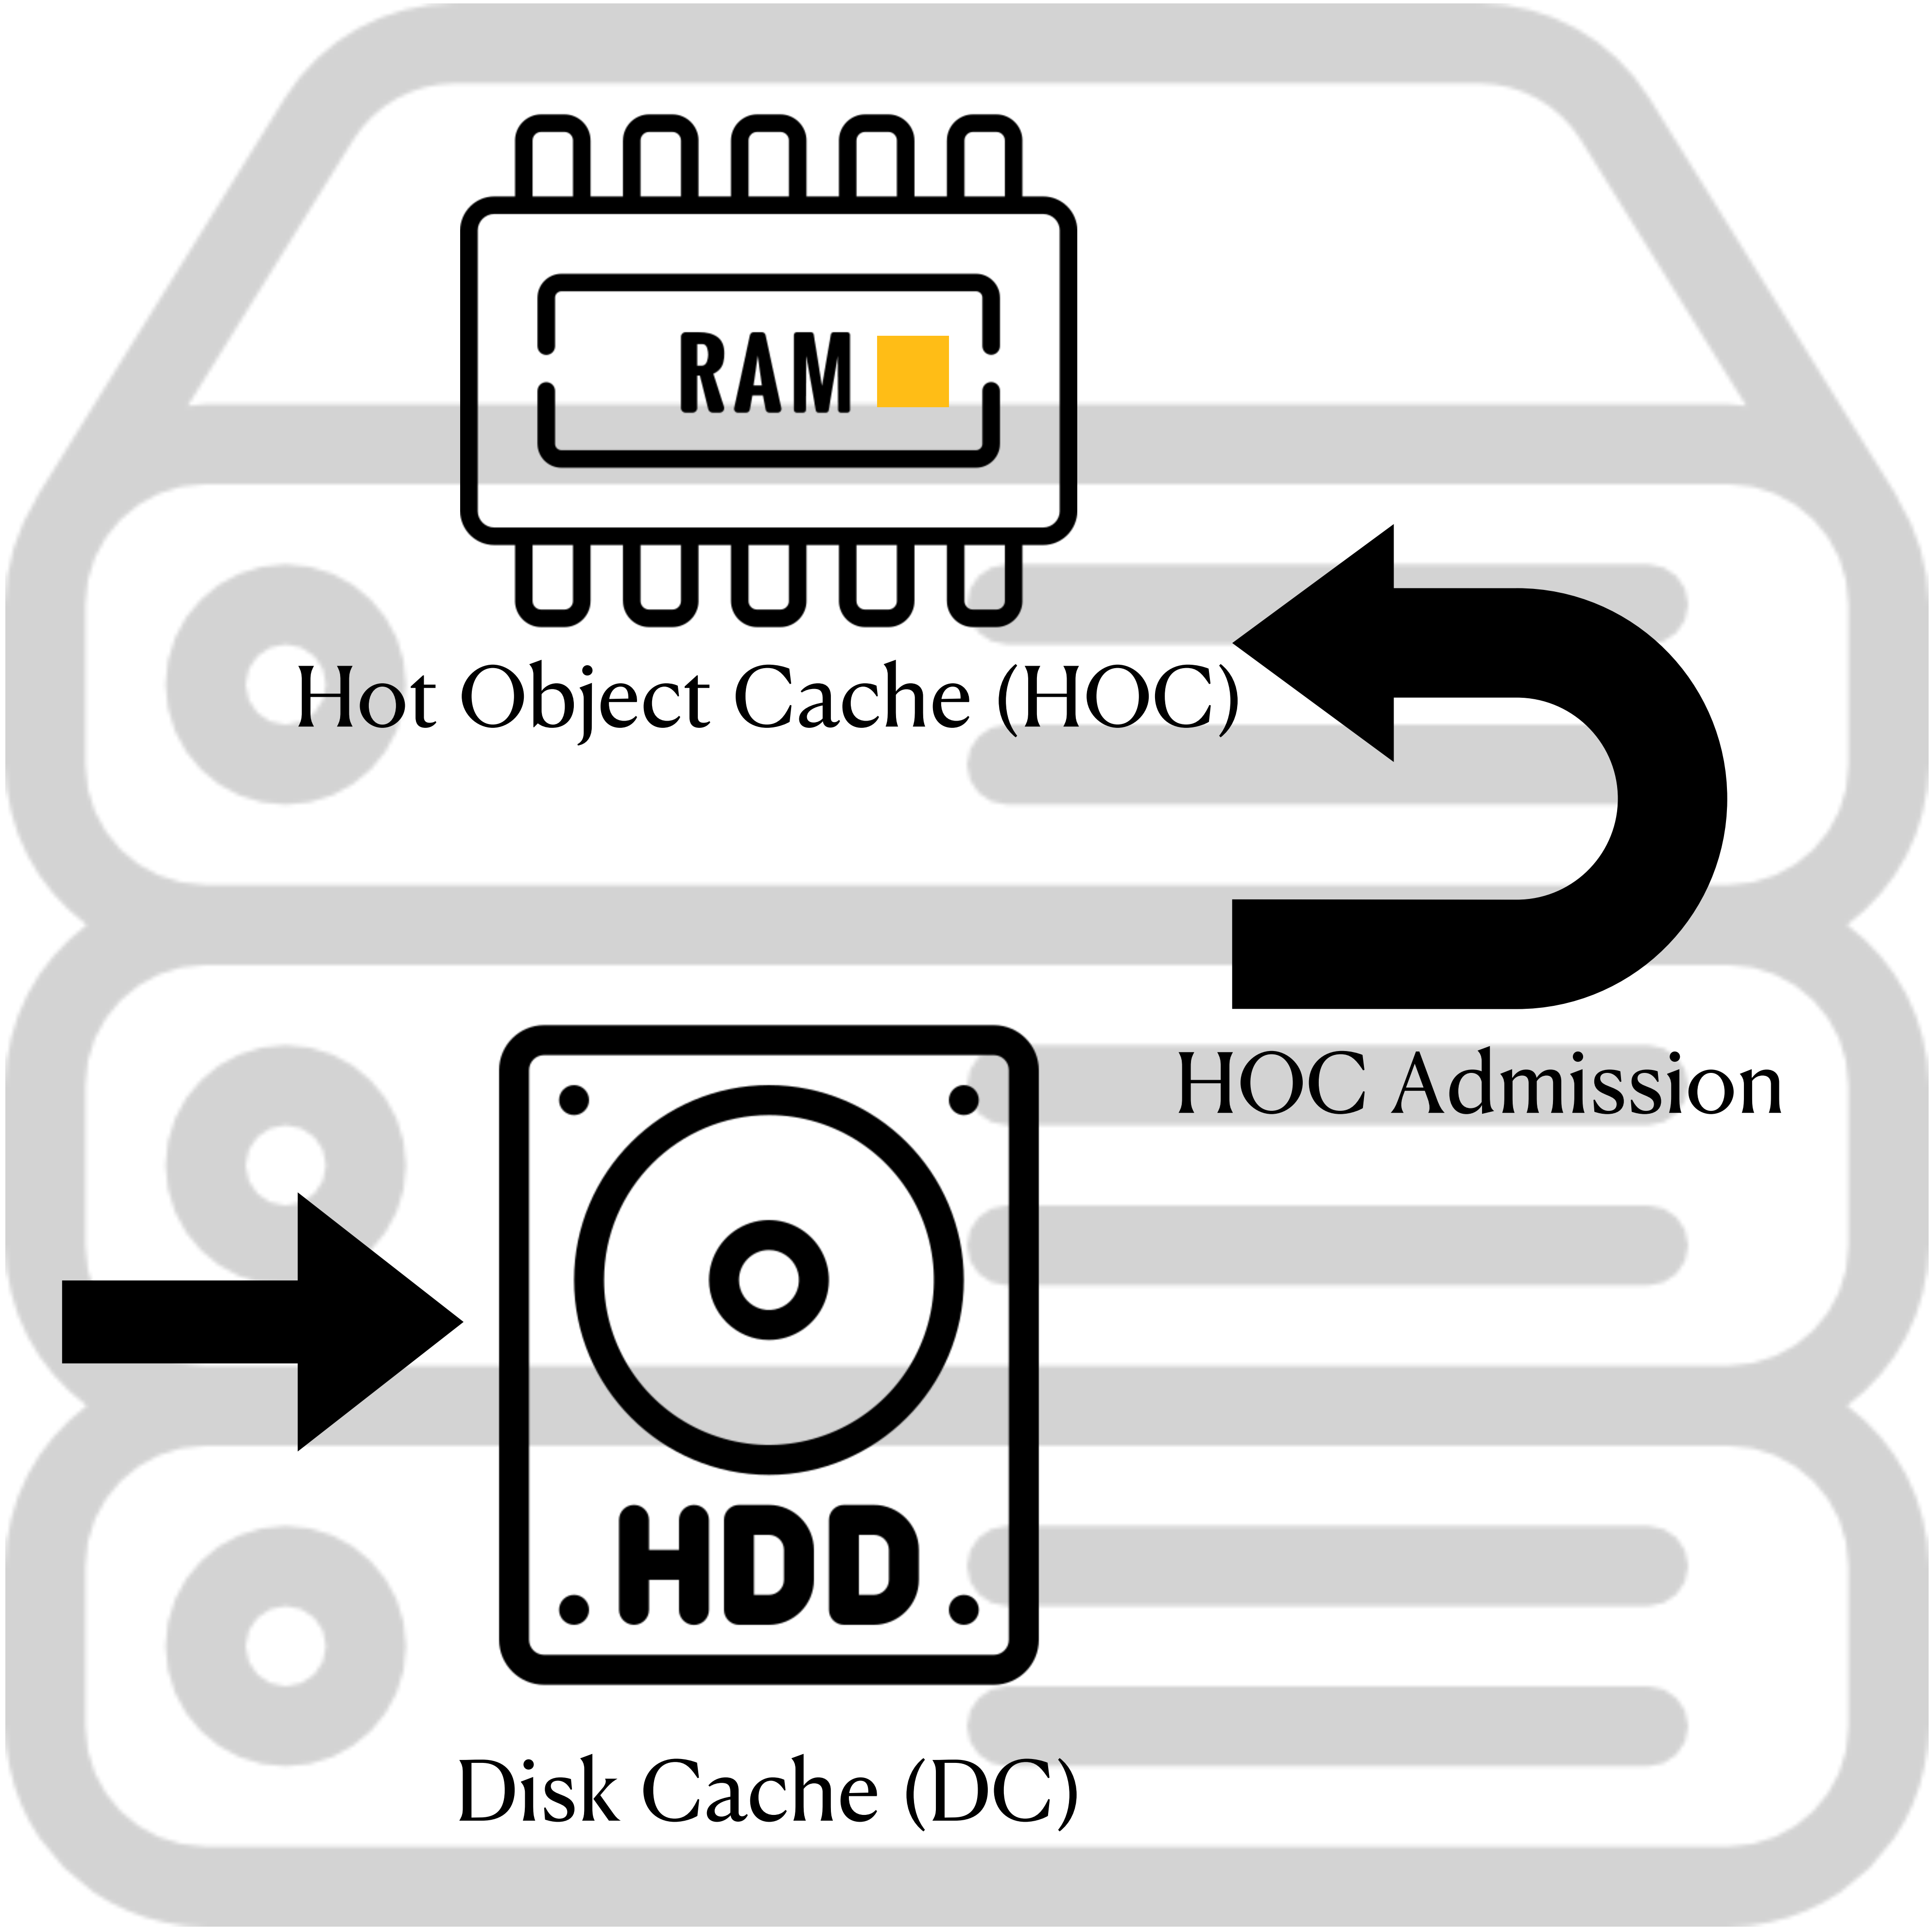

CDN Server

A common policy:
frequency $\geq$ f, size $\leq$ s

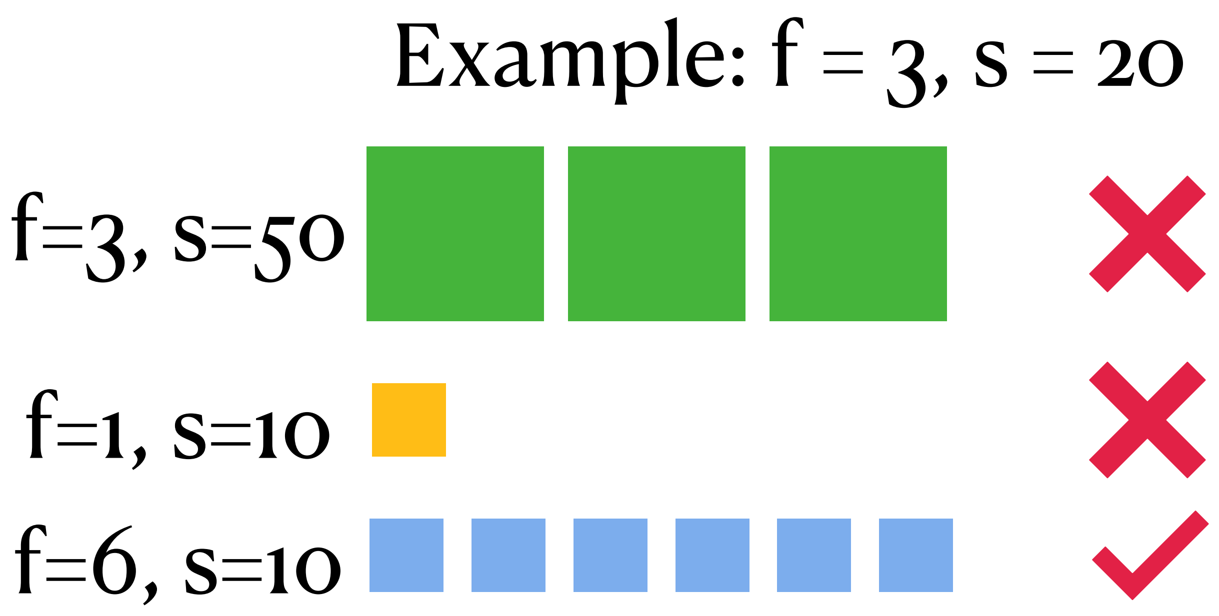

Hot Object Cache (HOC) Admission Policy

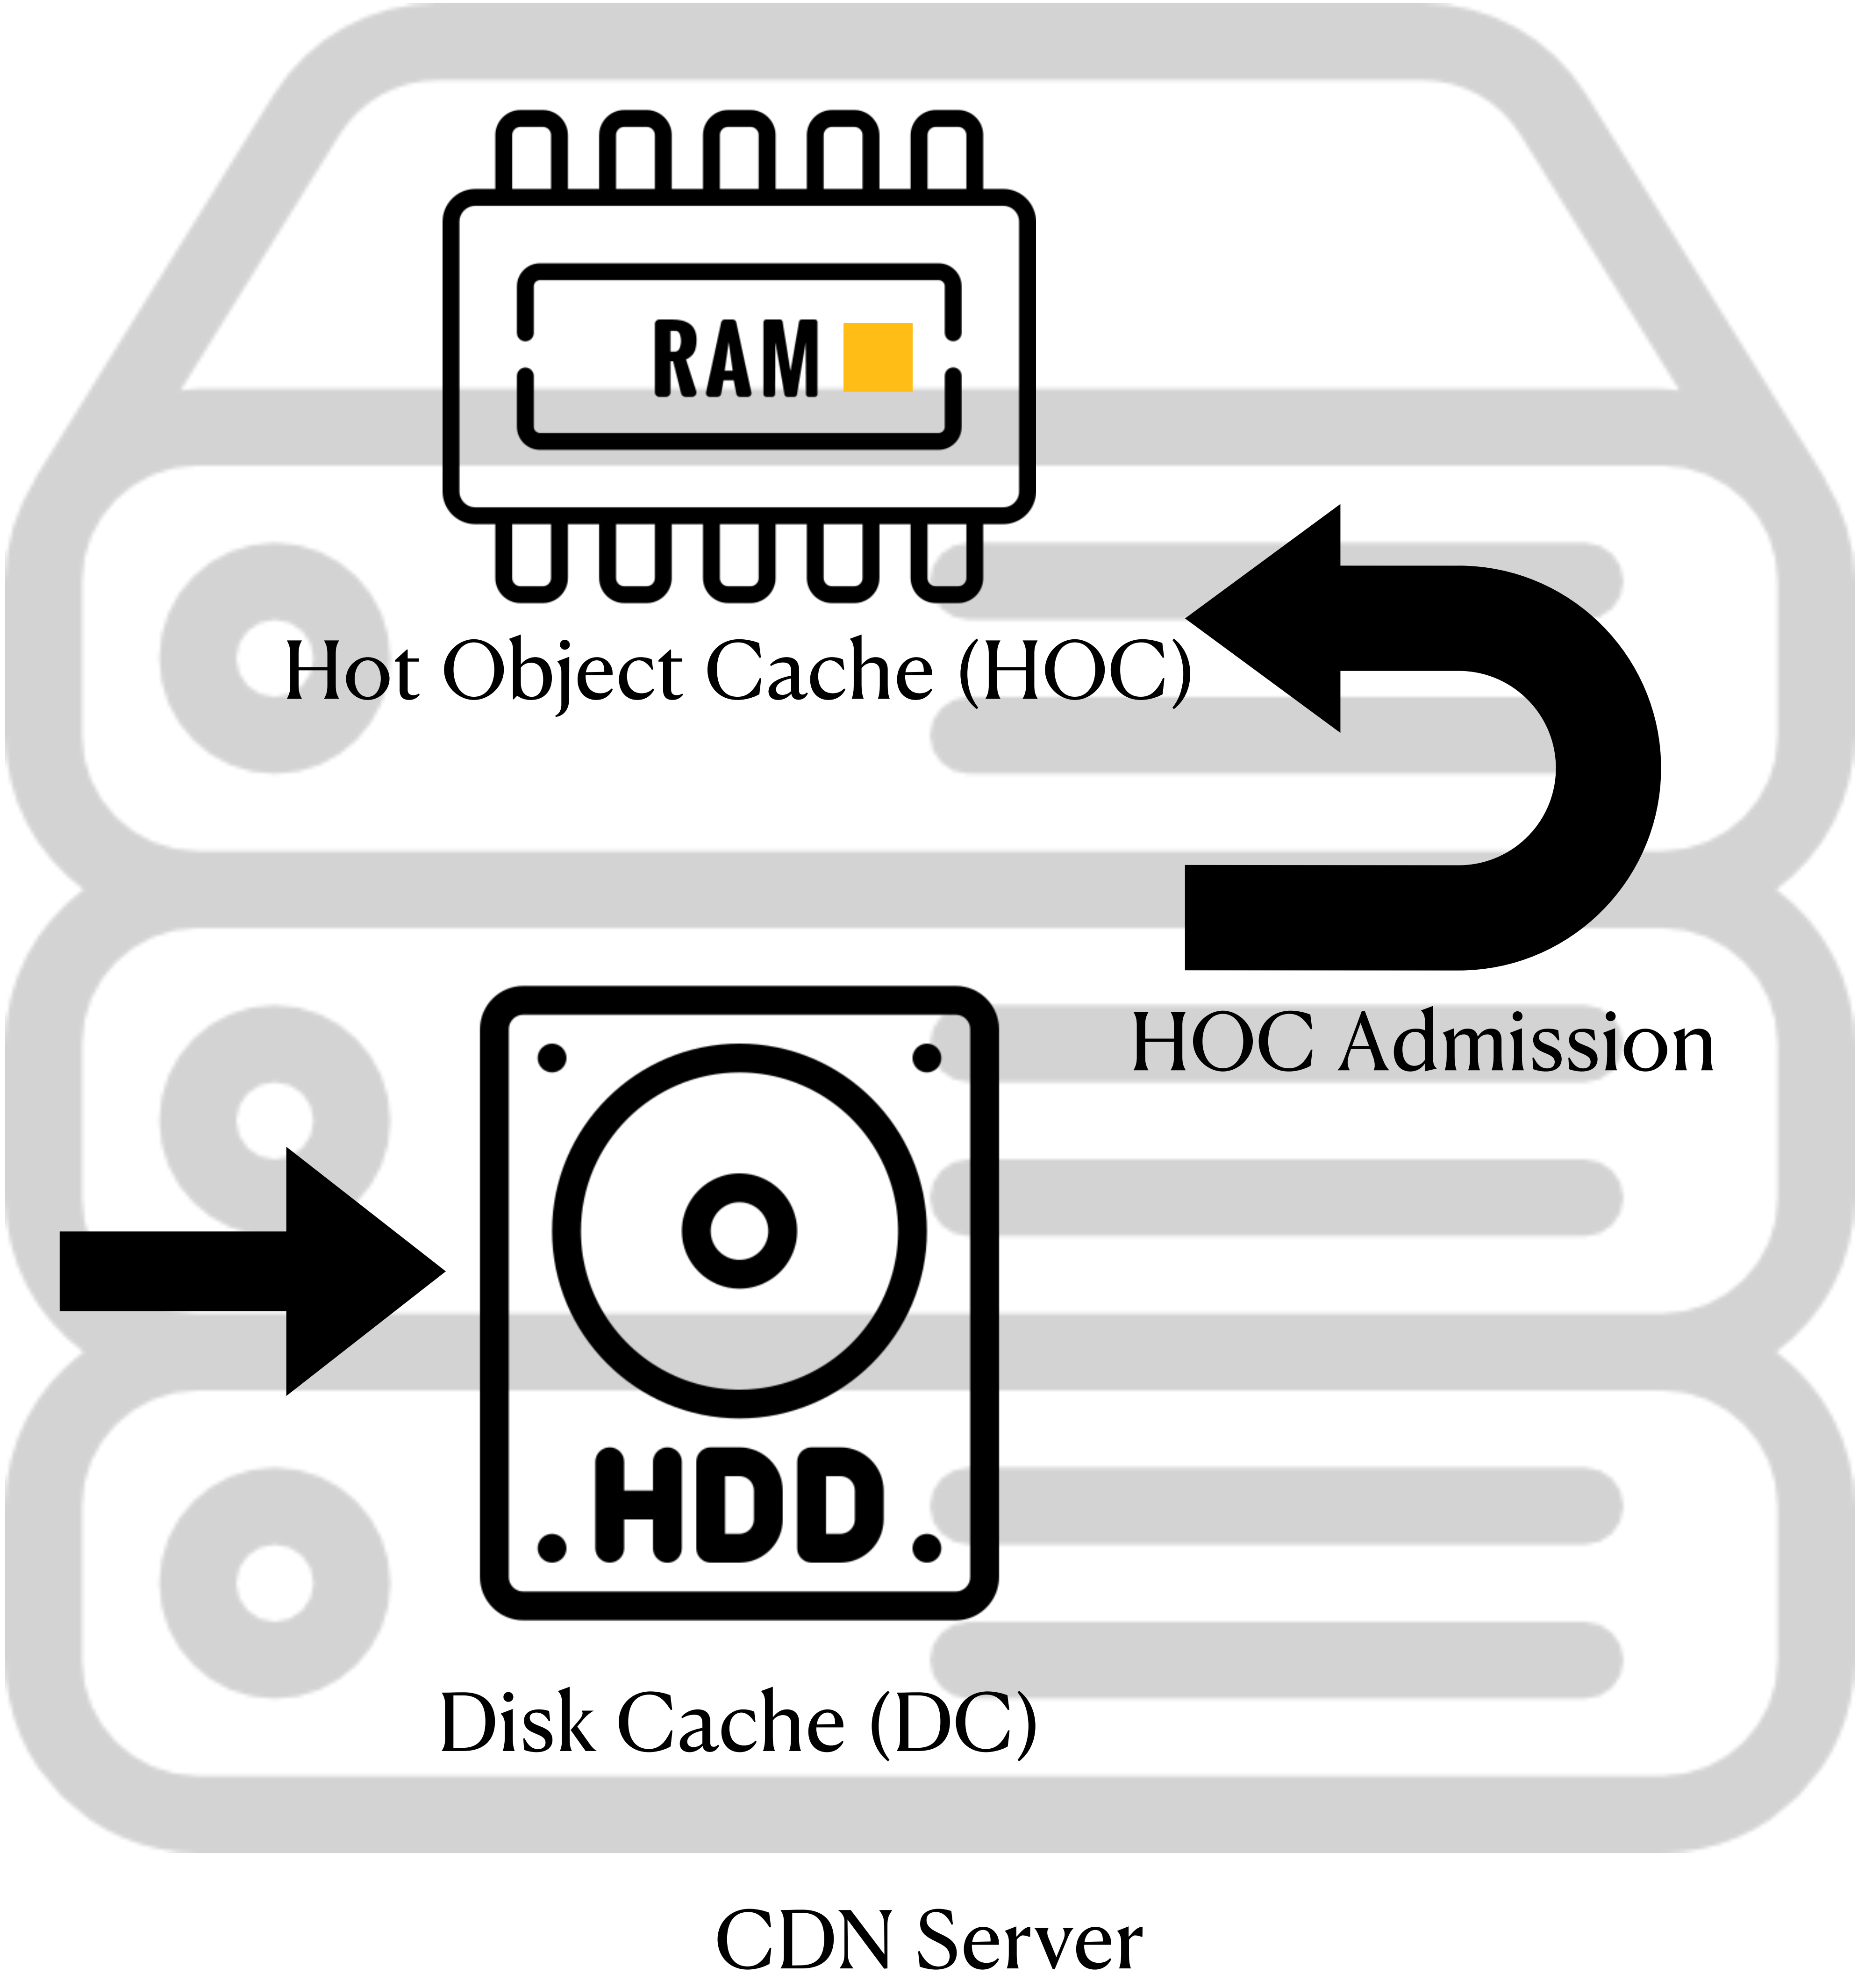

A common policy:
frequency $\geq$ f, size $\leq$ s

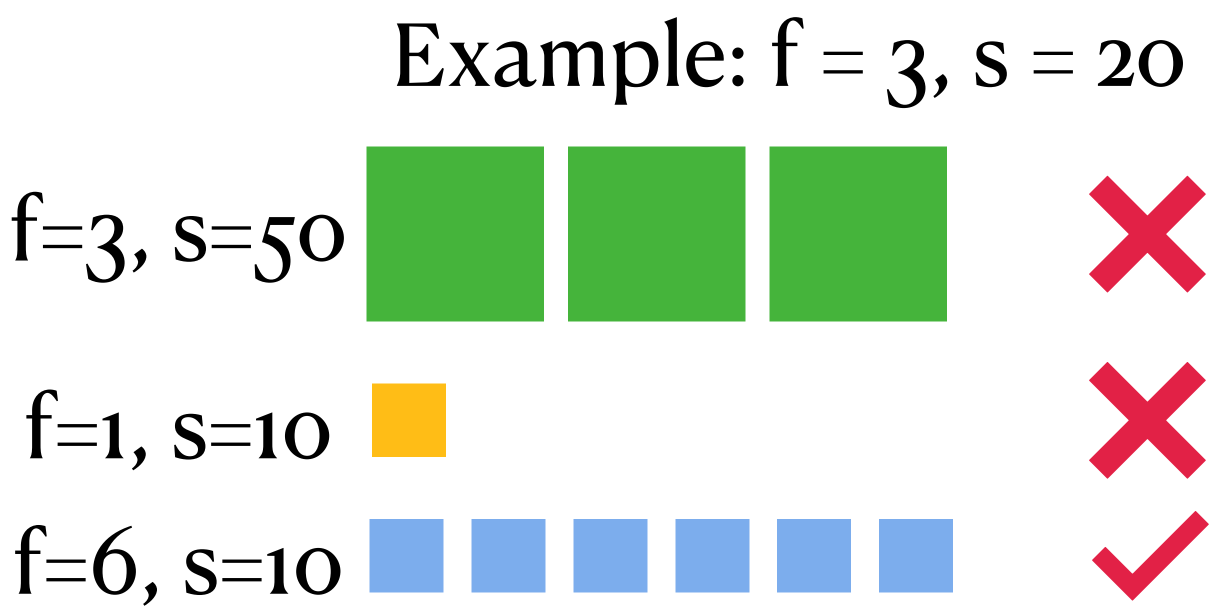

Metric: Object Hit Rate (OHR)

$$\text{HOC OHR} = \frac{\#\text{HOC Hits}}{\#\text{Requests}}$$

Static HOC admission policies fall short

Static HOC admission policies fall short

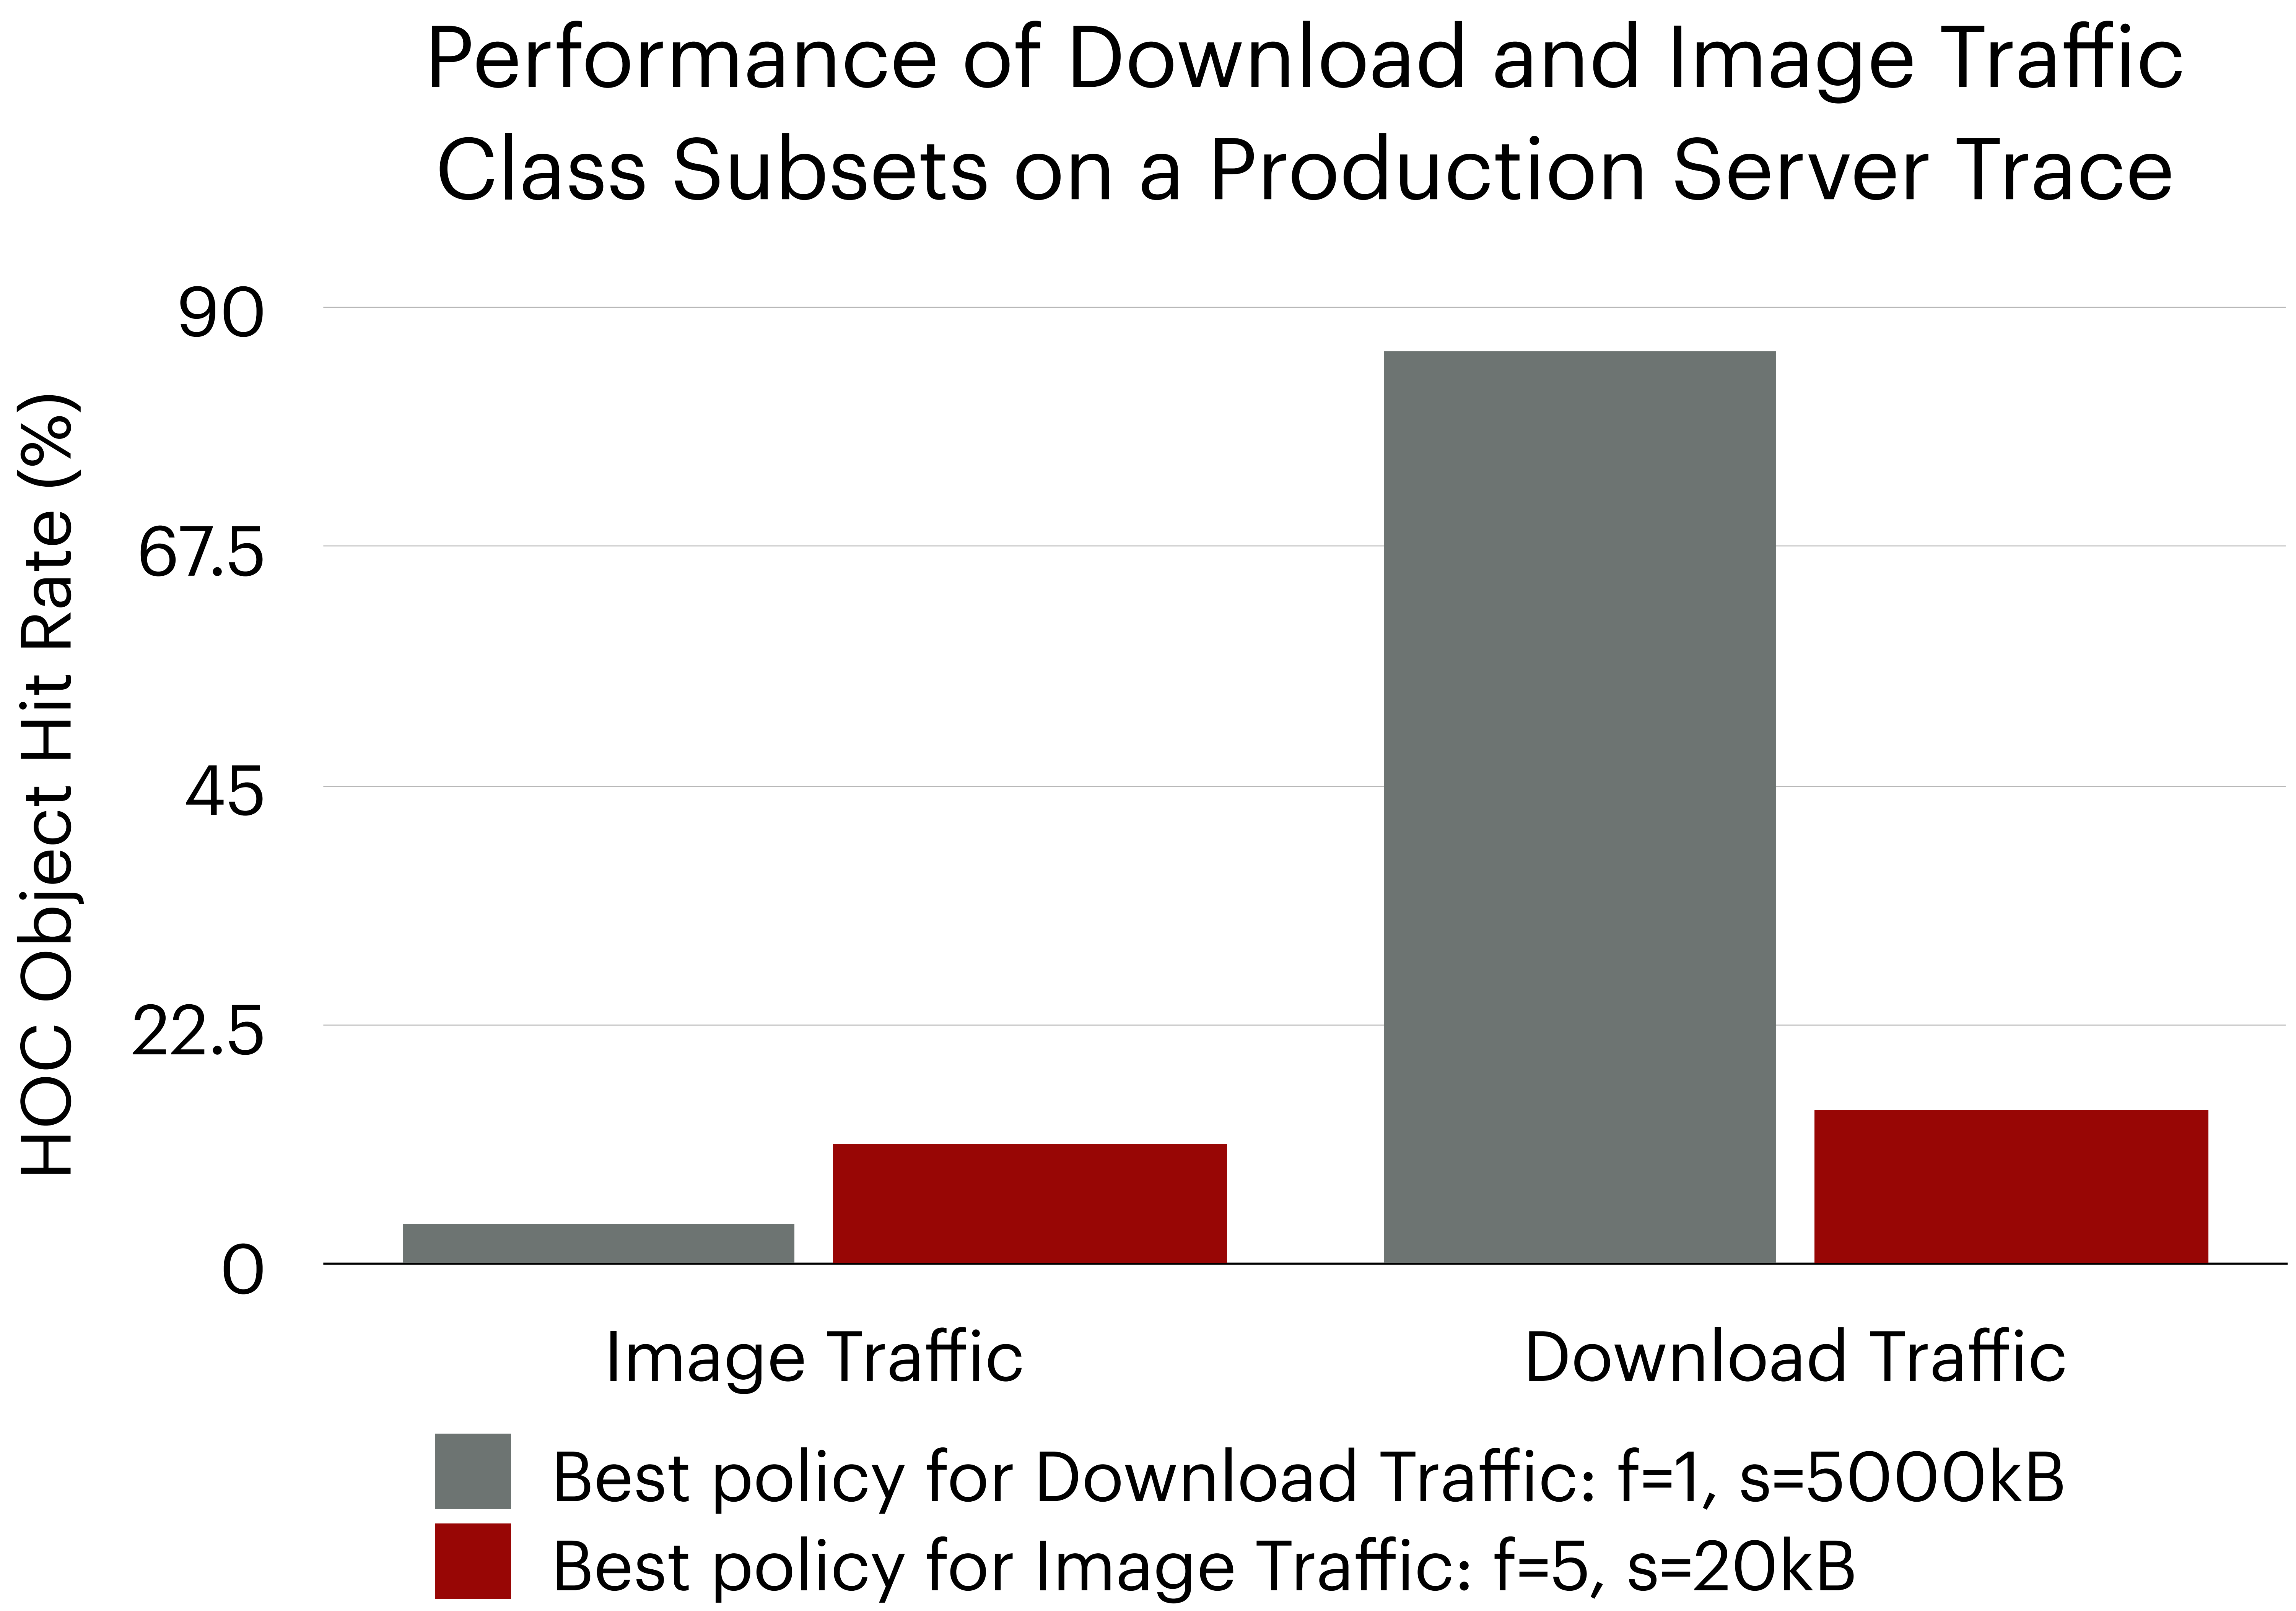

Static HOC admission policies fall short

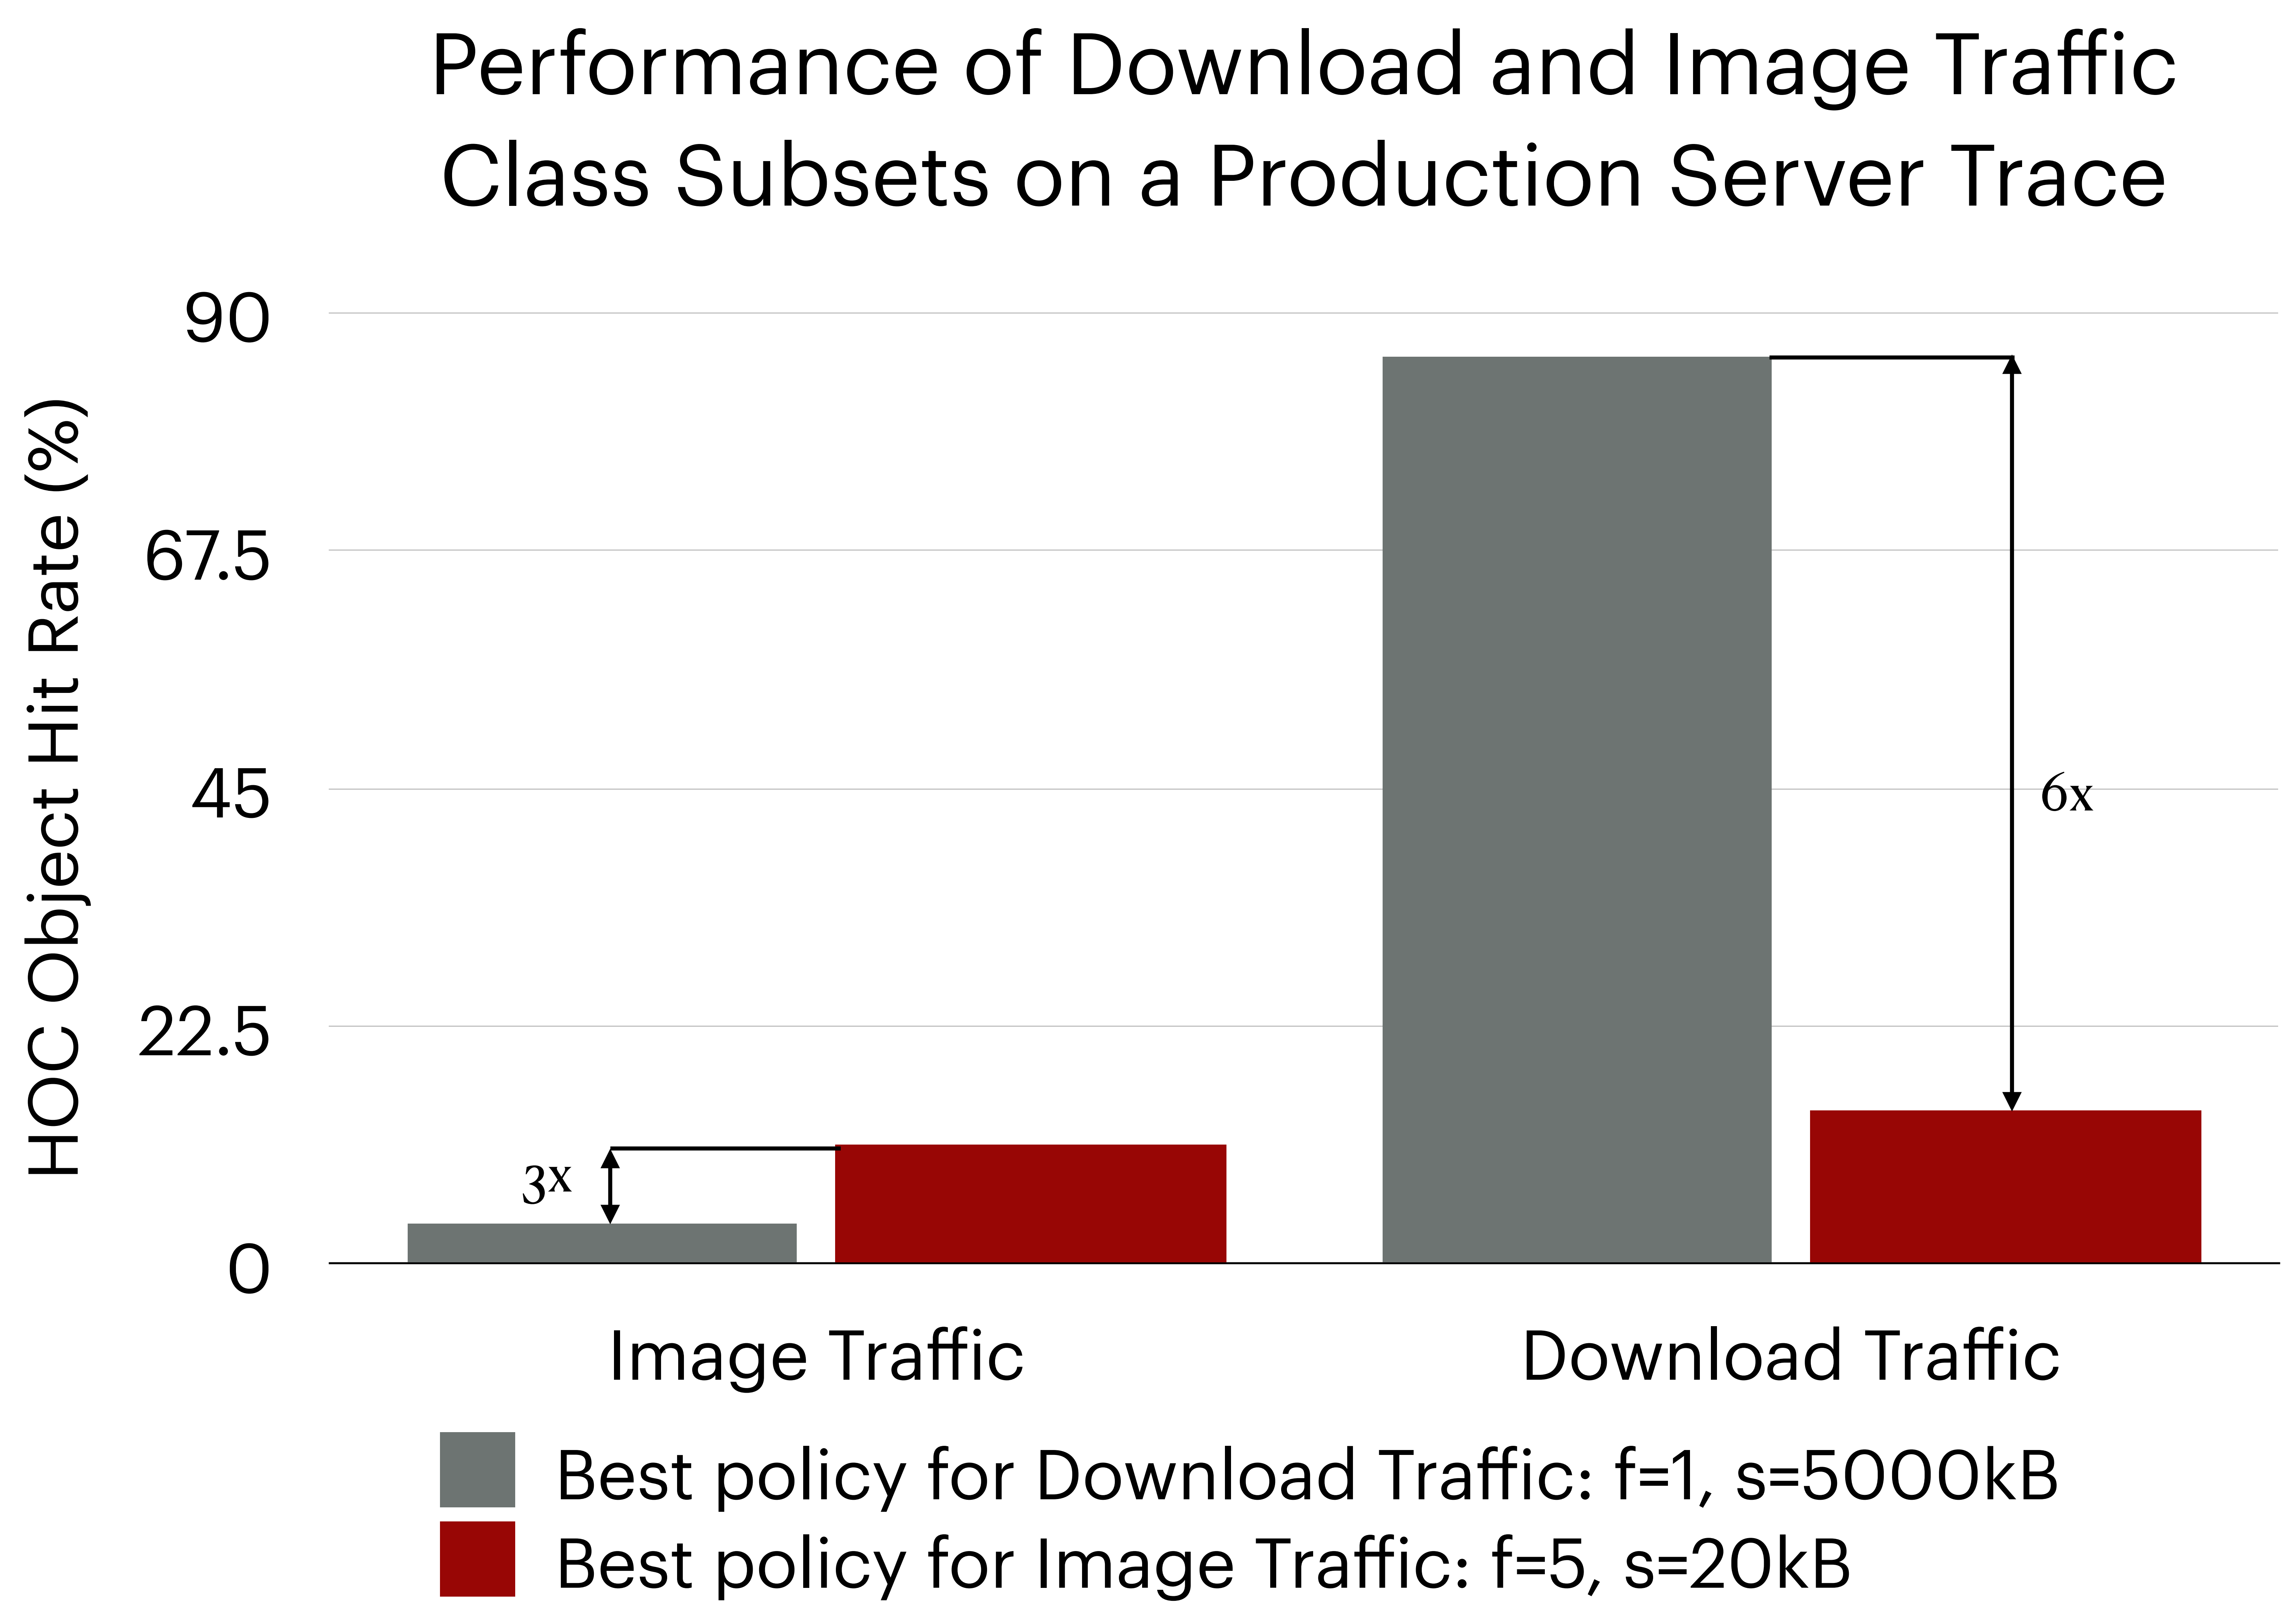

No one-size-fits-all static policy.

No one-size-fits-all static policy.
Can we learn the optimal policy for the current traffic?

Issues with Prior Adaptive Admission Schemes

Issues with Prior Adaptive Admission Schemes

- Restrict the **policy decision knobs**
    - *AdaptSize@NSDI'17* can only adapt size threshold.

Issues with Prior Adaptive Admission Schemes

- Restrict the **policy decision knobs**
  - *AdaptSize@NSDI'17* can only adapt size threshold.
- Don't accommodate **hardware-dependent metrics**
  - Shadow cache-based approaches (e.g. *HillClimbing@NSDI'17)* cannot model Disk Ops.

Issues with Prior Adaptive Admission Schemes

- Restrict the **policy decision knobs**
  - *AdaptSize@NSDI'17* can only adapt size threshold.
- Don't accommodate **hardware-dependent metrics**
  - Shadow cache-based approaches (e.g. *HillClimbing@NSDI'17)* cannot model disk Ops.
- Impose high **overhead**
  - *RL-Cache@NetAI'19* performs per-request inference.

Issues with Prior Adaptive Admission Schemes

- Restrict the **policy decision knobs**
  - *AdaptSize@NSDI'17* can only adapt size threshold.
- Don't accommodate **hardware-dependent metrics**
  - Shadow cache-based approaches (e.g. *HillClimbing@NSDI'17)* cannot model disk Ops.
- Impose high **overhead**
  - *RL-Cache@NetAI'19* performs per-request inference.

Issues with Prior Adaptive Admission Schemes

Darwin

- Restrict the **policy decision knobs**
  - *AdaptSize@NSDI'17* can only adapt size threshold.
- Don't accommodate **hardware-dependent metrics**
  - Shadow cache-based approaches (e.g. *HillClimbing@NSDI'17)* cannot model disk Ops.
- Impose high **overhead**
  - *RL-Cache@NetAI'19* performs per-request inference.

Issues with Prior Adaptive Admission Schemes

Darwin

Unrestricted Knobs ✓

- Restrict the **policy decision knobs**
  - *AdaptSize@NSDI'17* can only adapt size threshold.
- Don't accommodate **hardware-dependent metrics**
  - Shadow cache-based approaches (e.g. *HillClimbing@NSDI'17)* cannot model disk Ops.
- Impose high **overhead**
  - *RL-Cache@NetAI'19* performs per-request inference.

Issues with Prior Adaptive Admission Schemes

- Restrict the **policy decision knobs**
  - *AdaptSize@NSDI'17* can only adapt size threshold.
- Don't accommodate **hardware-dependent metrics**
  - Shadow cache-based approaches (e.g. *HillClimbing@NSDI'17)* cannot model disk Ops.
- Impose high **overhead**
  - *RL-Cache@NetAI'19* performs per-request inference.

Darwin

Unrestricted Knobs ✓

Hardware-dependent Metrics ✓

Issues with Prior Adaptive Admission Schemes

- Restrict the **policy decision knobs**
  - *AdaptSize@NSDI'17* can only adapt size threshold.
- Don't accommodate **hardware-dependent metrics**
  - Shadow cache-based approaches (e.g. *HillClimbing@NSDI'17)* cannot model disk Ops.
- Impose high **overhead**
  - *RL-Cache@NetAI'19* performs per-request inference.

Darwin

Unrestricted Knobs ✓

Hardware-dependent Metrics ✓

Low Overhead ✓

Darwin Overview

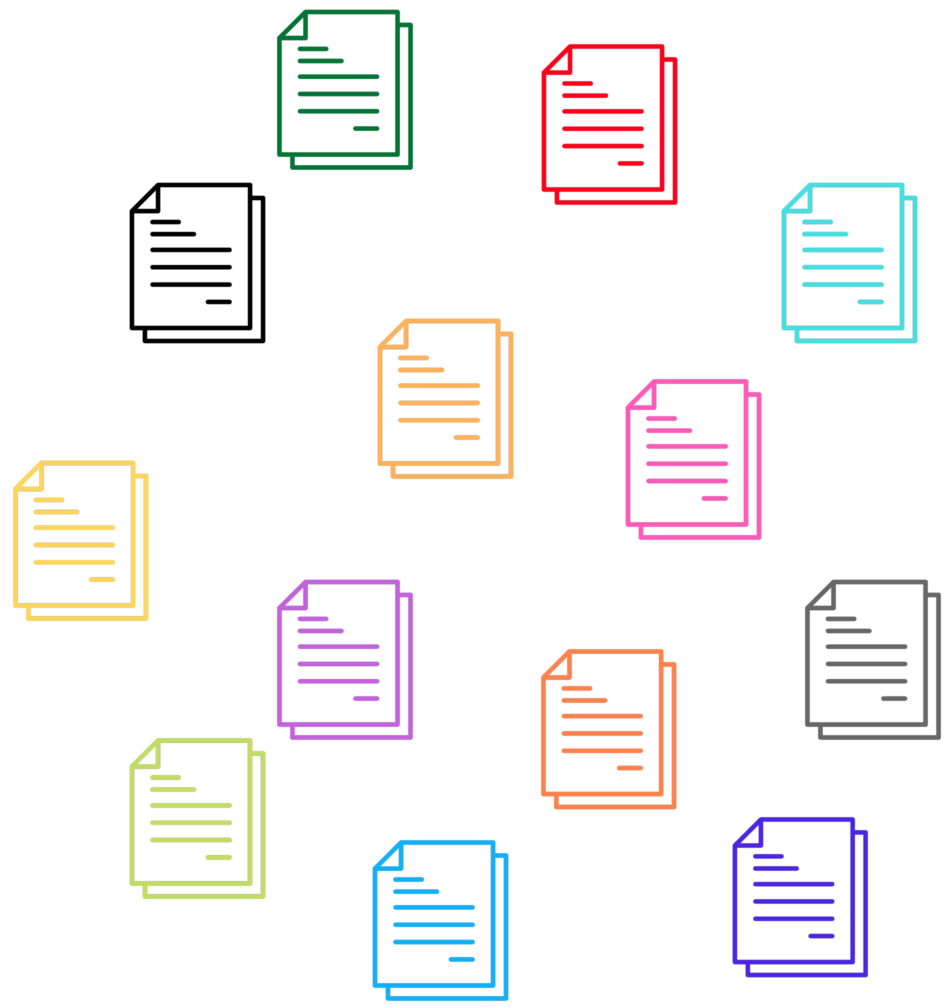

Policies (f, s)

Darwin Overview

Performance Evaluation

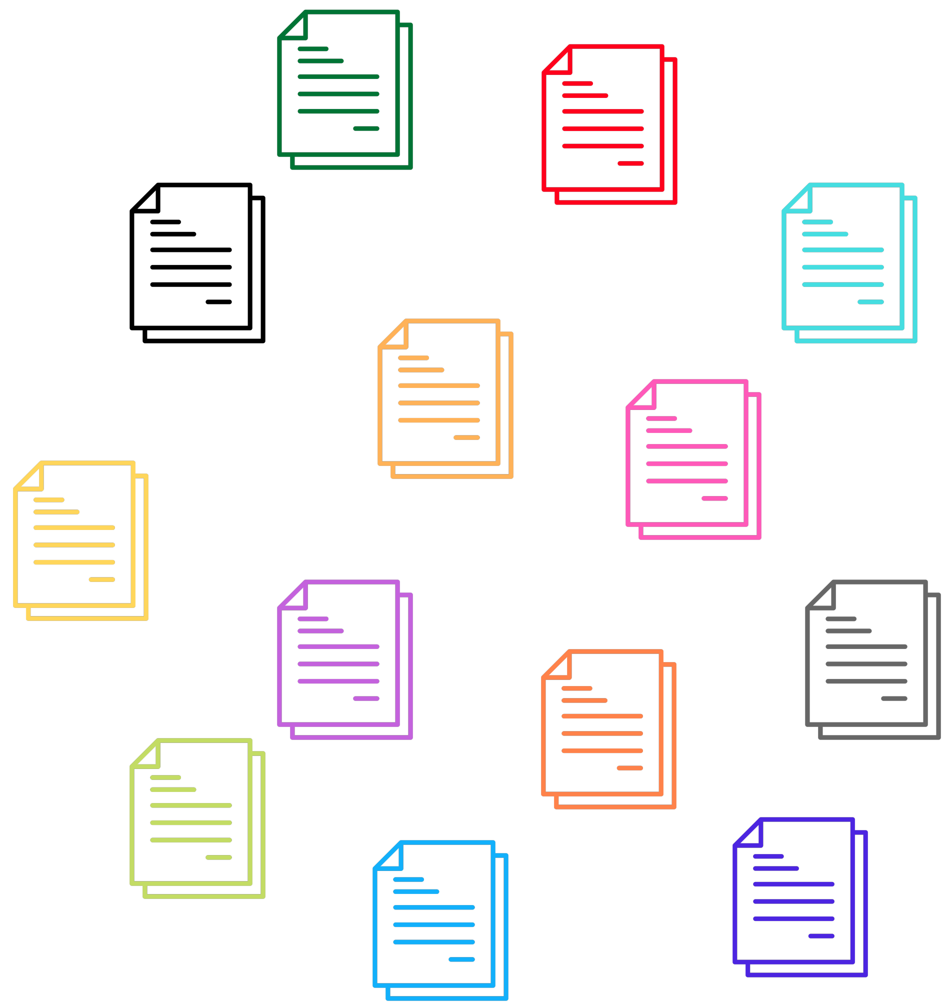

Policies (f, s)

Performance a

Policy a: f = 2, s = 20

Performance b

Policy b: f = 2, s = 500

Performance c

Policy c: f = 3, s = 20

⋮

Darwin Overview

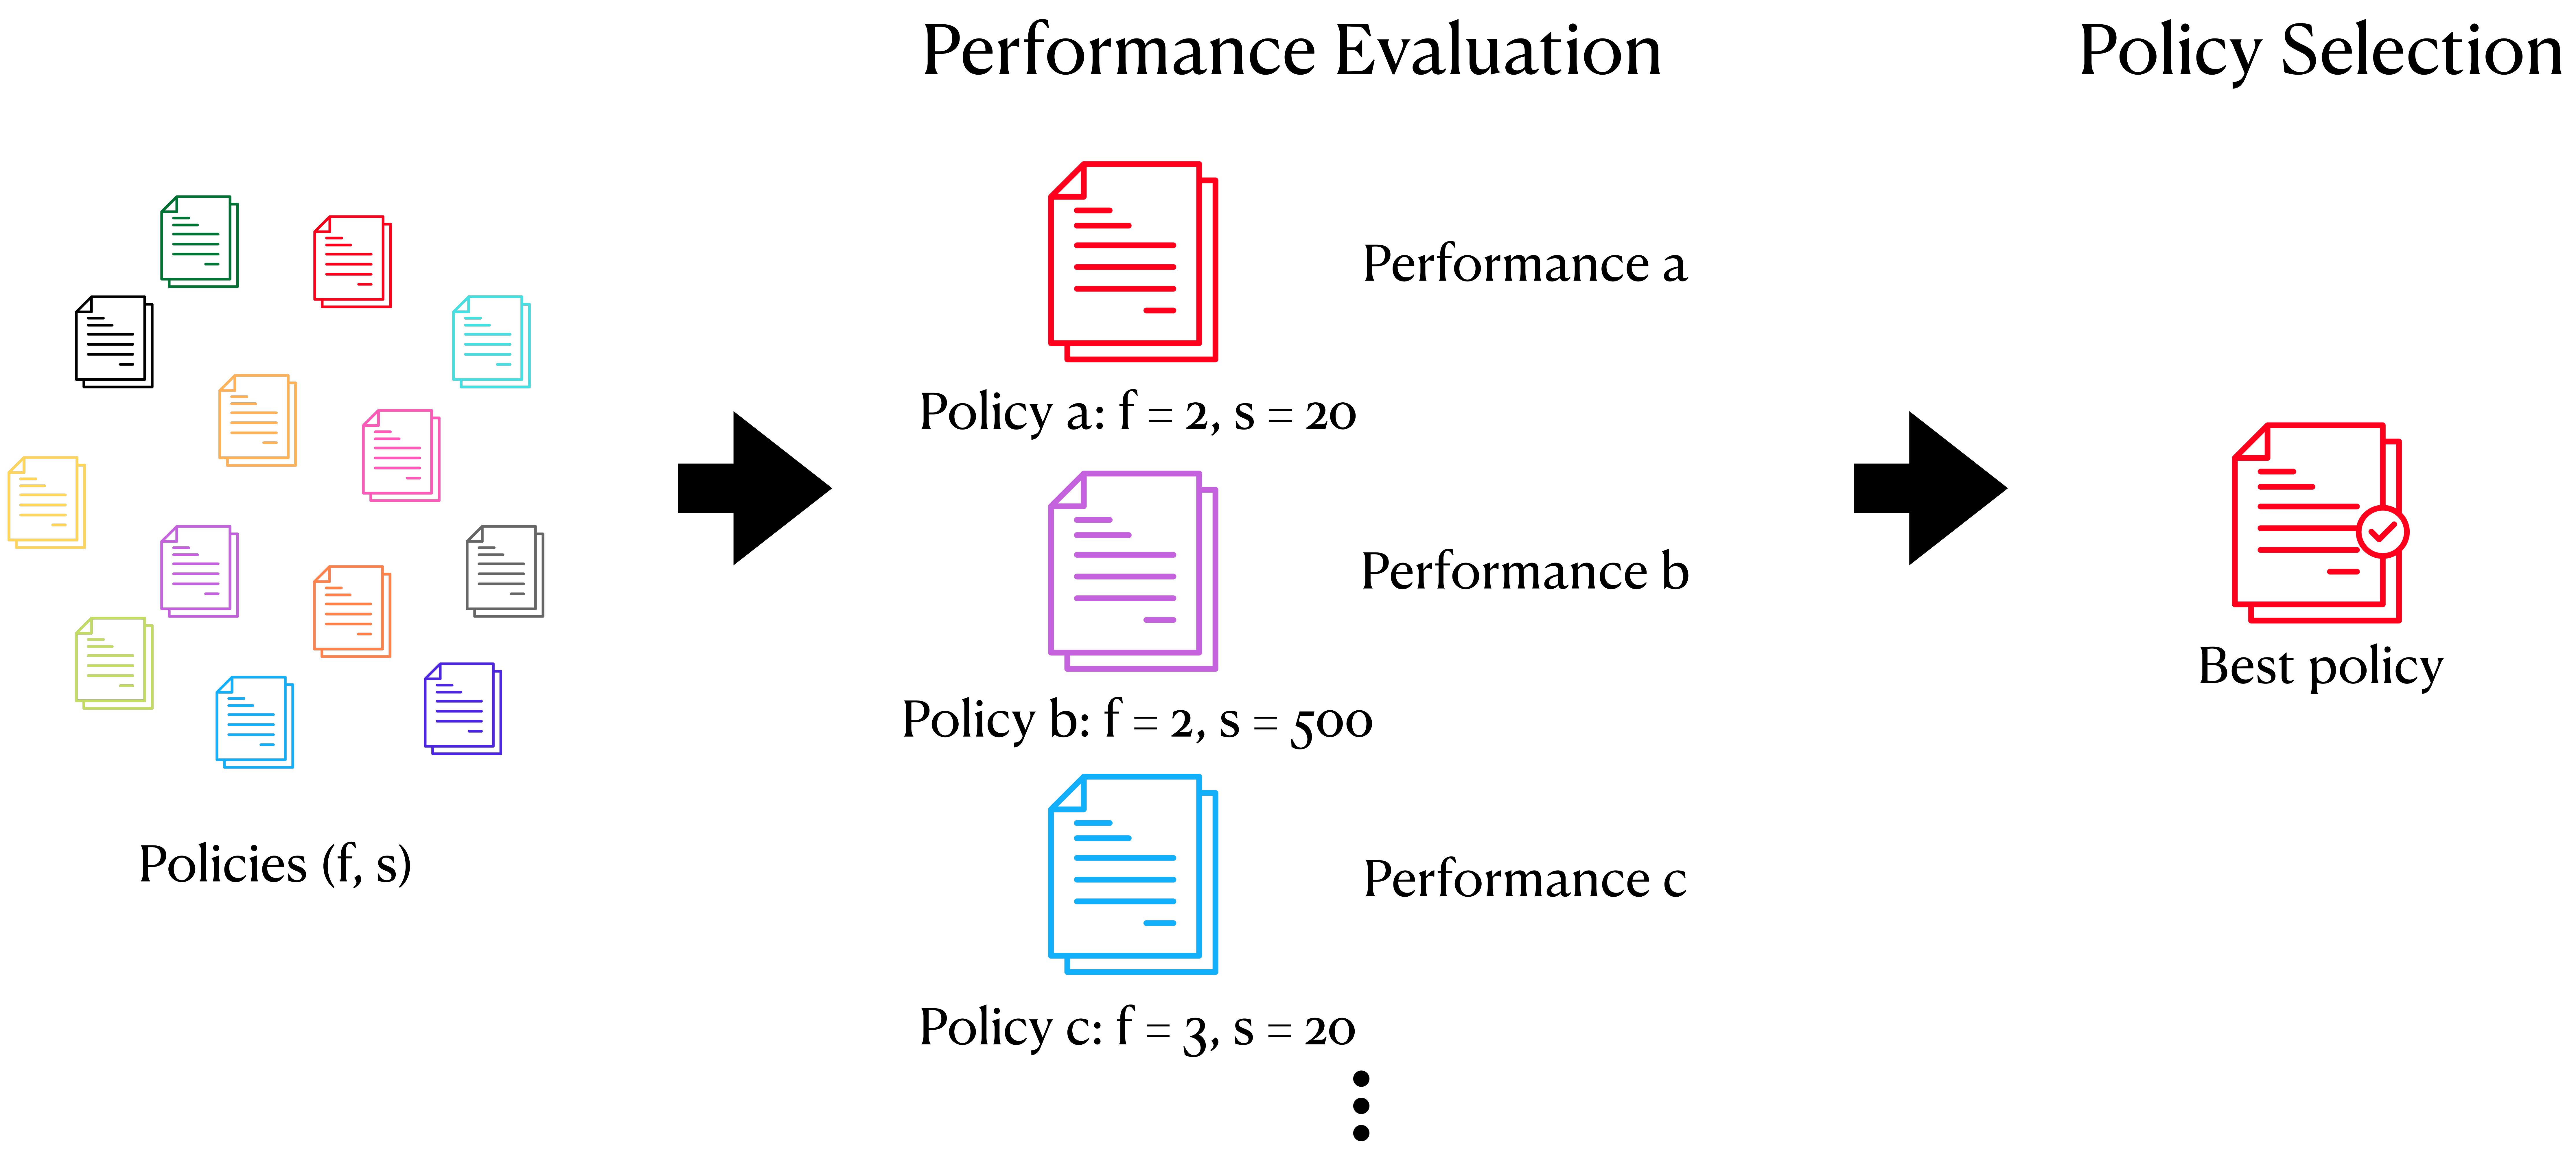

Darwin Overview

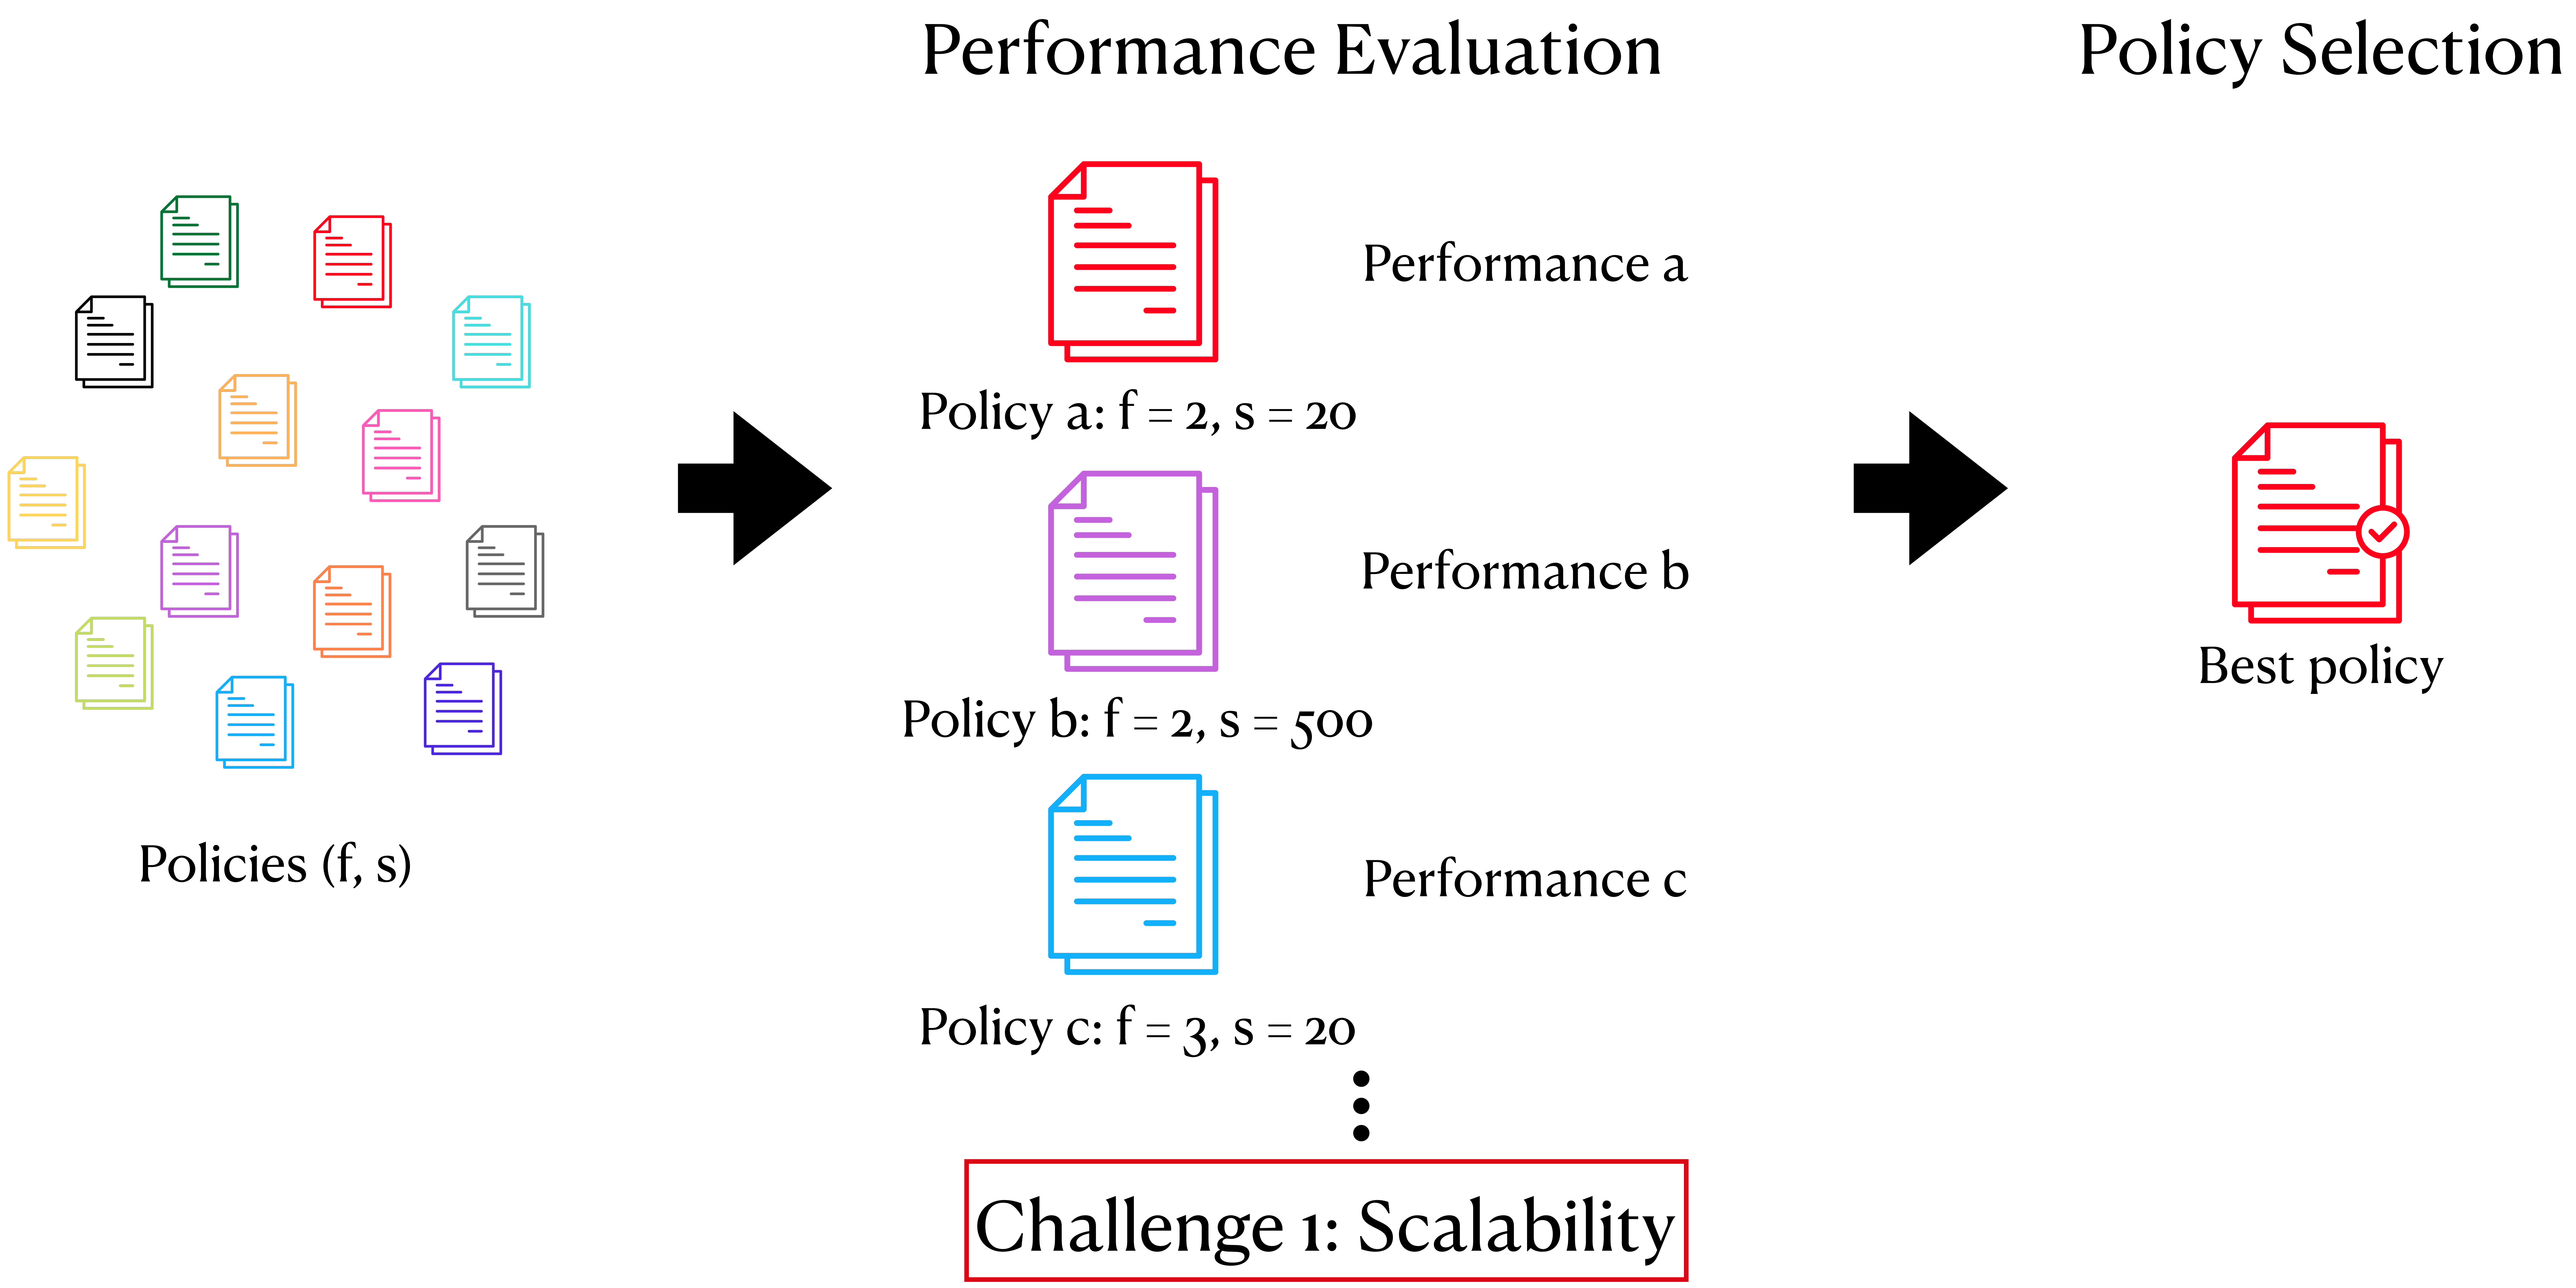

Darwin Overview

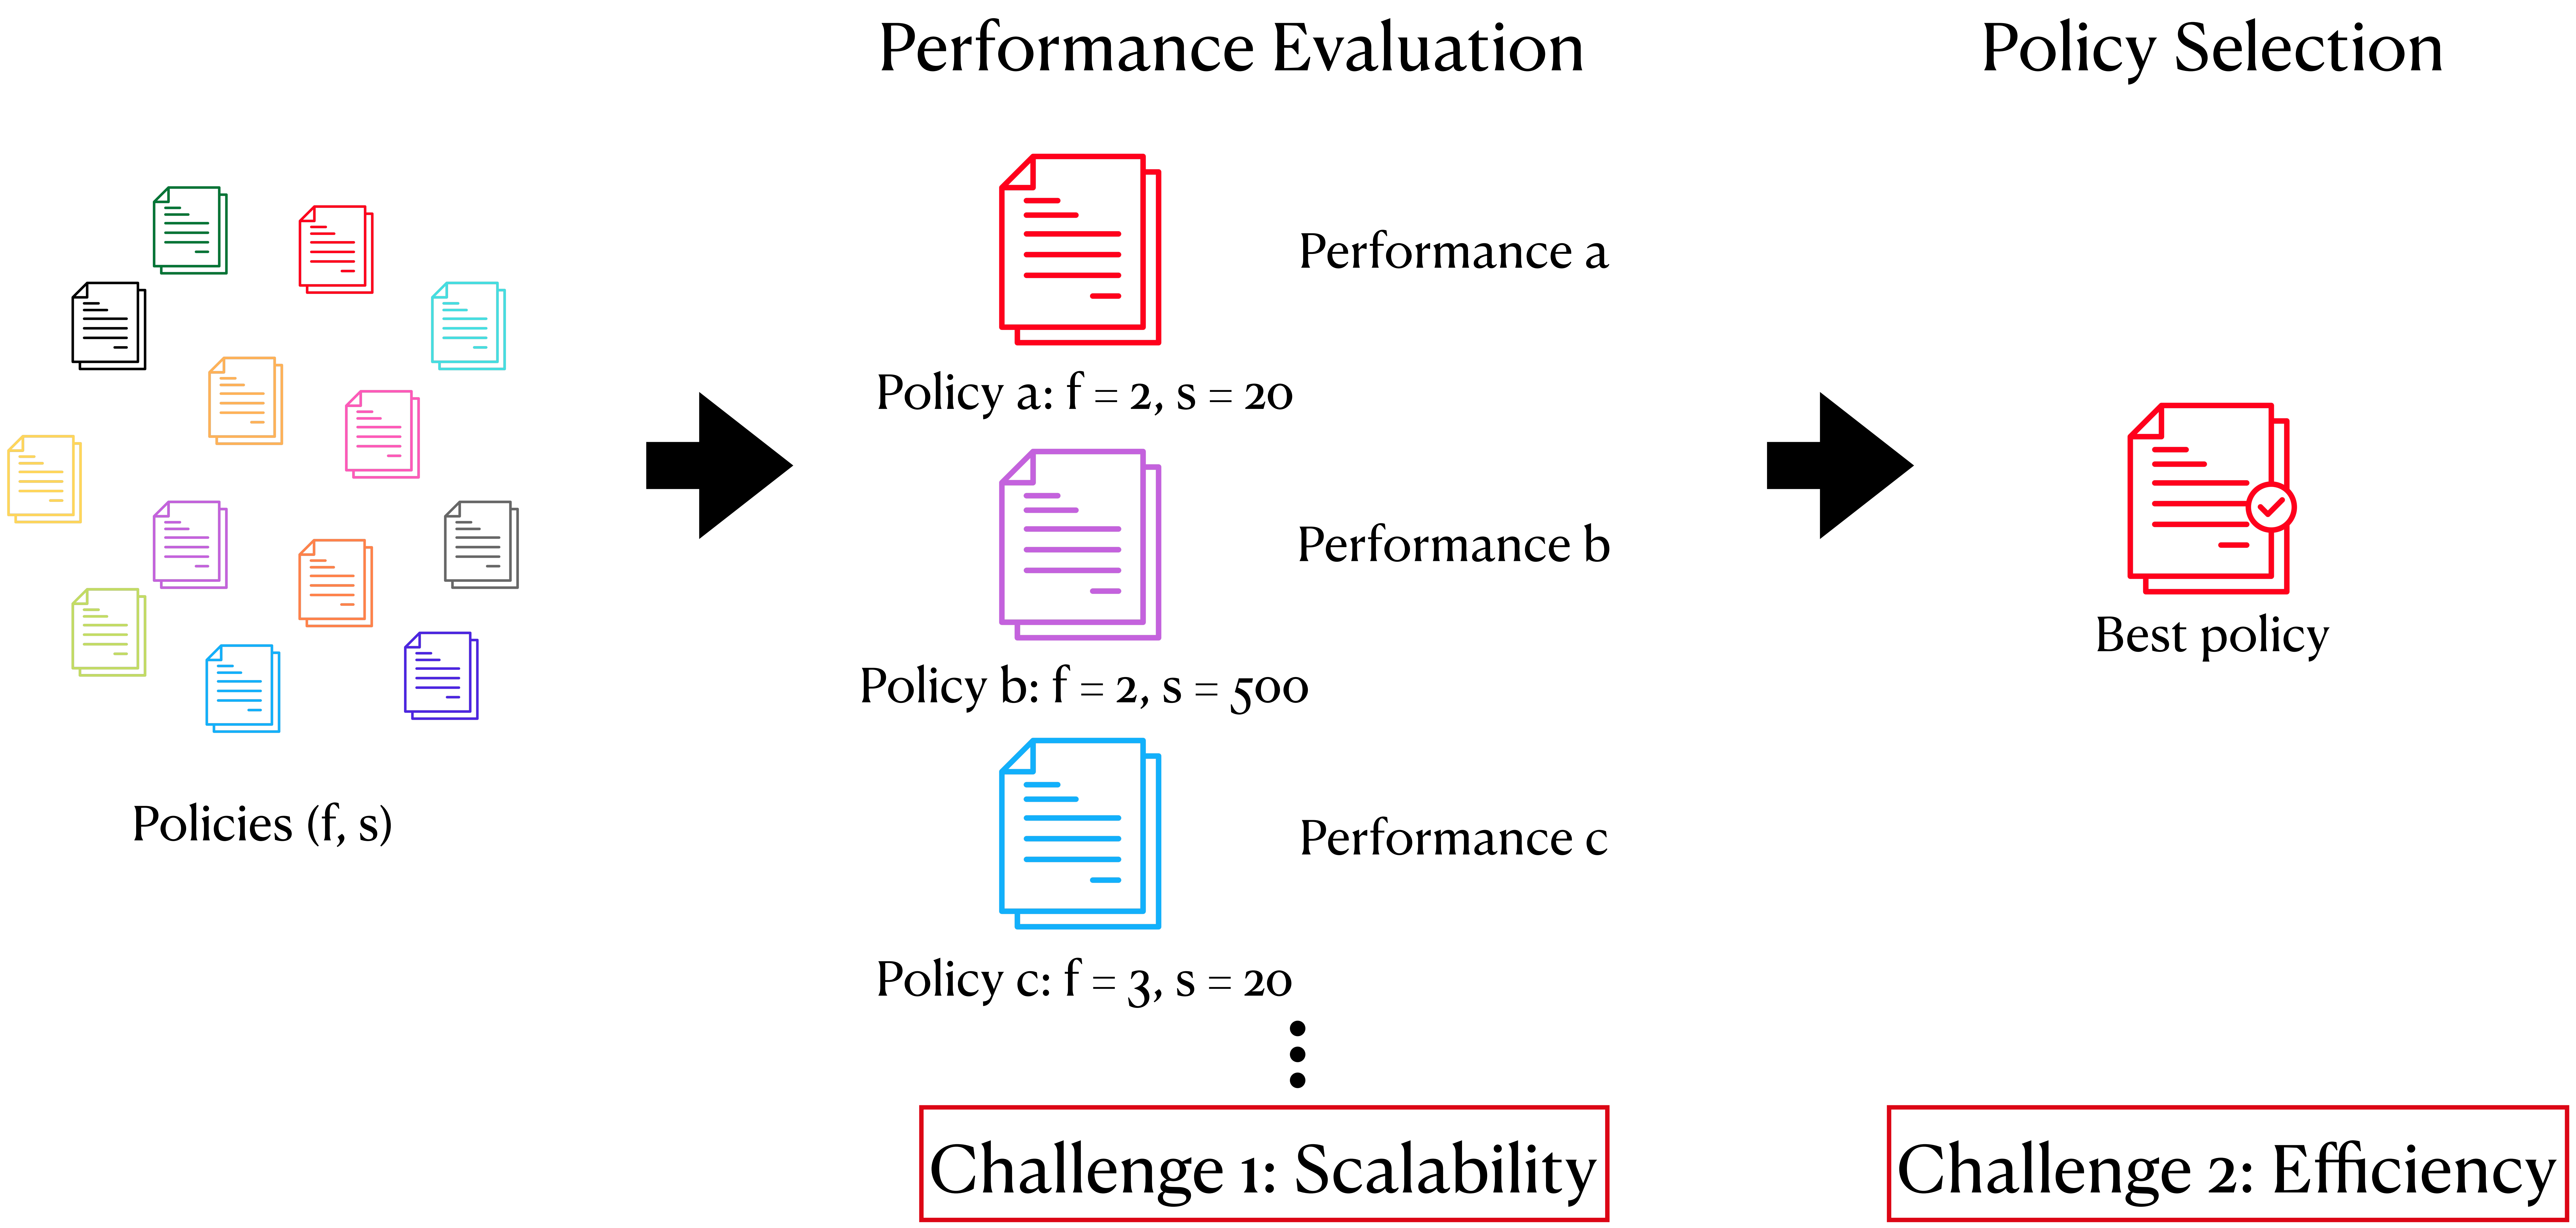

Challenge 1: Scalable Performance Observation

Challenge 1: Scalable Performance Observation

Sequential

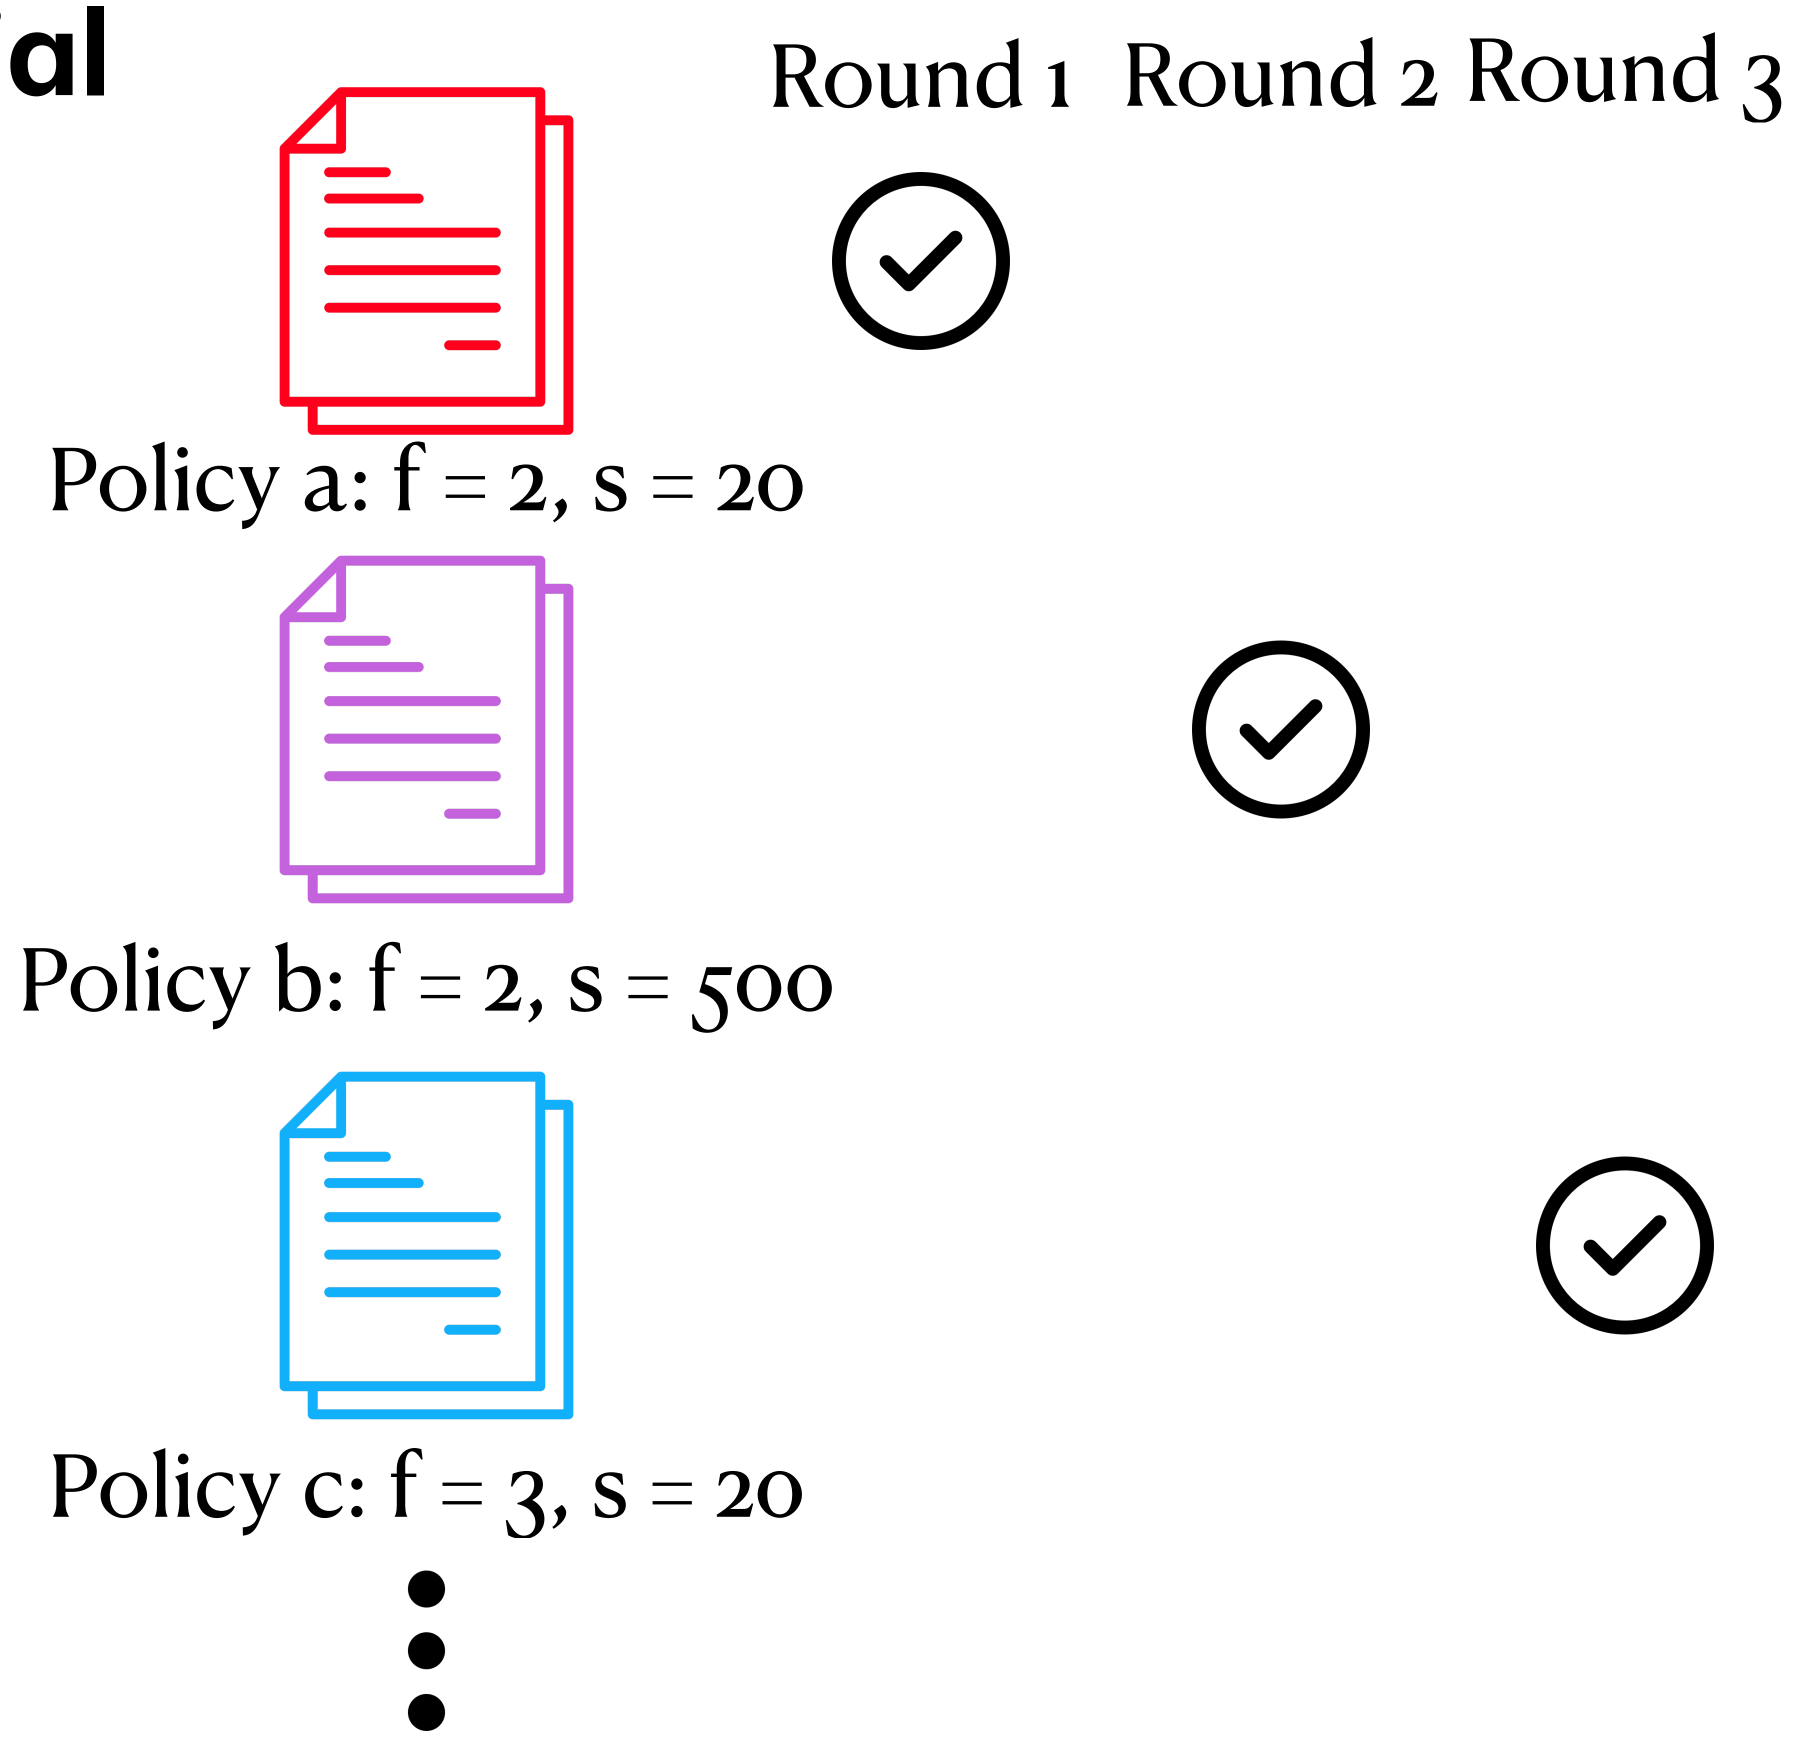

Challenge 1: Scalable Performance Observation

Sequential

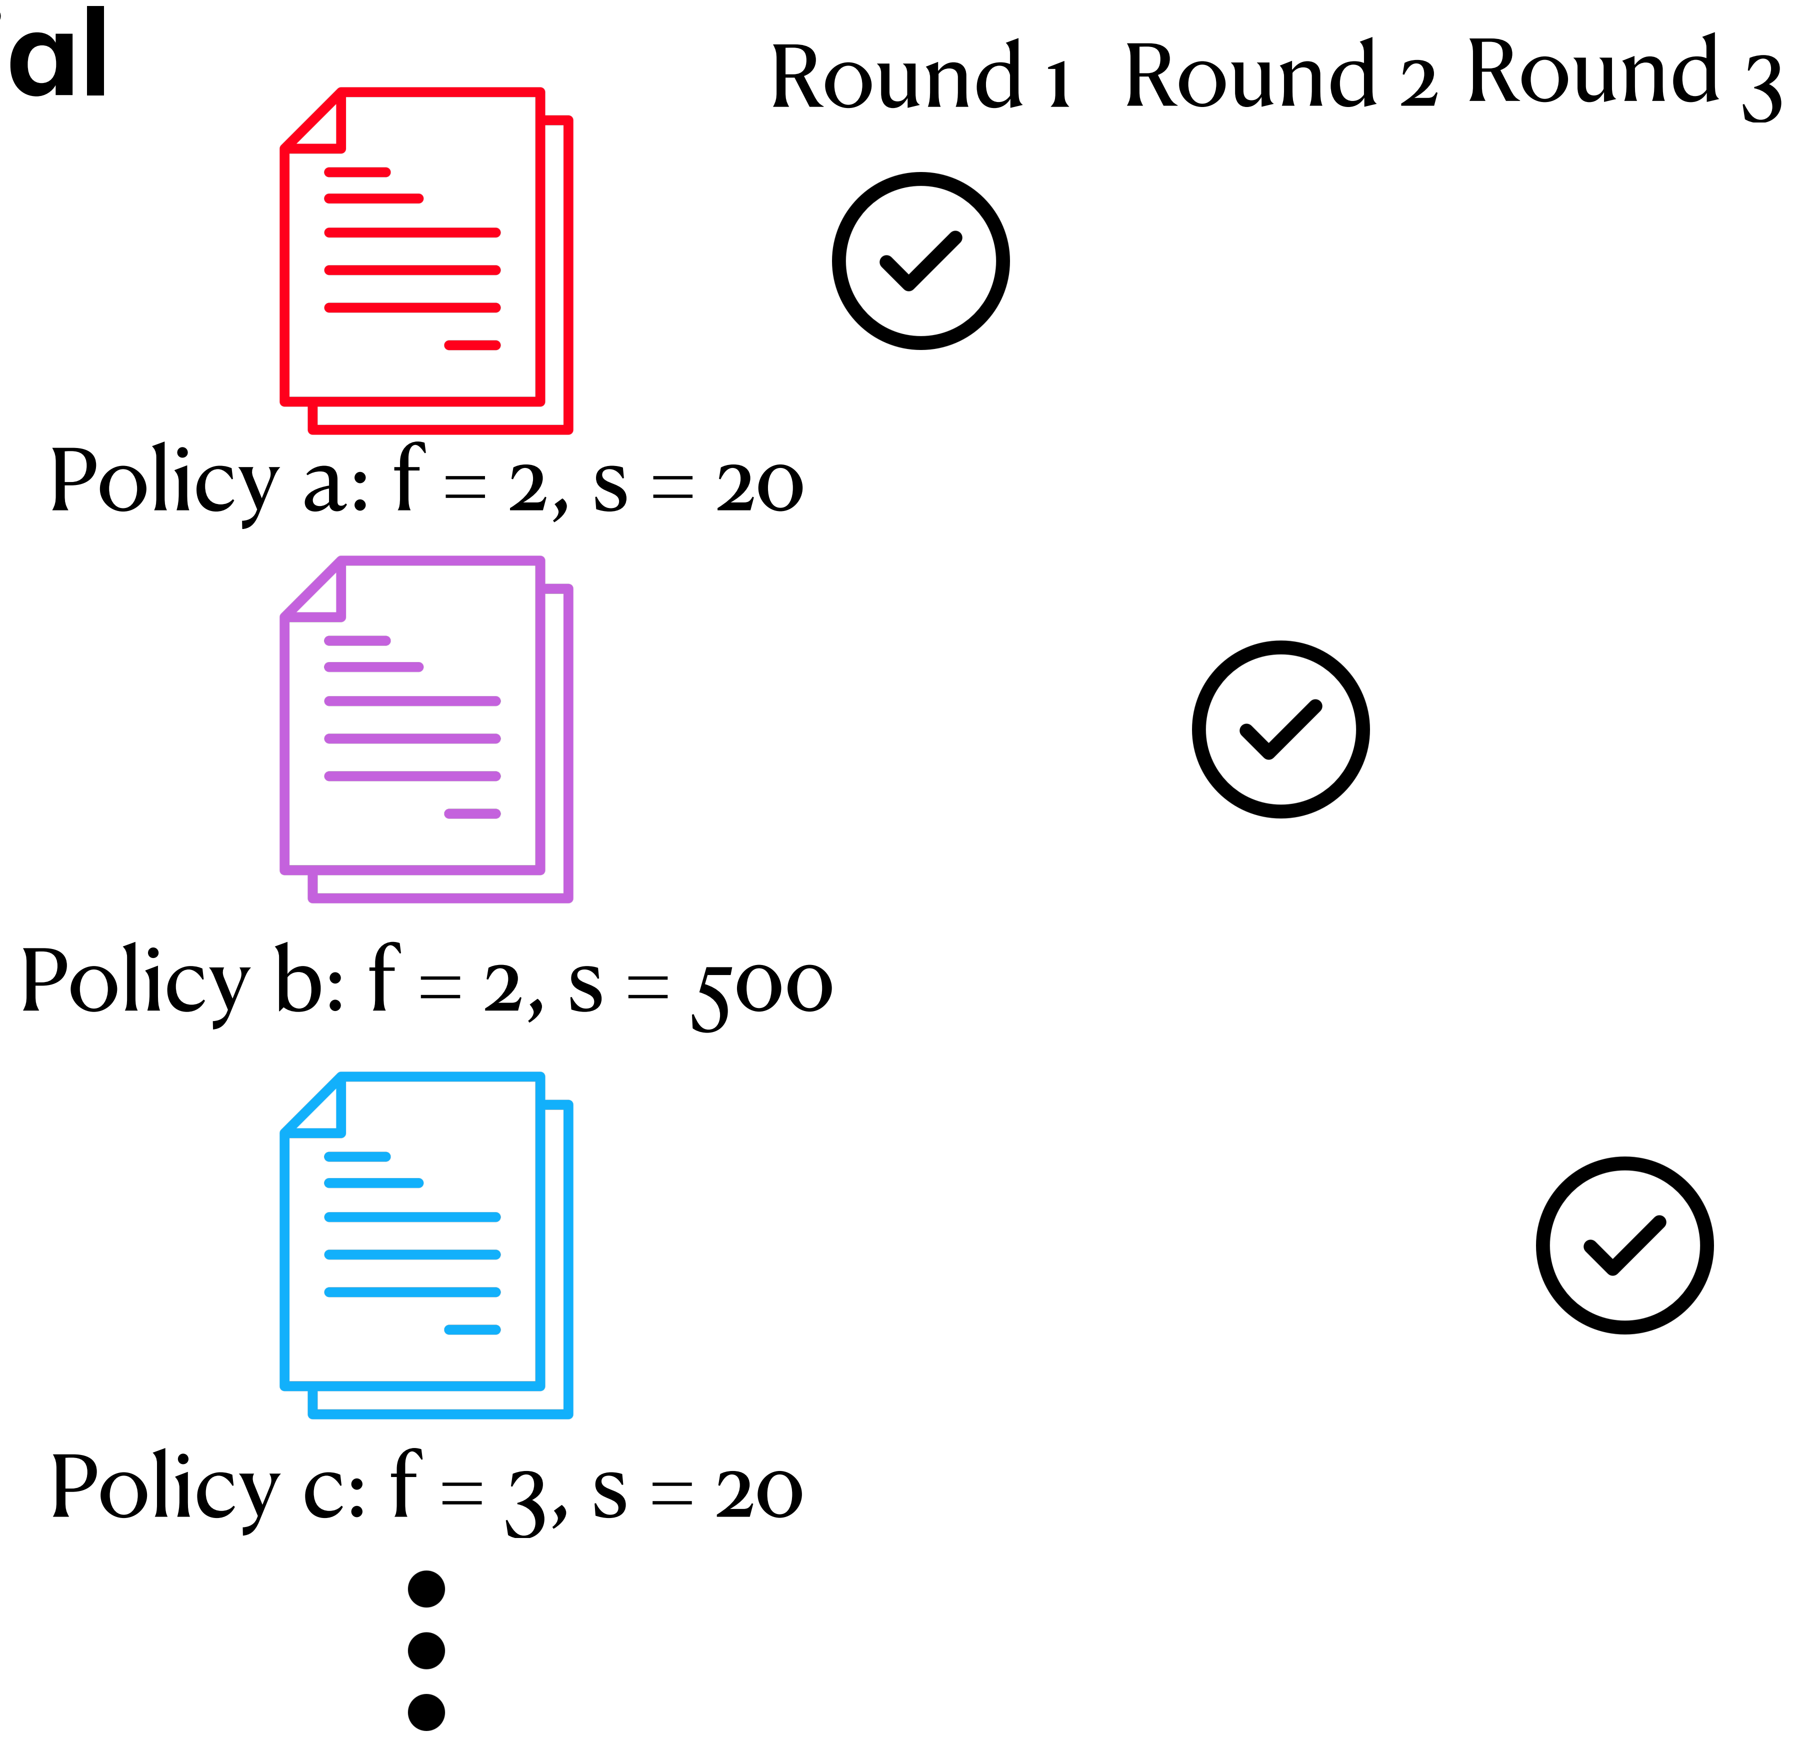

Problem: Observation Rounds

Challenge 1: Scalable Performance Observation

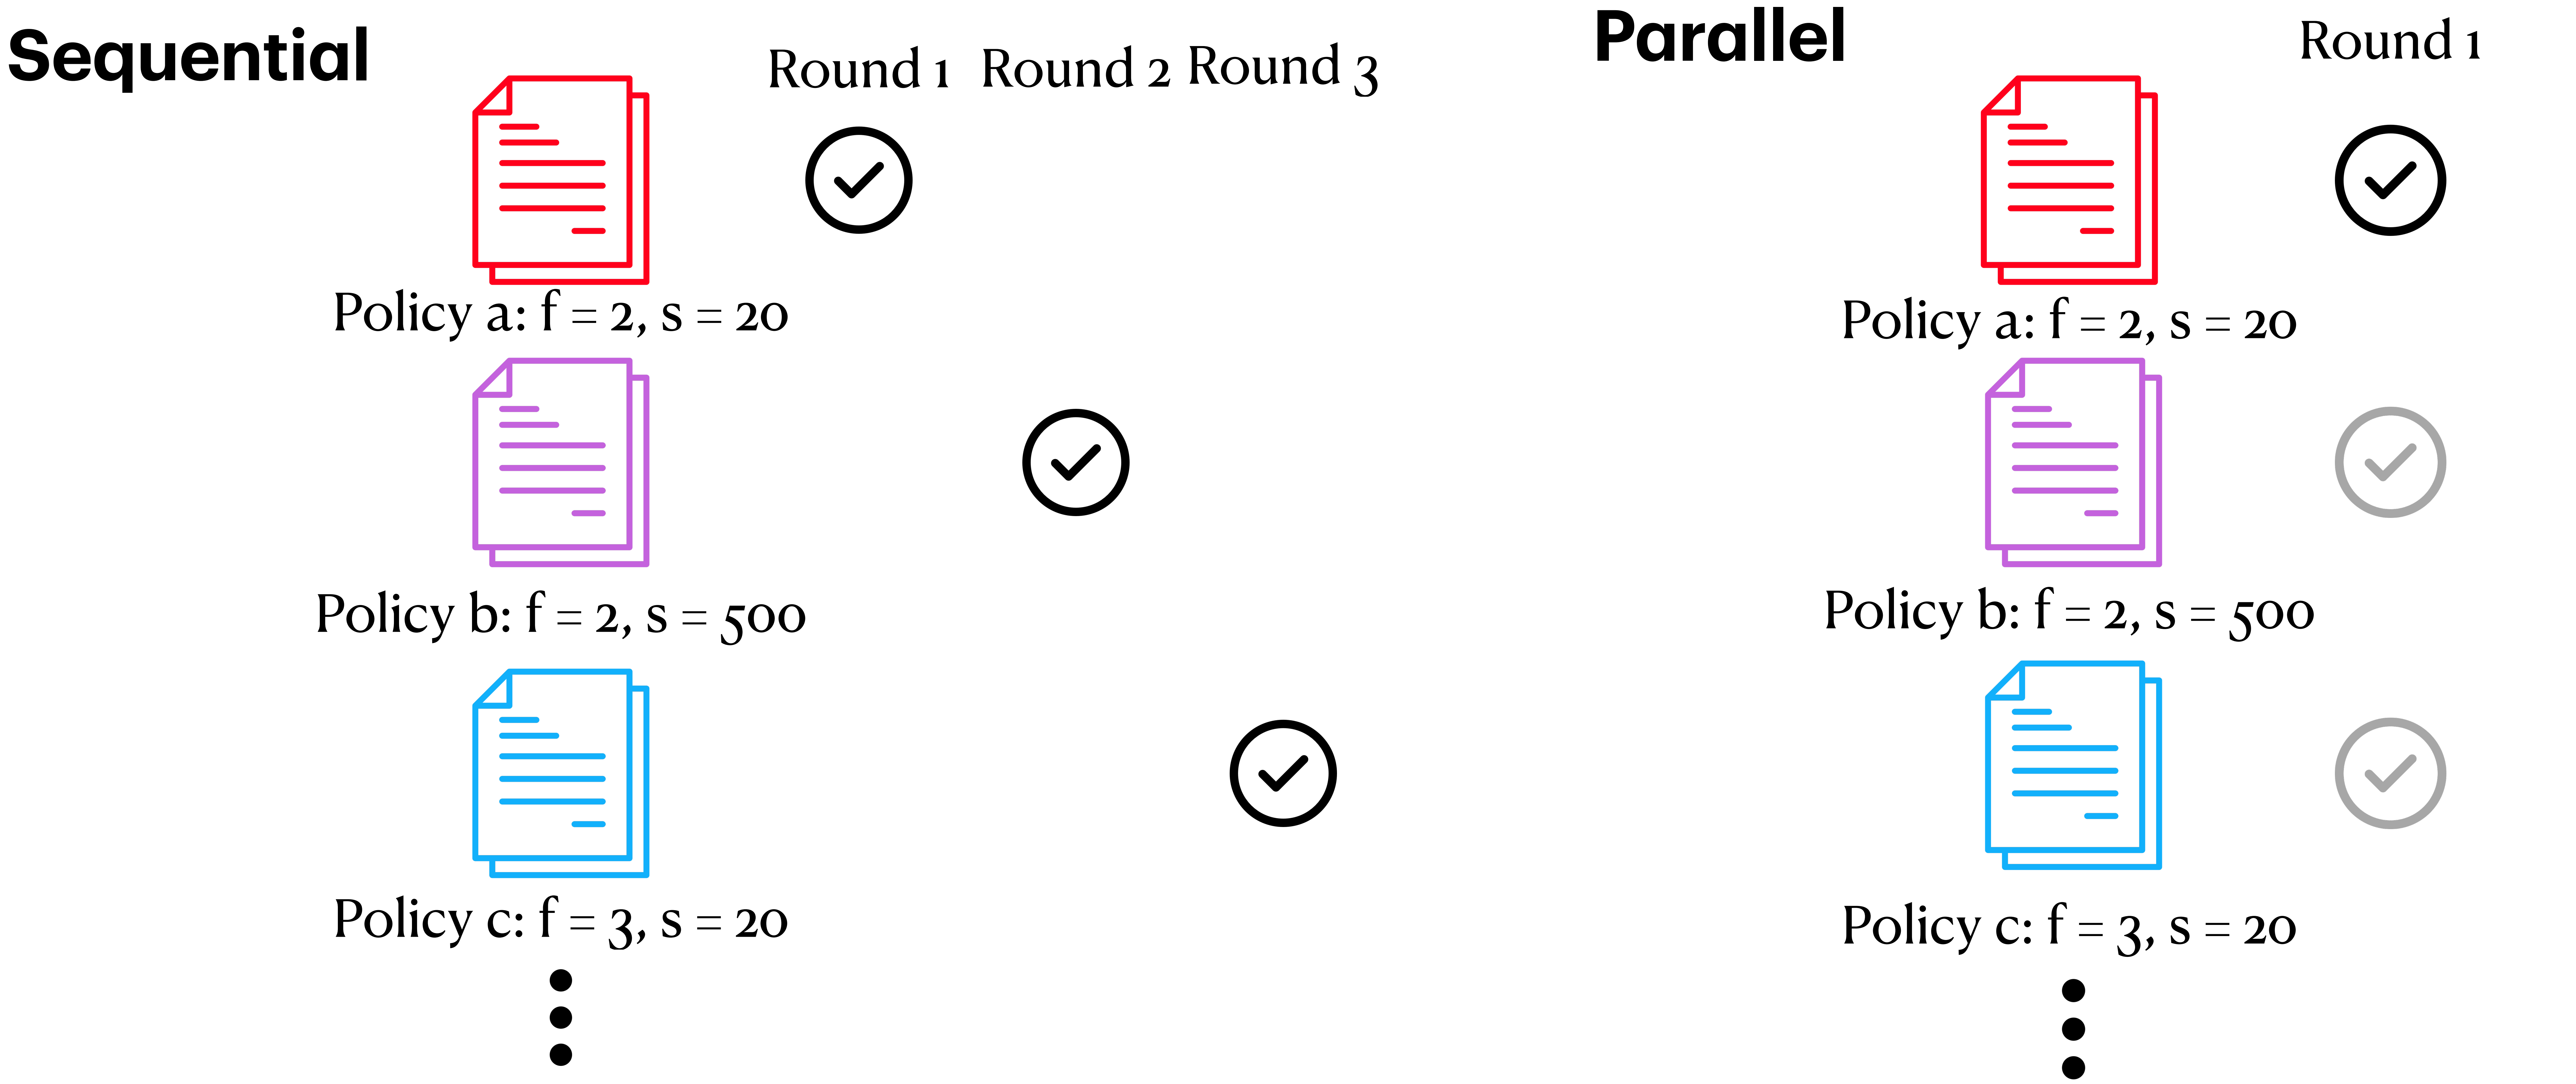

Problem: Observation Rounds

Challenge 1: Scalable Performance Observation

Sequential

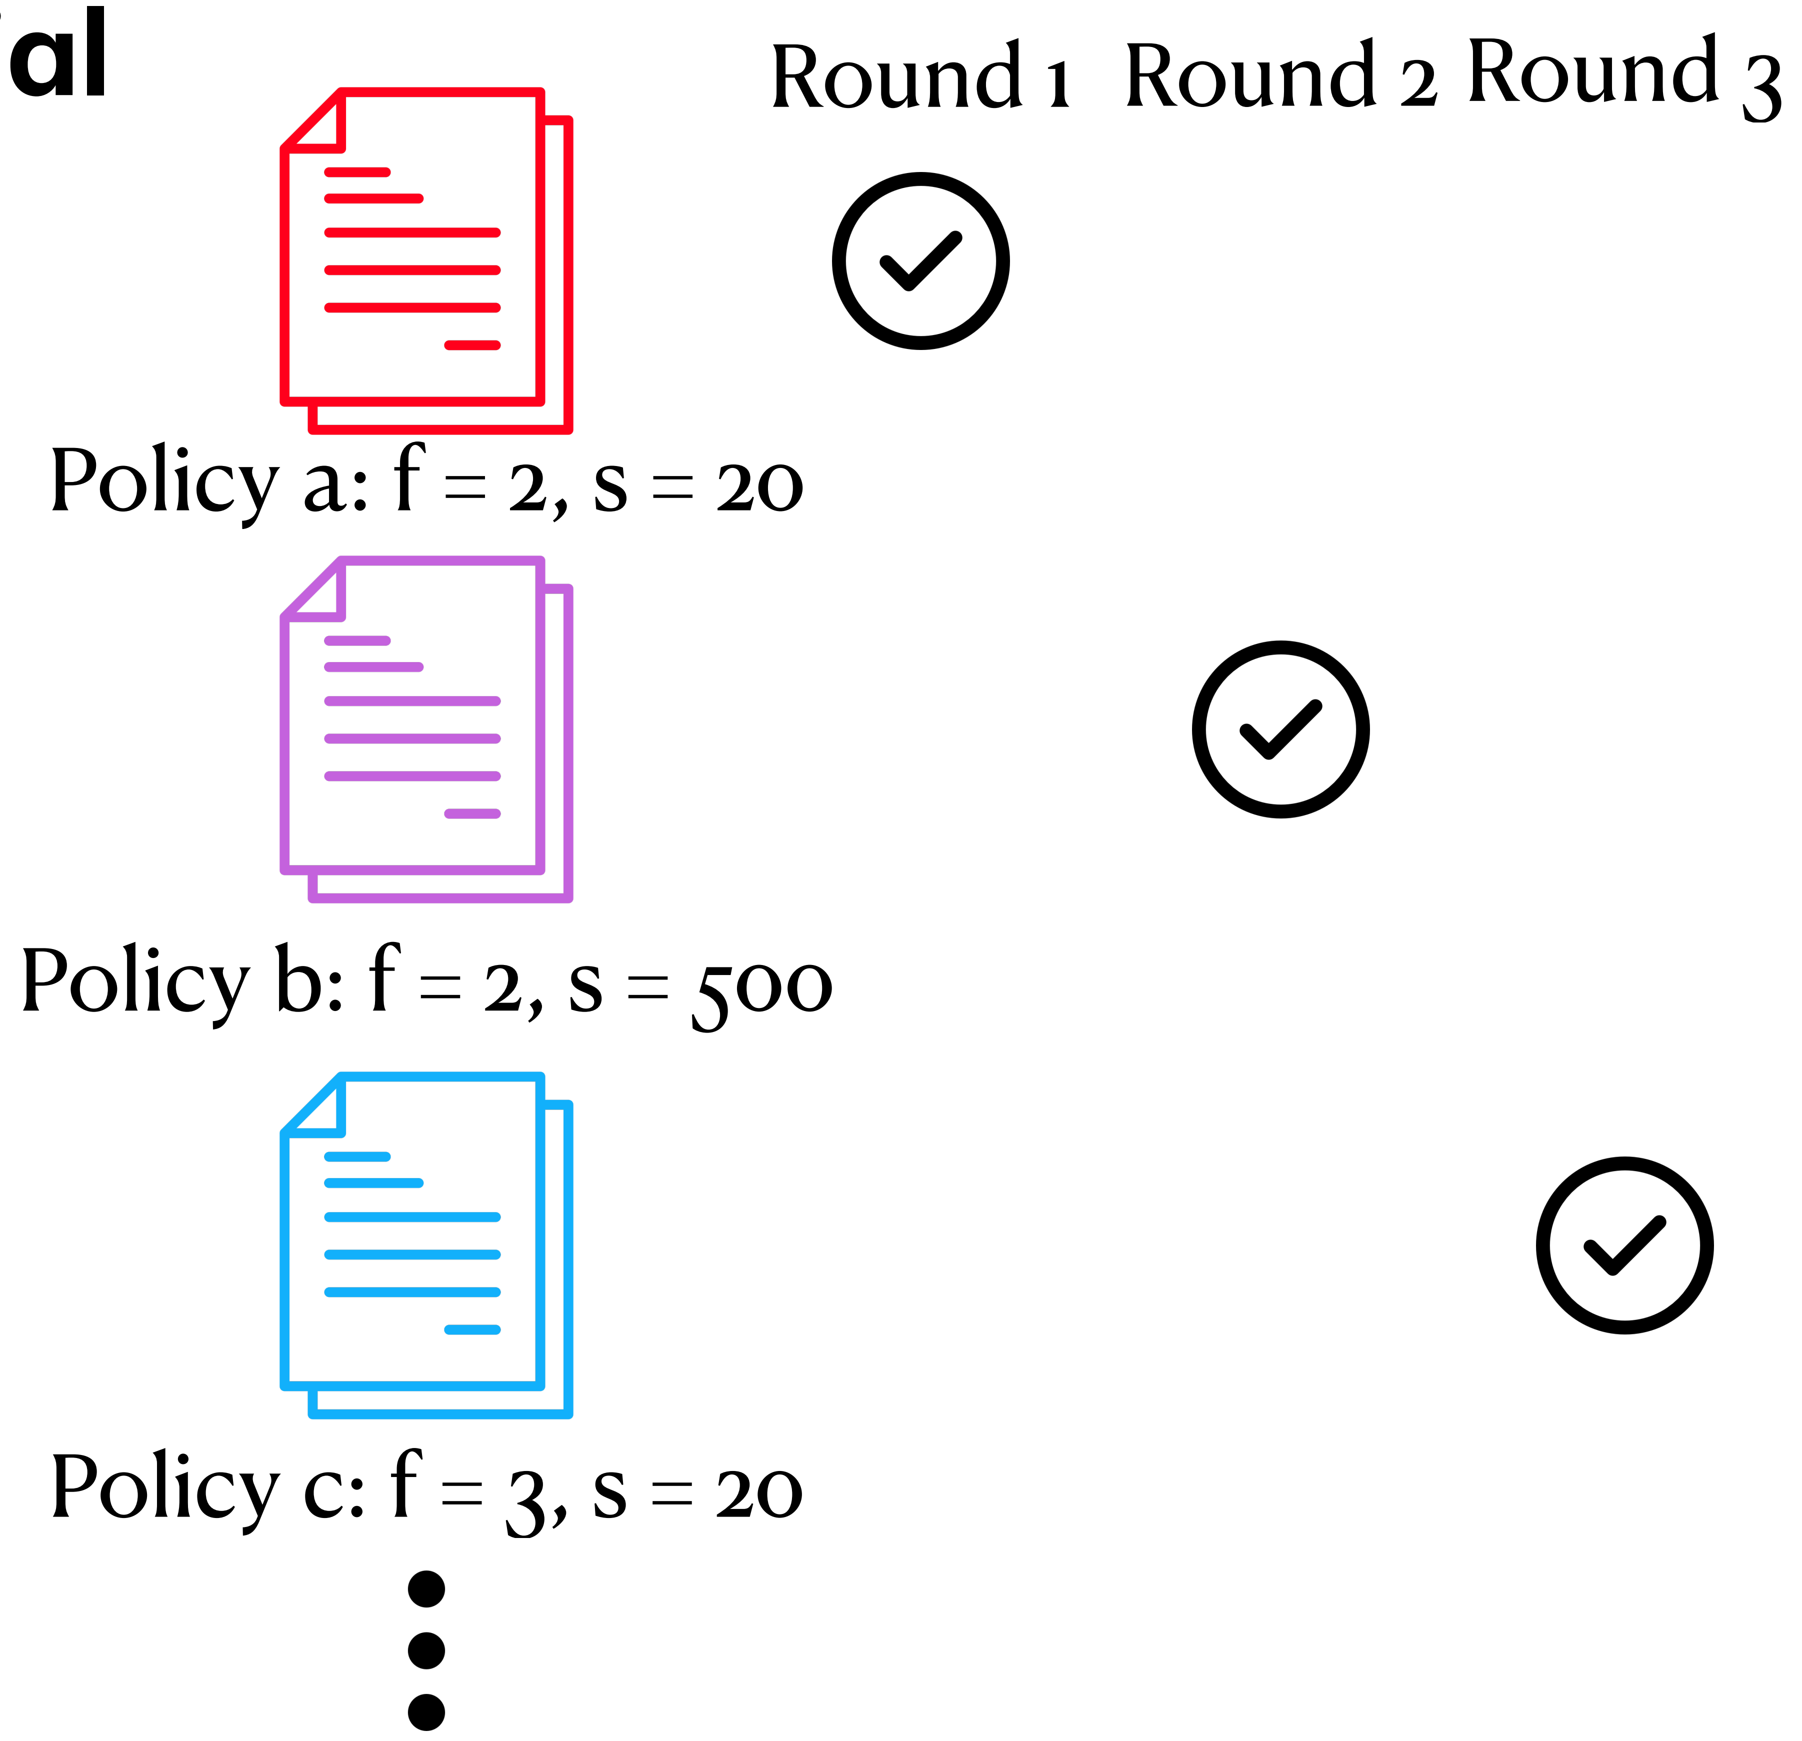

Problem: Observation Rounds

Parallel

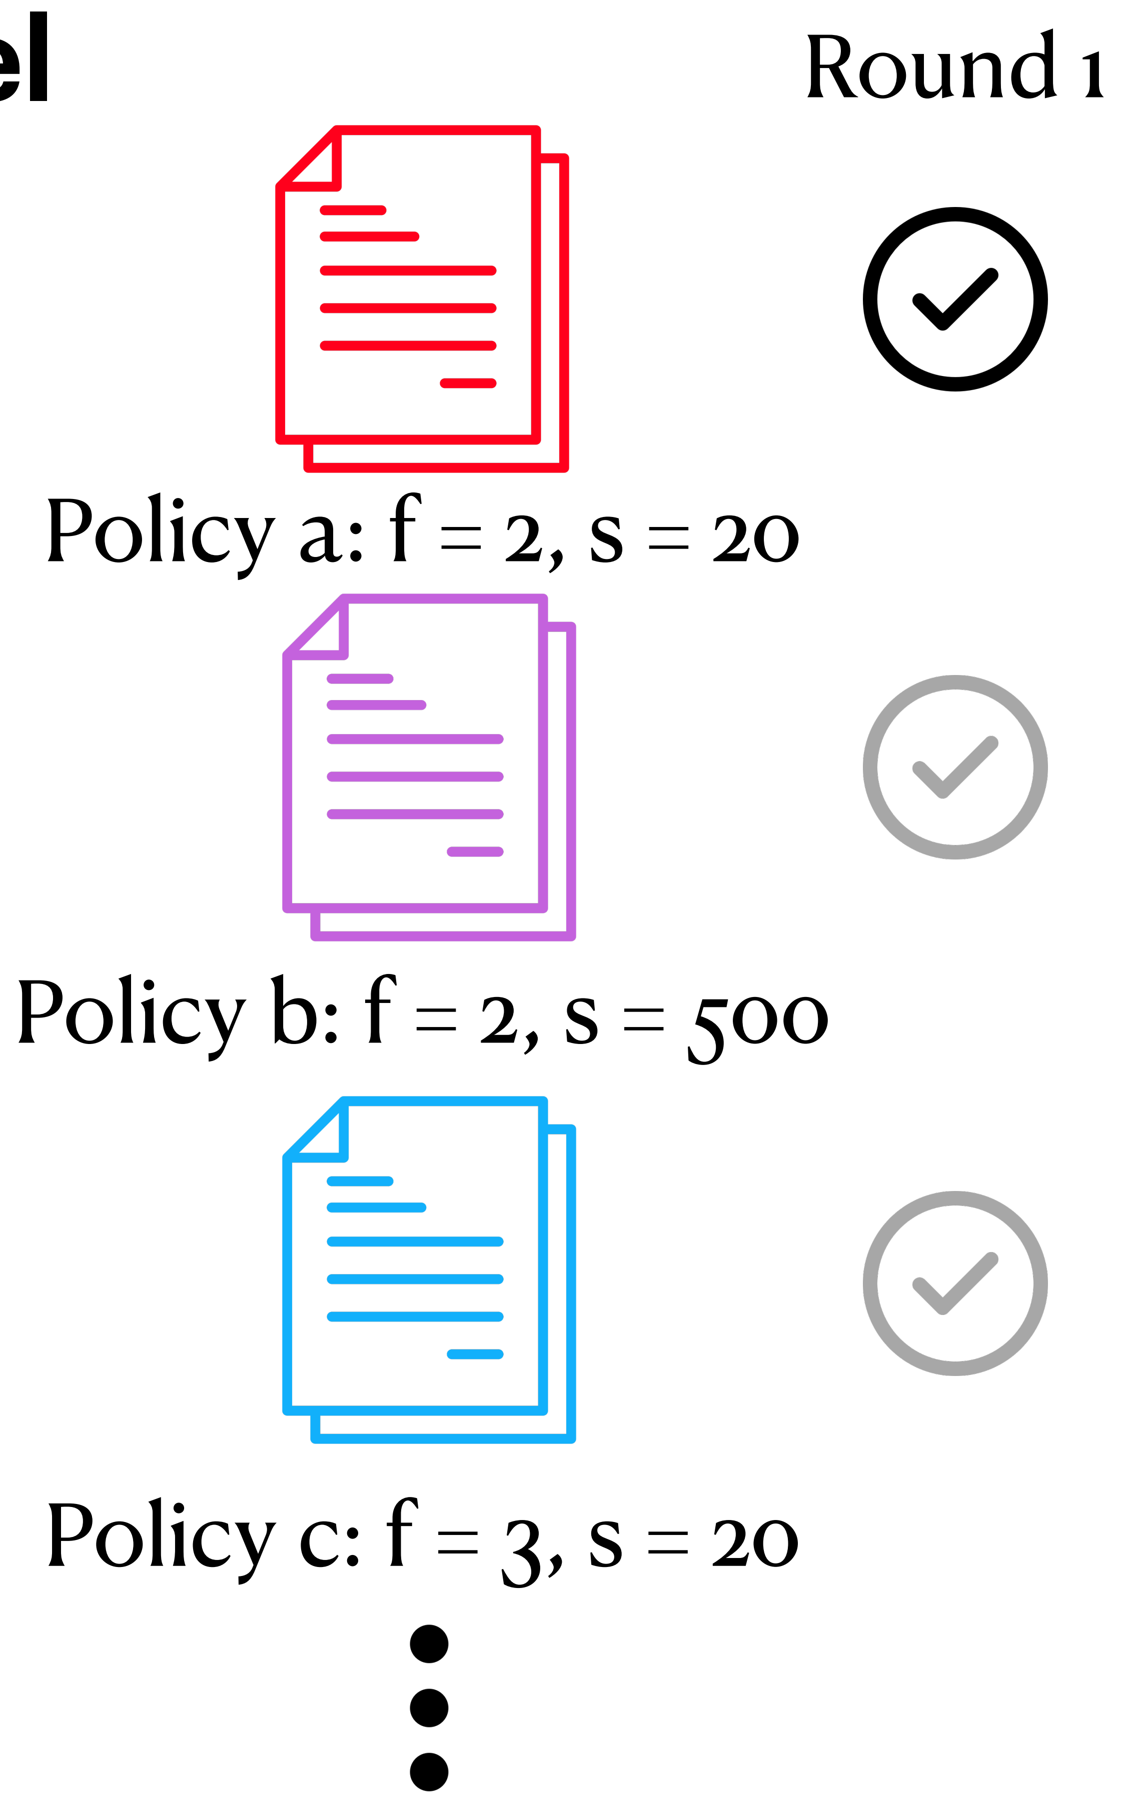

Problem: Resource Overhead

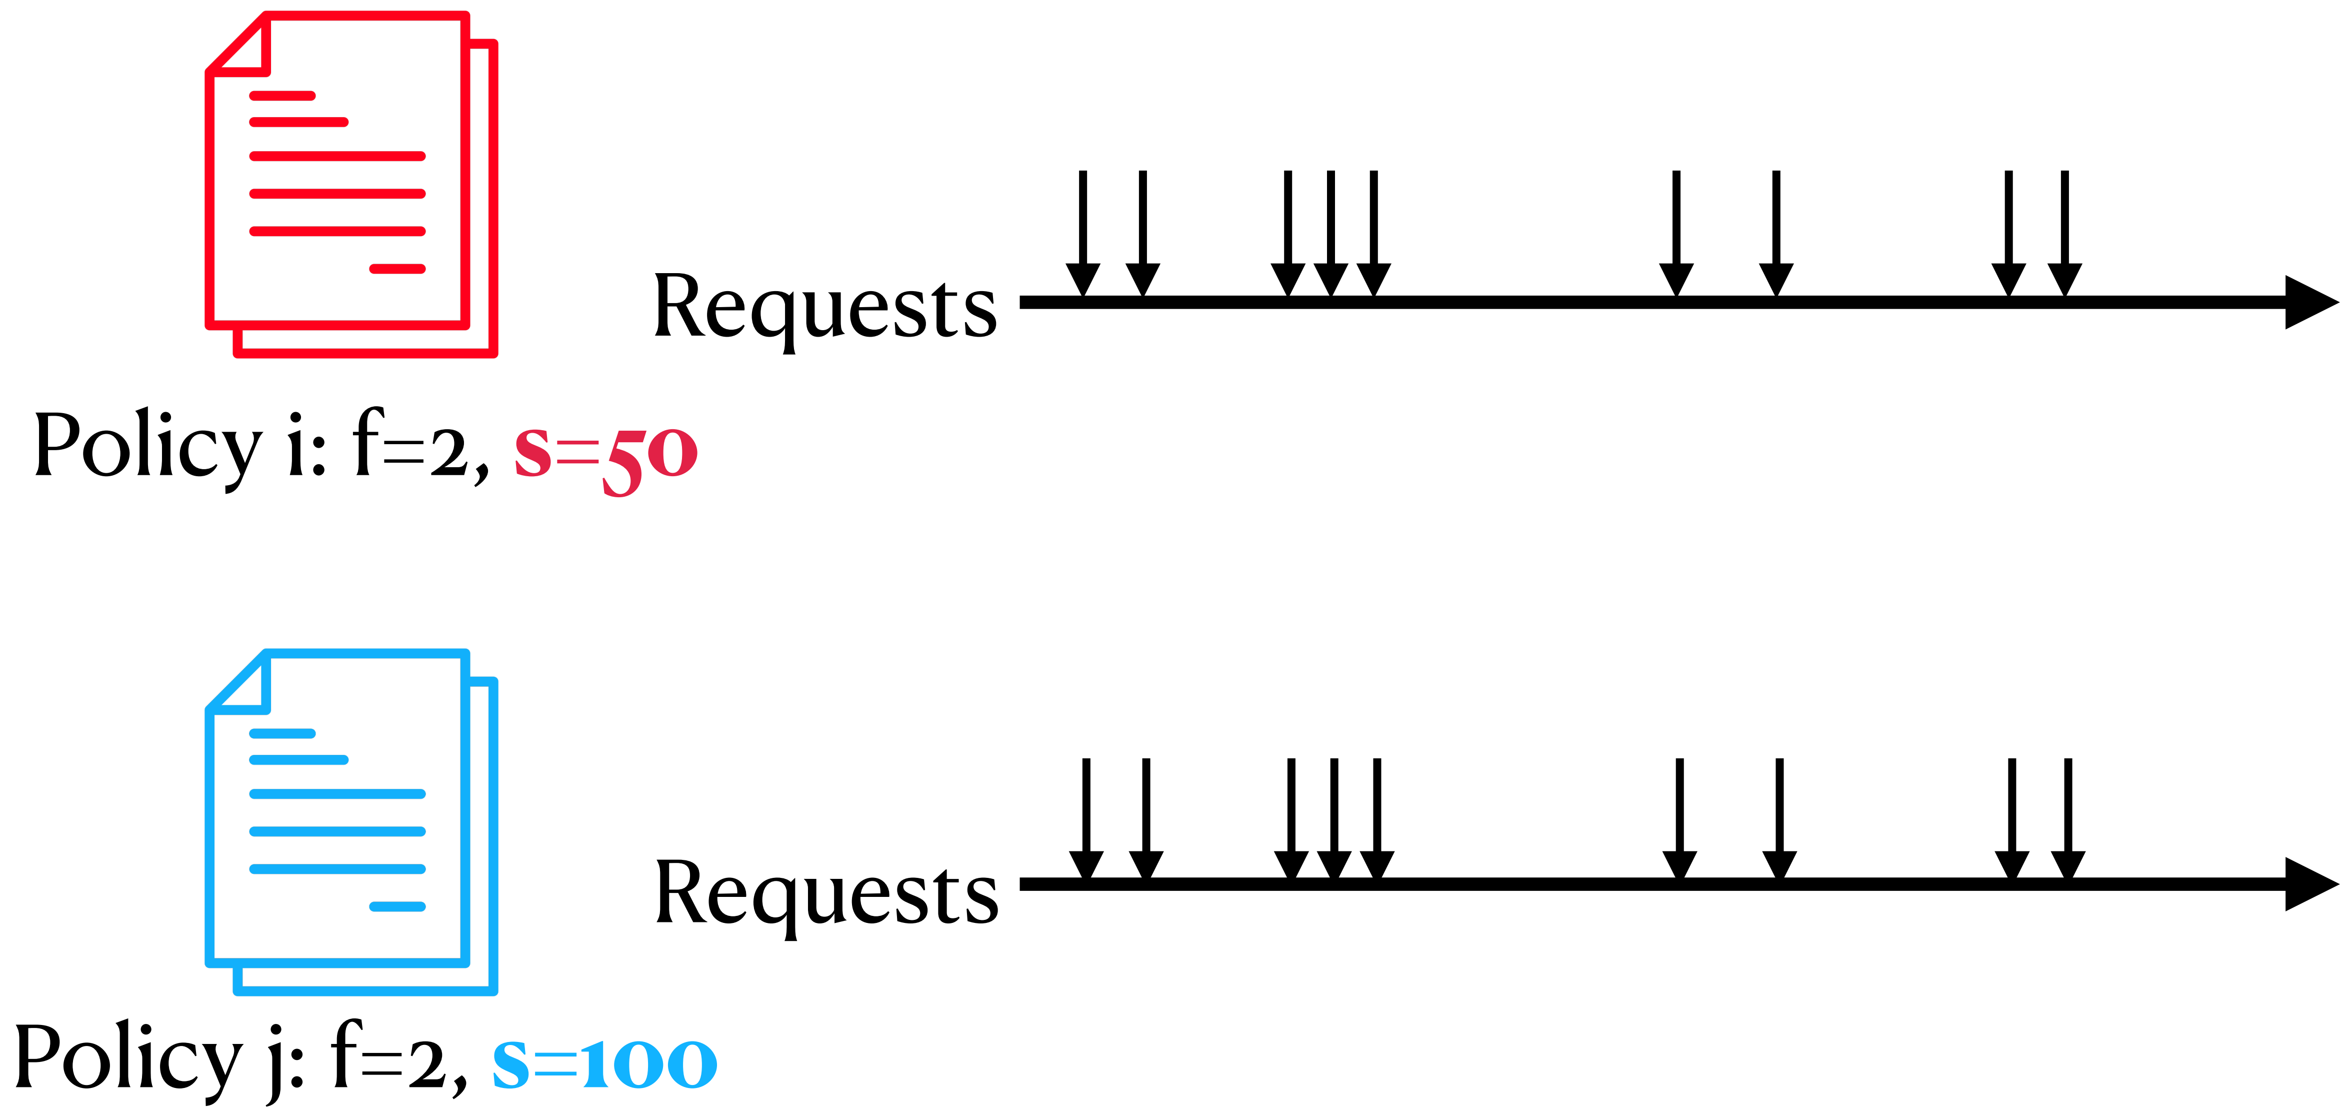
Requests
Policy i: f=2, s=50
Requests
Policy j: f=2, s=100

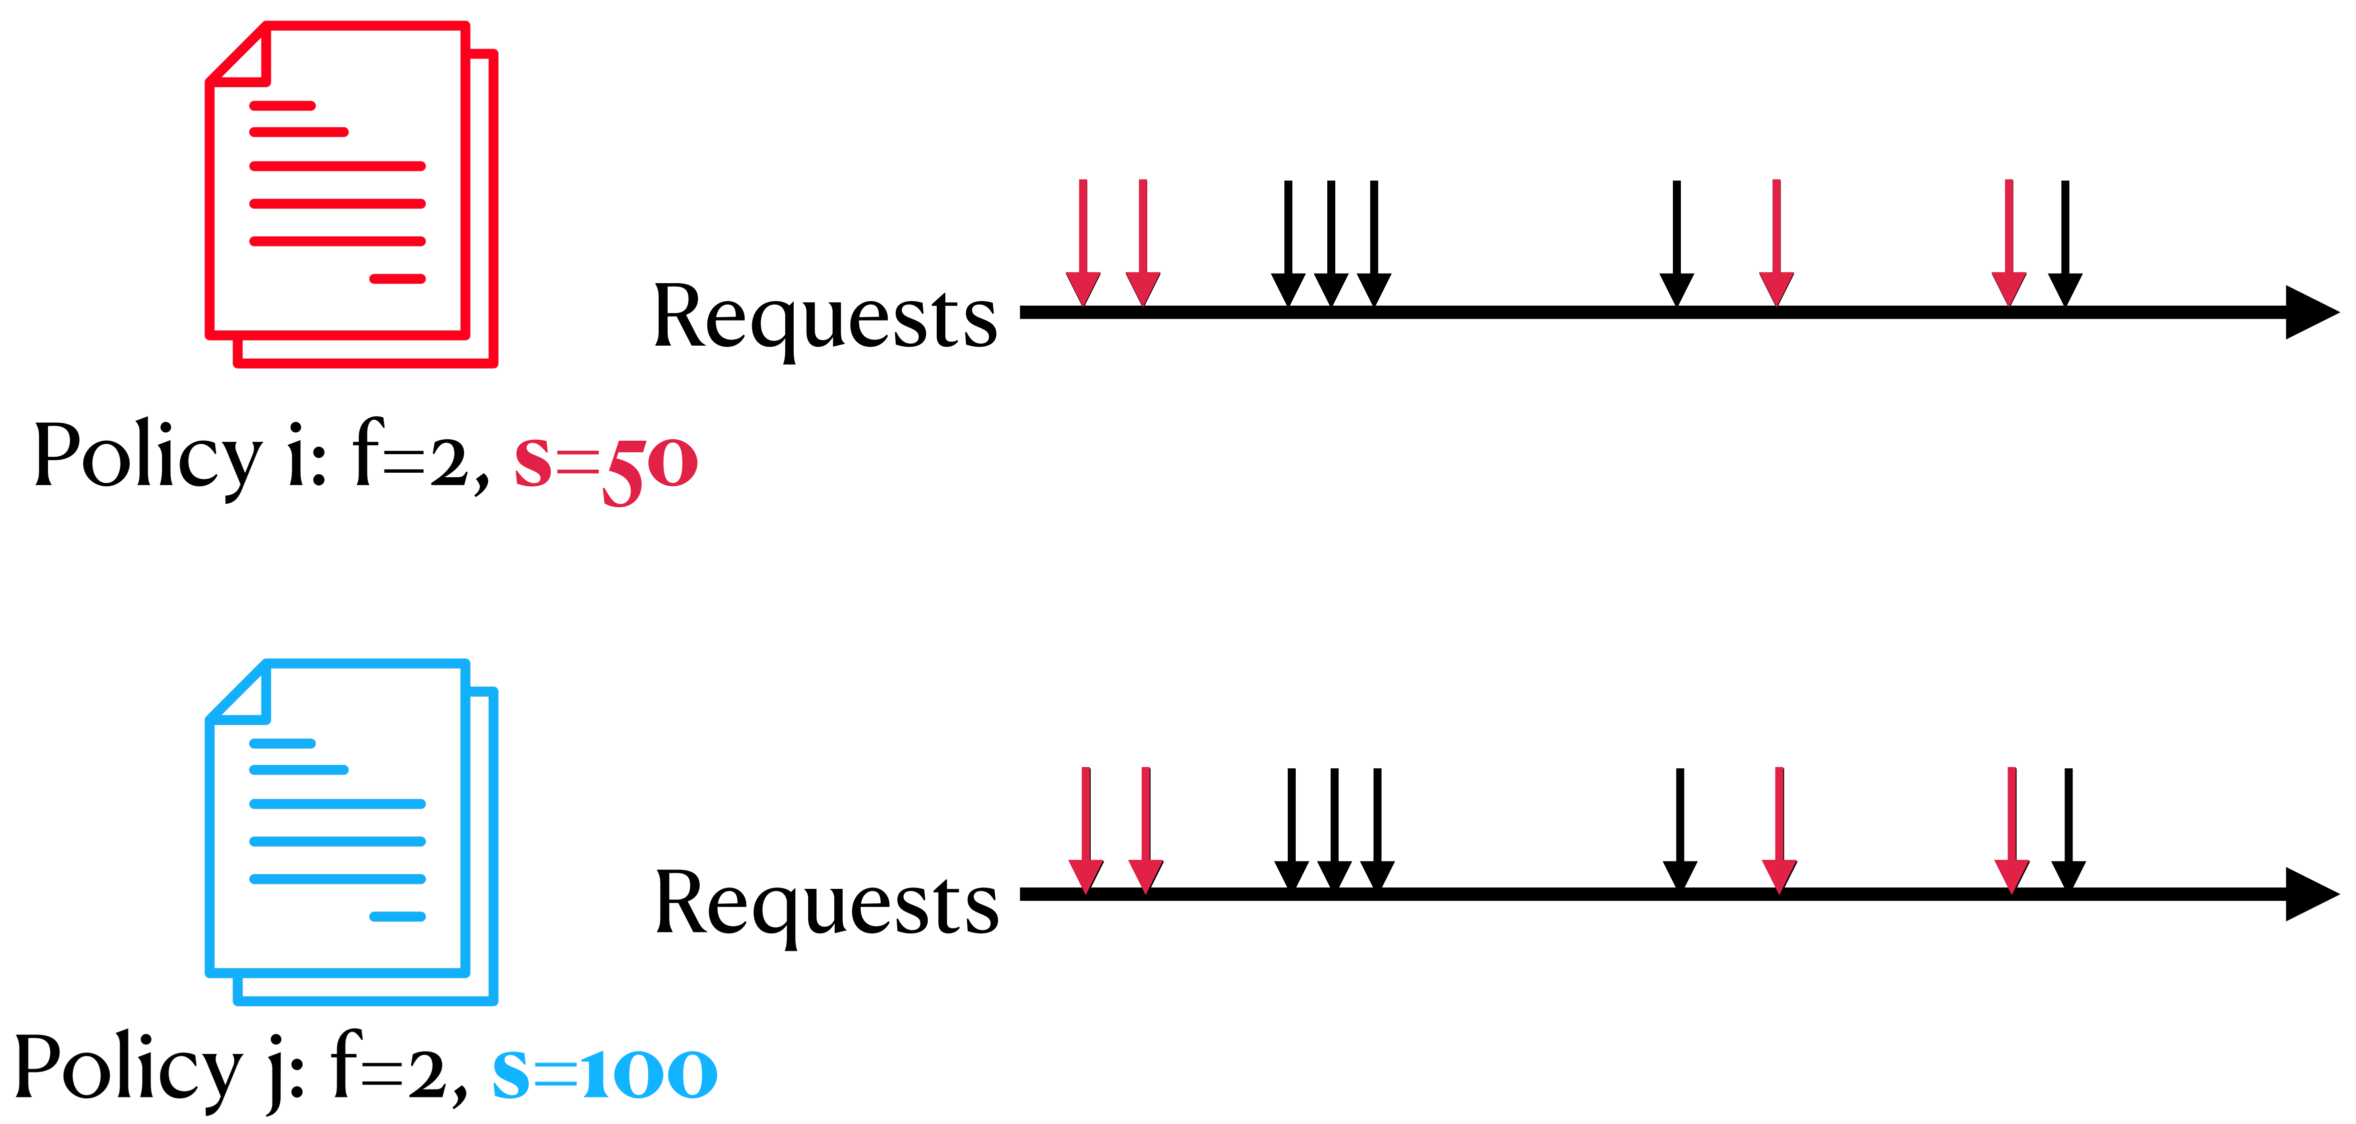
Requests
Policy i: f=2, s=50
Requests
Policy j: f=2, s=100

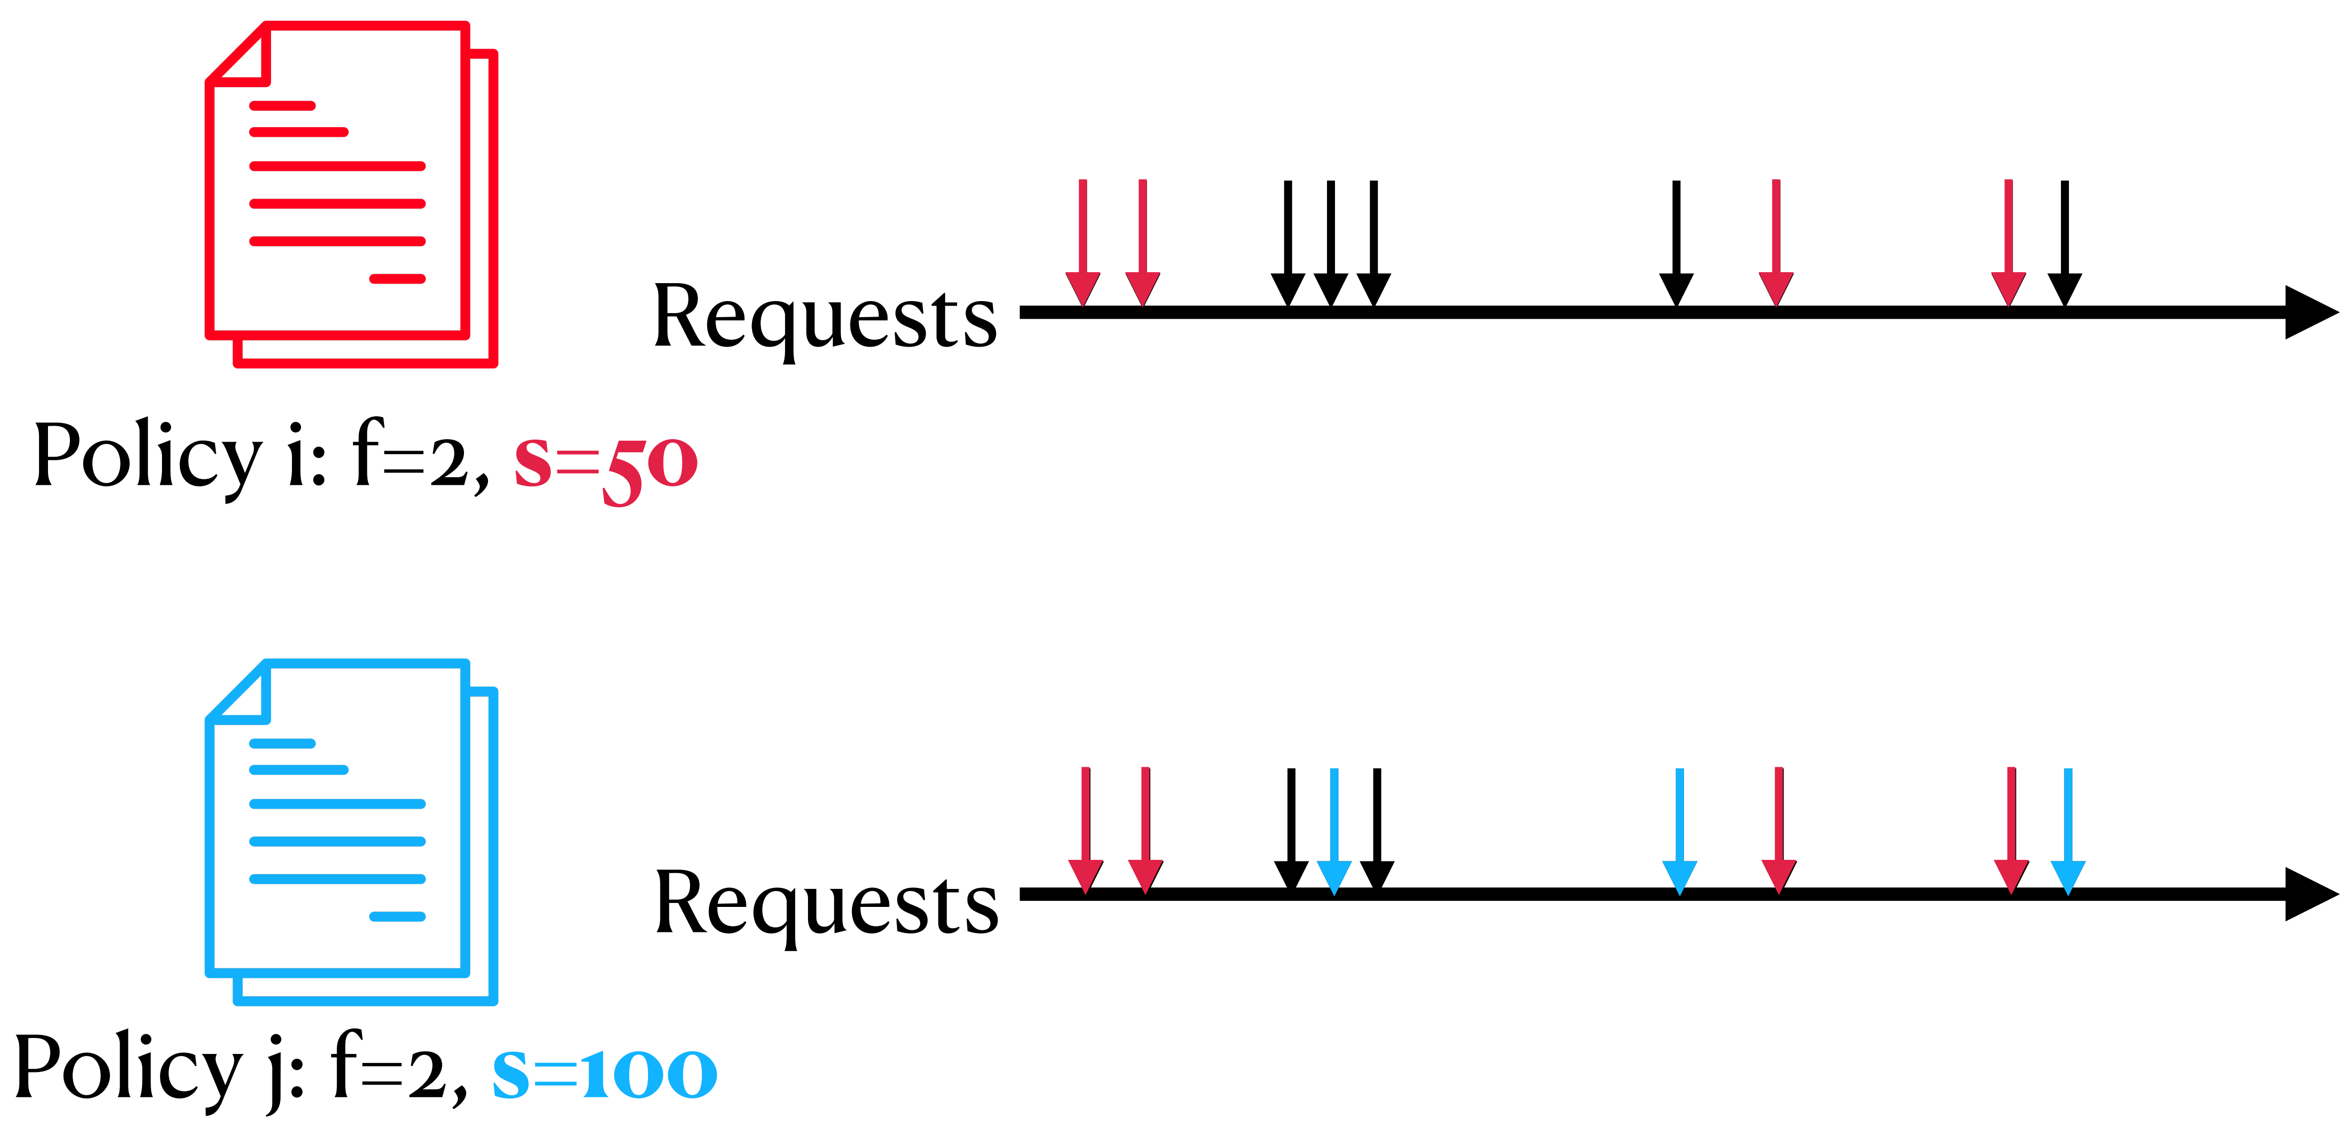
Requests
Policy i: f=2, s=50
Requests
Policy j: f=2, s=100

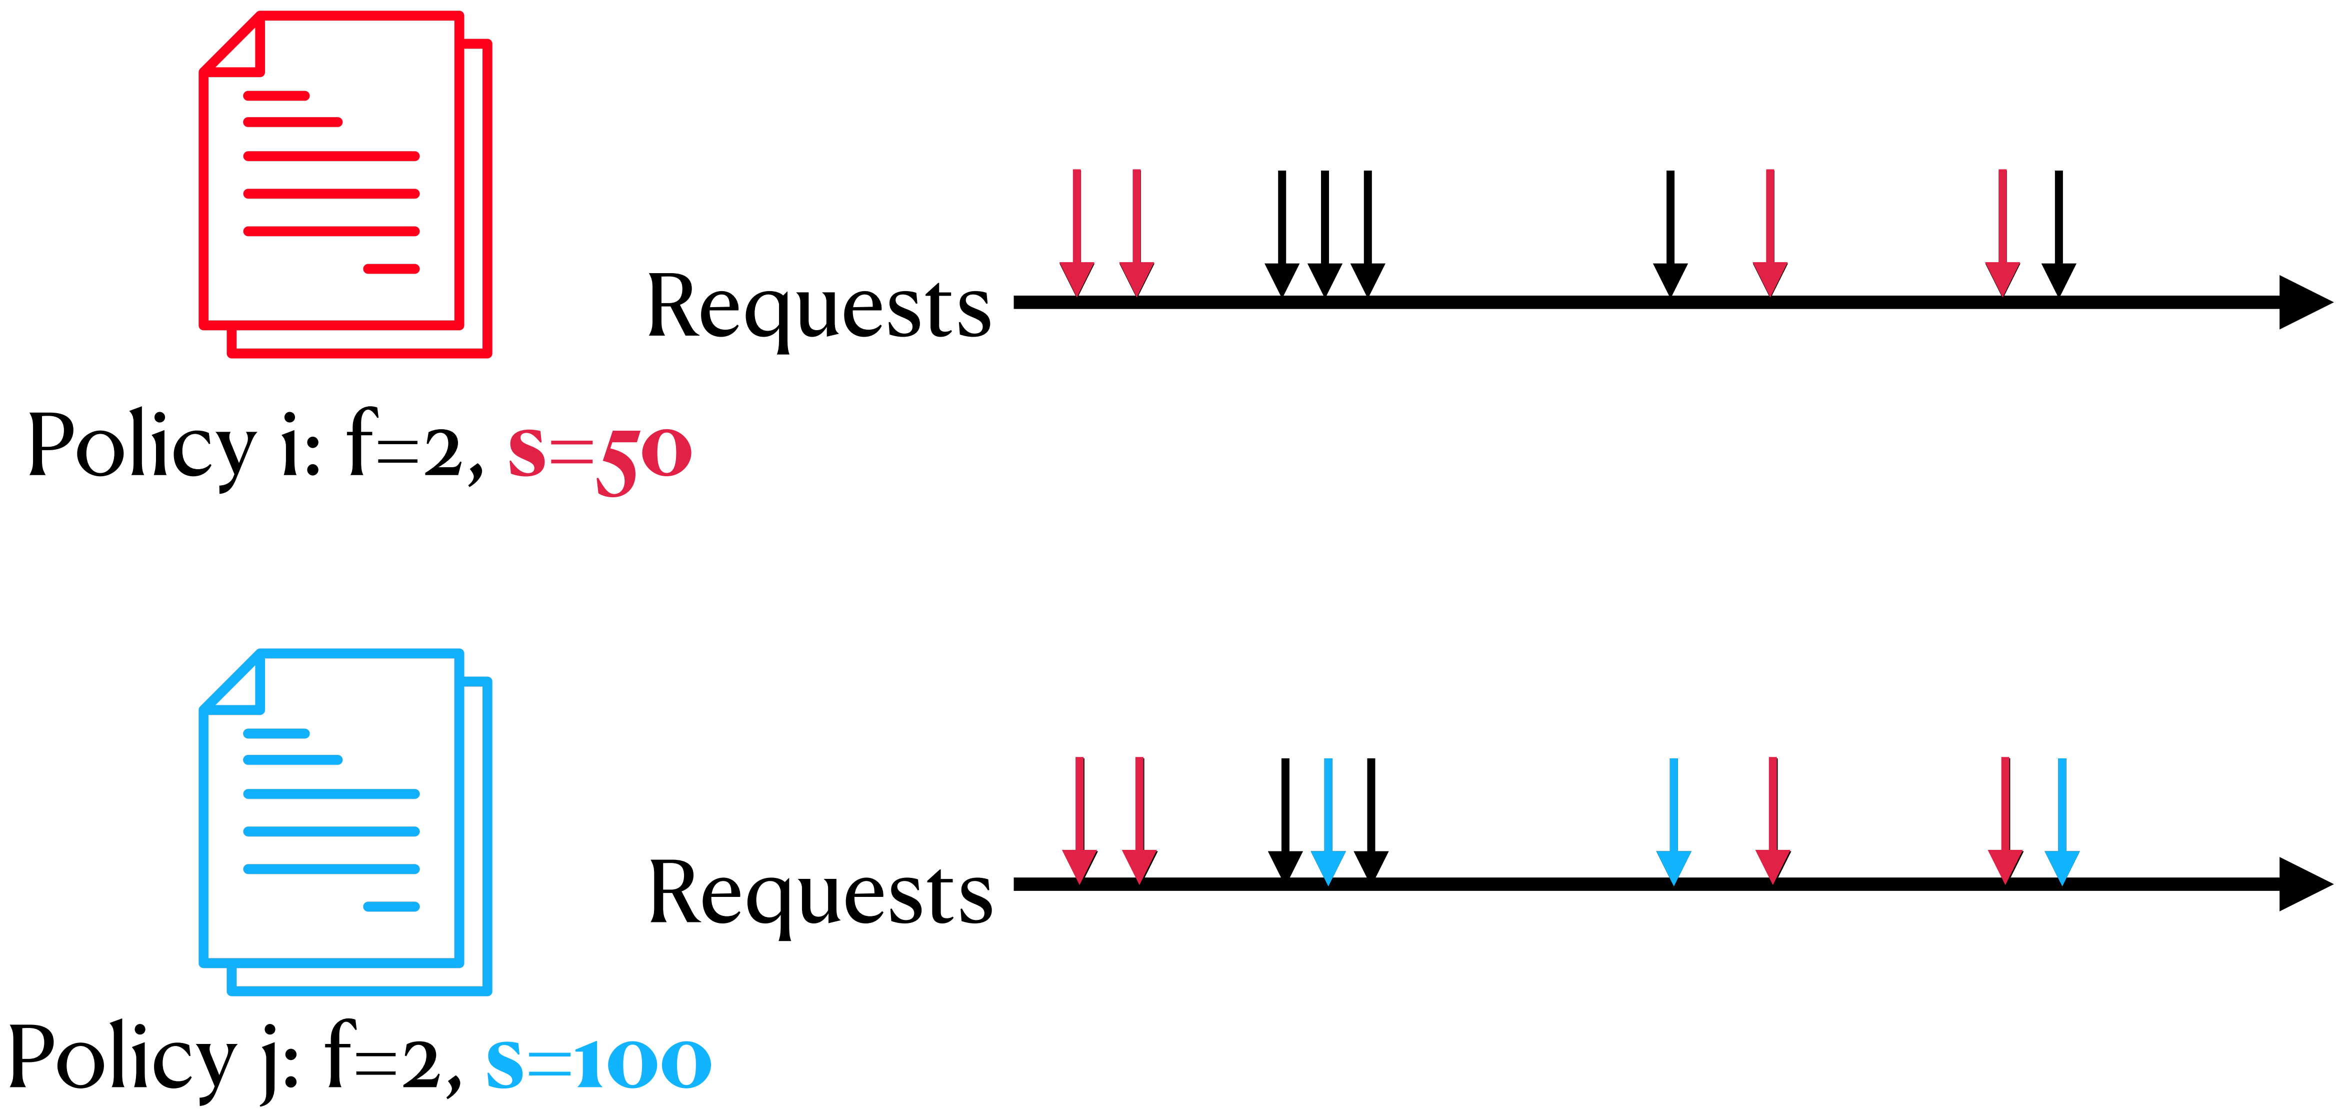

Policy performances are correlated

Cross-policy Prediction Models

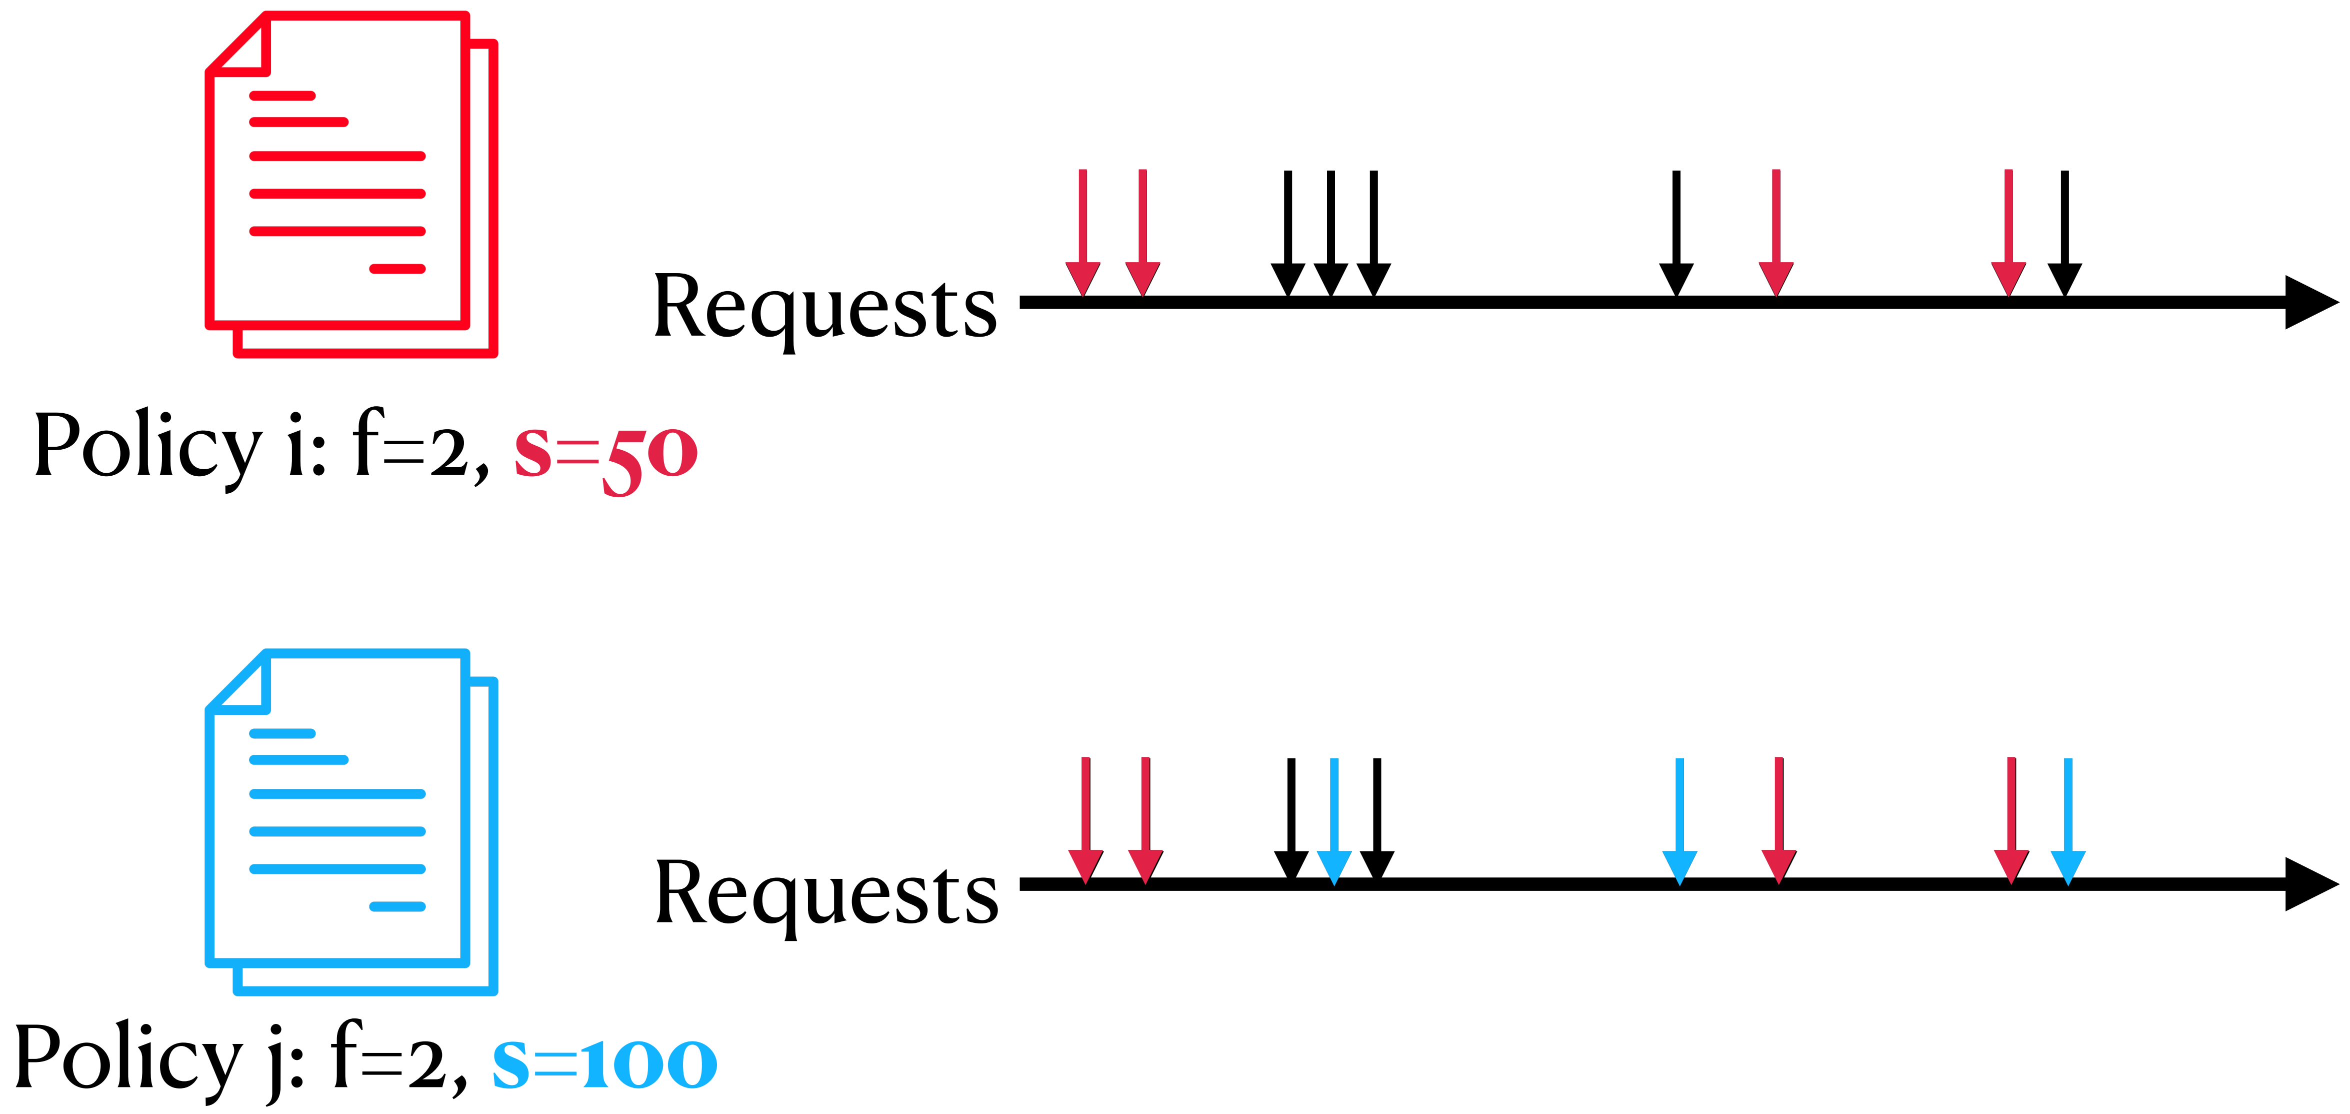

Policy performances are correlated

Cross-policy Prediction Models

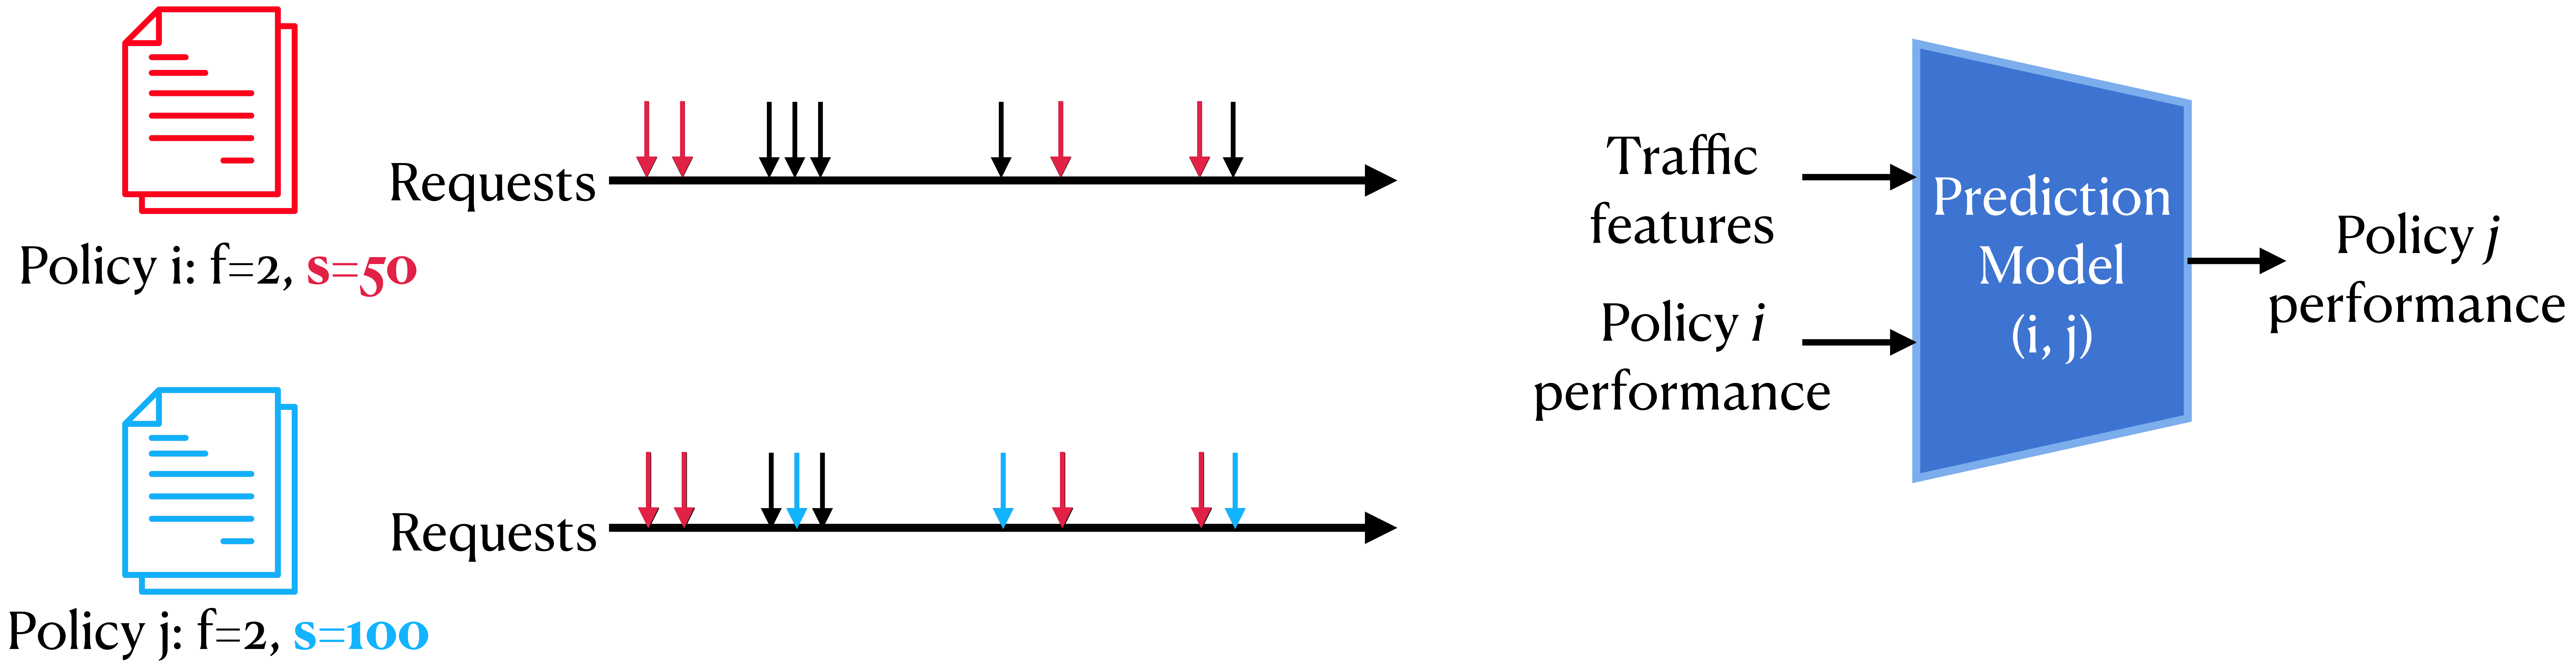

Policy performances are correlated

Cross-policy Prediction Models

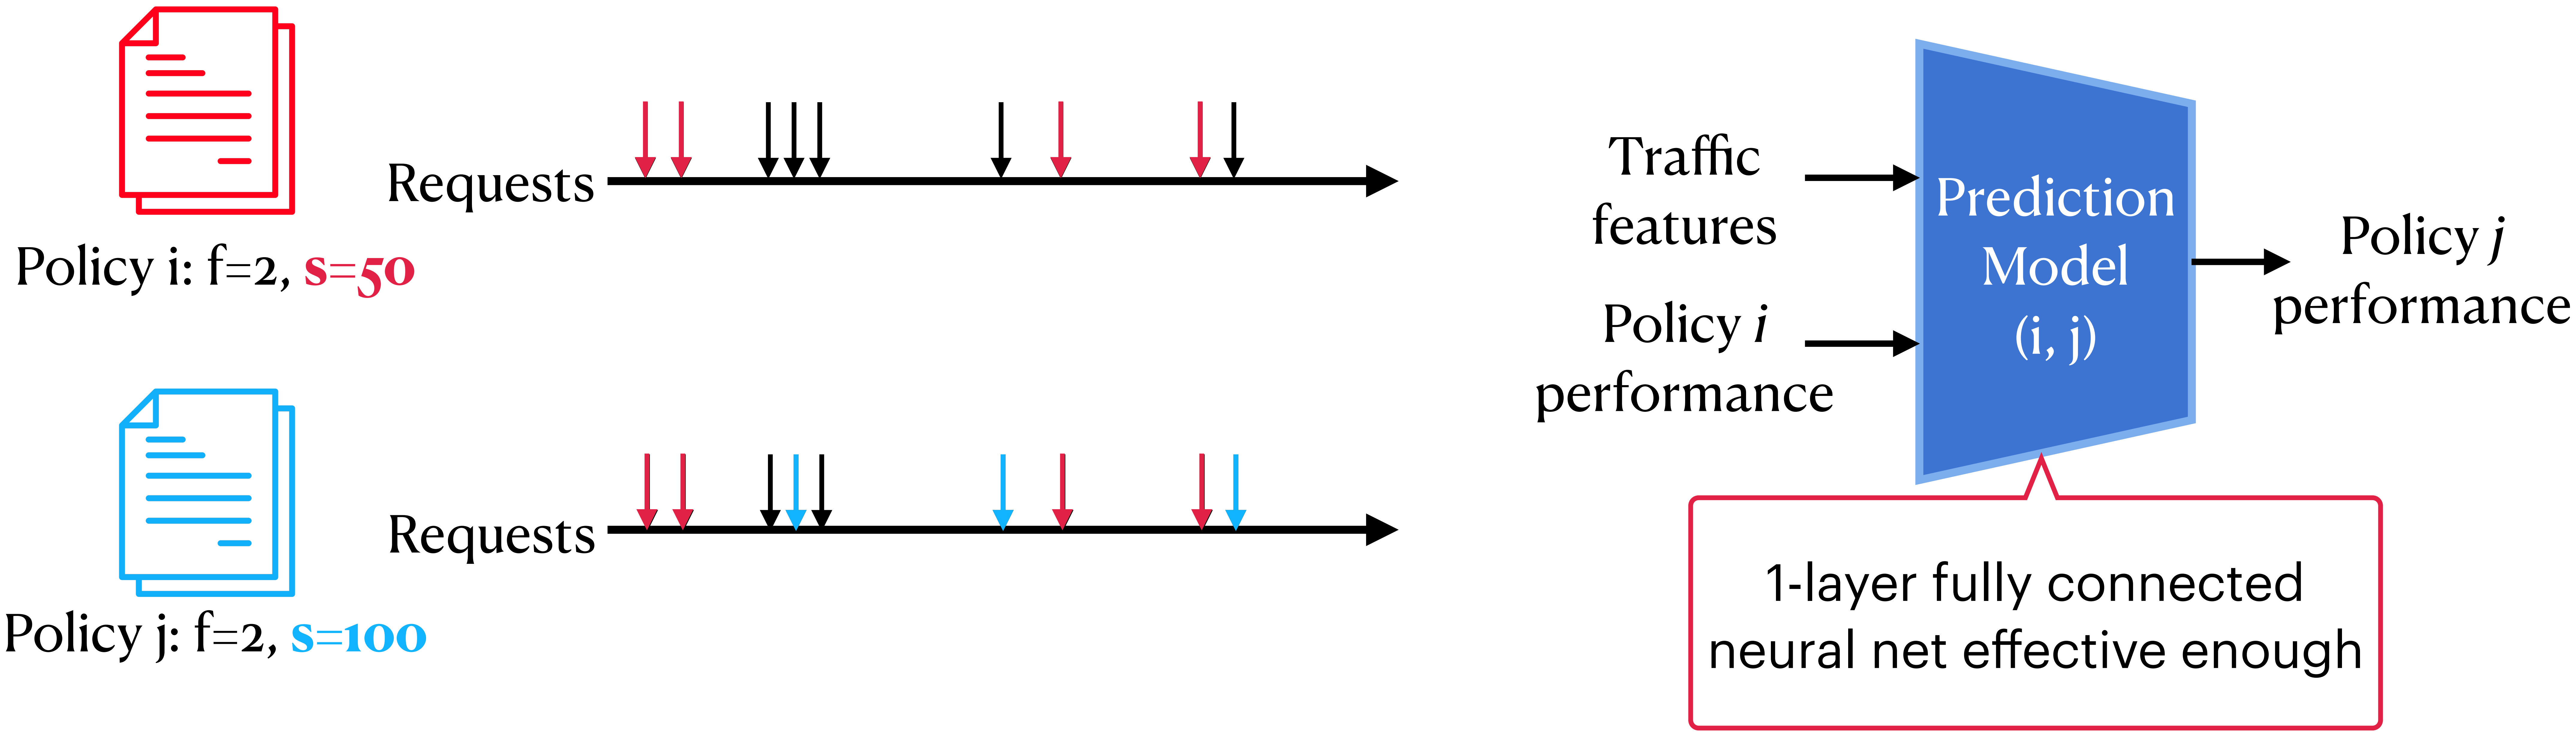

Policy performances are correlated

How to reduce the policy space?

Feature Clustering and Policy Association

How to reduce the policy space?

Feature Clustering and Policy Association

<u>Features</u>:
Average object size
Inter-arrival time
Stack distance
......

How to reduce the policy space?

Feature Clustering and Policy Association

<u>Features</u>:
Average object size
Inter-arrival time
Stack distance
……

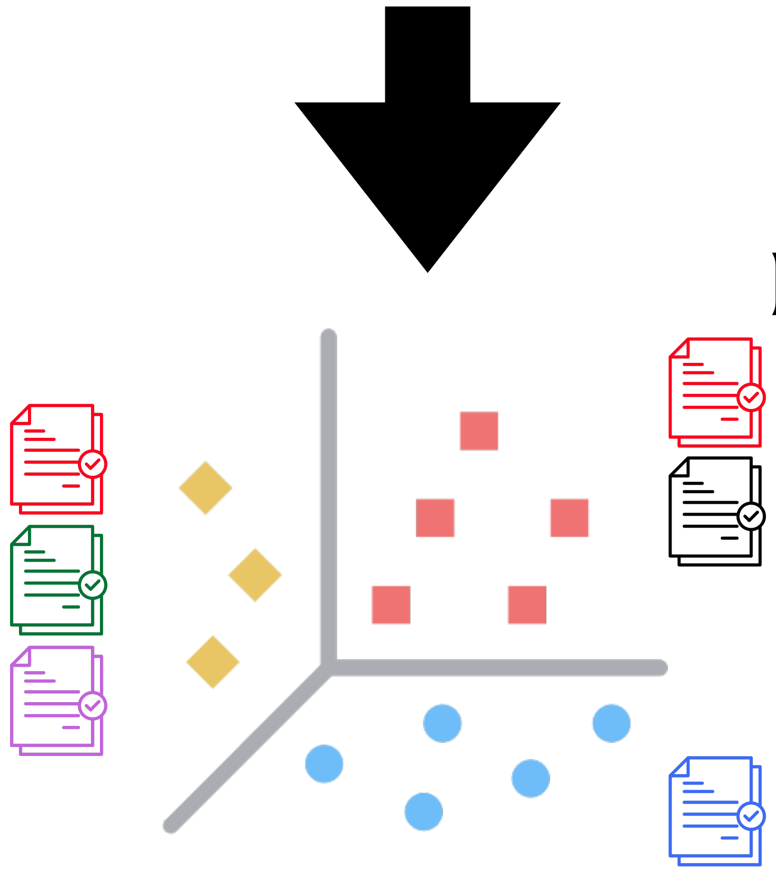

Clusters of features

How to reduce the policy space?

Feature Clustering and Policy Association

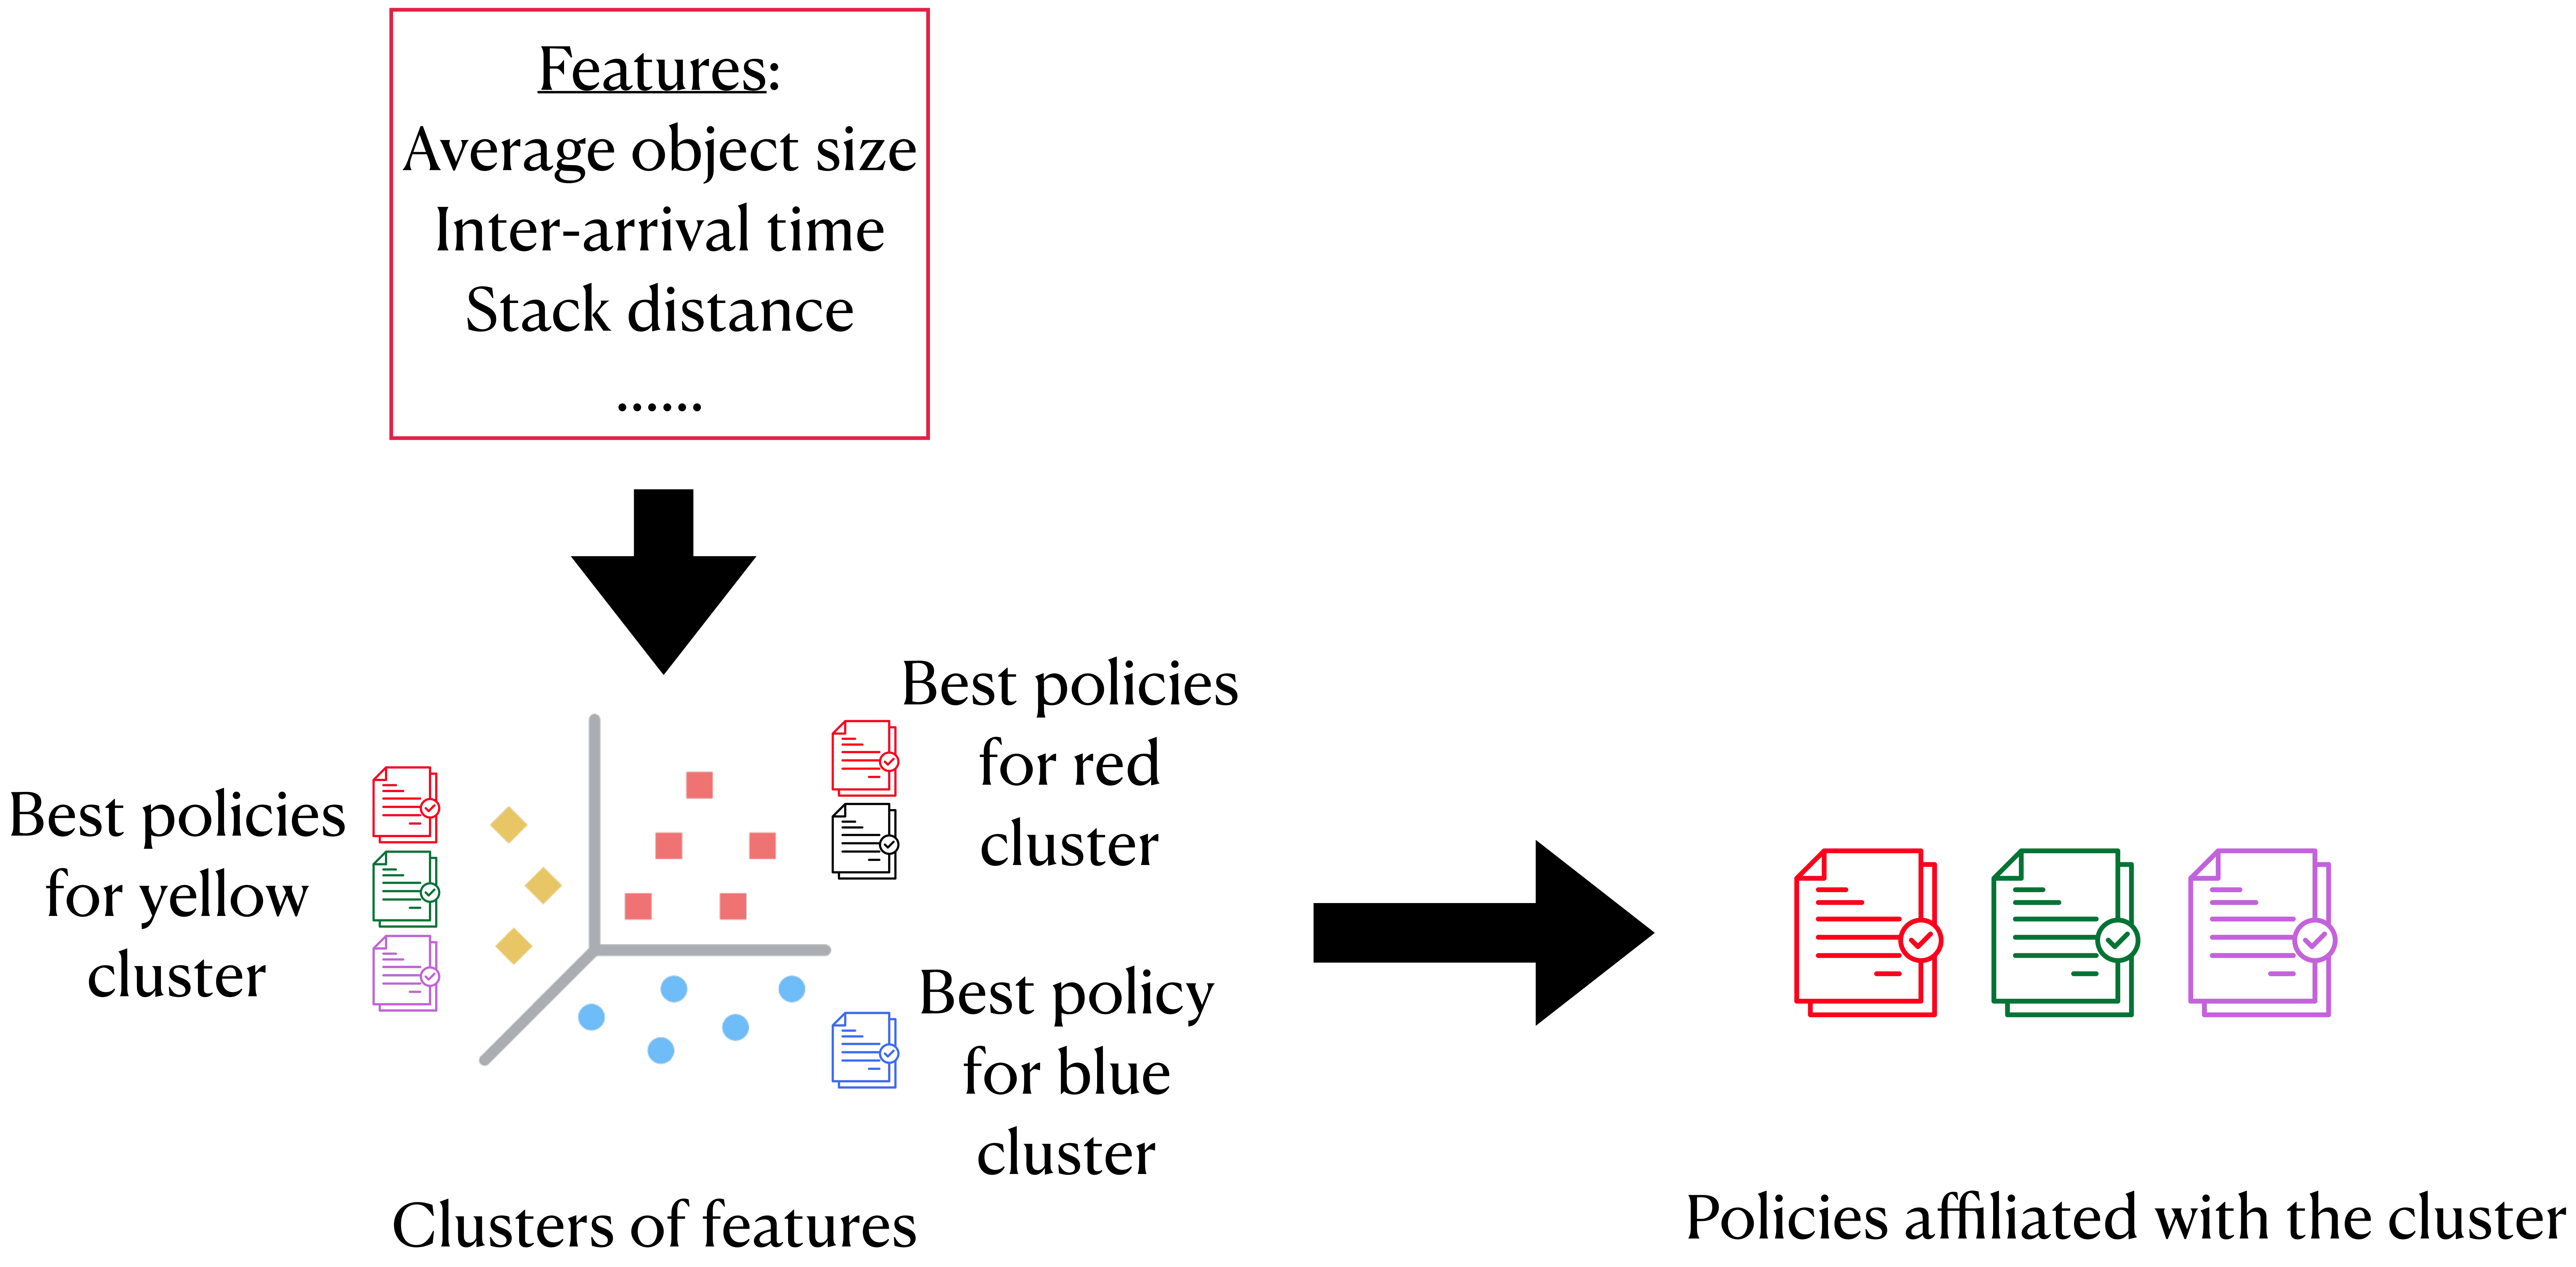

Challenge 2: Efficient Policy Selection

Challenge 2: Efficient Policy Selection

Track and Stop with Side Info (Bandit Best-arm Identification)

Challenge 2: Efficient Policy Selection

Track and Stop with Side Info (Bandit Best-arm Identification)

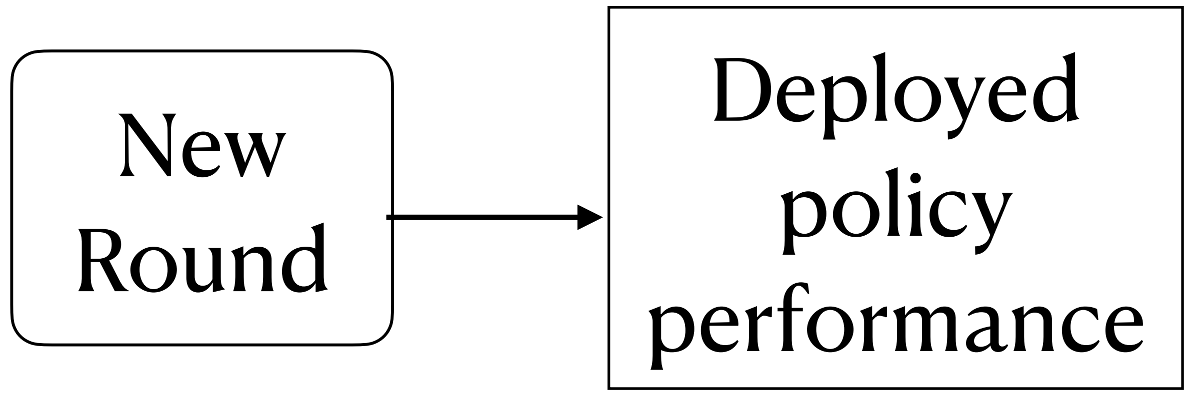

Challenge 2: Efficient Policy Selection

Track and Stop with Side Info (Bandit Best-arm Identification)

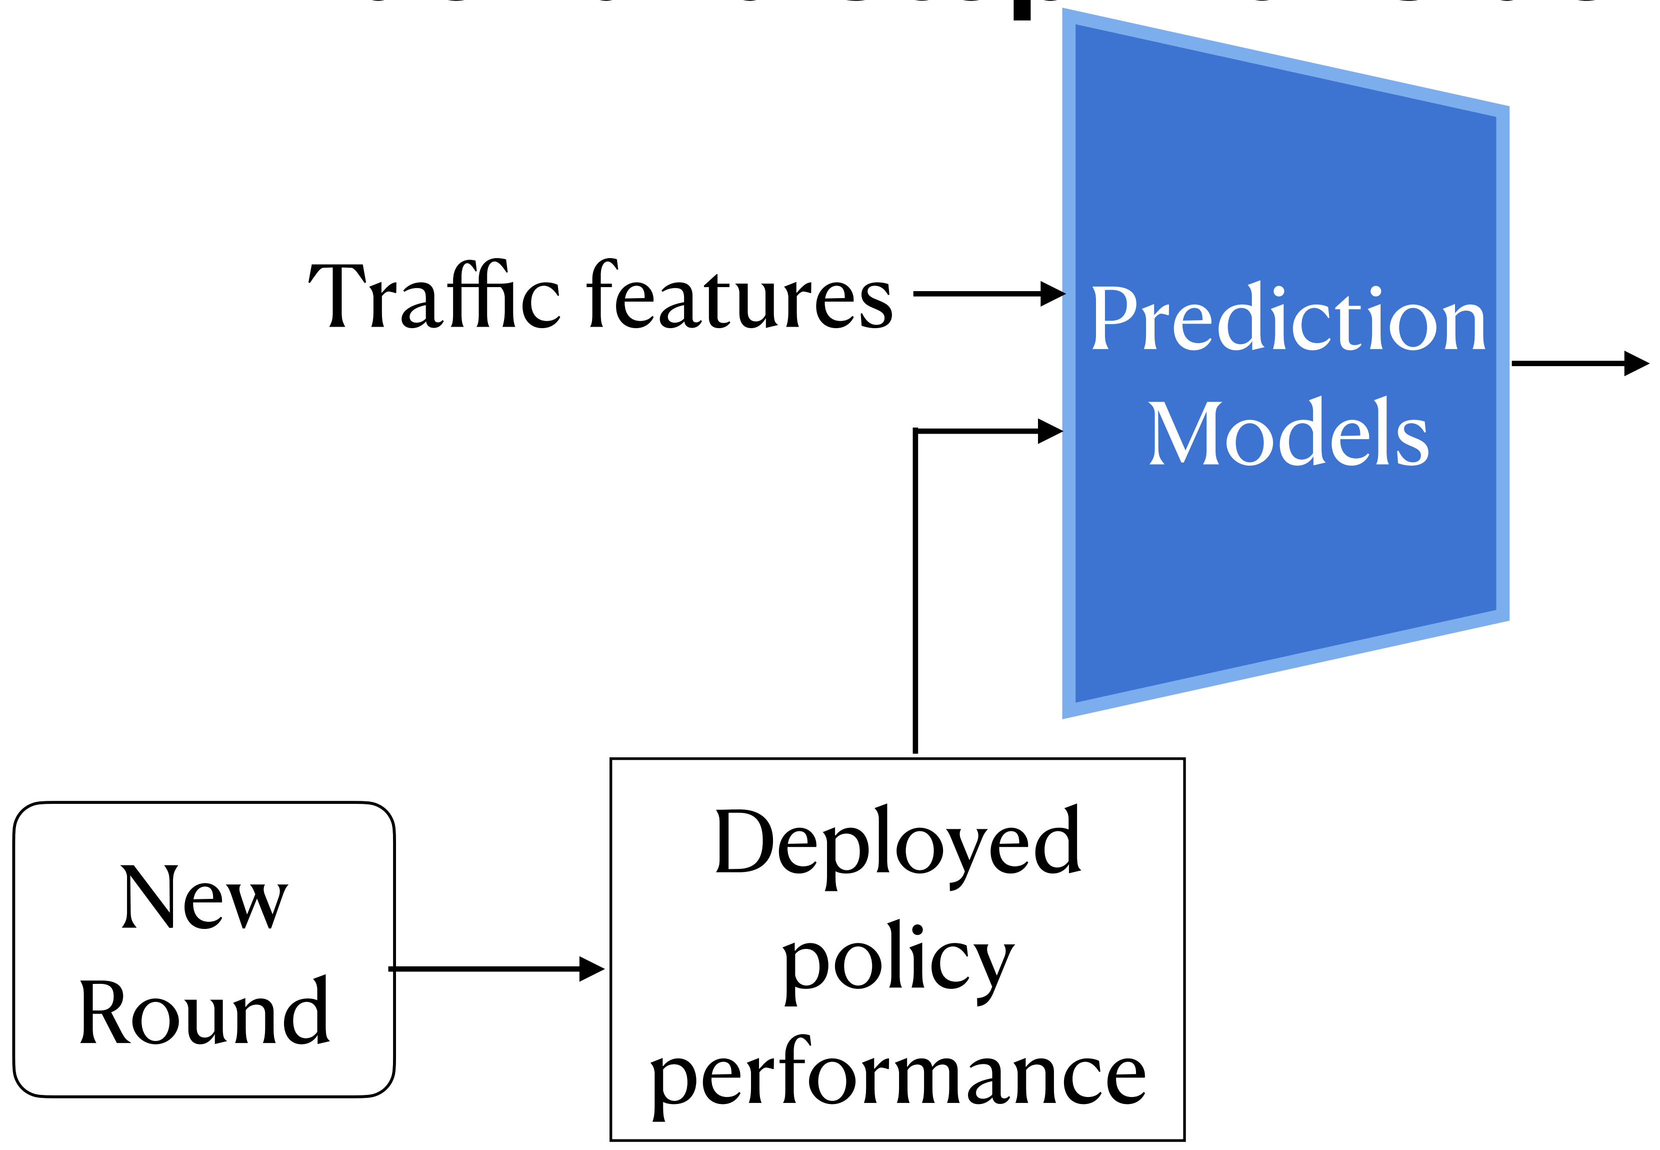

Challenge 2: Efficient Policy Selection

Track and Stop with Side Info (Bandit Best-arm Identification)

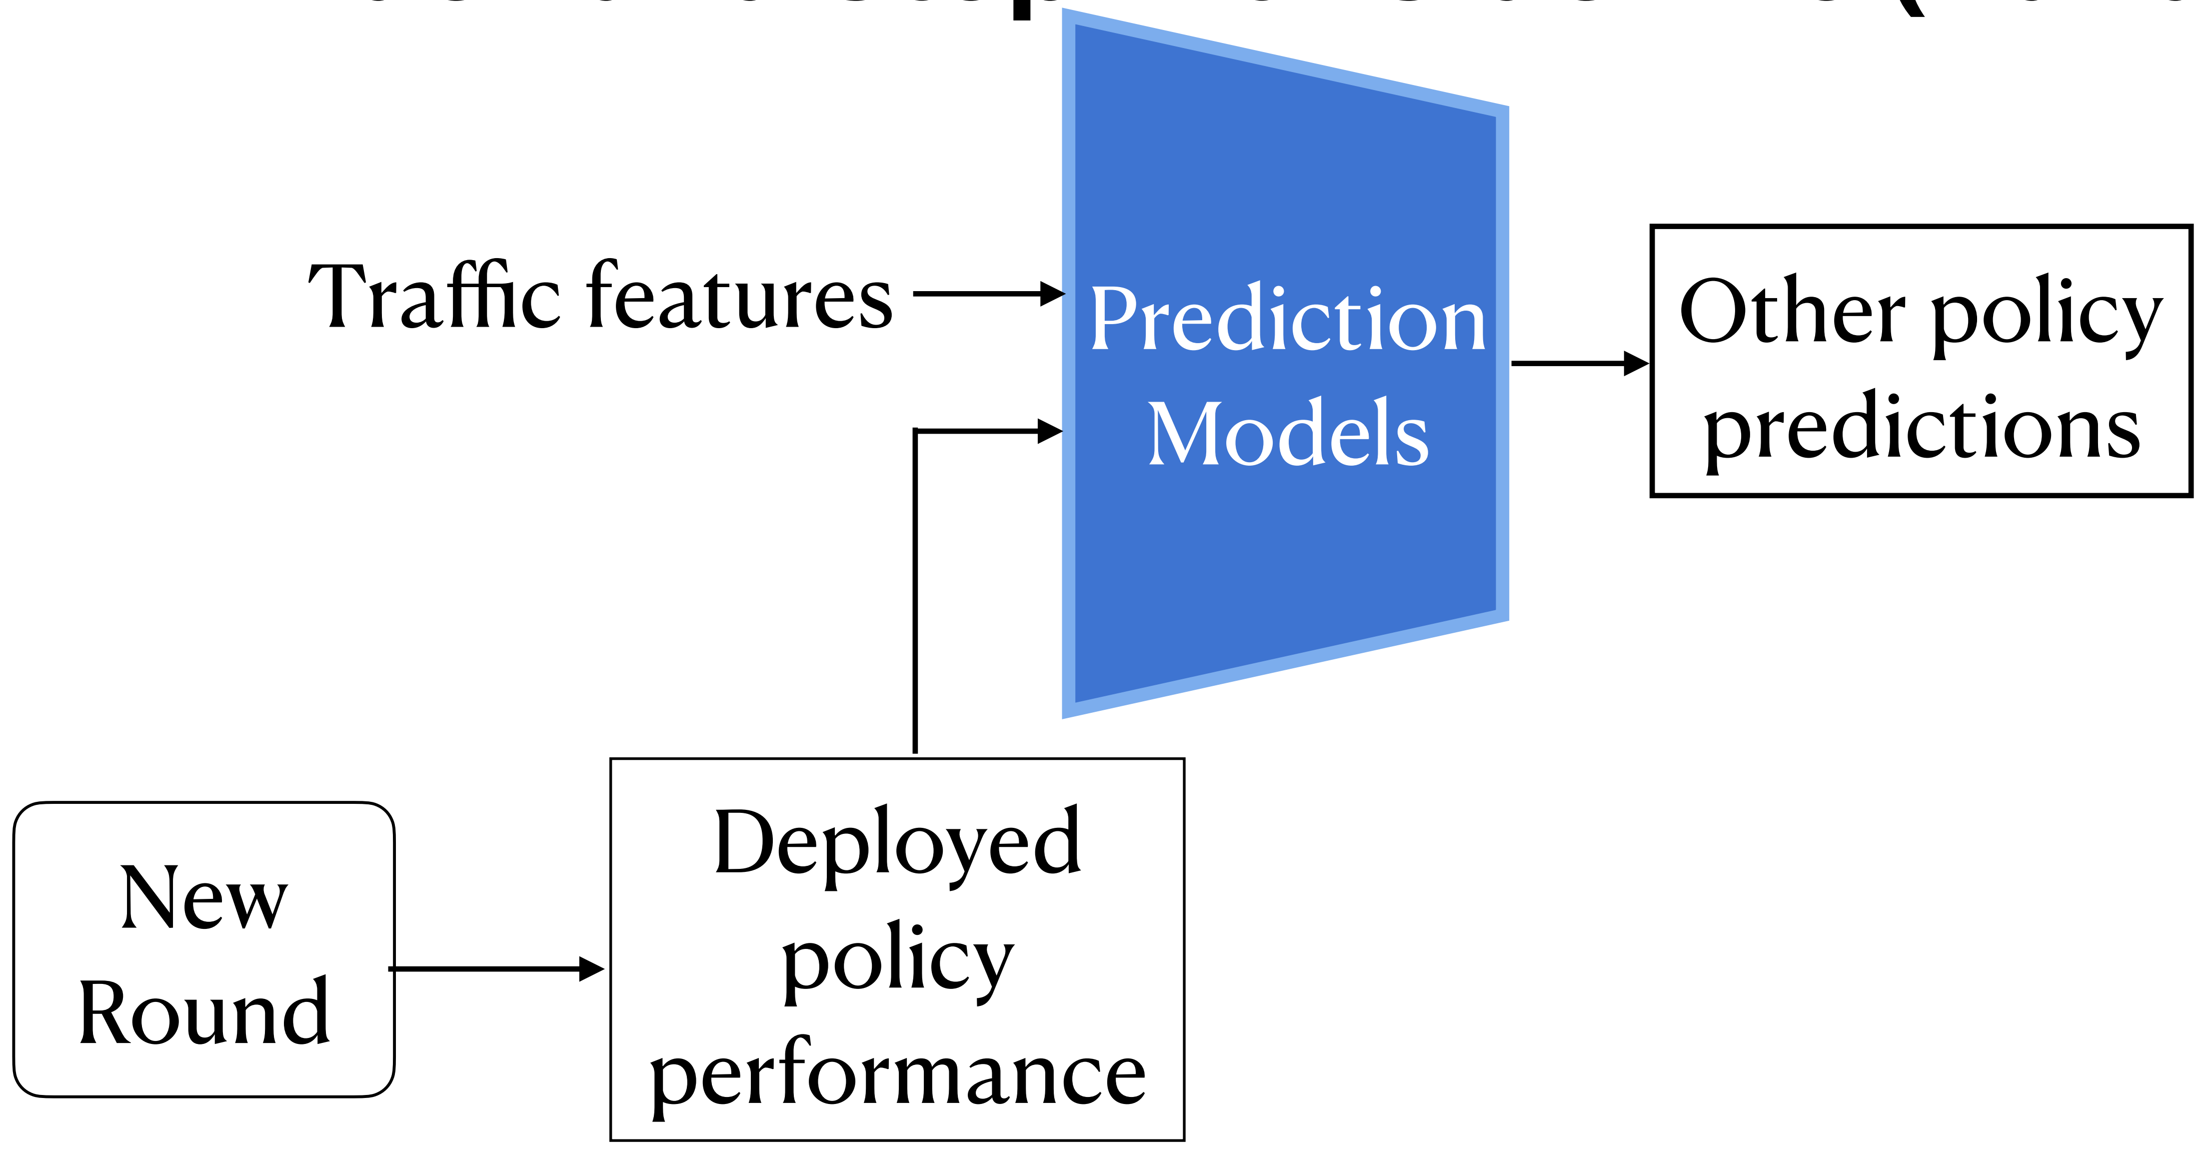

Challenge 2: Efficient Policy Selection

Track and Stop with Side Info (Bandit Best-arm Identification)

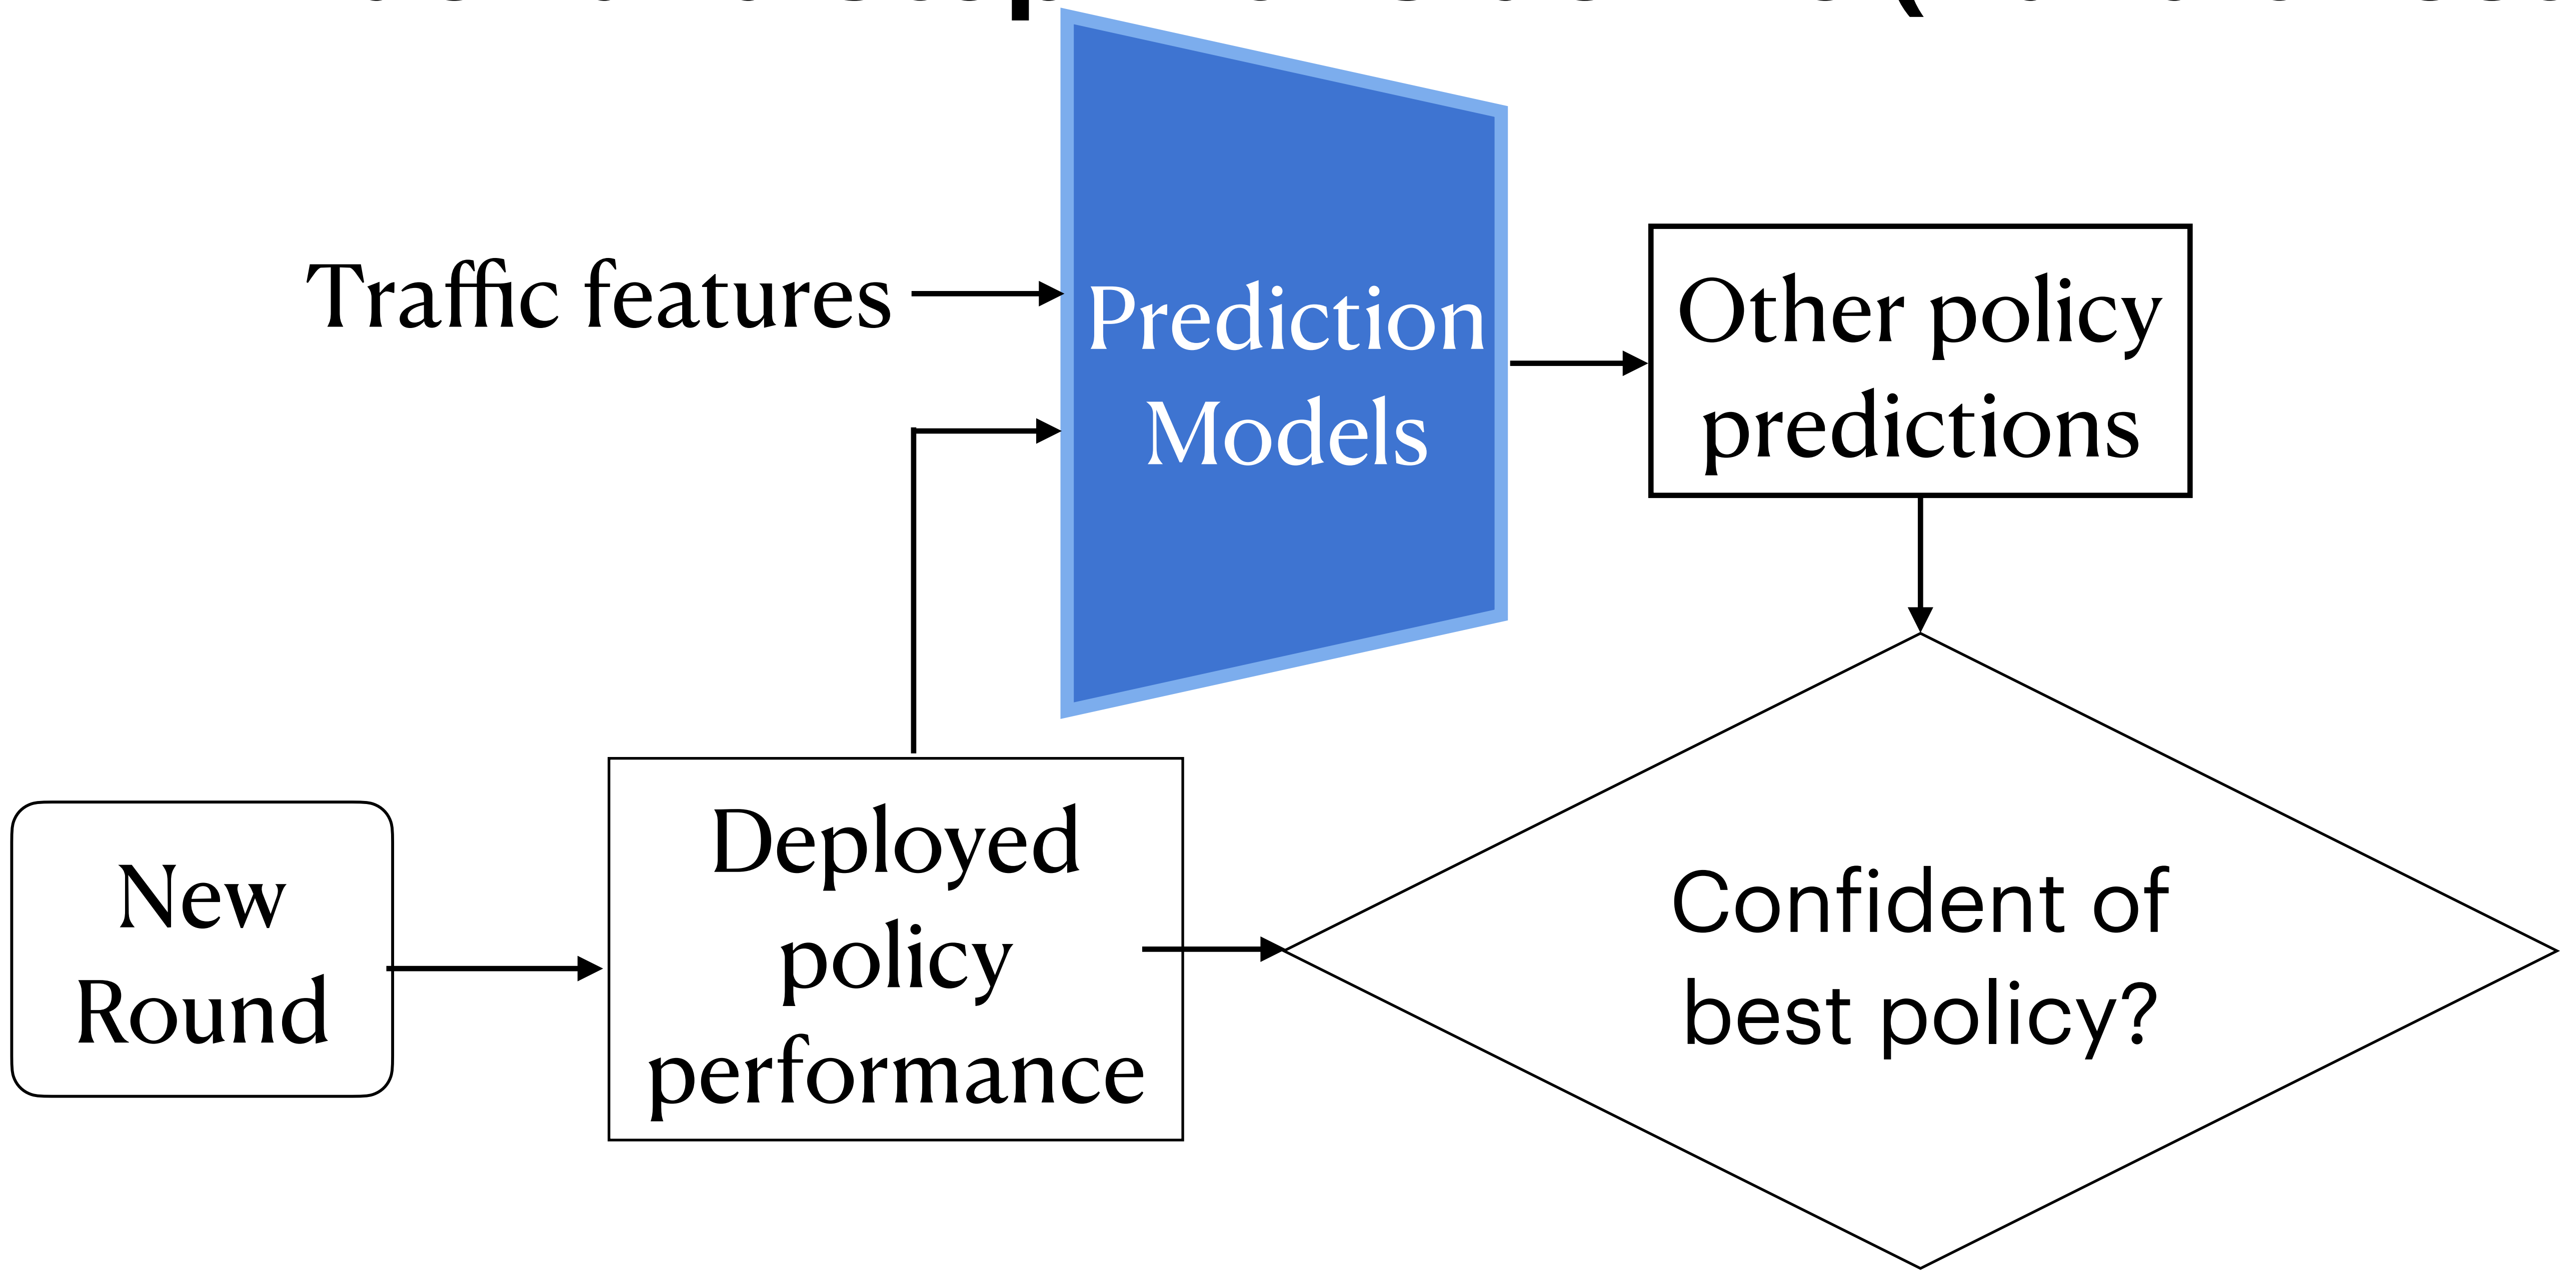

Challenge 2: Efficient Policy Selection

Track and Stop with Side Info (Bandit Best-arm Identification)

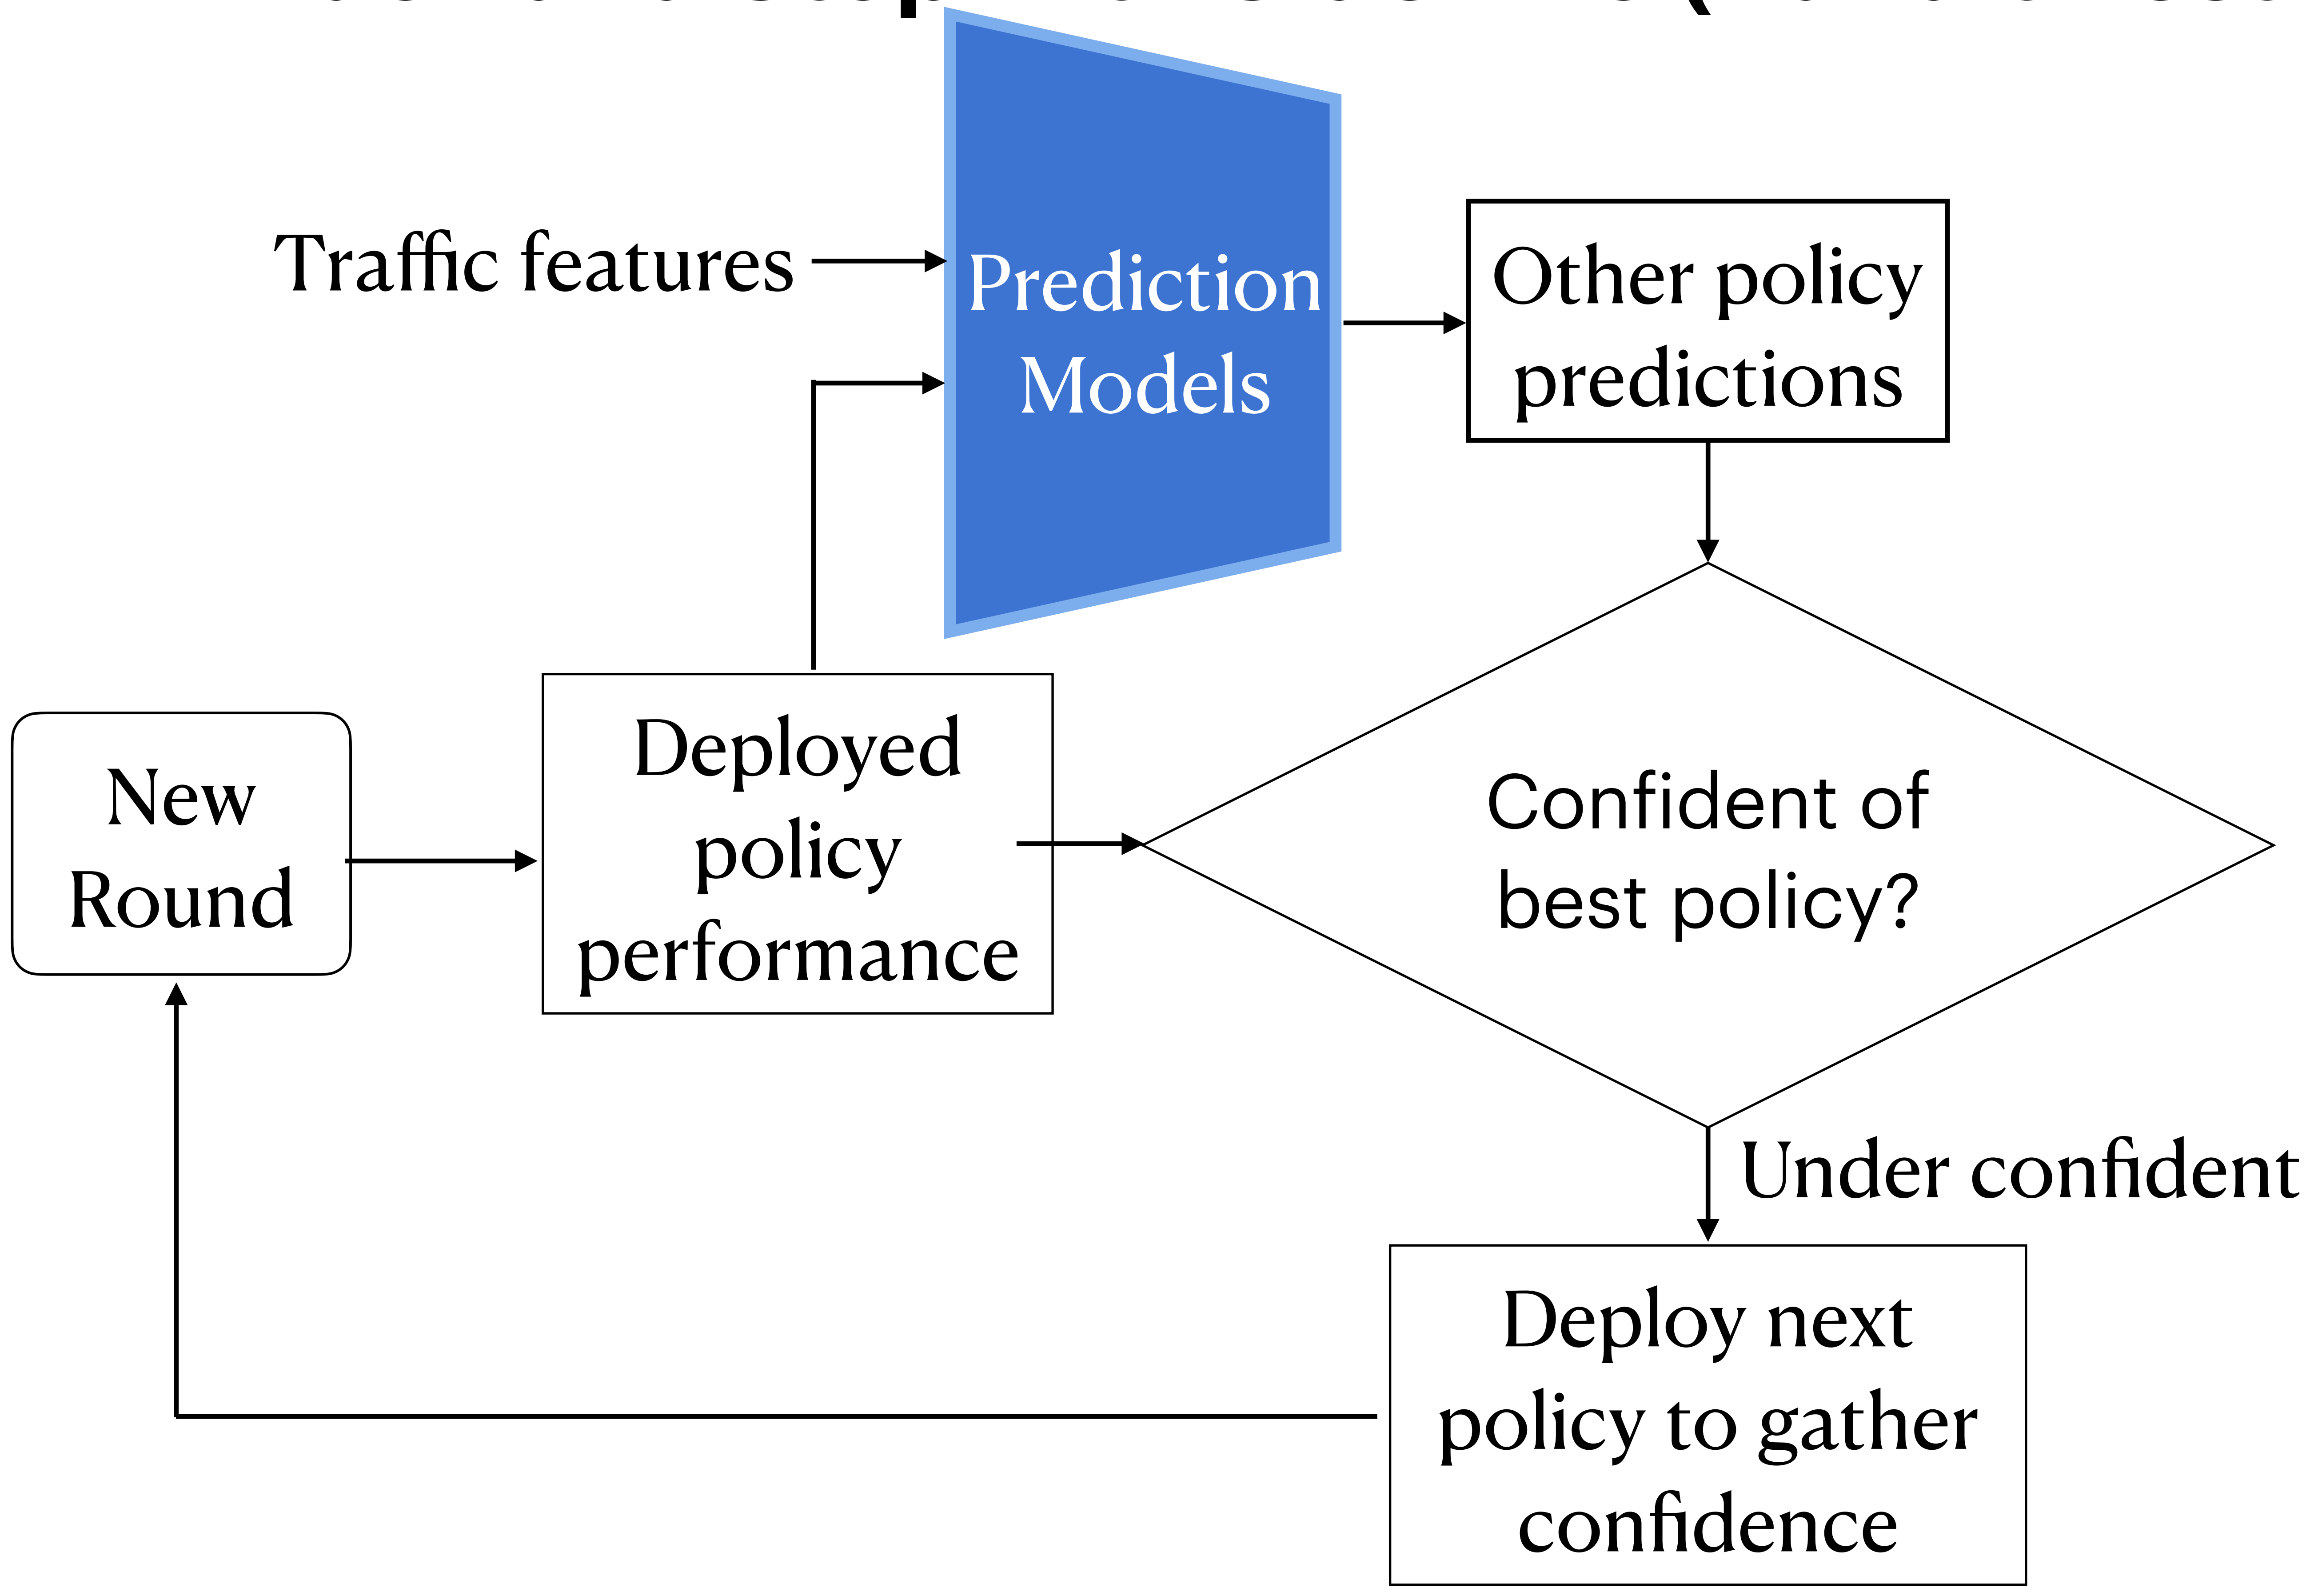

Challenge 2: Efficient Policy Selection

Track and Stop with Side Info (Bandit Best-arm Identification)

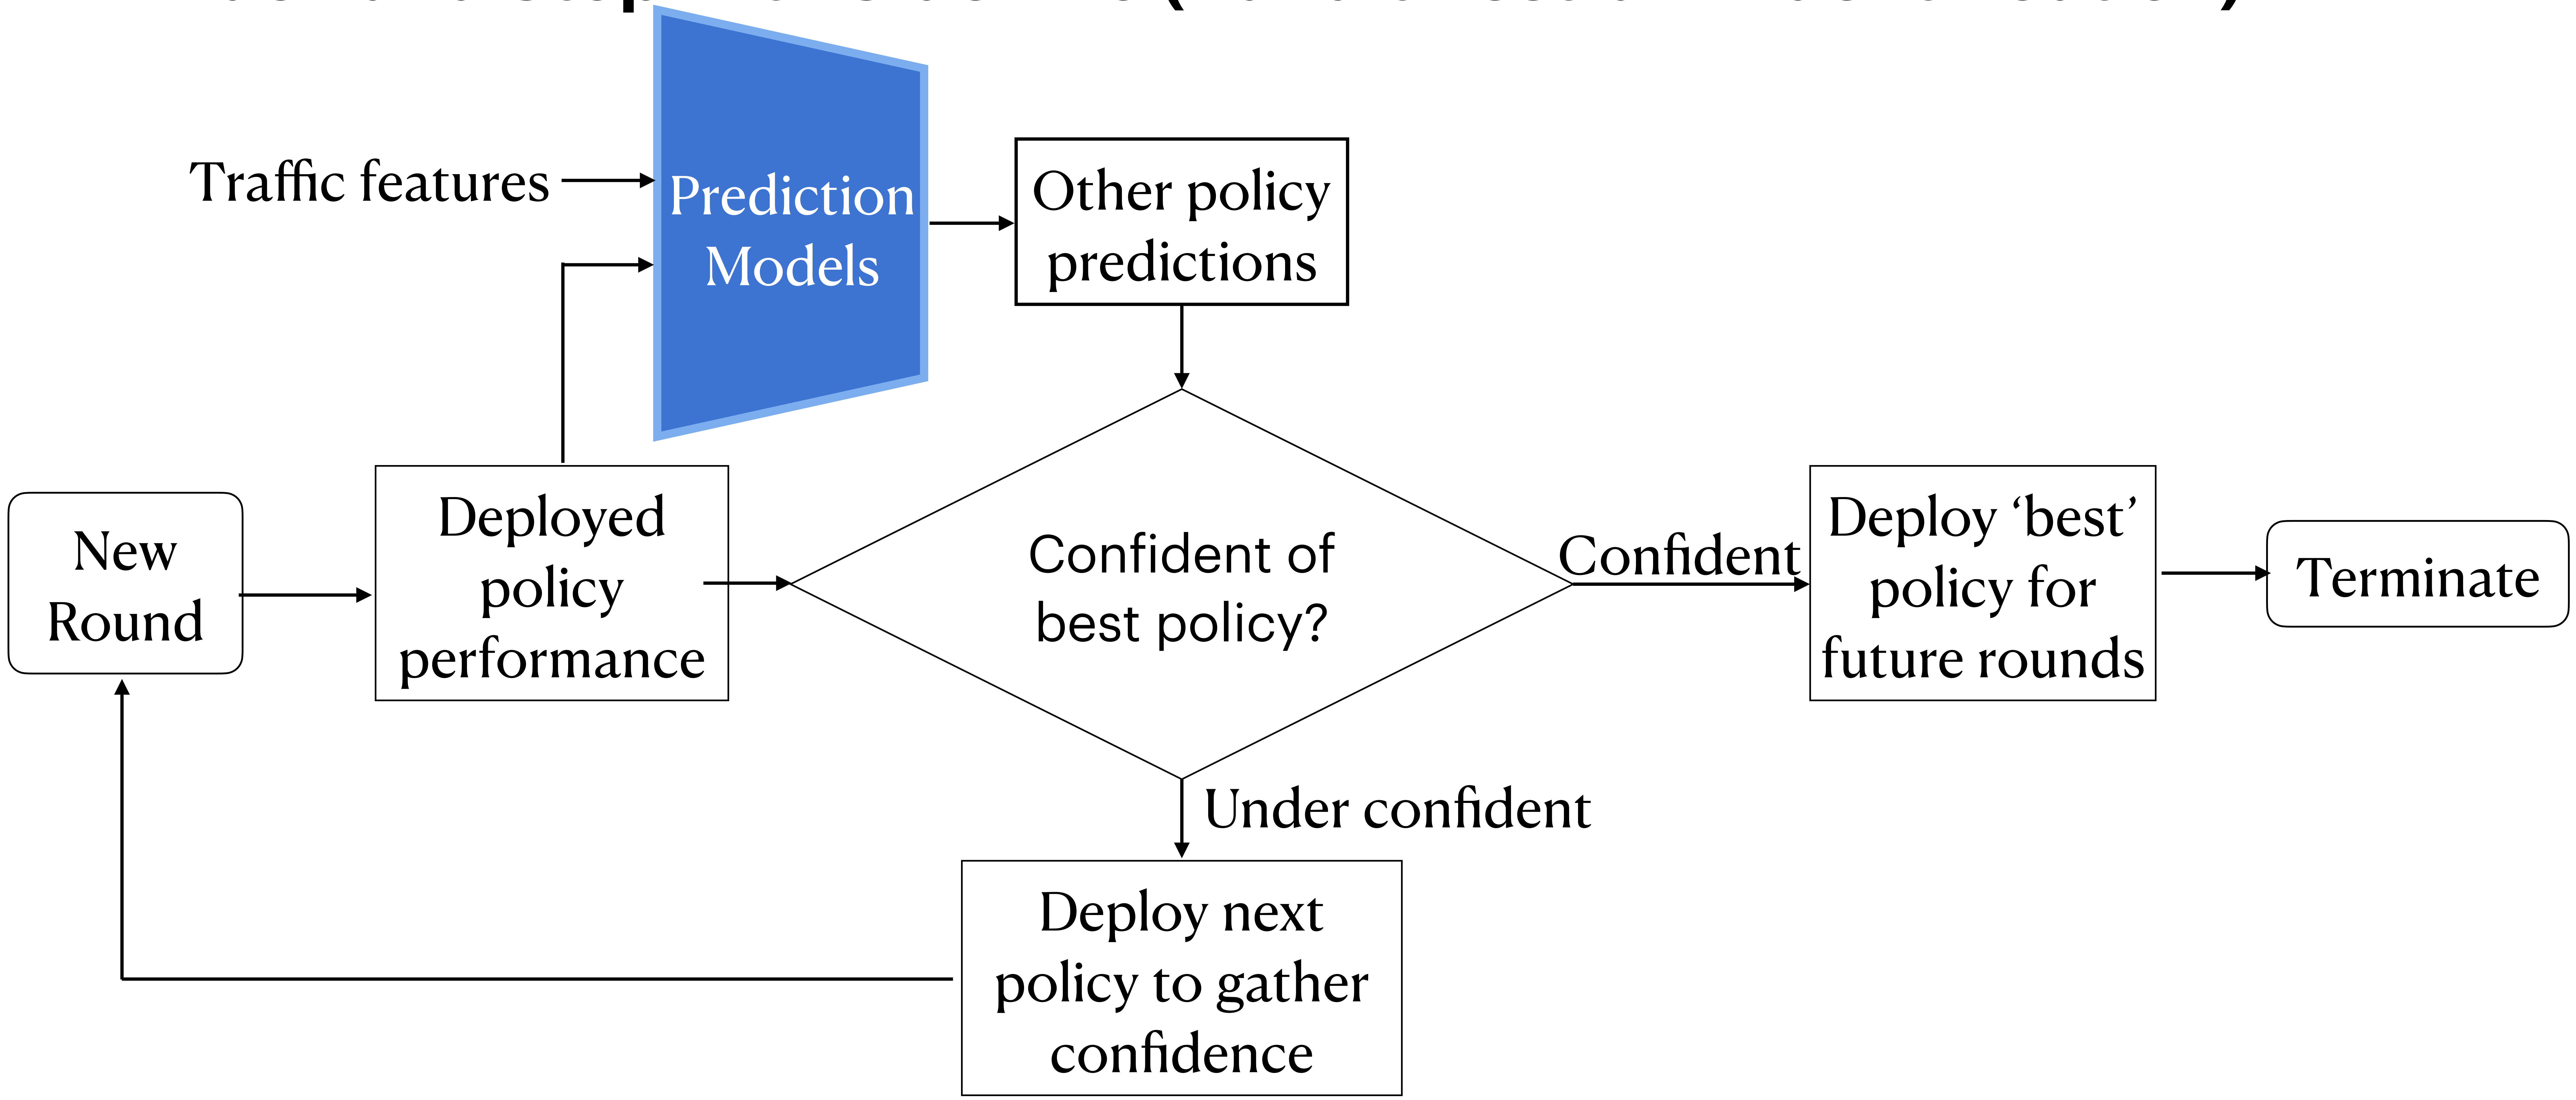

Challenge 2: Efficient Policy Selection

Track and Stop with Side Info (Bandit Best-arm Identification)

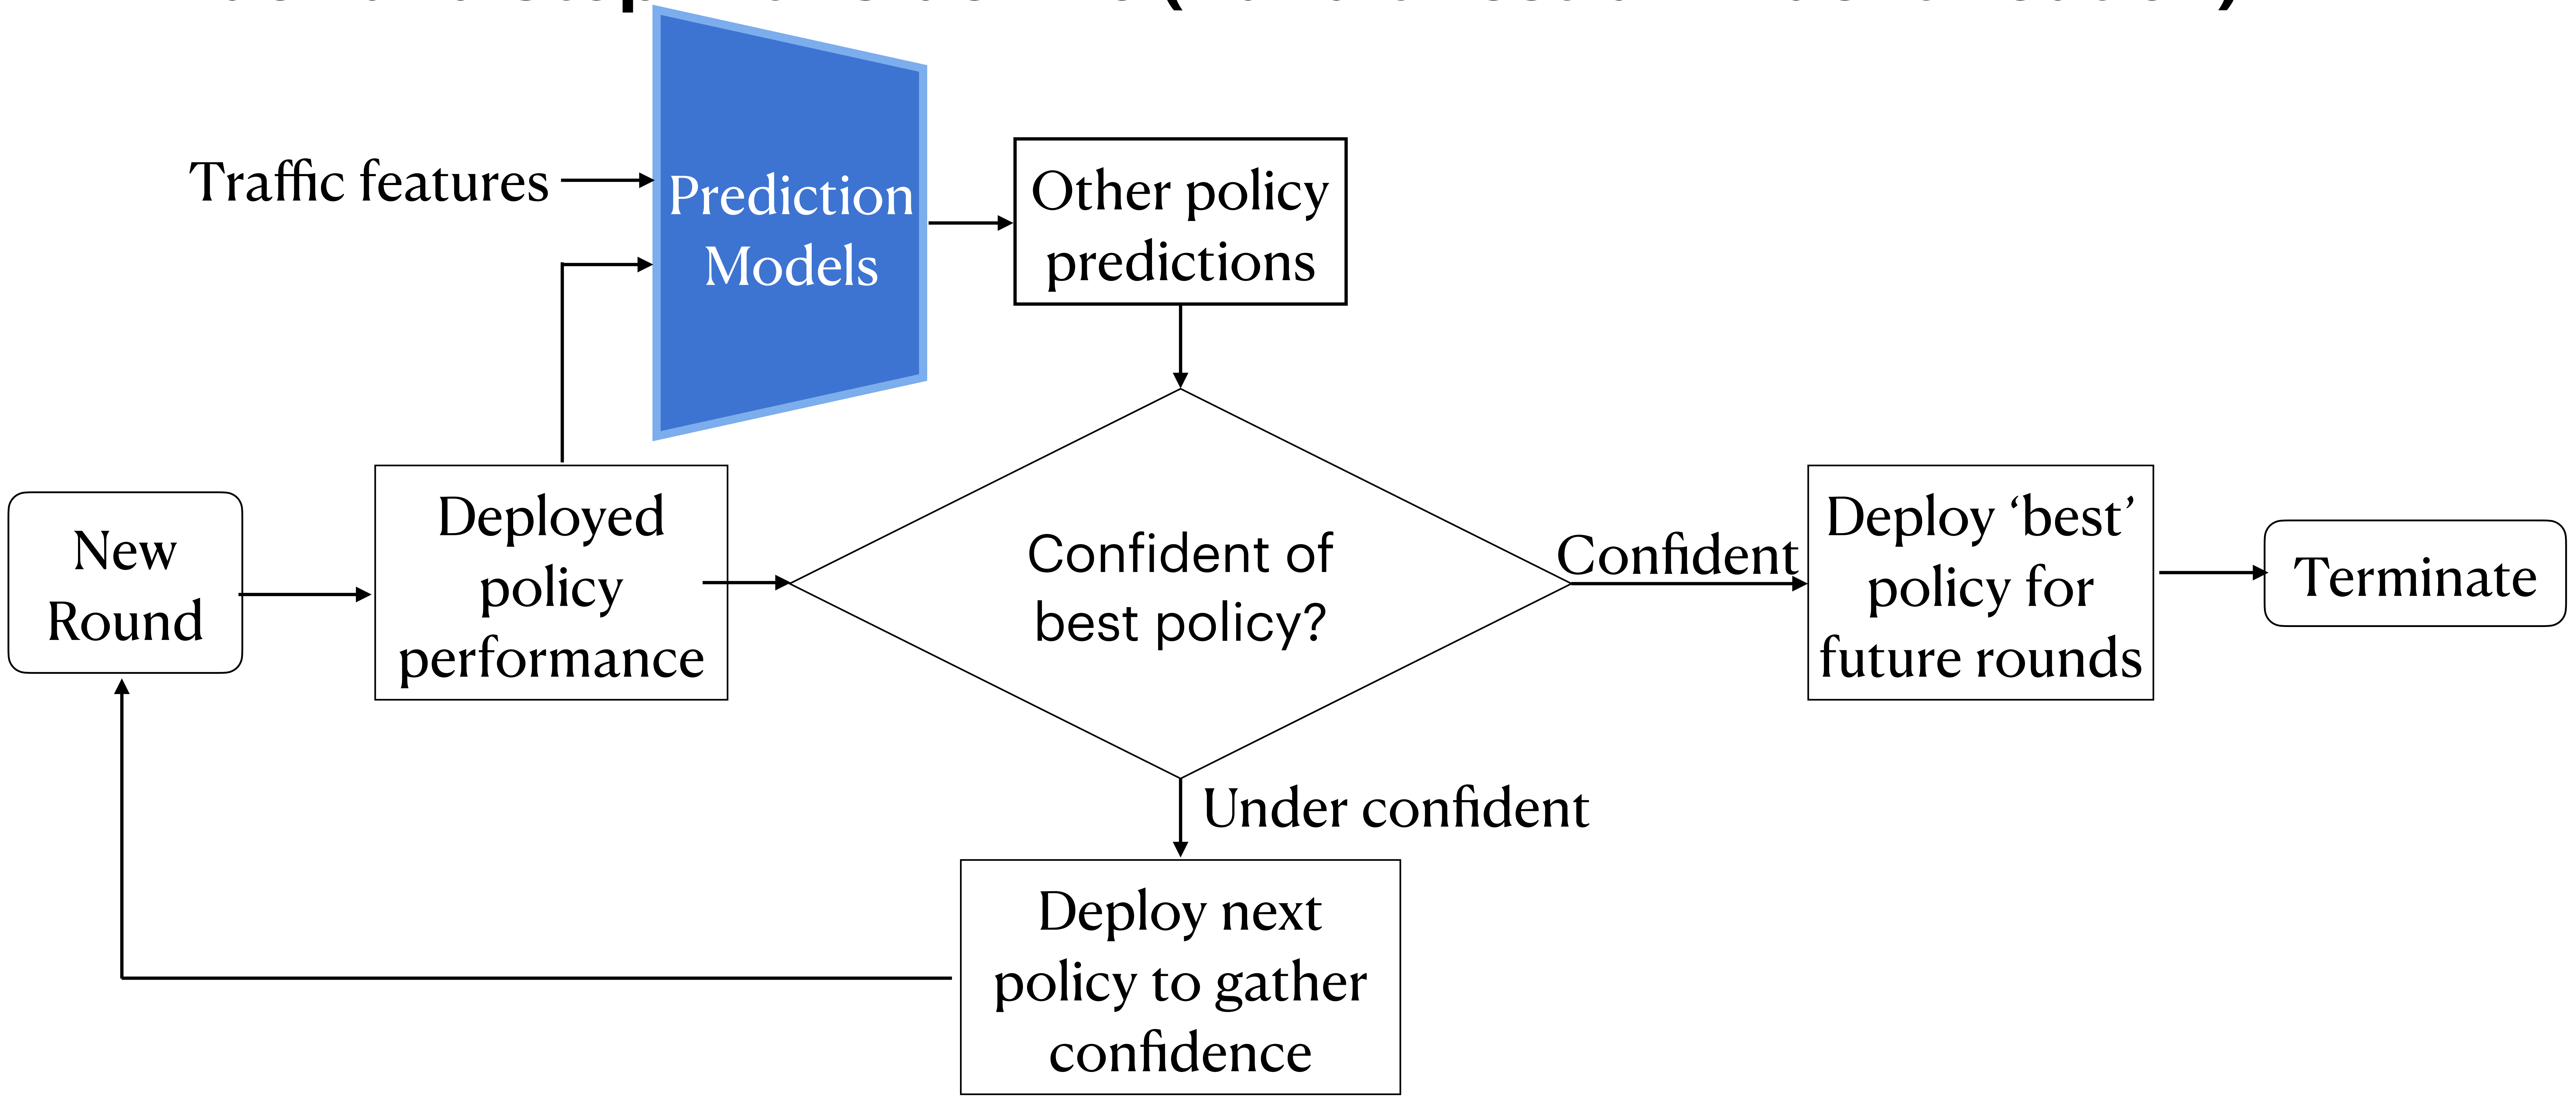

Theorem: Convergence time is bounded by a constant that is **independent of the number of policies**

Darwin Design Overview

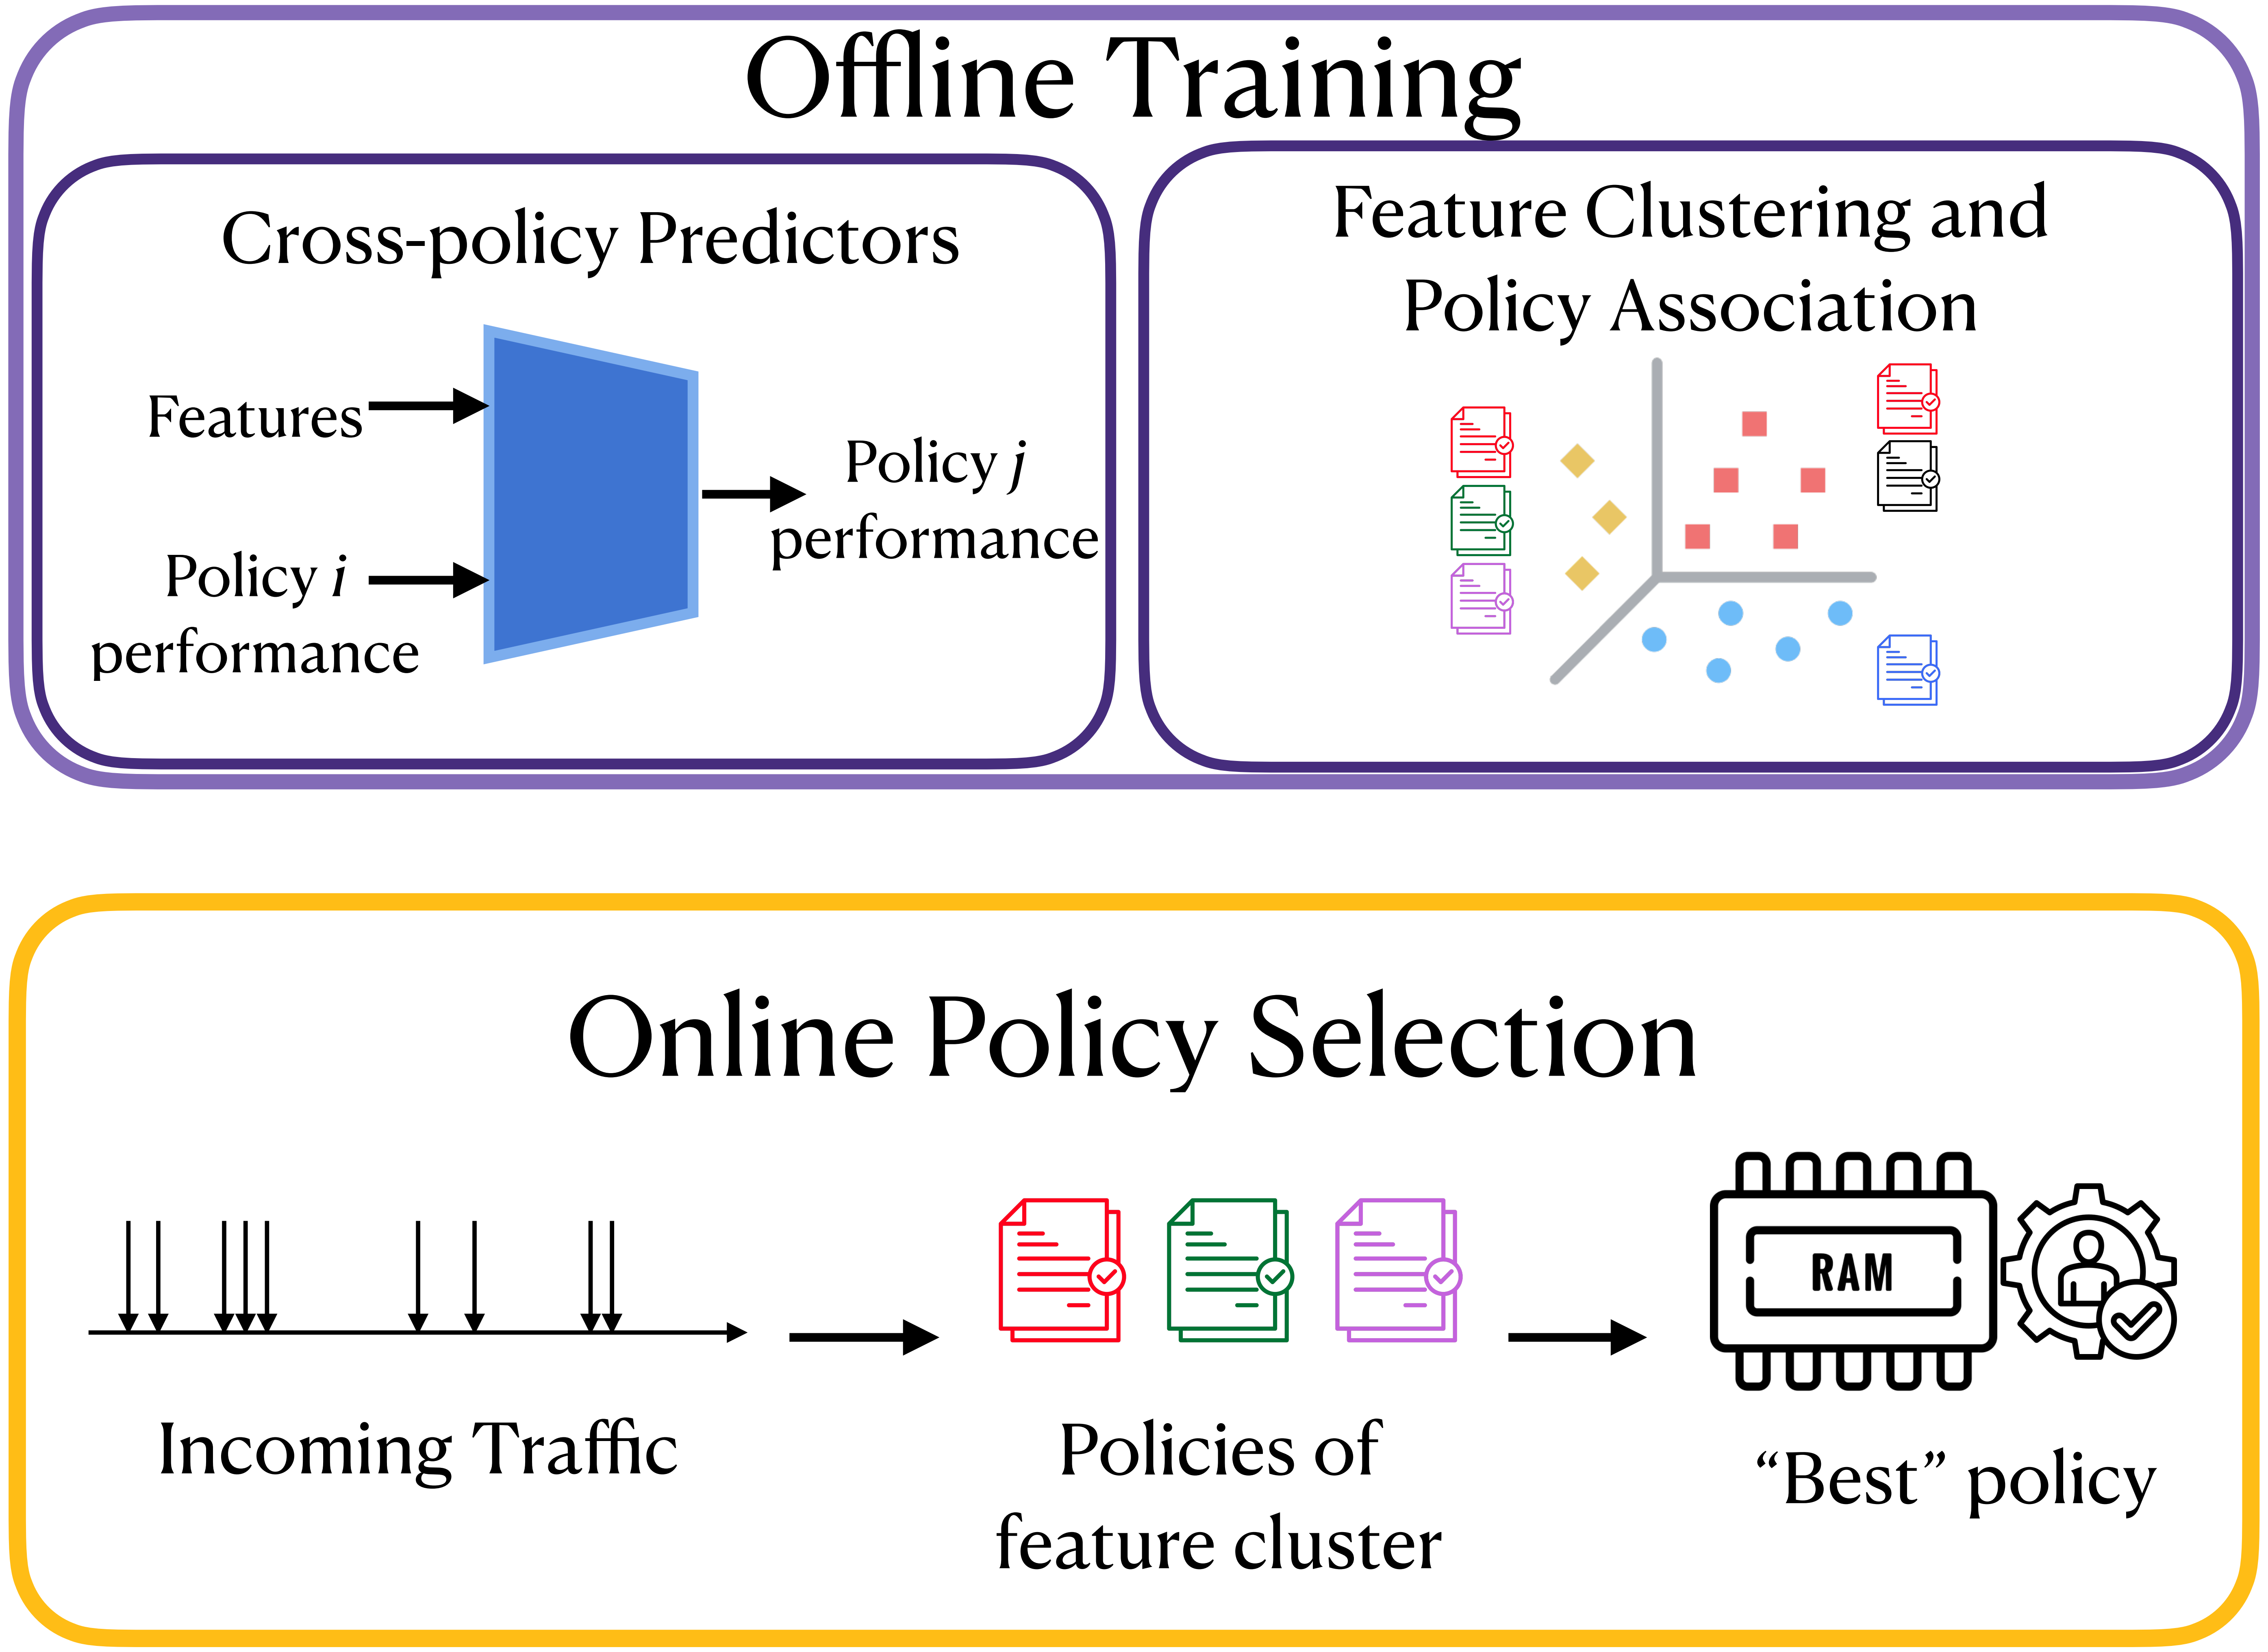

Darwin Design Overview

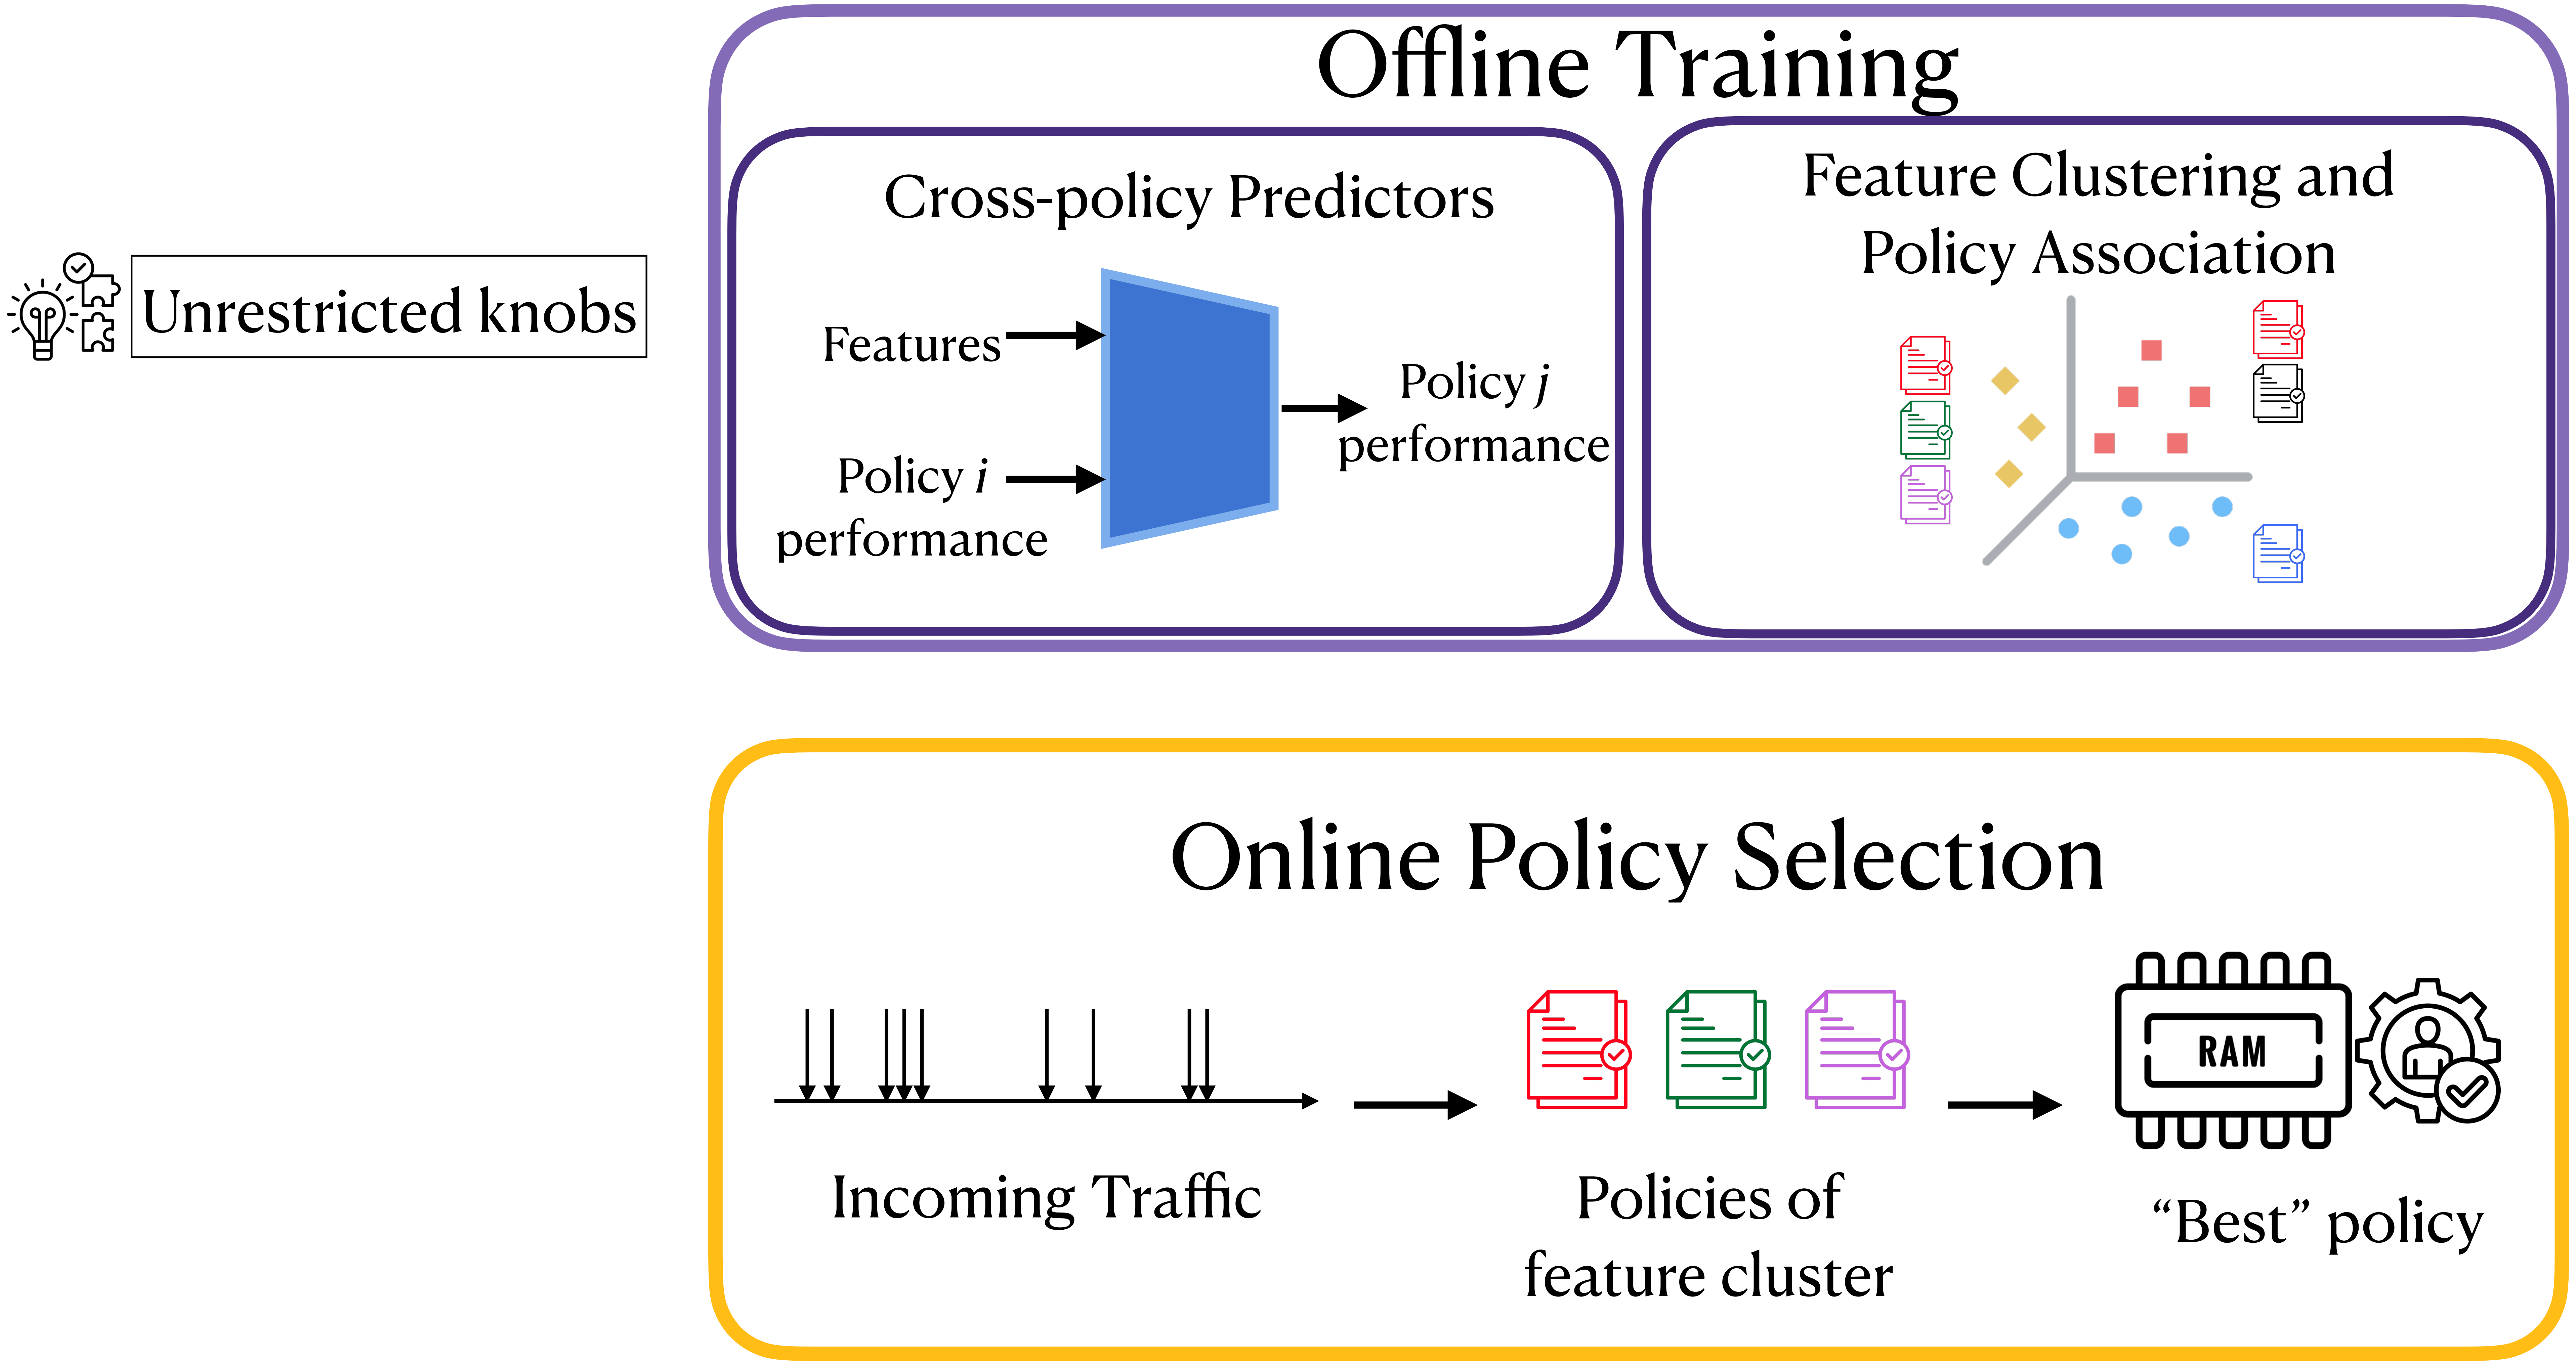

Darwin Design Overview

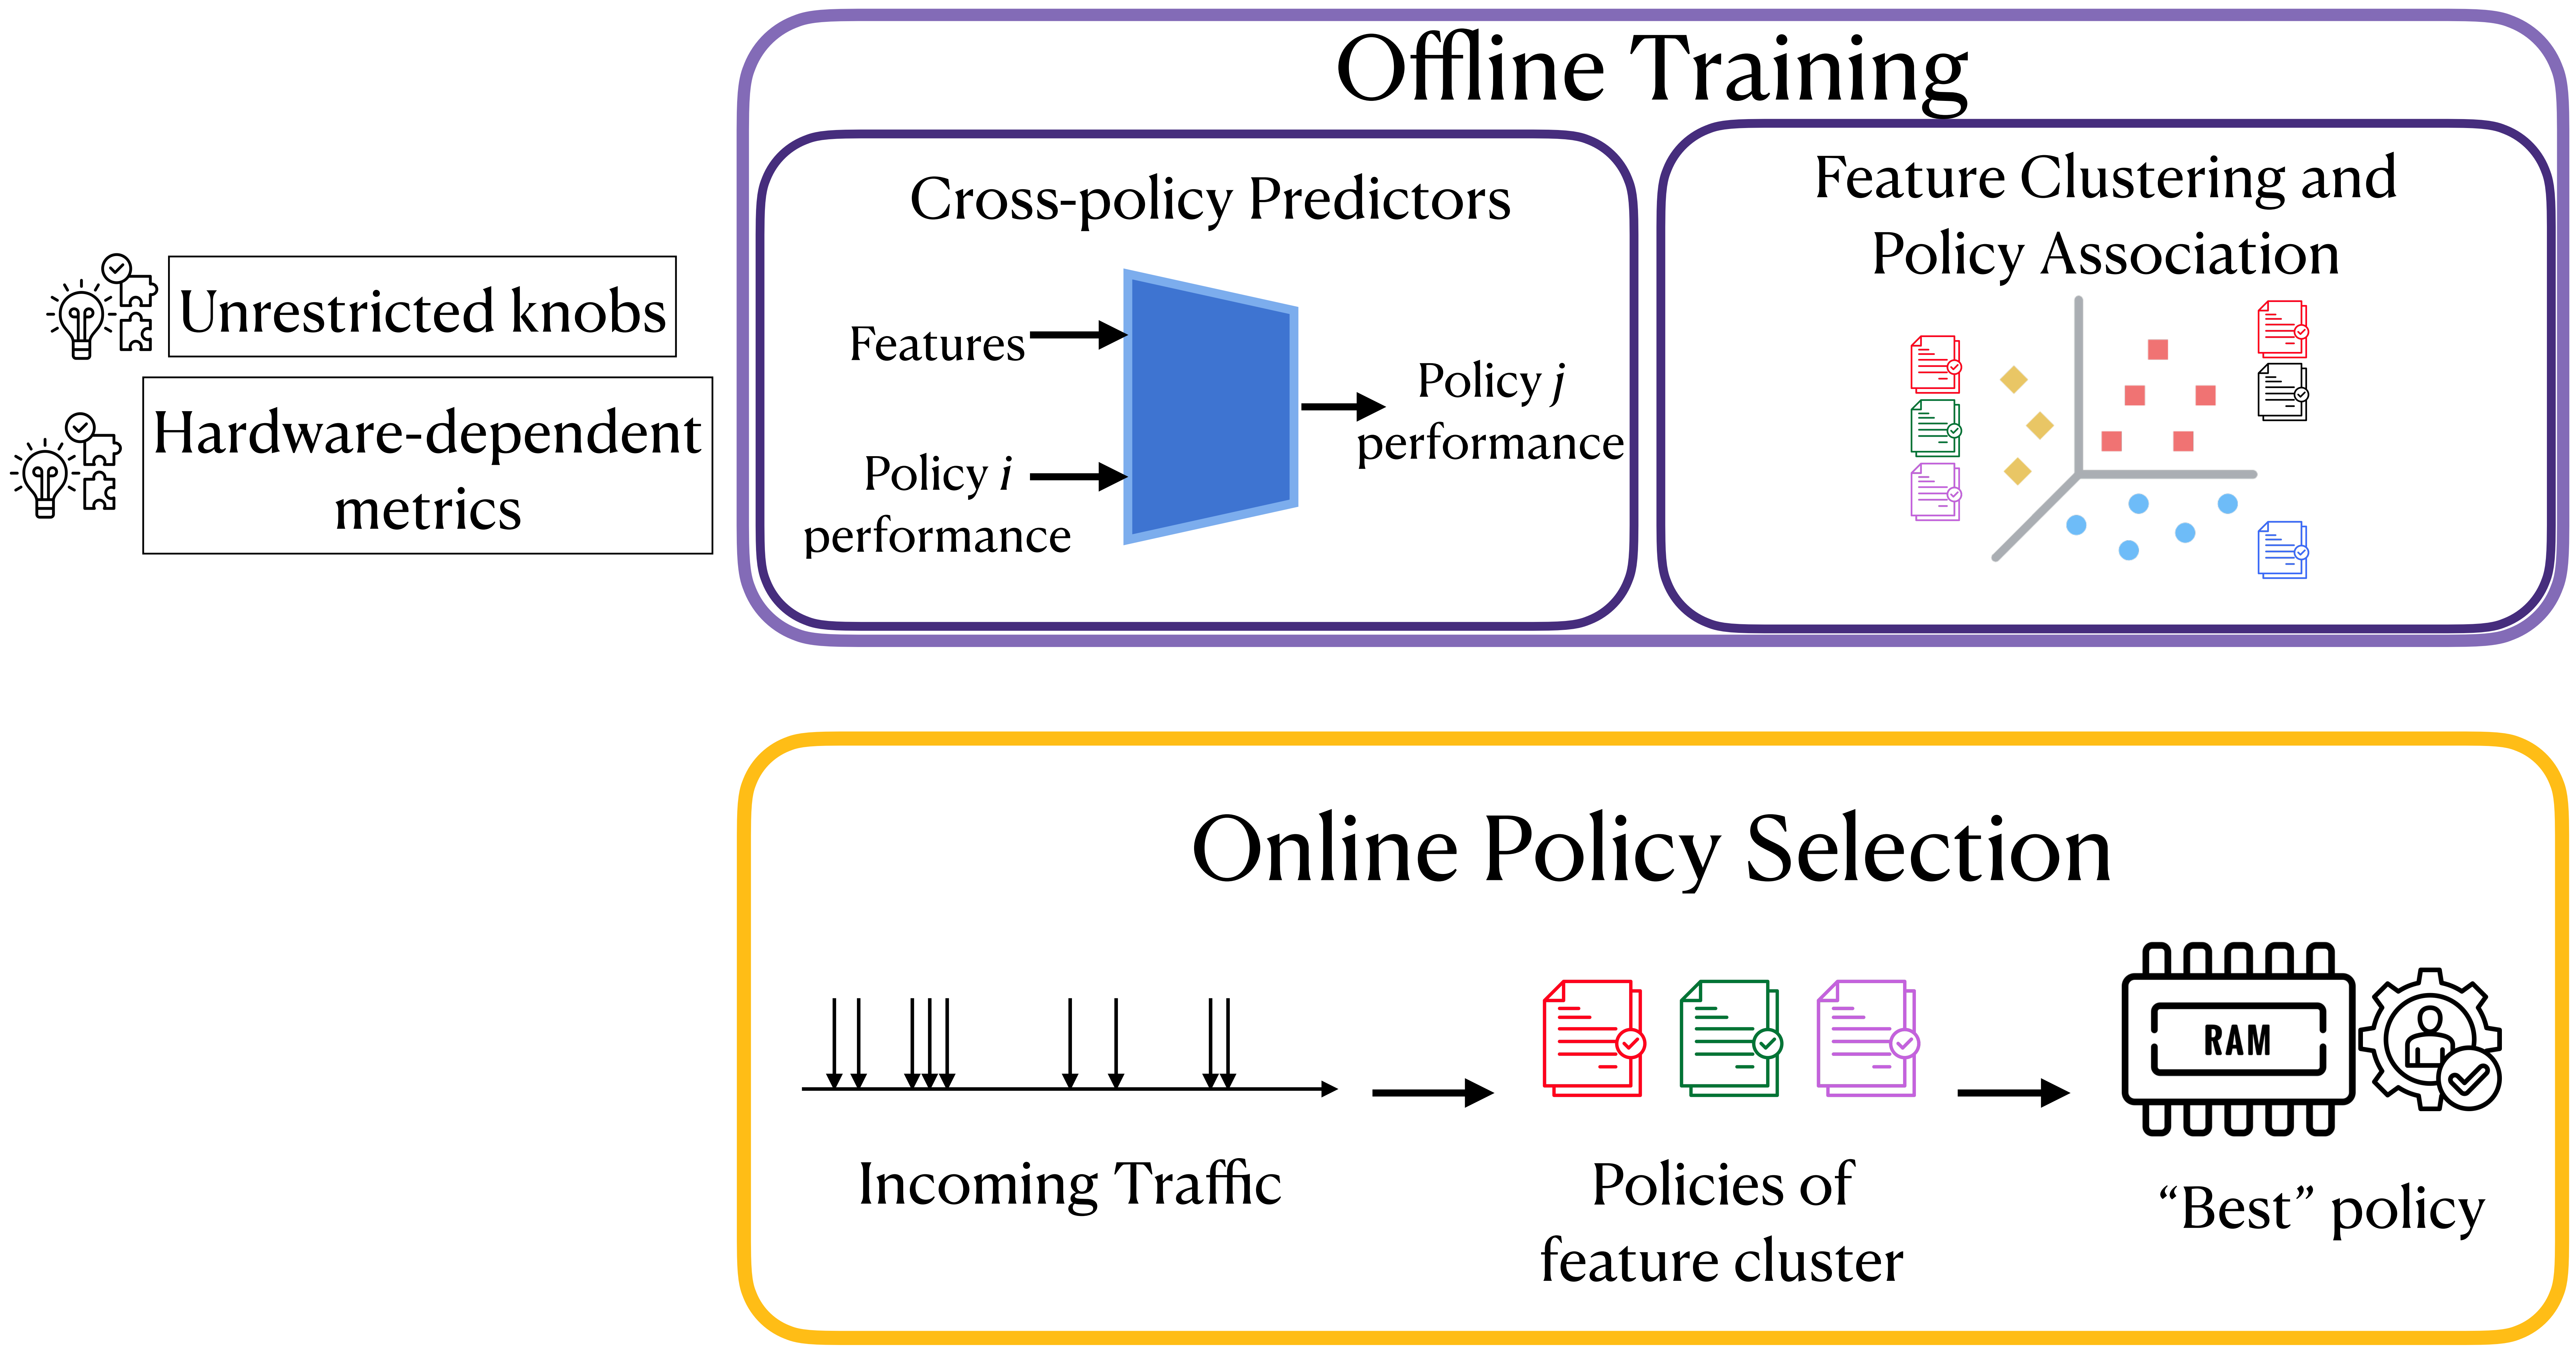

Darwin Design Overview

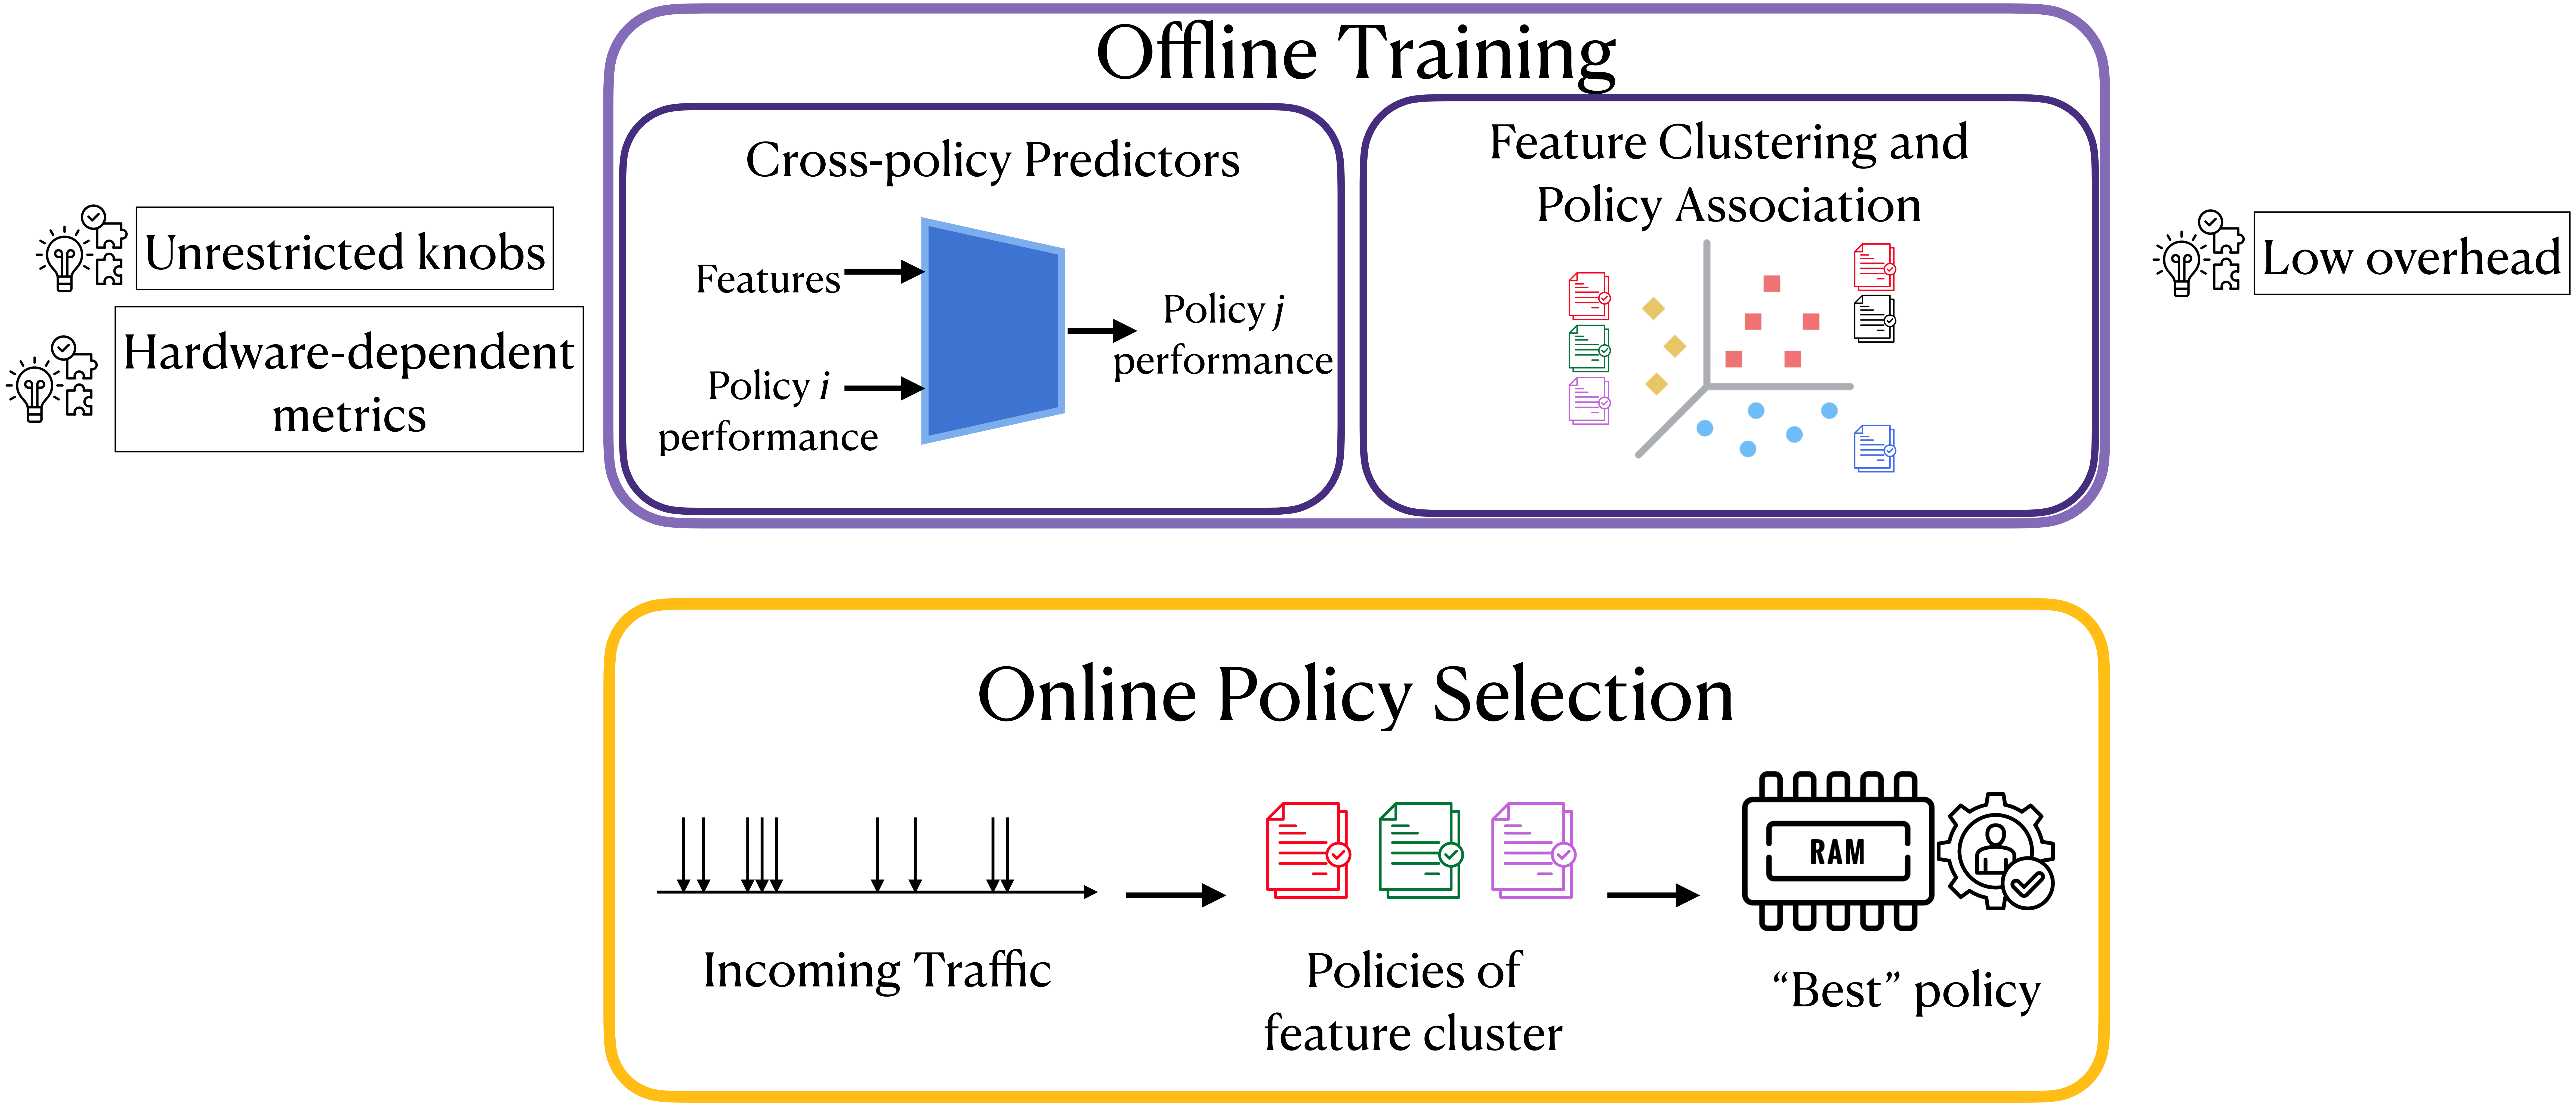

Evaluation Setup

Darwin Simulator [1] and Apache Traffic Server (ATS)-based Prototype

- HOC Cache Size
  - 100MB, 200MB, 500MB
- CDN Traces
  - 100 mixed configurations for two traffic classes
- Baselines
  - Static policies, AdaptSize, Percentile, HillClimbing

[1]: https://github.com/Janecjy/Darwin

Robustness to Traffic Changes

Darwin outperforms static baselines by 4.83%-28.16%

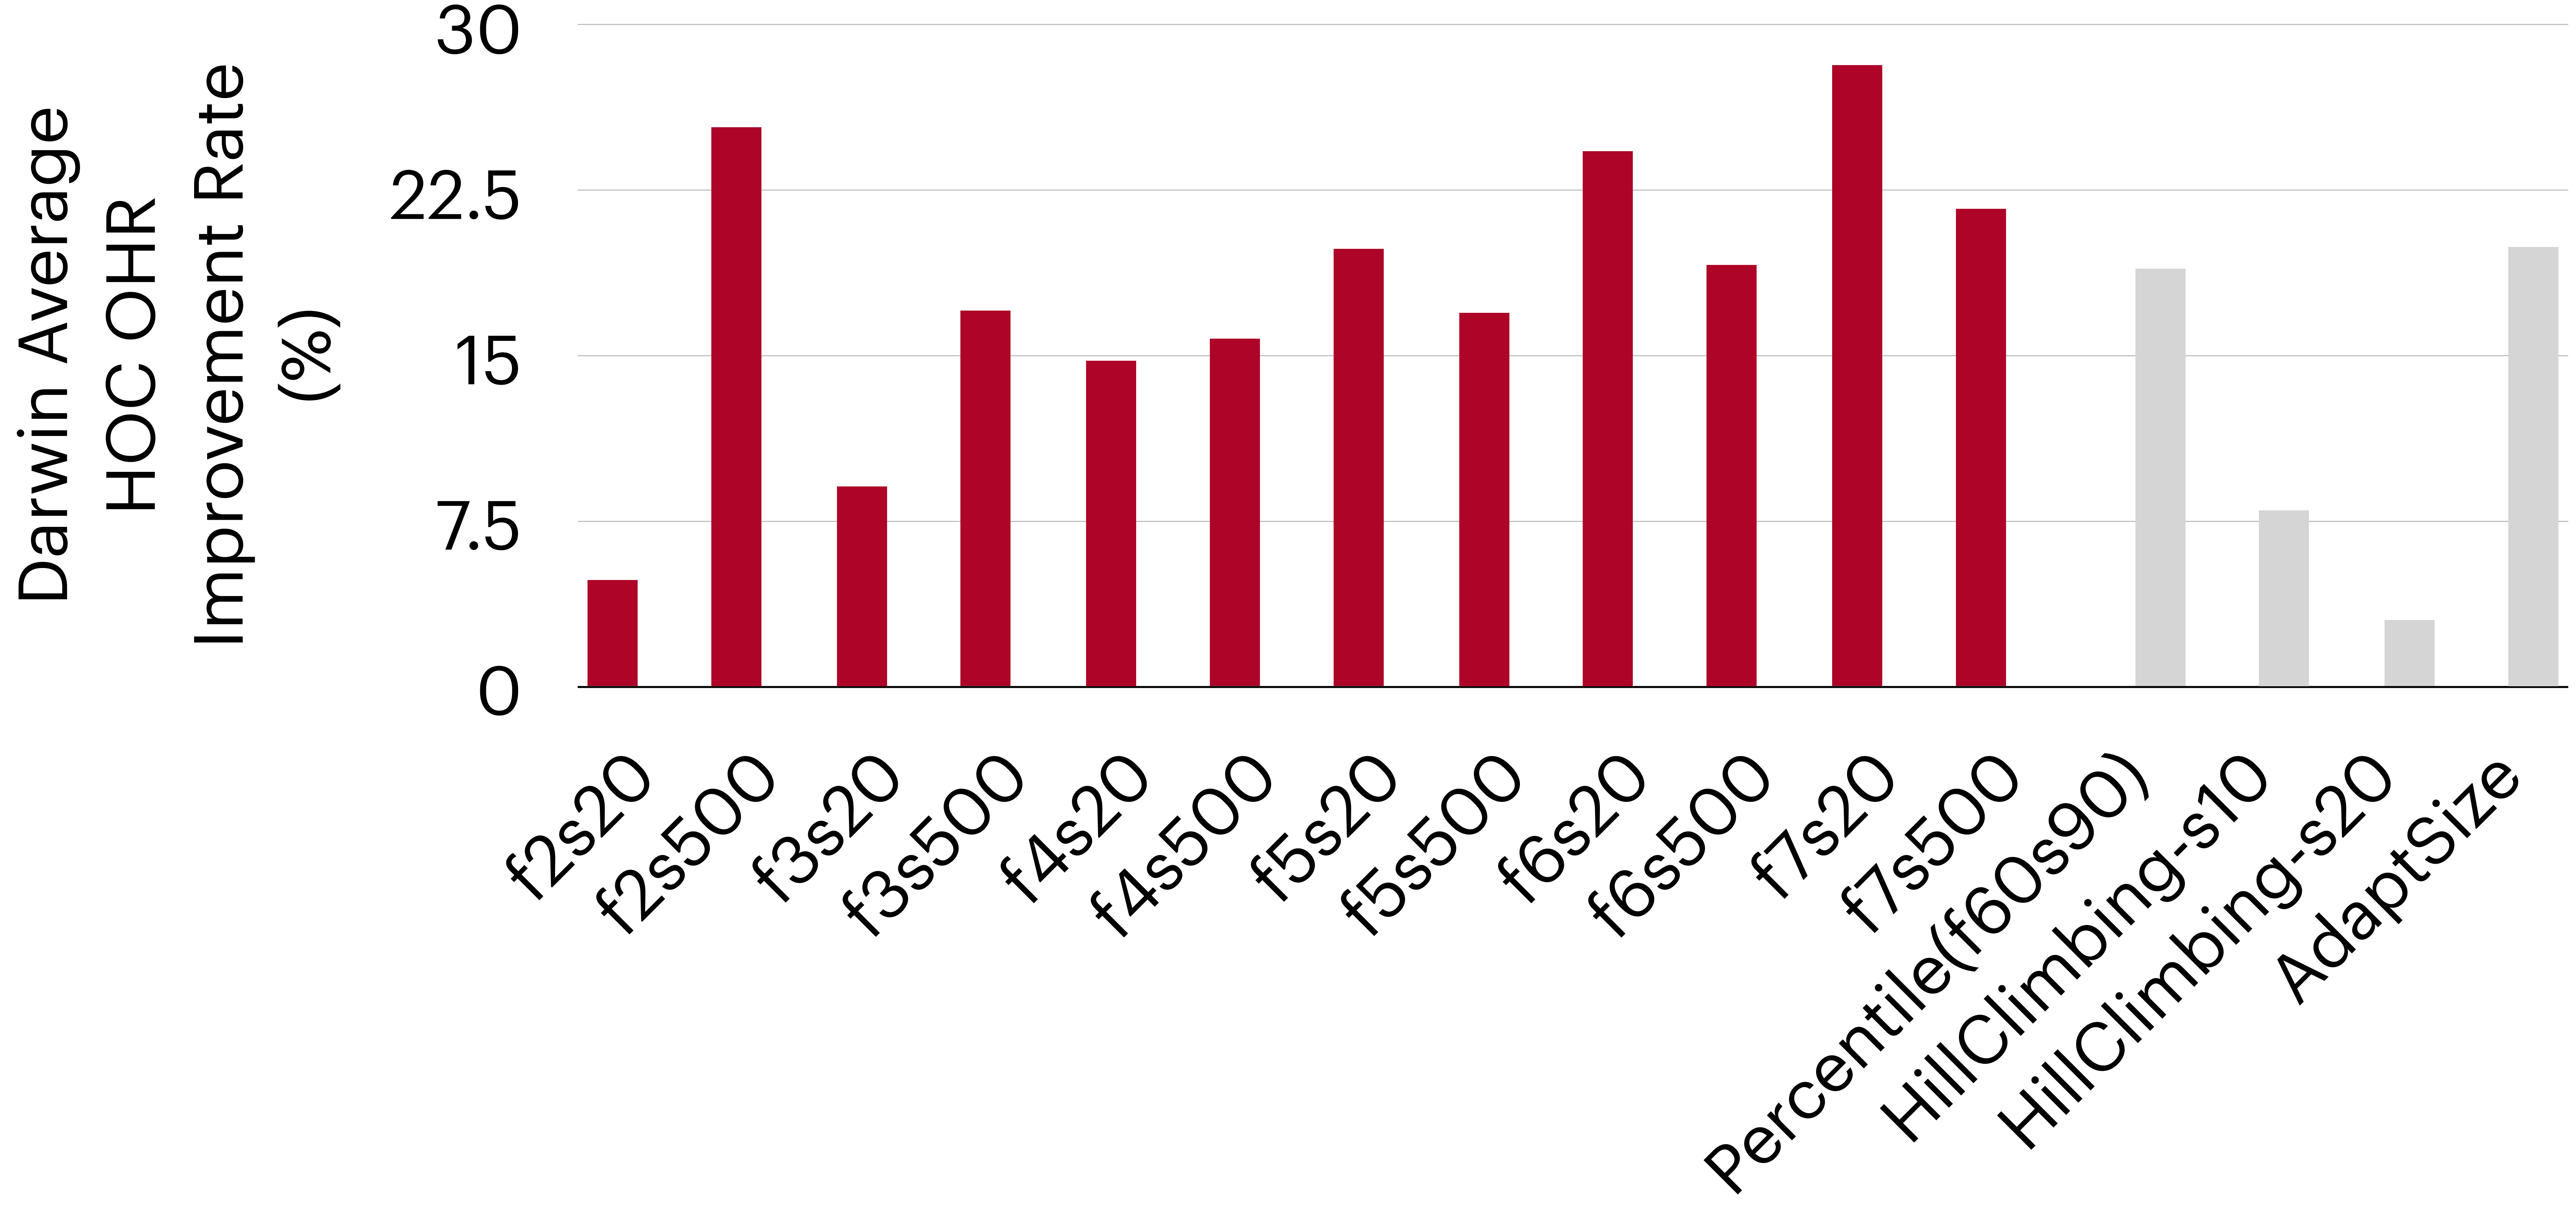

Robustness to Traffic Changes

Darwin outperforms static baselines by 4.83%-28.16%

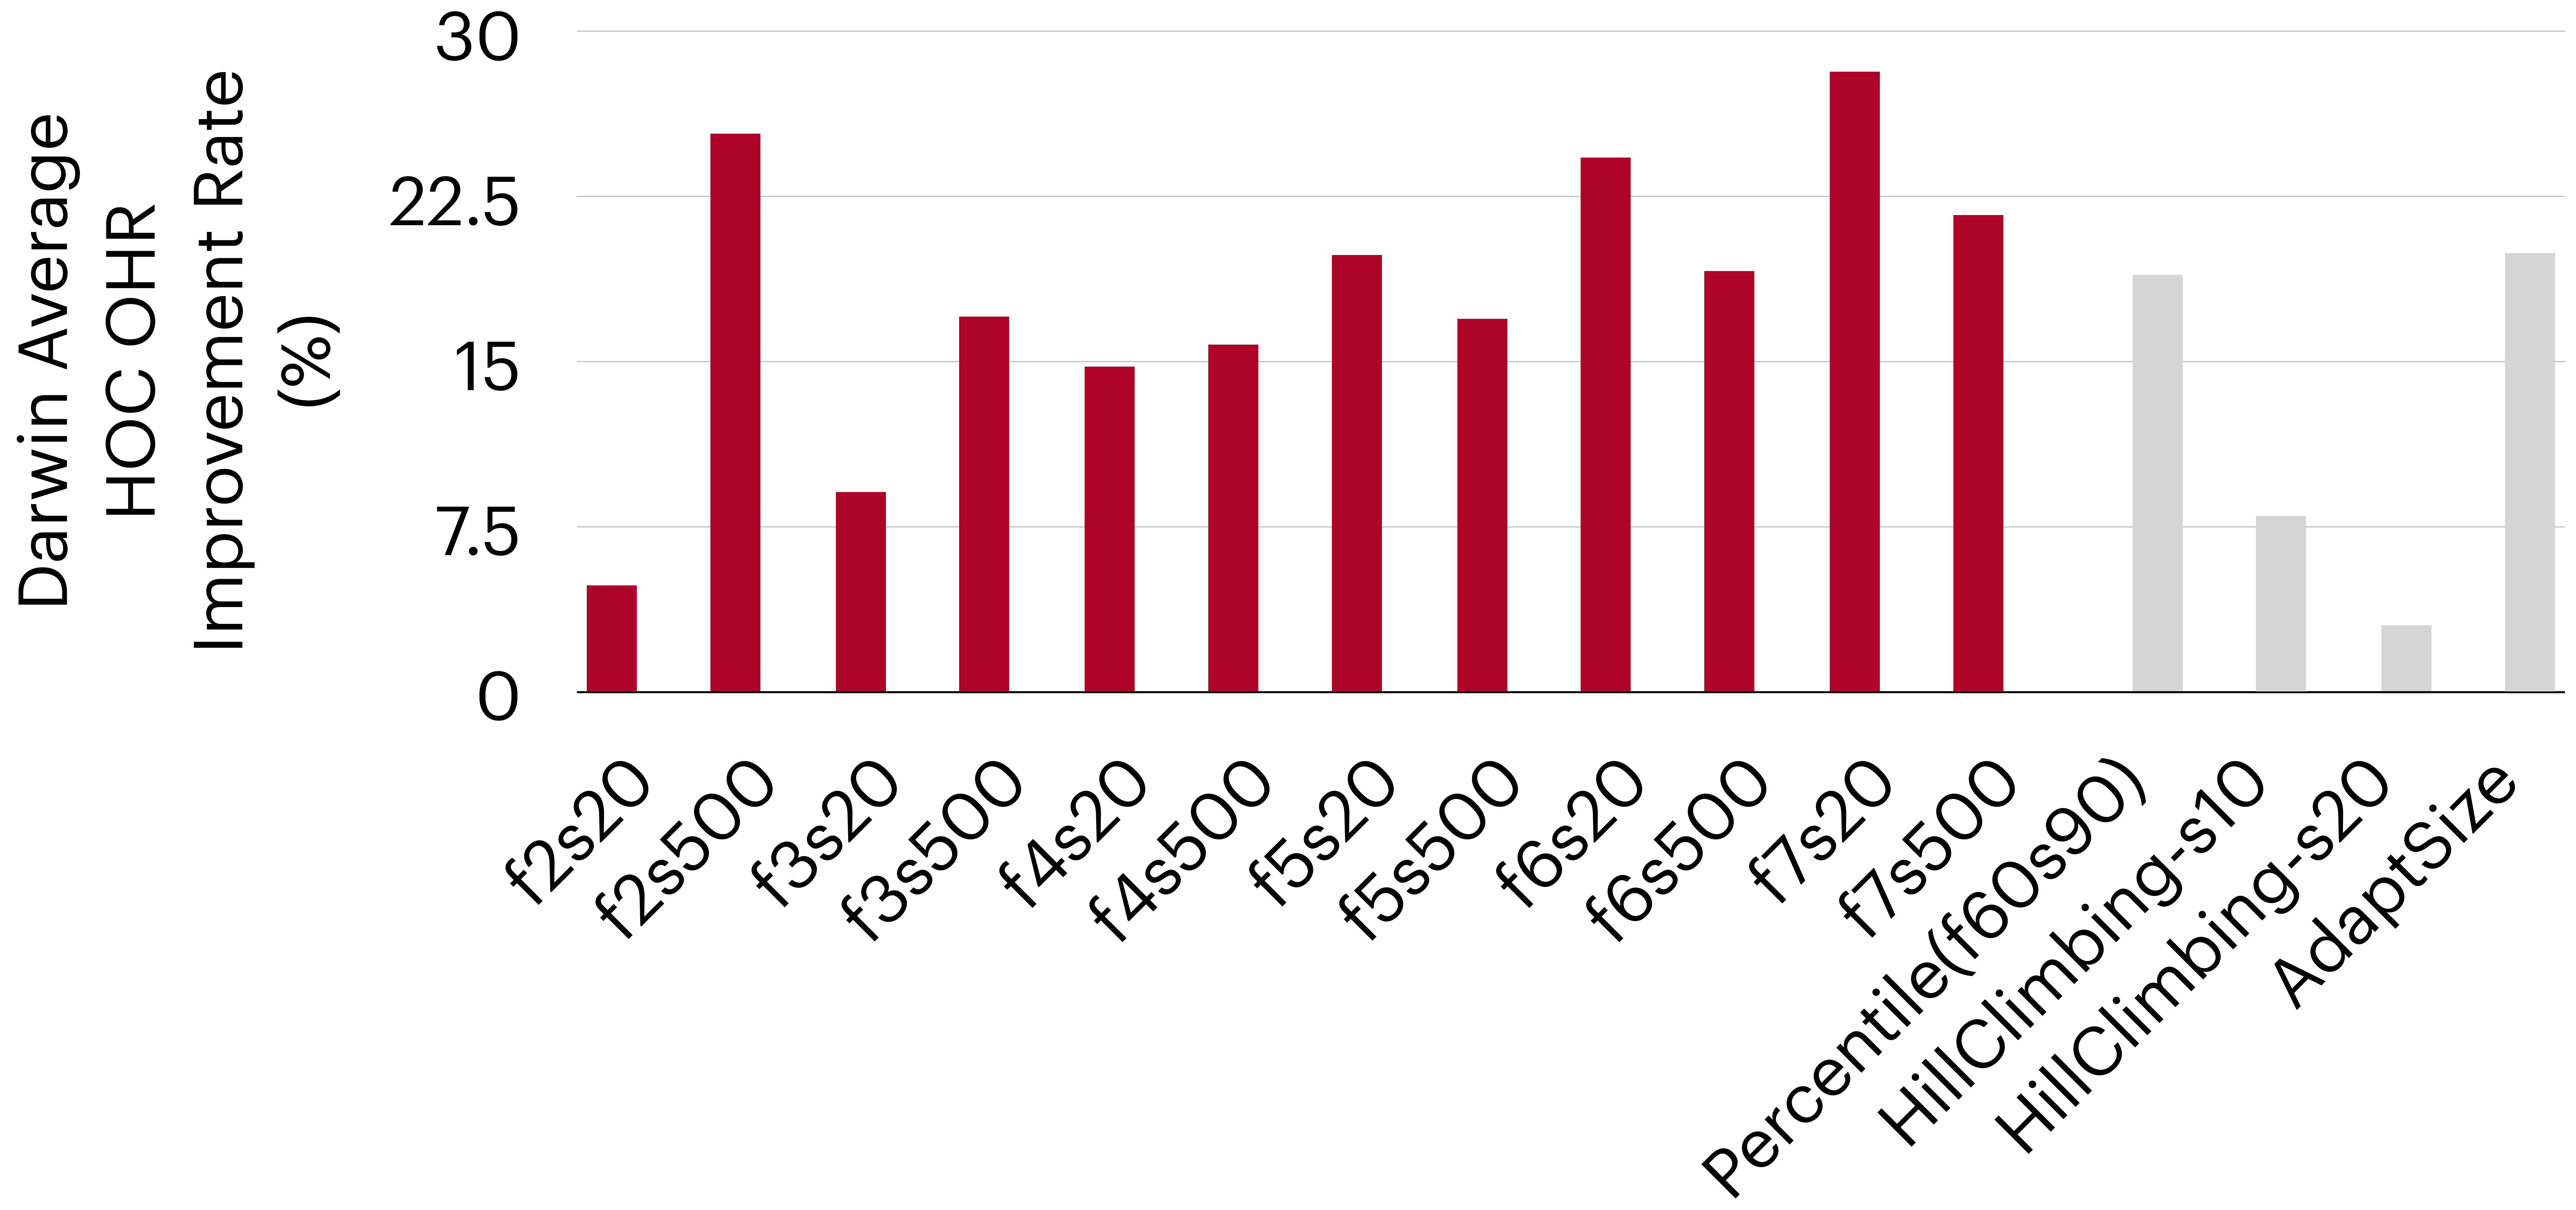

No static policy works well in all traces

Robustness to Traffic Changes

Darwin outperforms adaptive baselines by 3%-19.96%

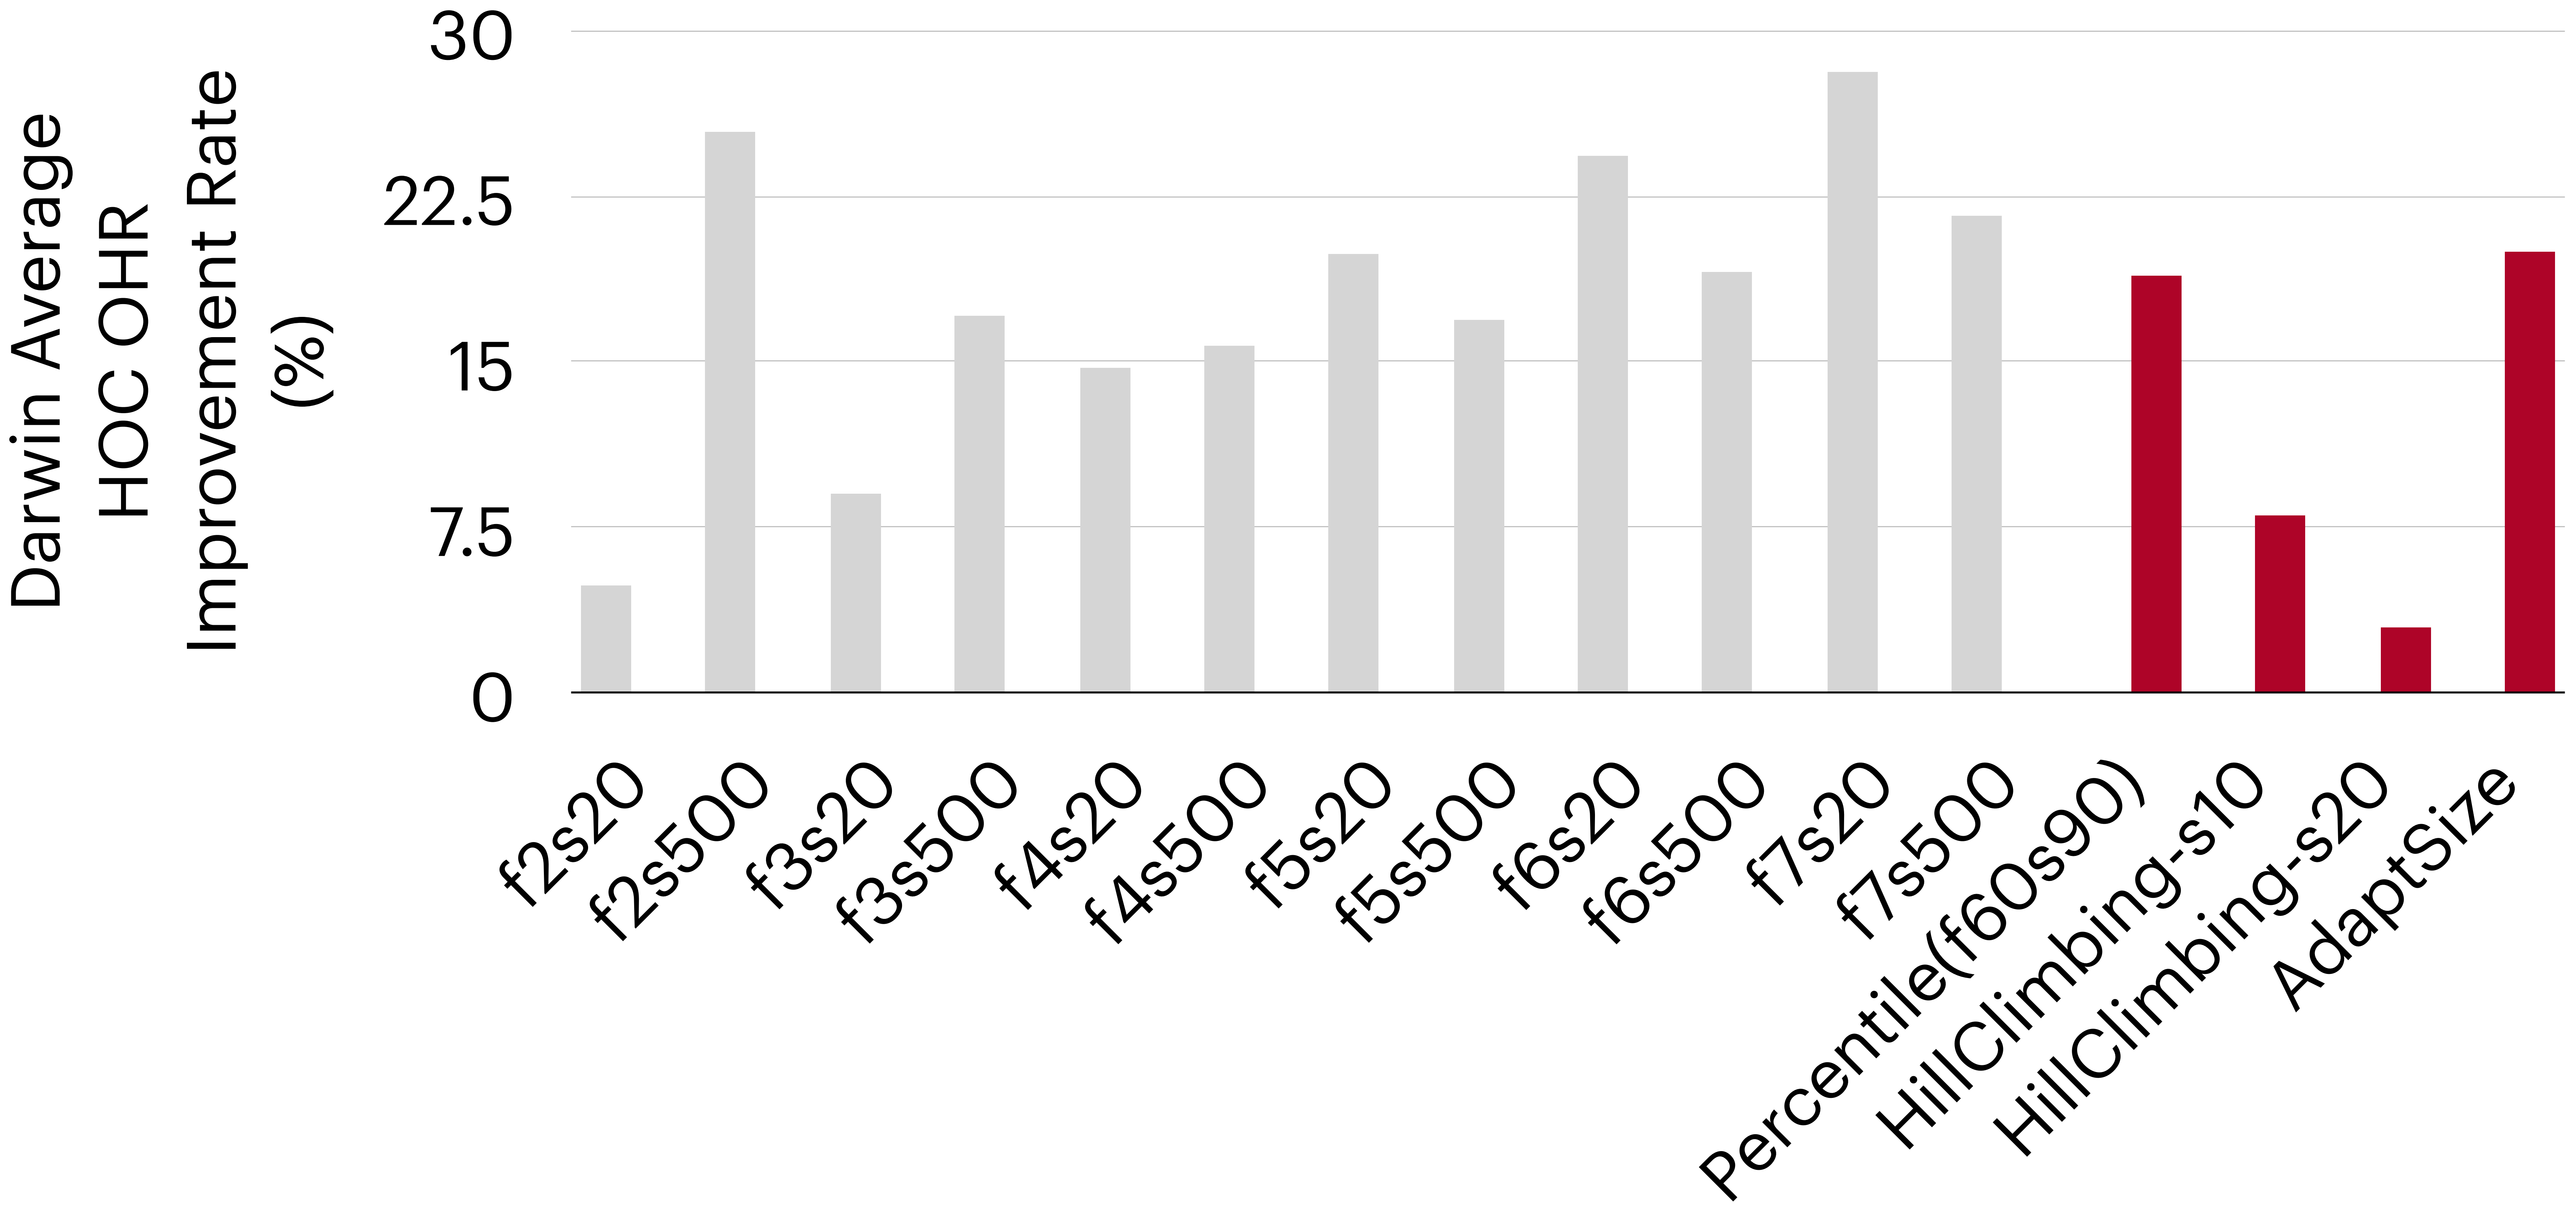

Robustness to Traffic Changes

Darwin outperforms adaptive baselines by 3%-19.96%

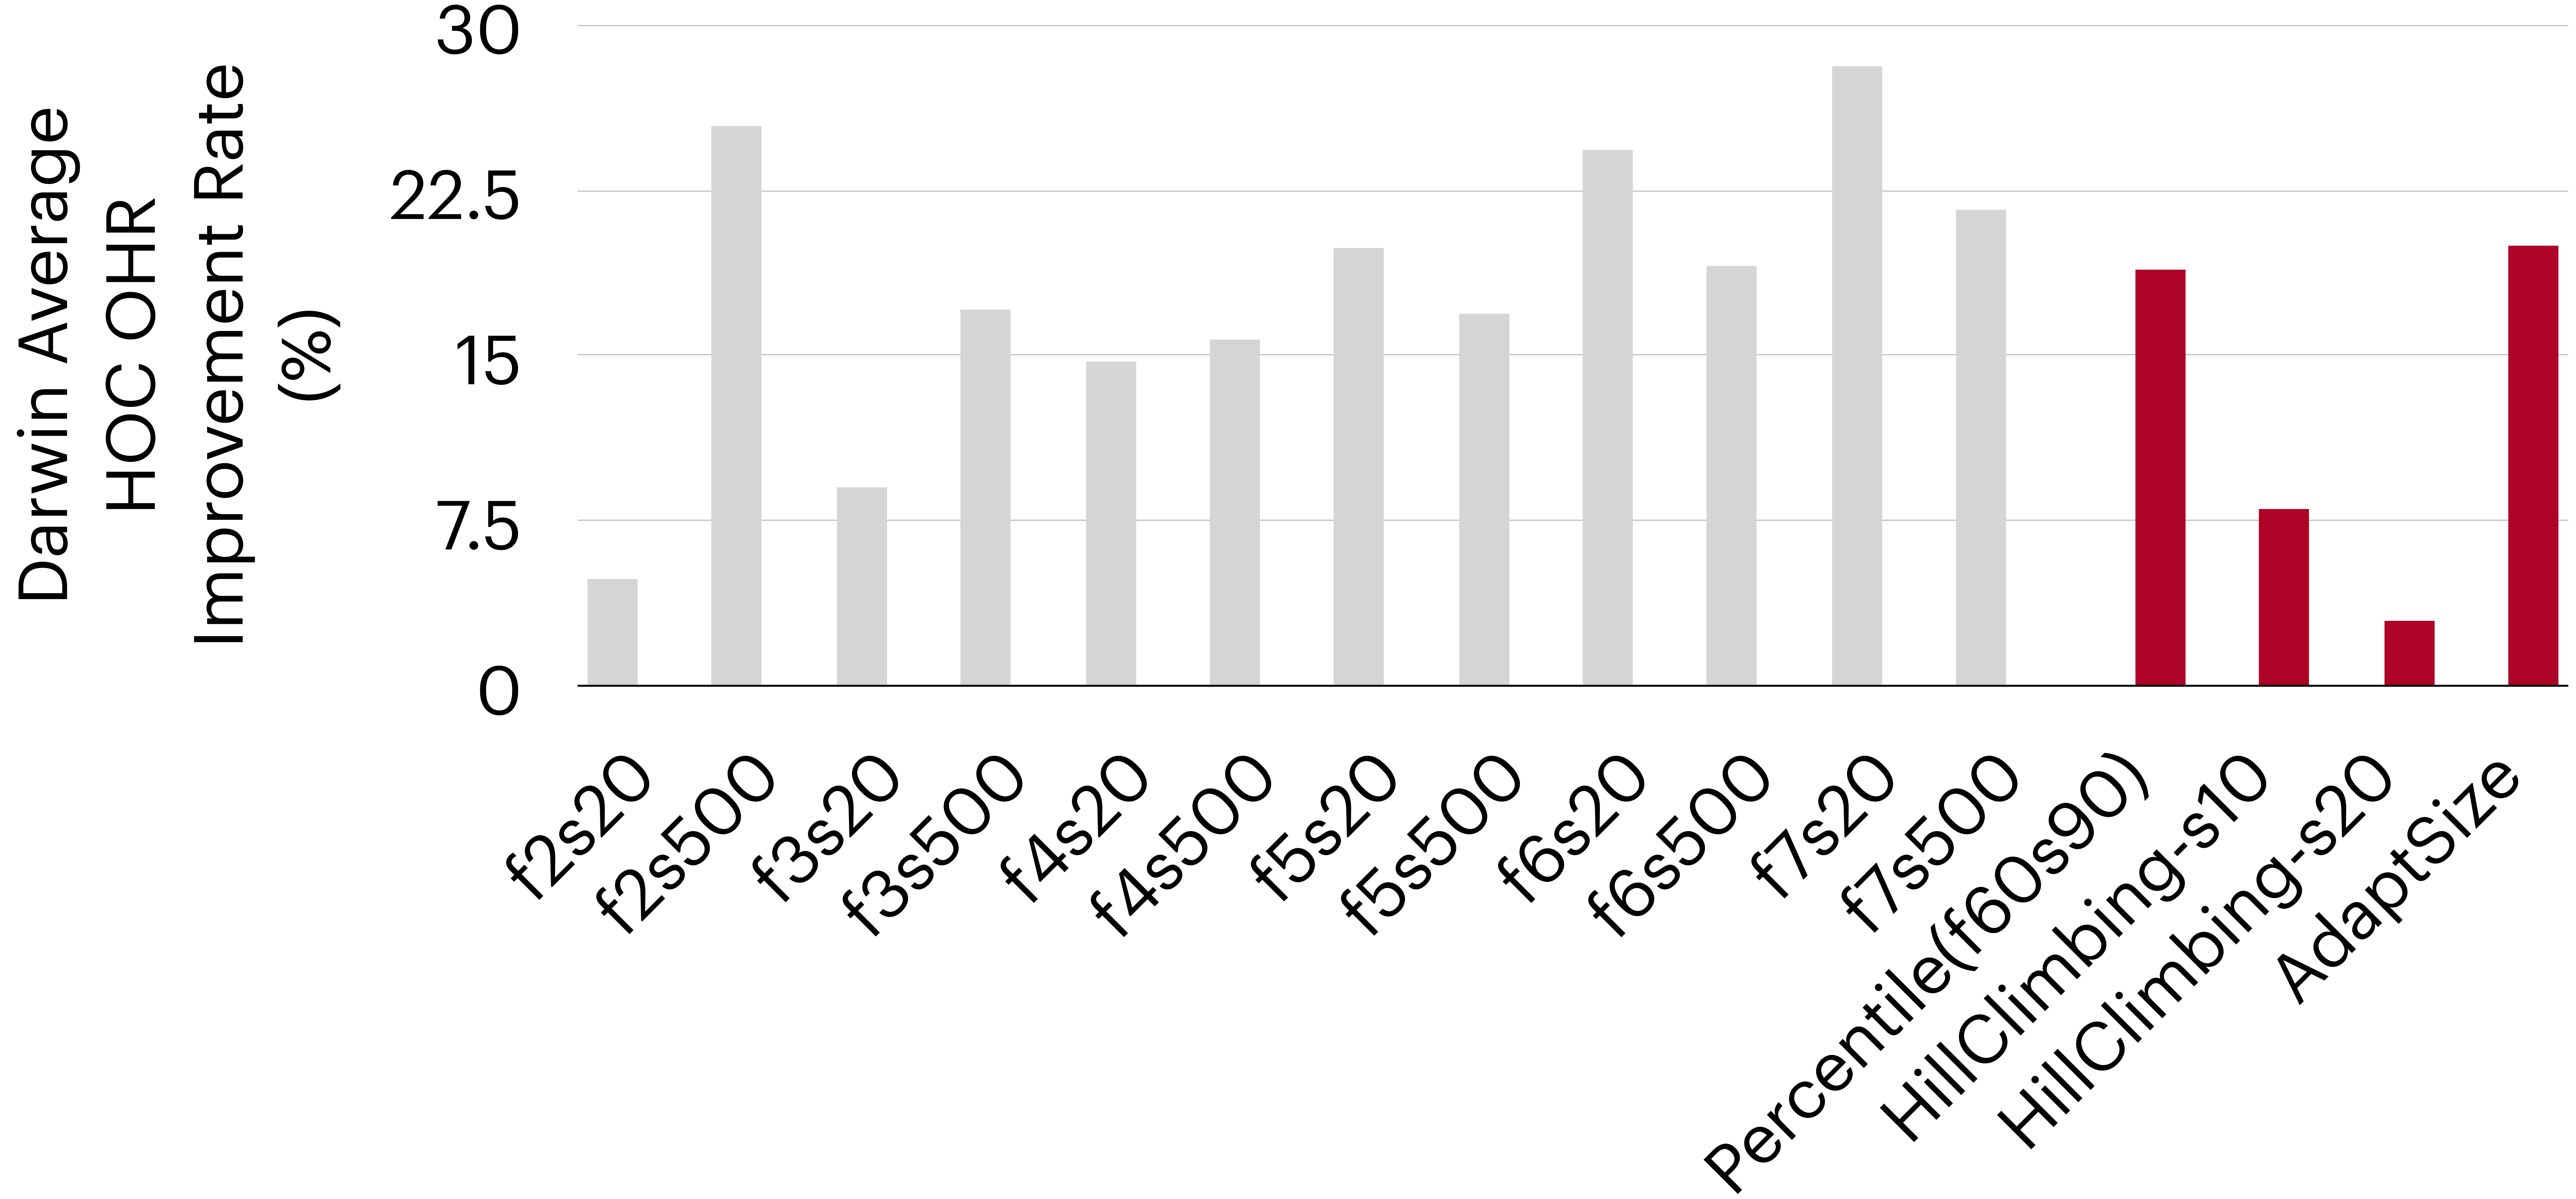

- Percentile: Use f-th, s-th percentile value as thresholds
- HillClimbing: Move toward the direction with better performance
- AdaptSize: Markov chain tuning of probabilistic size threshold

Robustness to Traffic Changes

Darwin outperforms adaptive baselines by 3%-19.96%

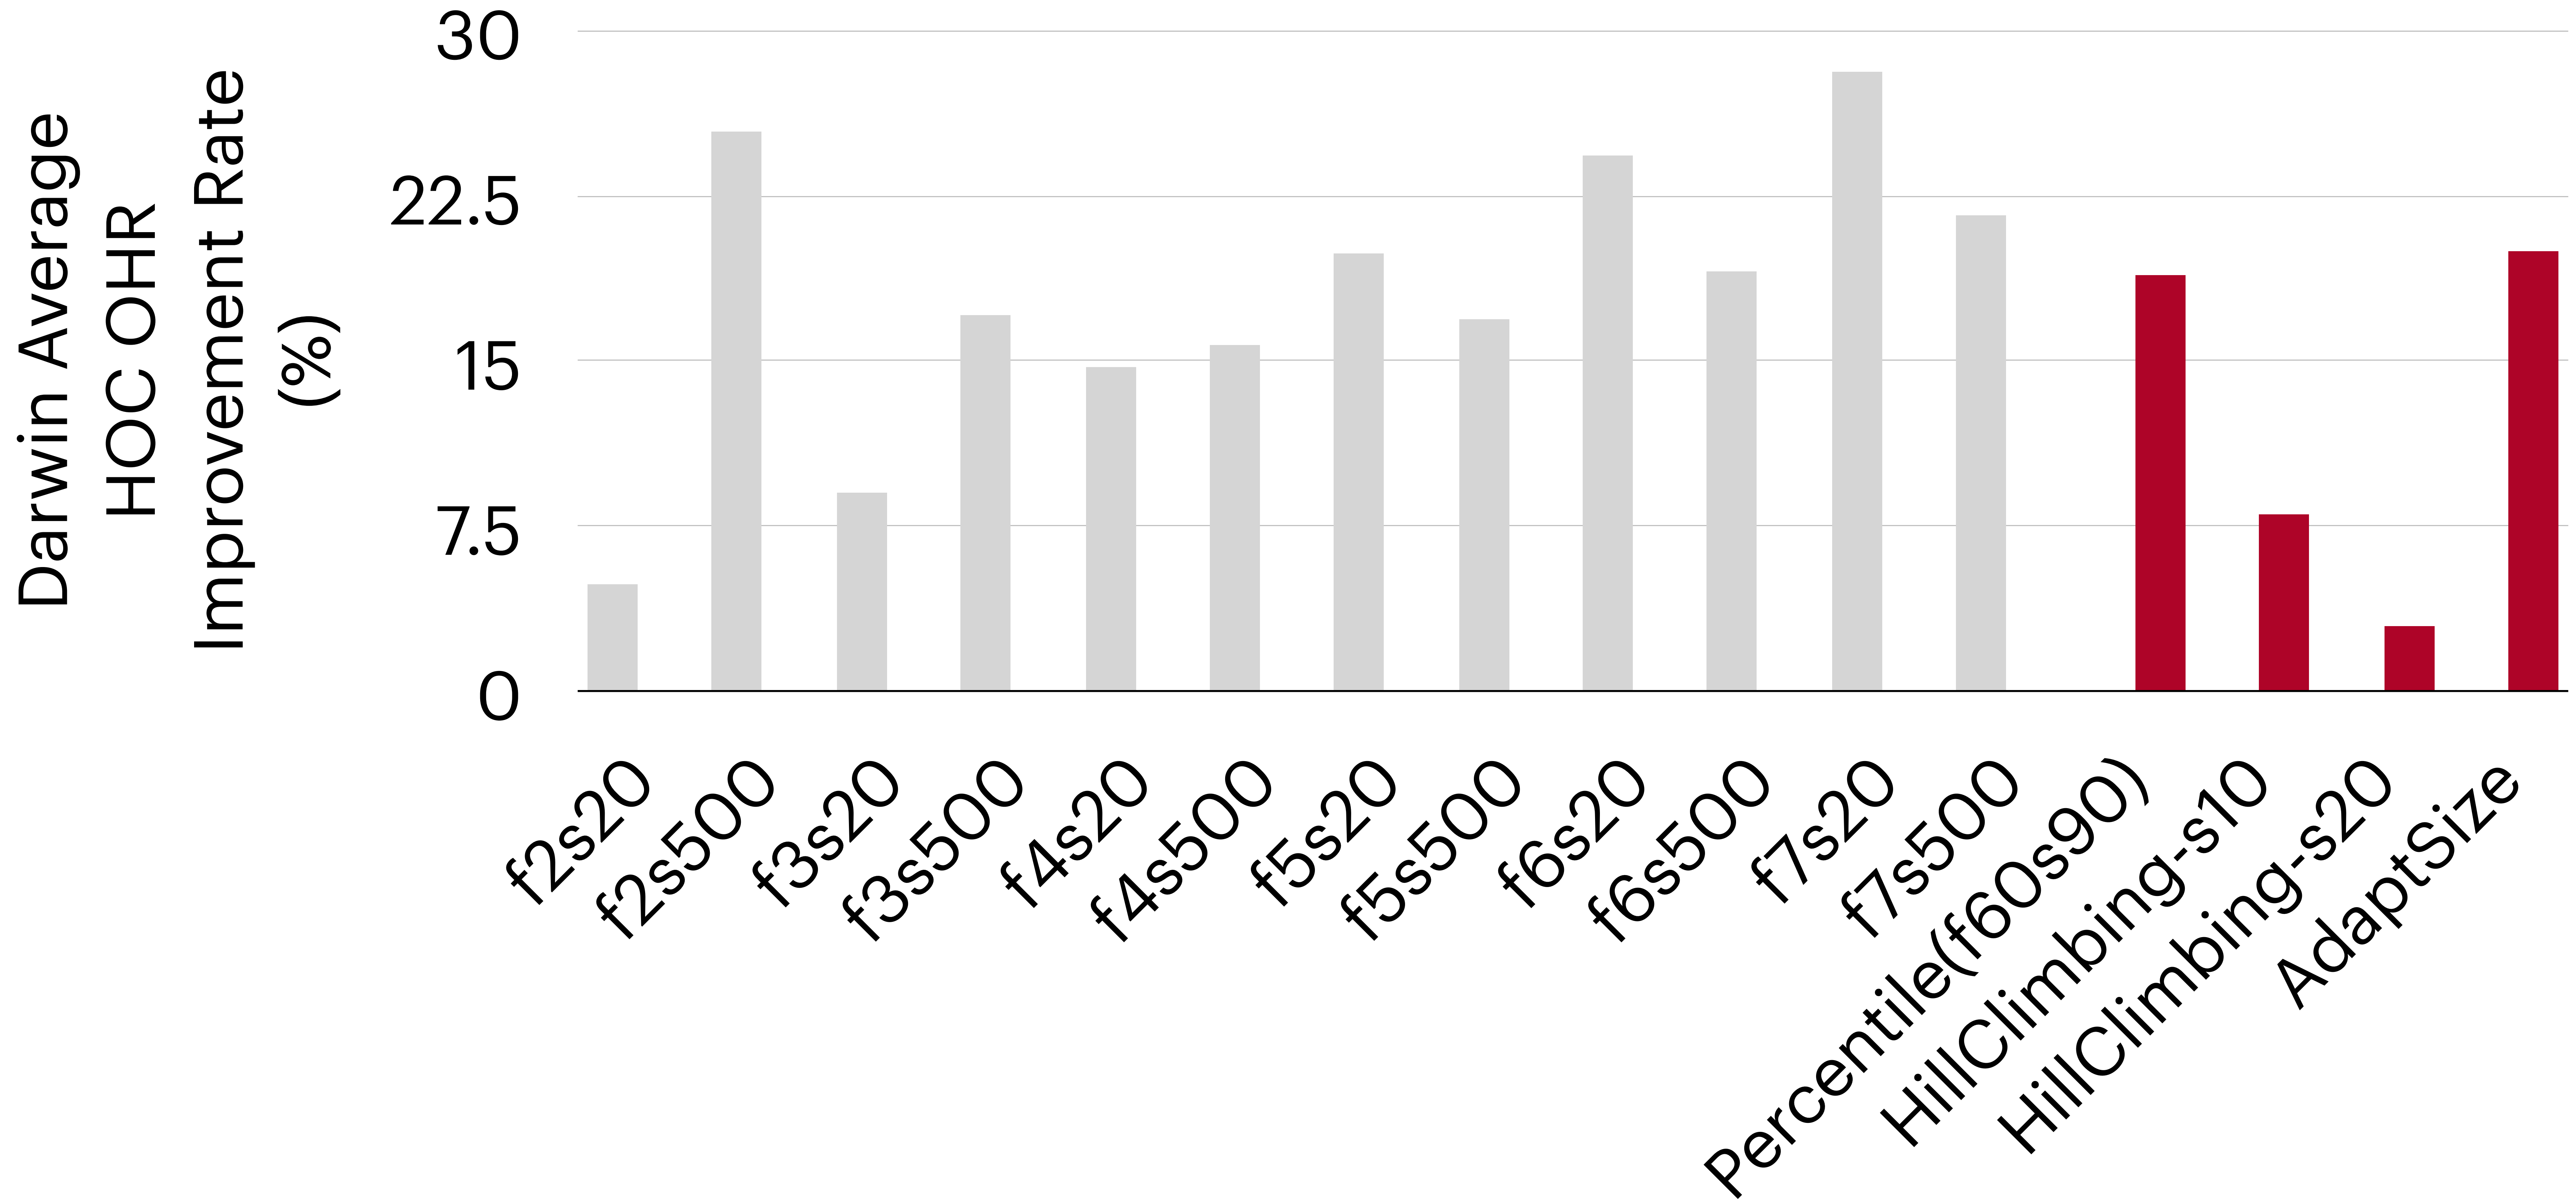

- Percentile: Use f-th, s-th percentile value as thresholds
- HillClimbing: Move toward the direction with better performance
- AdaptSize: Markov chain tuning of probabilistic size threshold

- Tuning of multiple parameters
- Access to finer granularity of policies

More Evaluation Results

- **Cross-policy prediction models are robust.**
  - >90% of the cross-policy predictors reach > 80% order prediction accuracy.
- **Darwin can be used to improve other metrics.**
  - e.g., improves $\left(\text{OHR} - \frac{\text{DiskWrite}}{\text{\#Requests}}\right)$ by 7.47%-96.67%
- **Darwin doesn't impose additional latency overhead and minimally impacts CPU and memory utilization.**

Conclusion

- Static HOC admission policies fall short when the workload shifts
- Darwin can learn the best CDN HOC admission policy flexibly with
  - Cross-policy prediction models
  - Feature clustering and policy association
  - Track and Stop with Side Info algorithm
- Darwin outperforms the state-of-the-art admission policies with respect to multiple metrics adding minimal overhead
- Darwin is a generally applicable policy selection approach.

Thank You!