Fault Tolerant Service Function Chaining

M. GHAZNAVI, E. JALALPOUR, B. WONG, R. BOUTABA, A. MASHTIZADEH

UNIVERSITY OF WATERLOO

Middleboxes and Service Function Chains

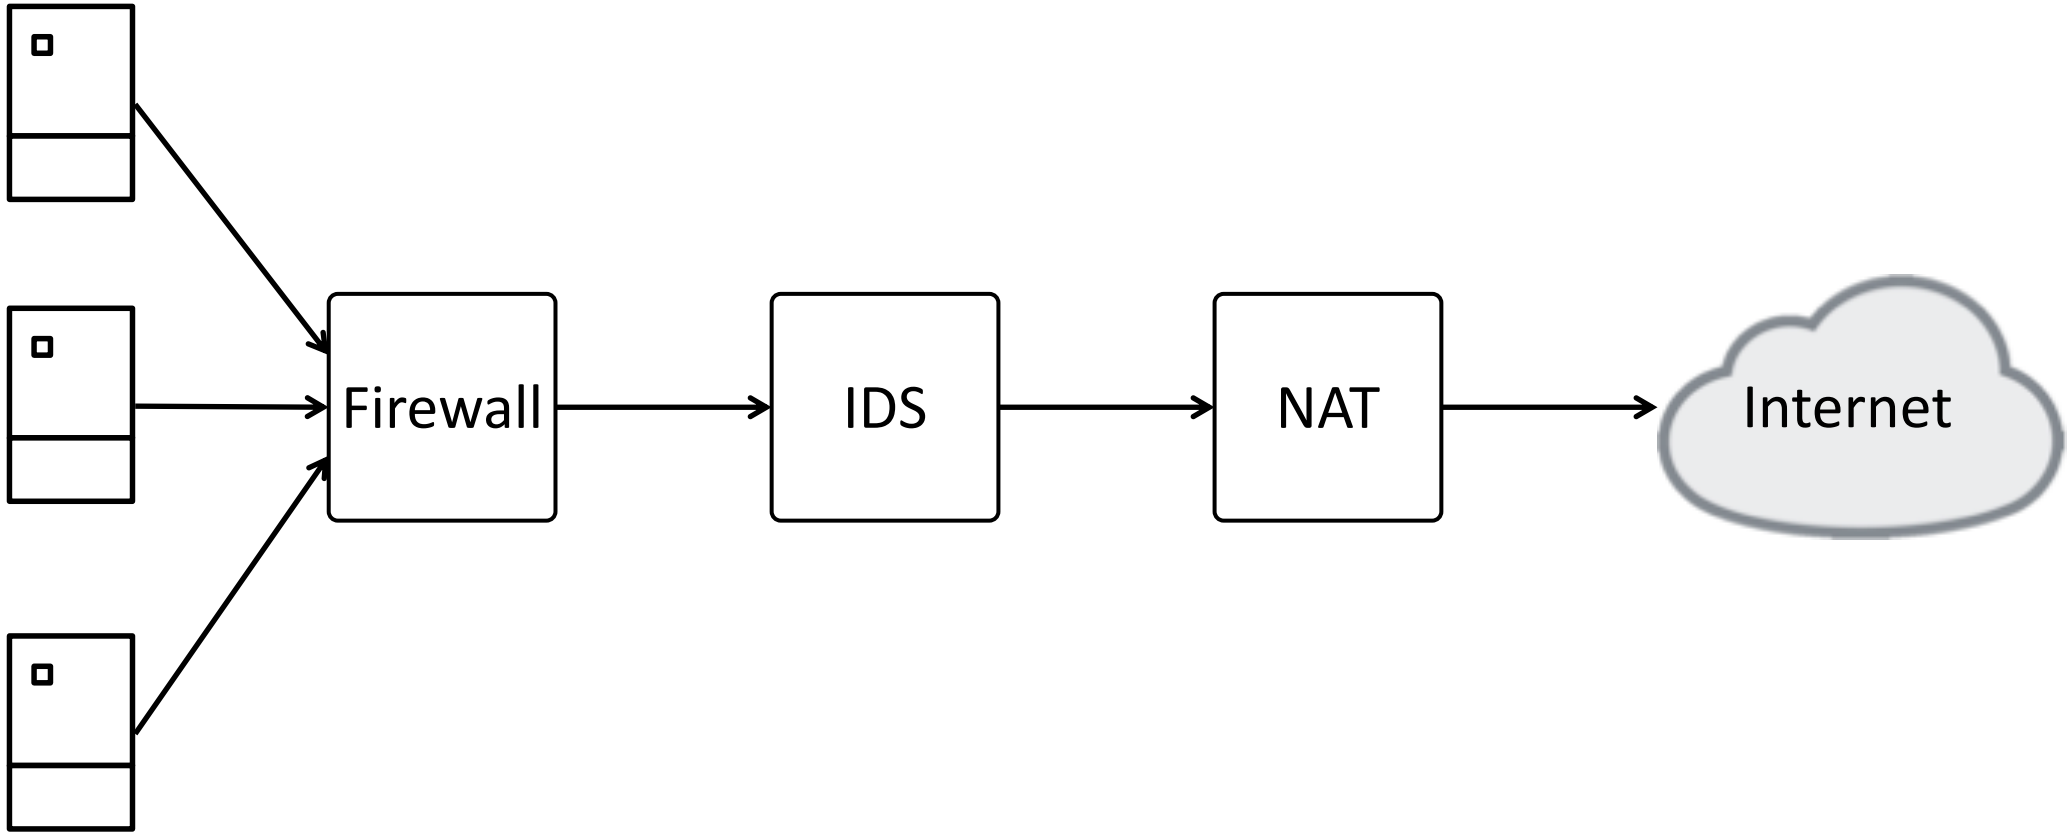

Middlebox Failures

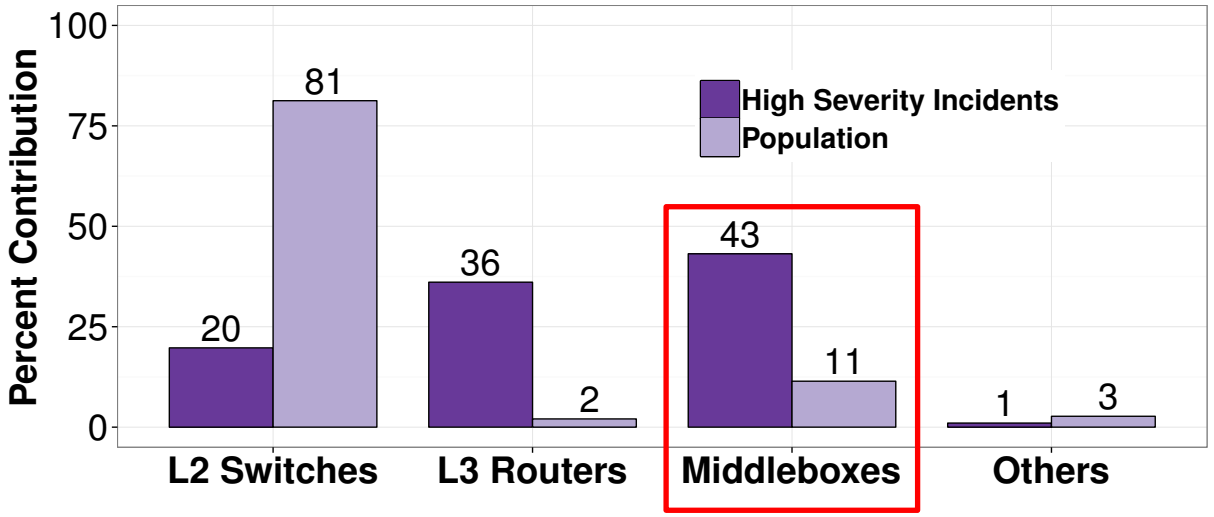

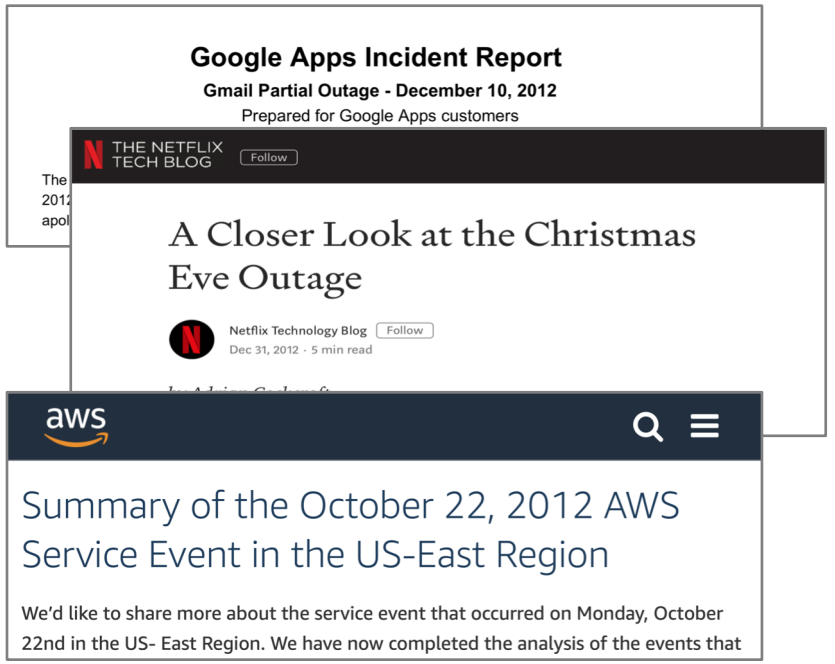

Demystifying the dark side of the middle:
A field study of middlebox failures in datacenters
IMC 2013

Middlebox Fault Tolerance

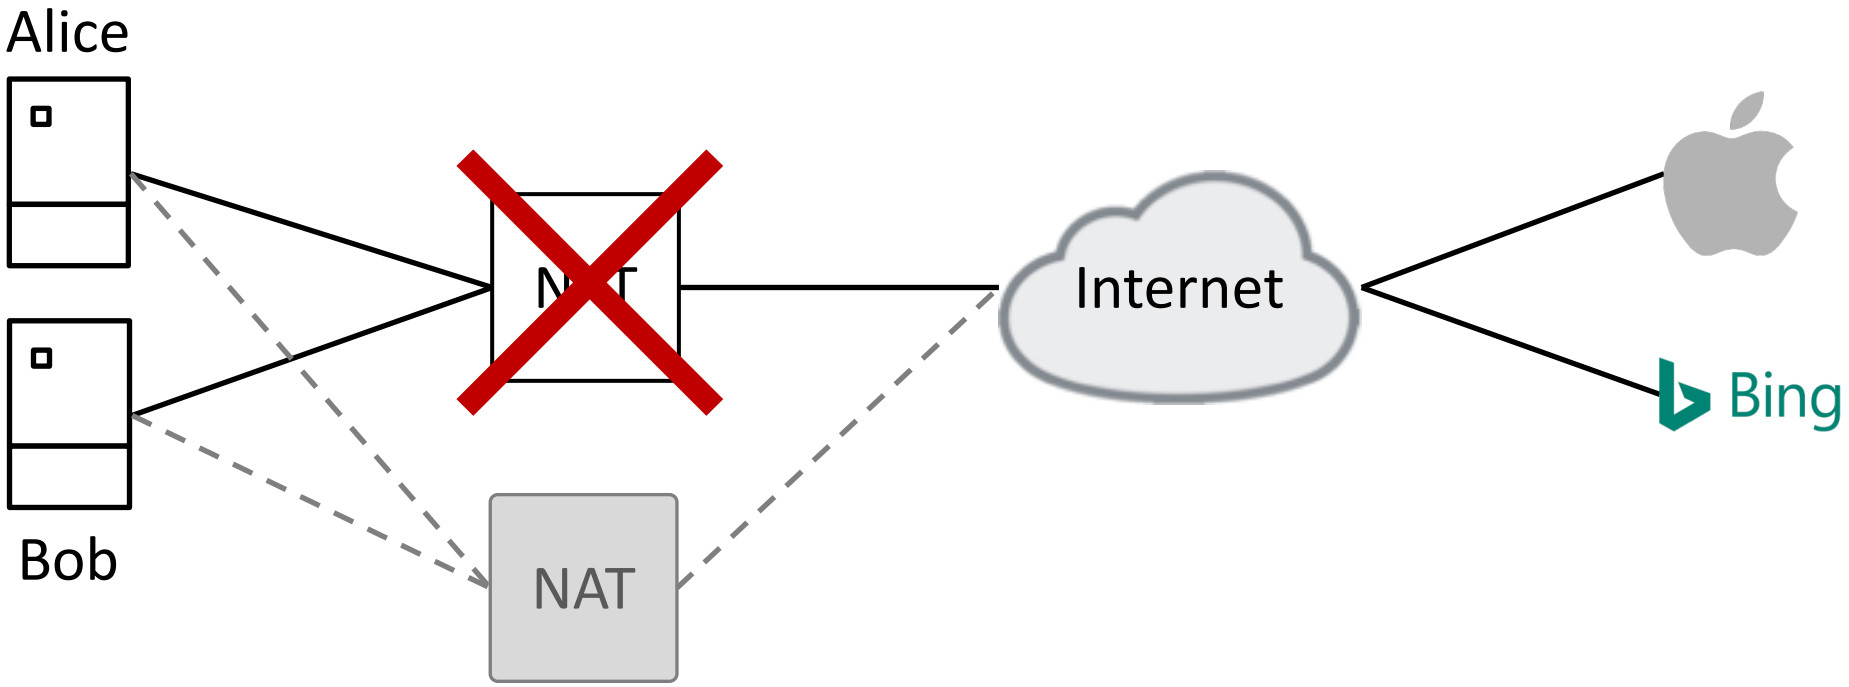

Consistent State Replication

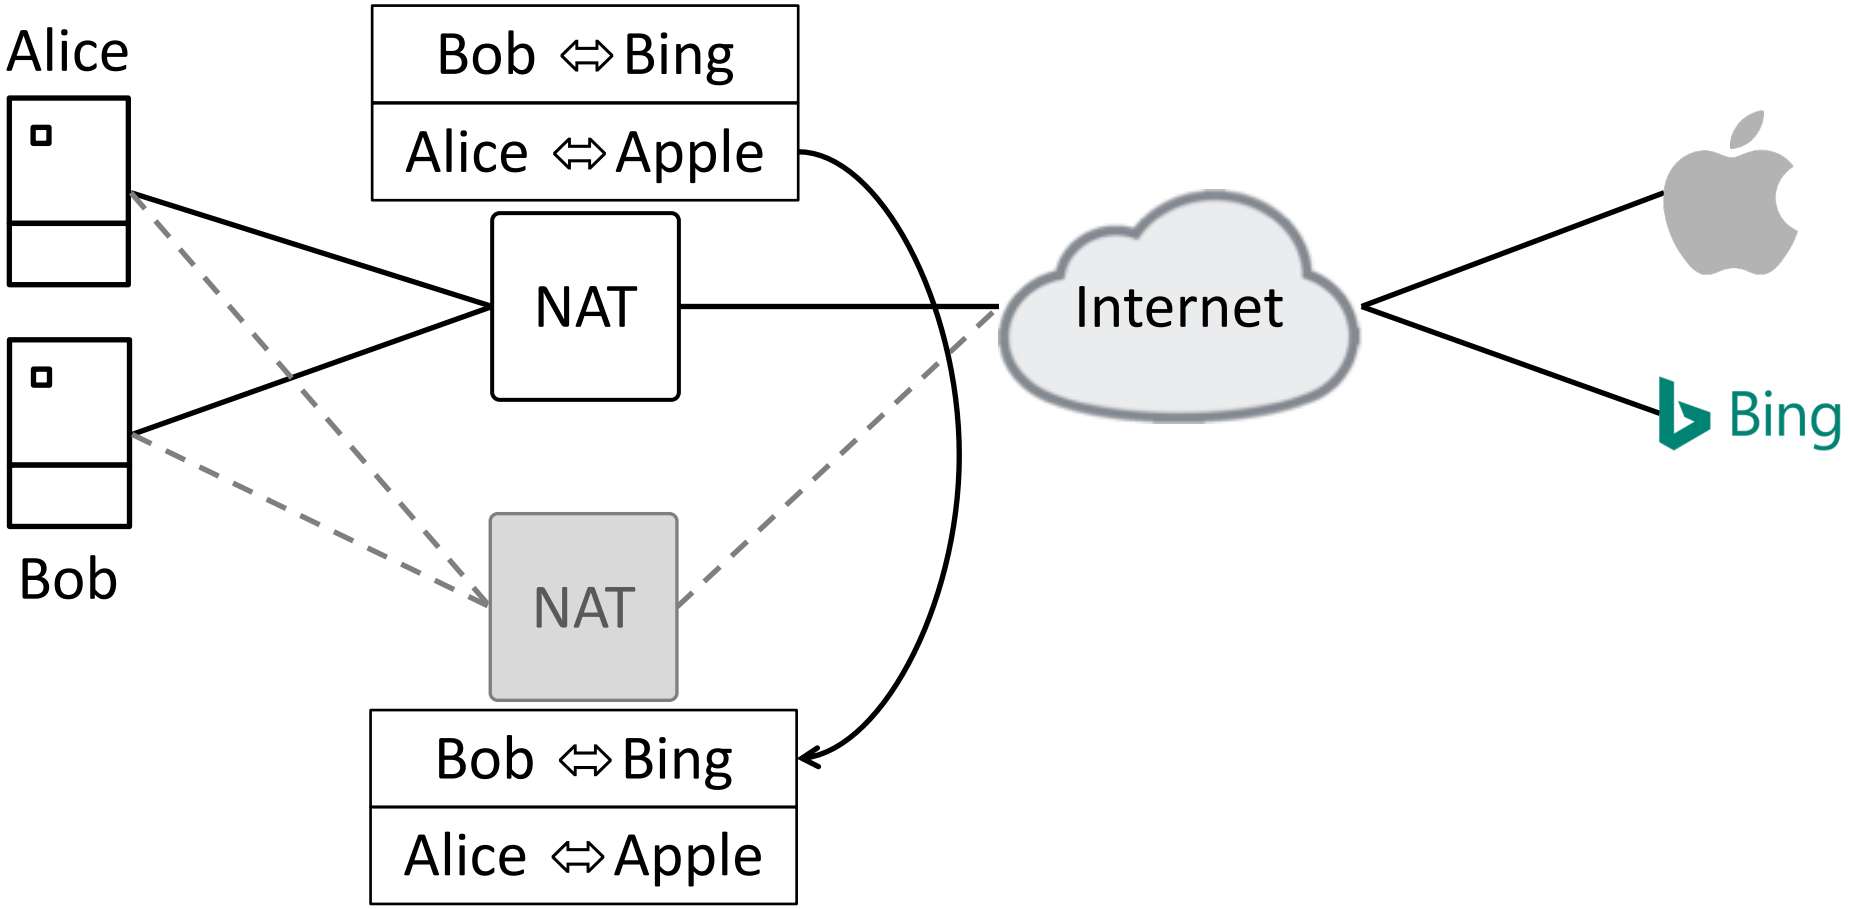

Consistent State Replication

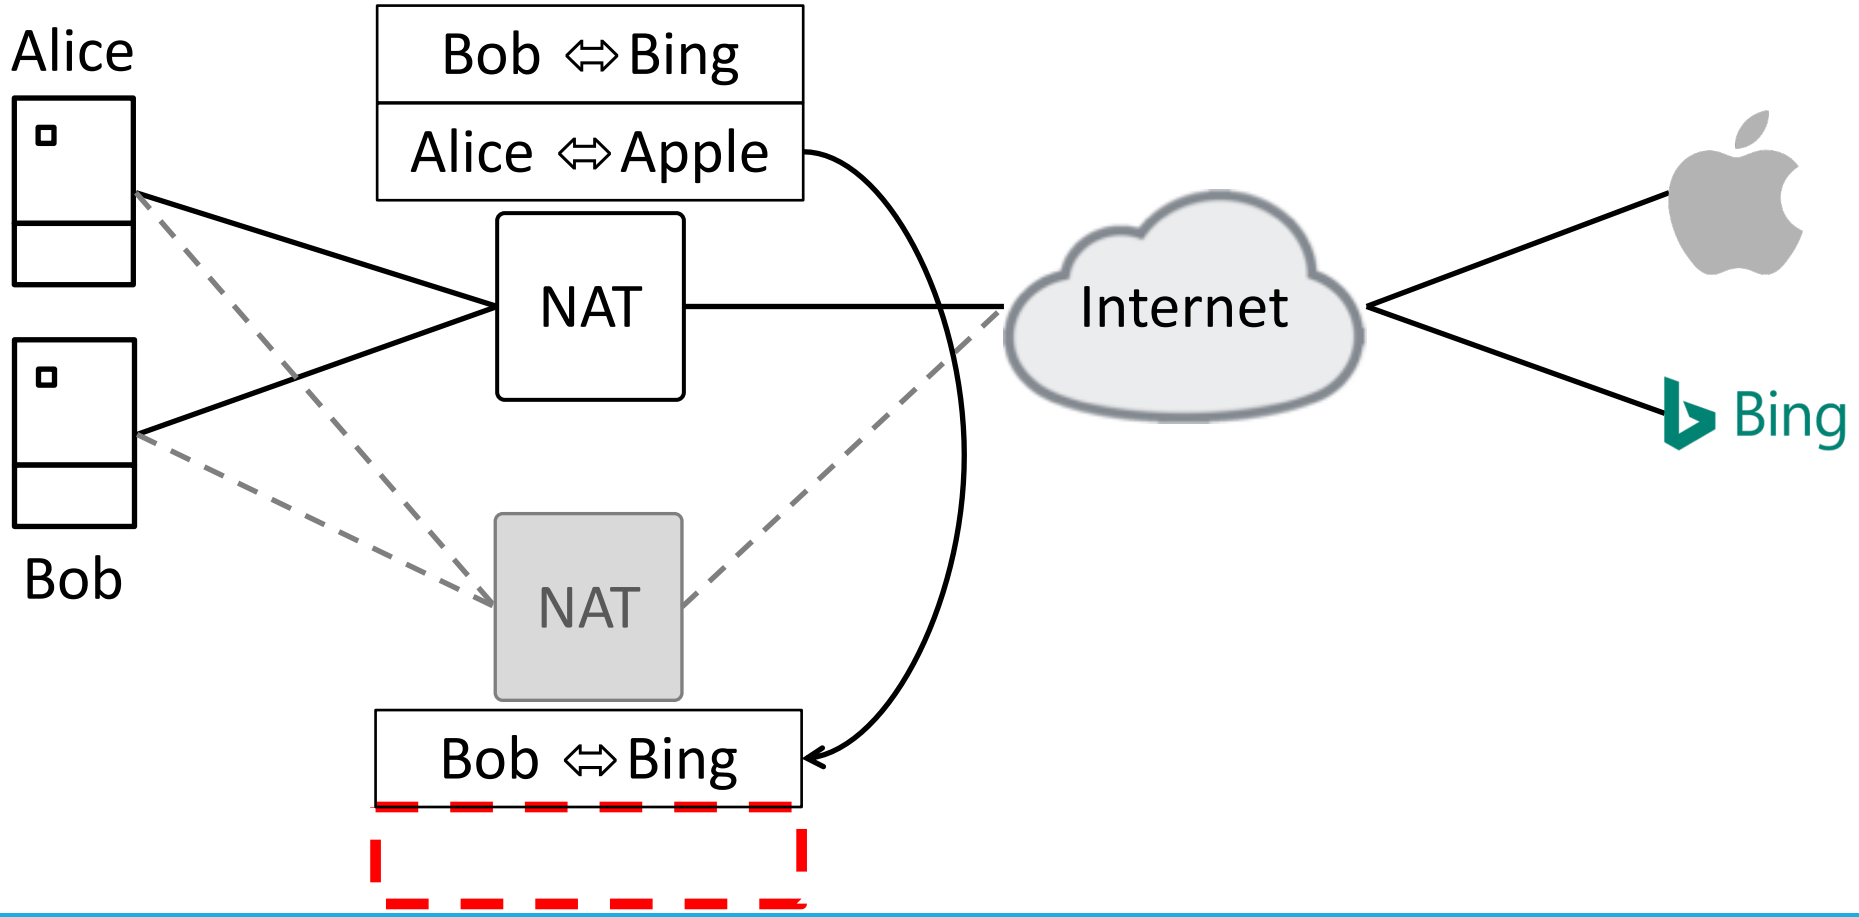

Previous Approaches

EXTERNALIZED STATE

StatelessNF, NSDI 2017

CHC, NSDI 2019

SNAPSHOT BASED

Pico Replication, SoCC 2013

FTMB, SIGCOMM 2015

REINFORCE, CoNEXT 2018

Externalized State Approach

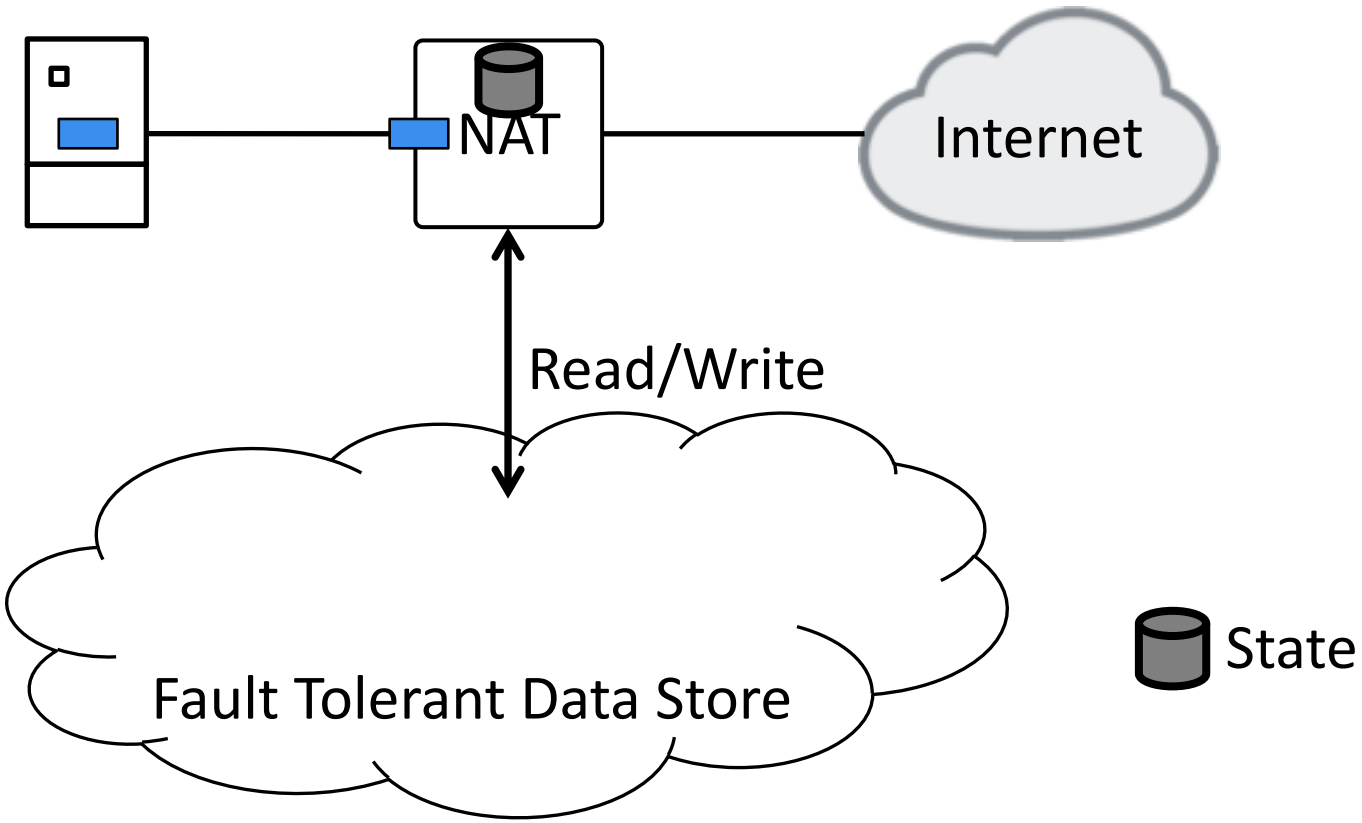

Snapshot Based Approaches

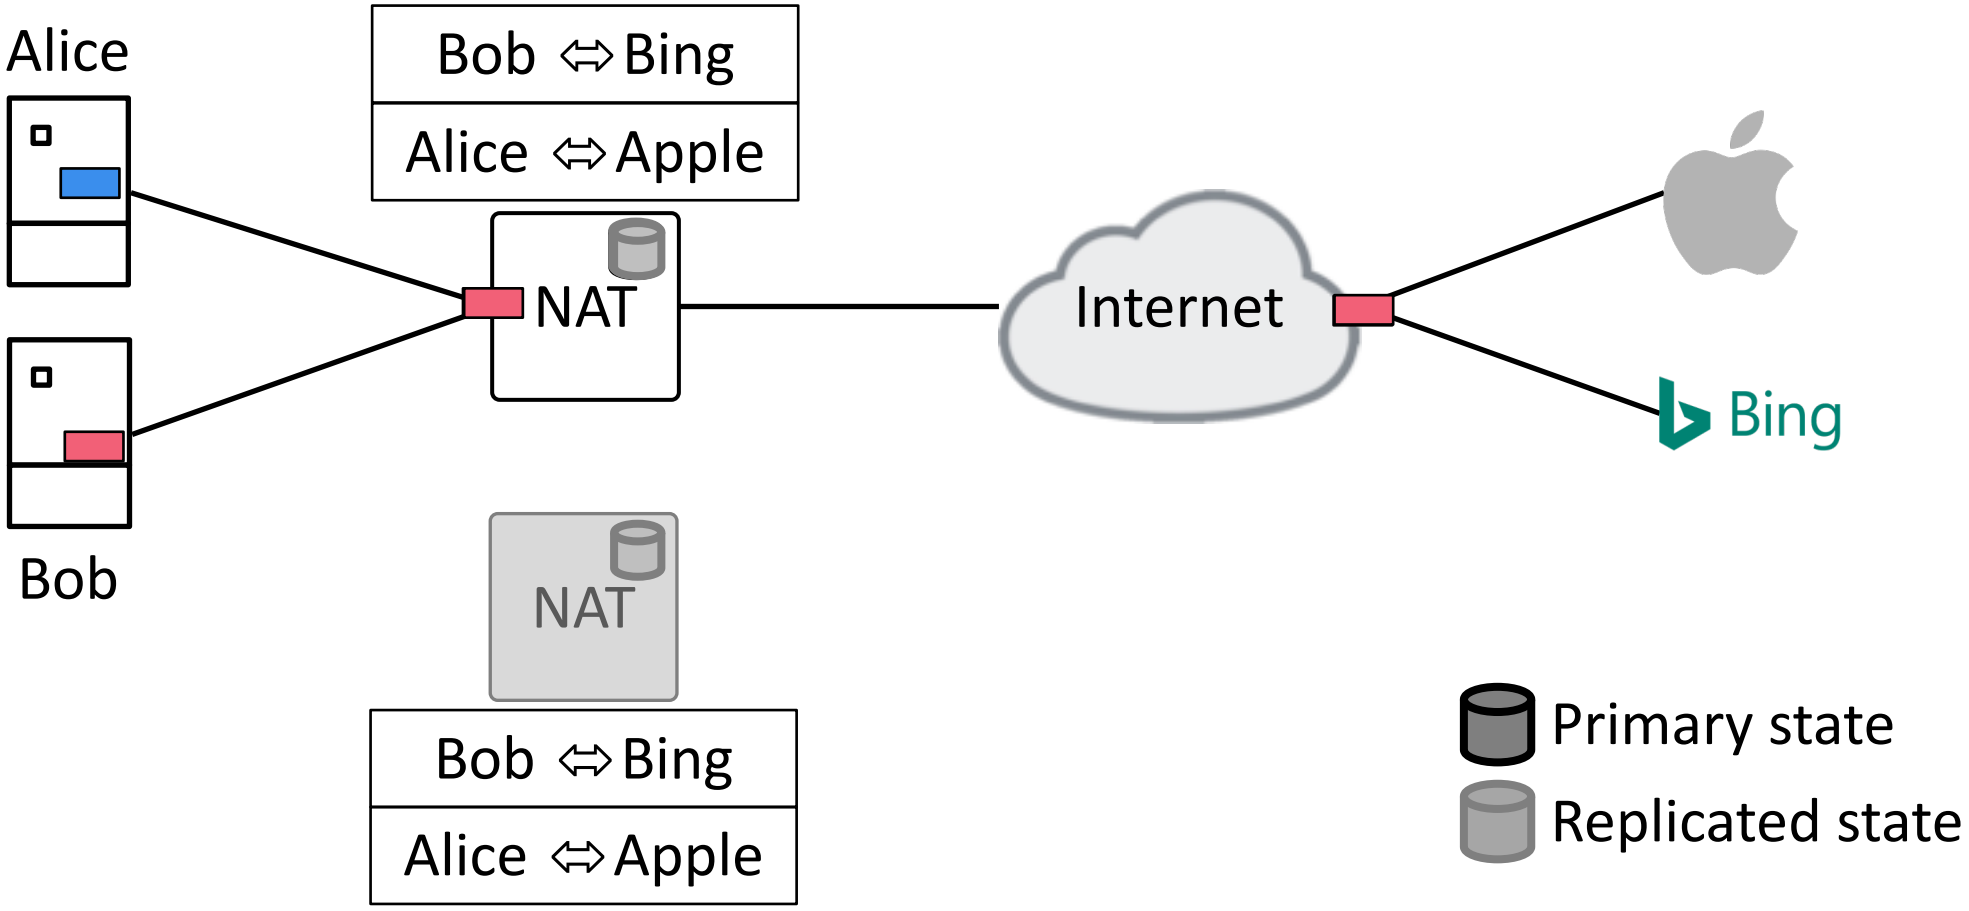

Snapshot Based Approaches for a Chain

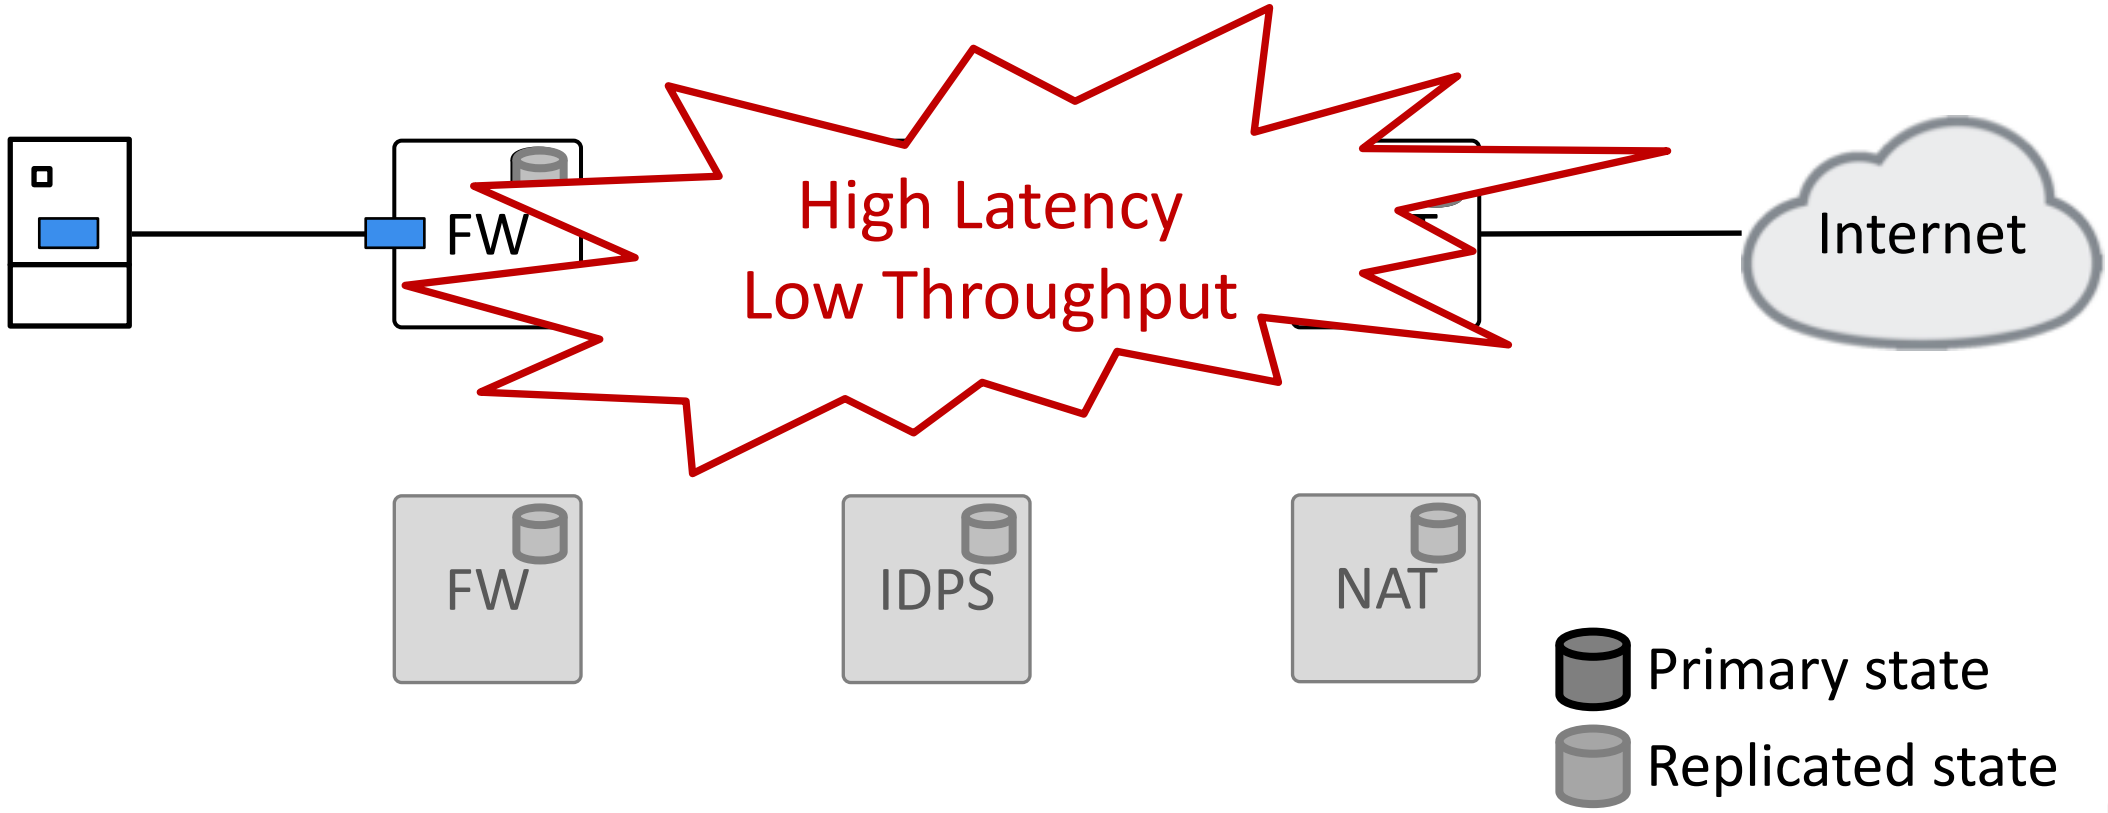

Our Approach

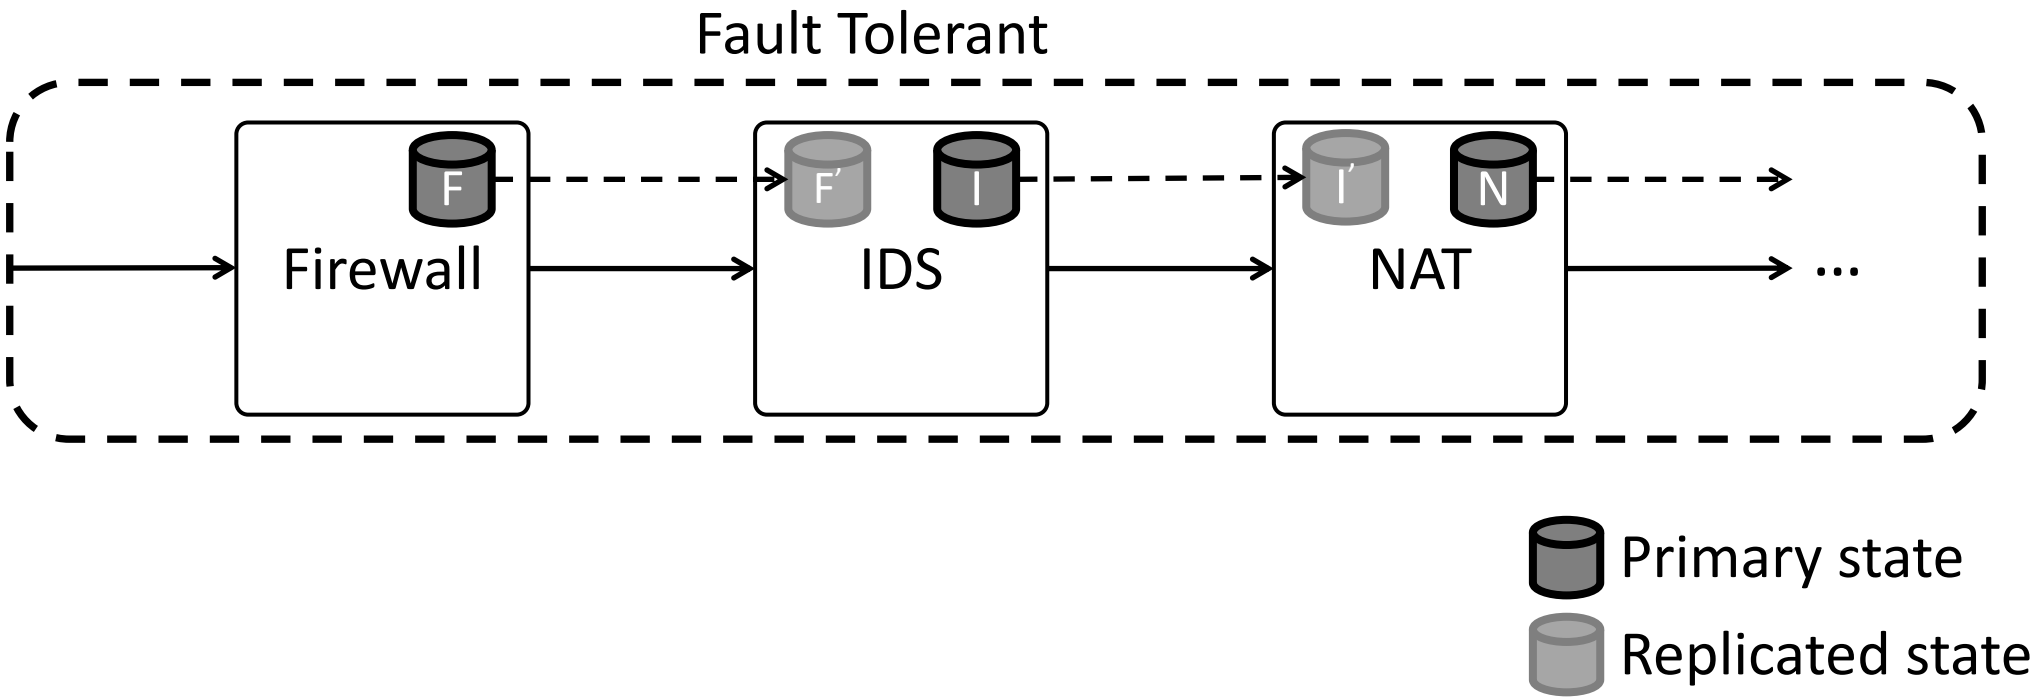

Goals

Consistent state replication to tolerate f middlebox failures

Minimizing performance overhead during normal operation

Minimizing disruption during middlebox failures

Fault Tolerant Chaining (FTC)

In-chain replication
- Replicates a chain's state instead of the state of individual middleboxes
- Each middlebox's state replicated to subsequent f middlebox servers

Transactional packet processing
- Simplifies the development of multi-threaded middleboxes
- Improves scalability and performance

Data dependency vectors
- Enables concurrent state replication

Normal Operation

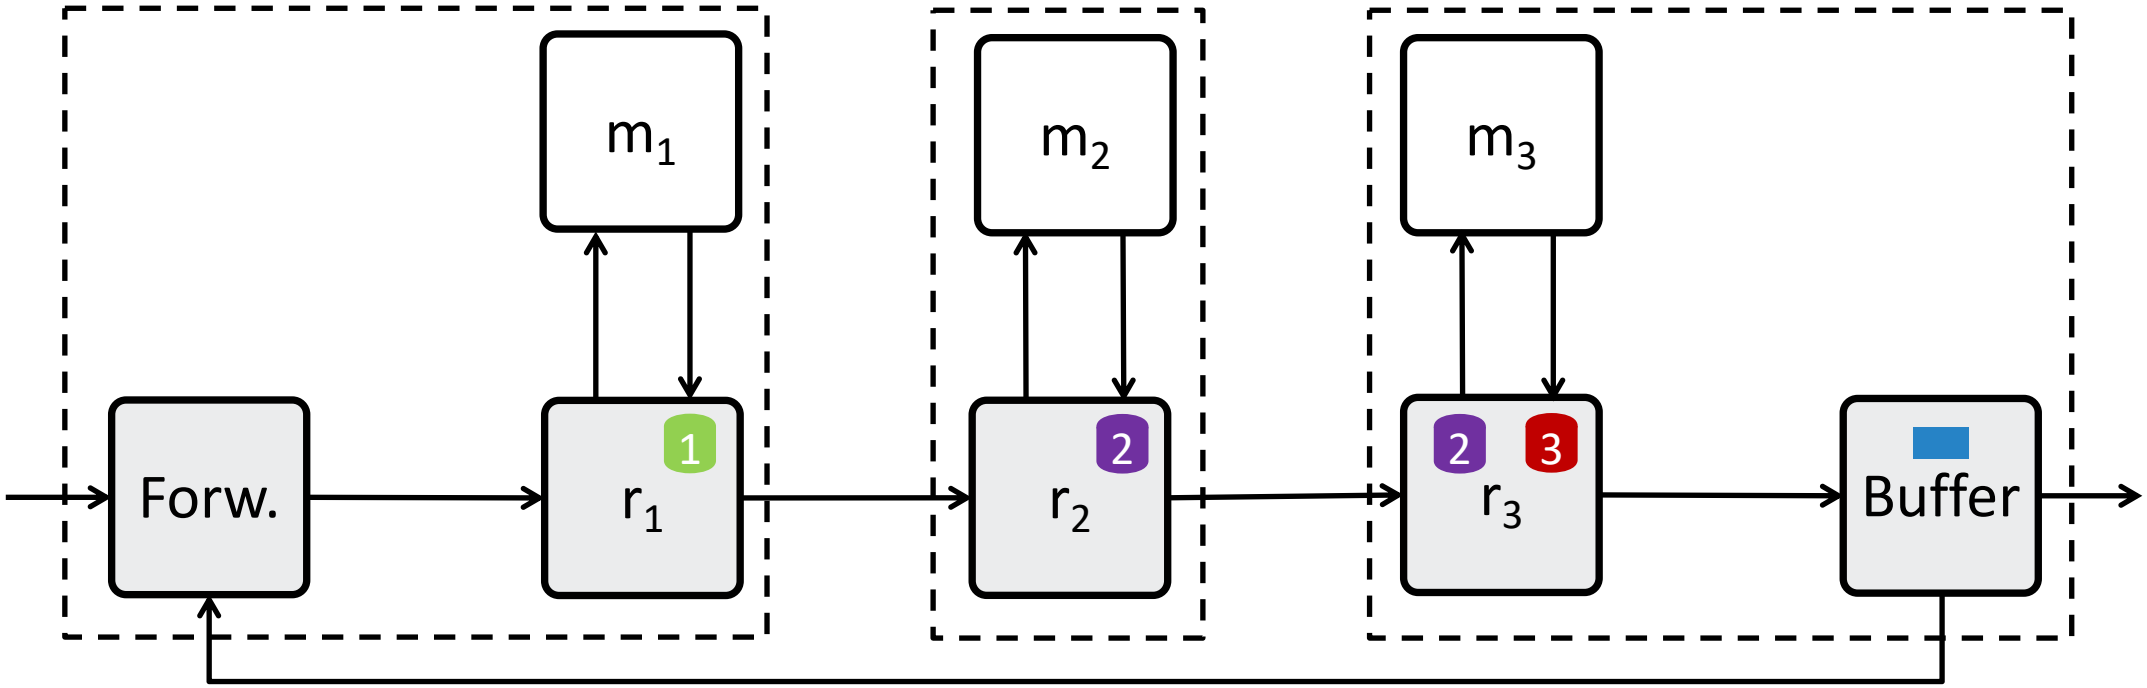

Normal Operation

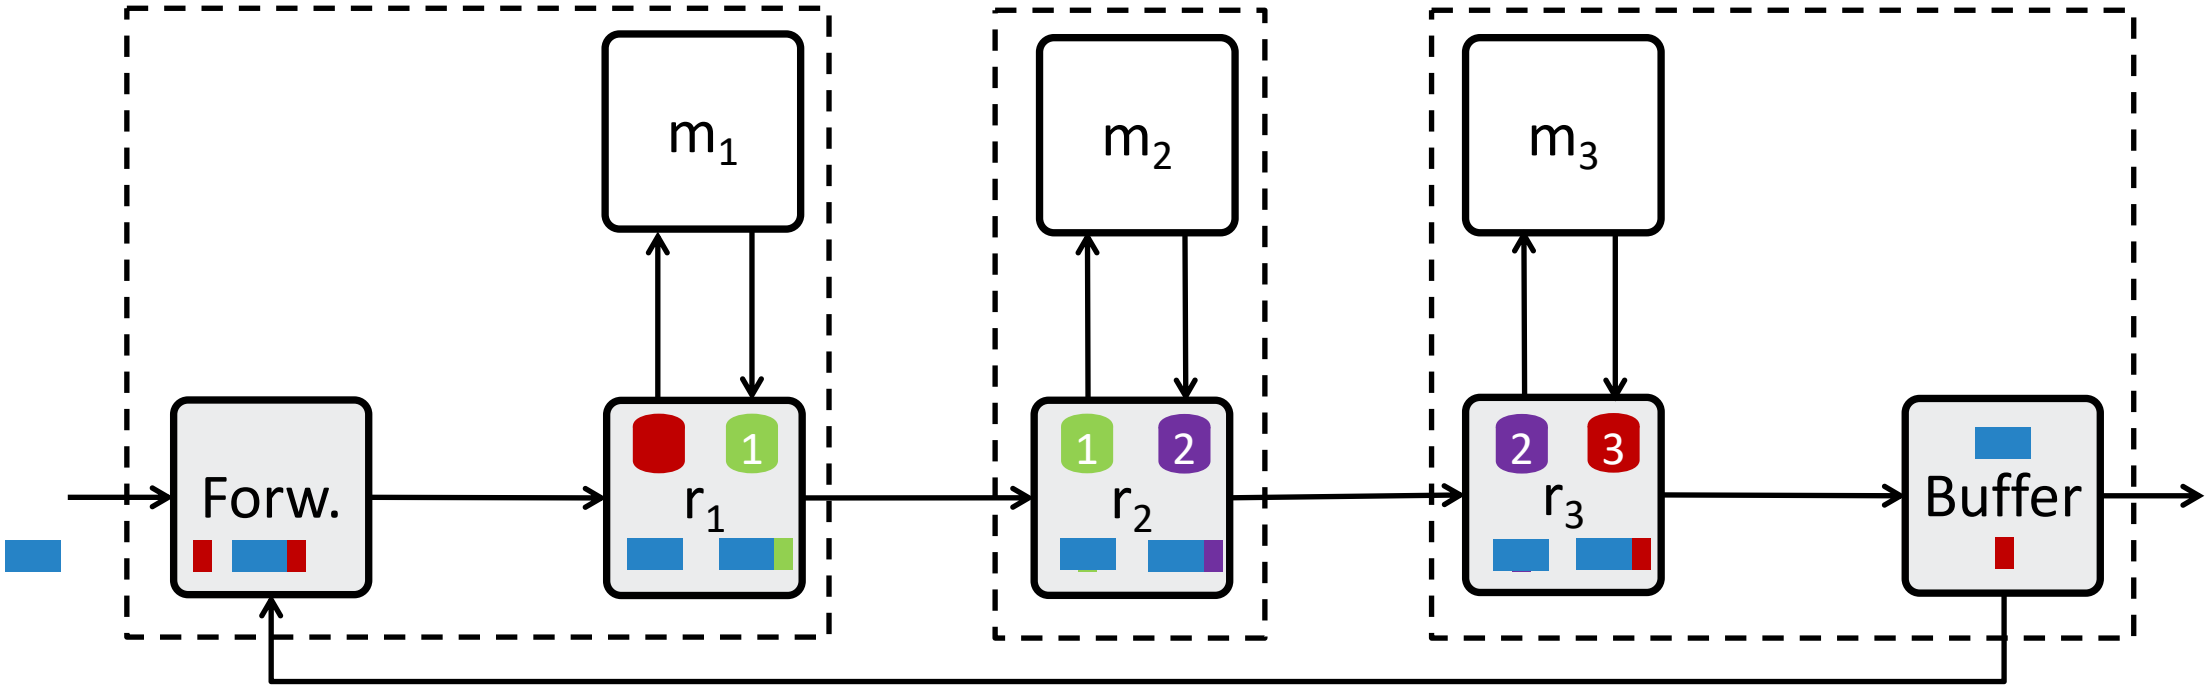

Normal Operation

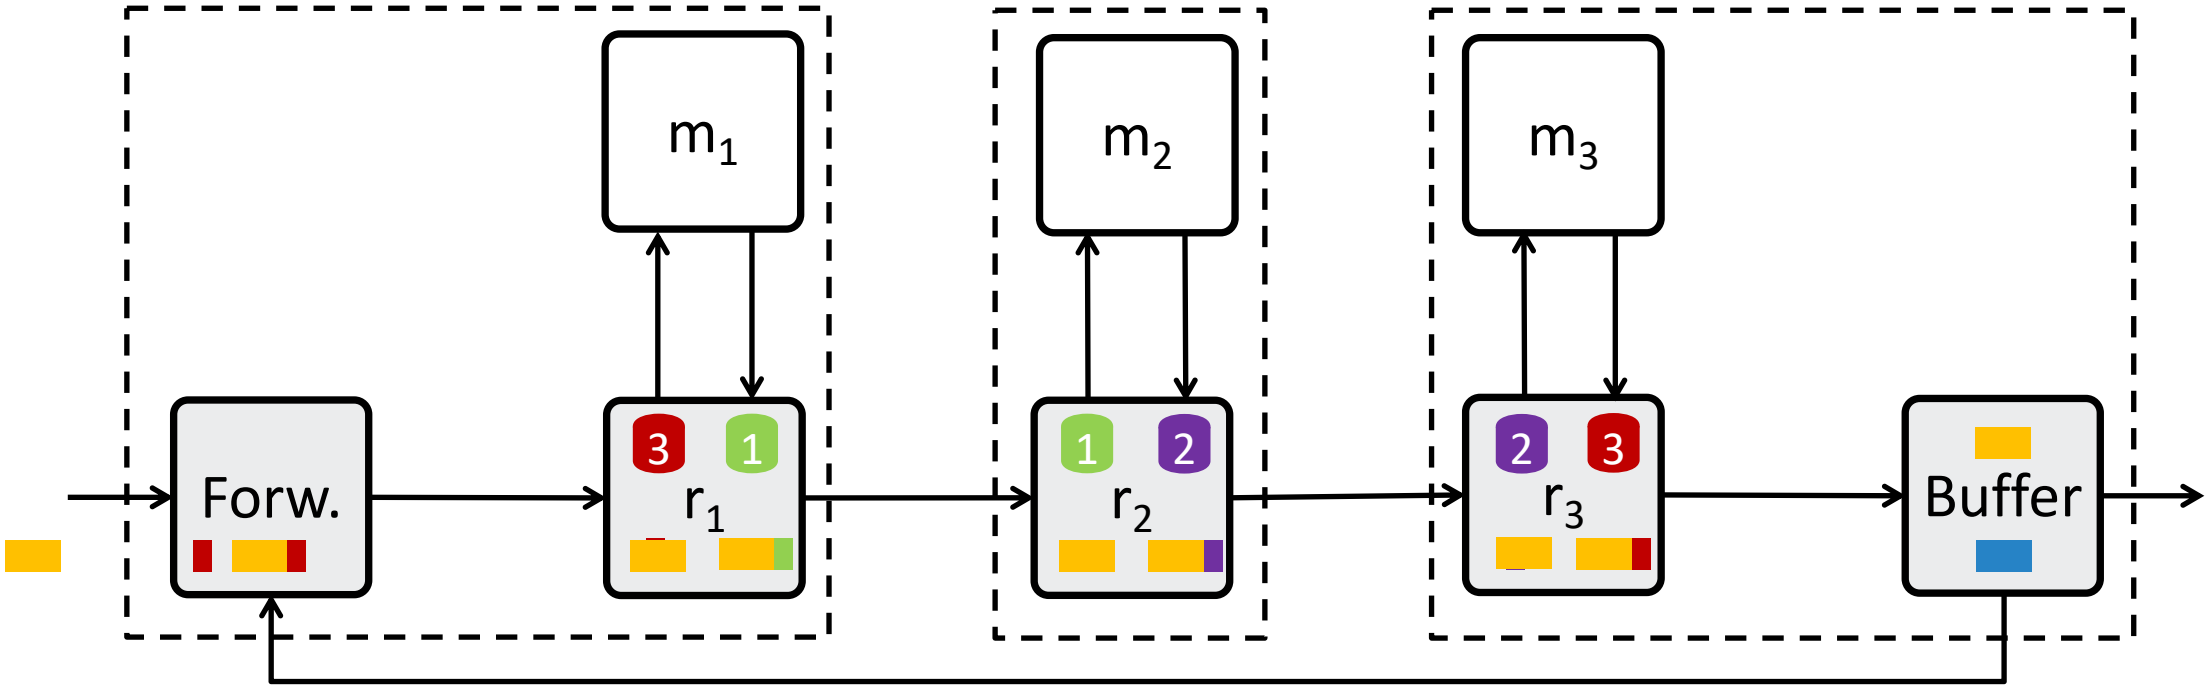

Failure Recovery

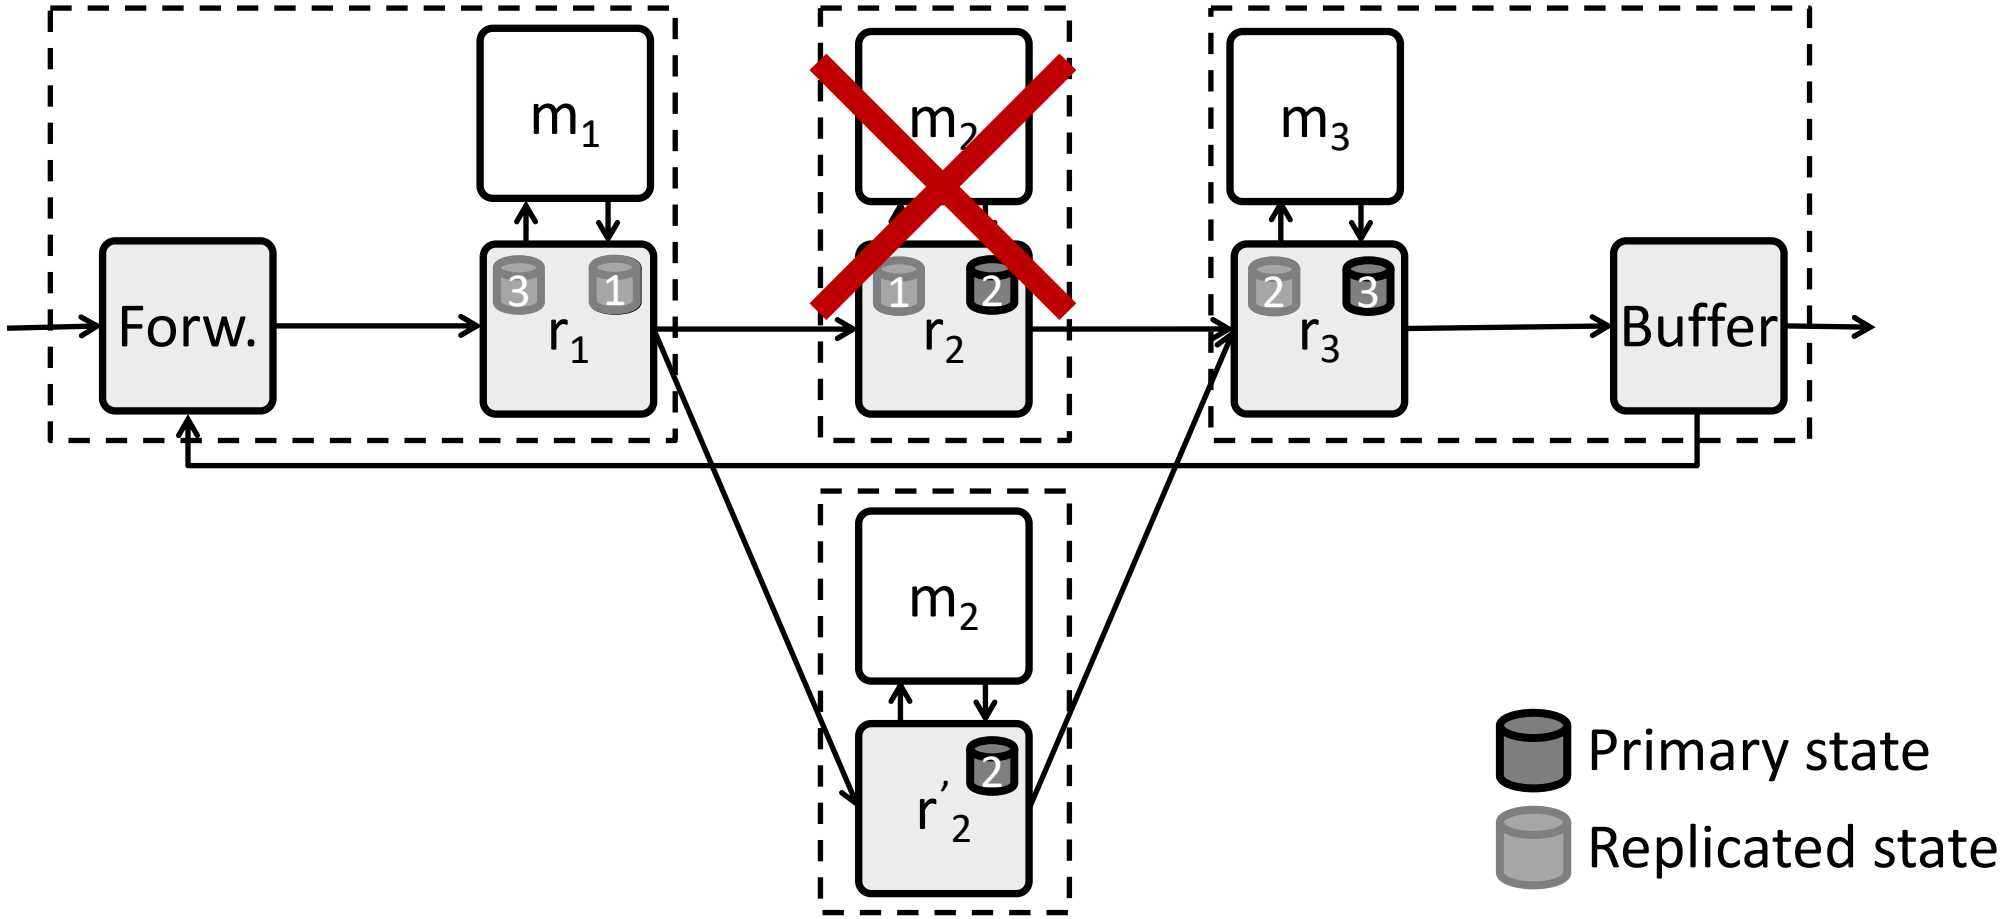

Transactional Packet Processing

Existing approaches

- Single thread or batched packet processing
- FTMB: multi threaded packet processing
  - Tracking state changes in granularity of each state variable read/write
  - Frequent periodic state snapshots

Our approach

- Packet transaction model for concurrent packet processing
- Using isolation property to tracking state changes in granularity of packet transactions

Data Dependency Vectors

Tracking data changes instead of thread operations

Enabling different number of threads at the middlebox and replicas

- Fail over to smaller machine
- Scale up to a larger machine

Middlebox	Product	Throughput	CPU Core
IPSec	HP VSR1001	268 Mbps	1
	HP VSR1008	926 Mbps	8
WAN Optimizer	STEELHEAD CCX770M	10 Mbps	2
	STEELHEAD CCX1555M	50 Mbps	4
WAF	Barracuda Level 1	100 Mbps	1
	Barracuda Level 5	200 Mbps	2

Data Dependency Vectors Example

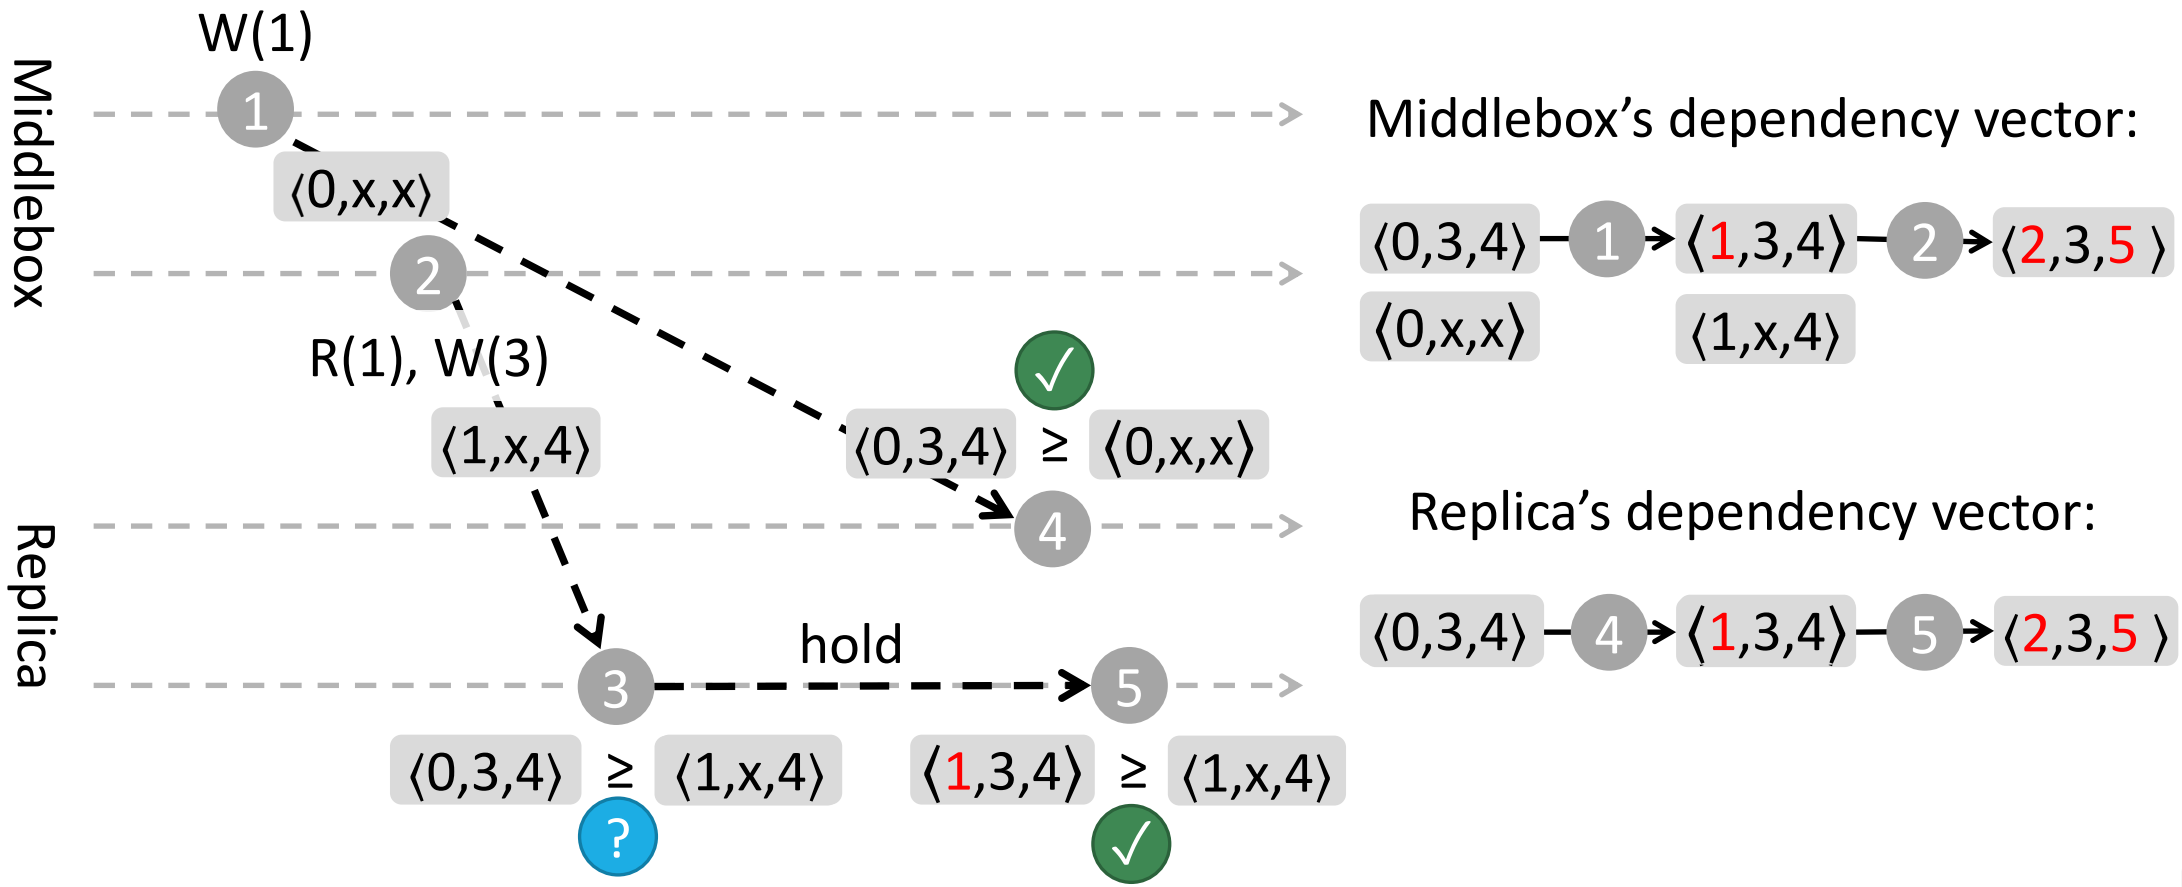

Evaluation

METHOD

Comparing FTC with:

NF, **N**on-**F**ault tolerant system
- Ideal performance

FTMB (SIGCOMM 2015)
- State logging ~~+ Snapshots~~

FTMB + Snapshot (SIGCOMM 2015)
- State logging + Snapshots

ENVIRONMENTS

A cluster of 12 servers
- 40 Gbps network

SAVI Cloud environment
- A virtual network of OVS switches

MoonGen and pktGen traffic generators
- UDP traffic
- Packet size: 256 B

Fault Tolerant NATs

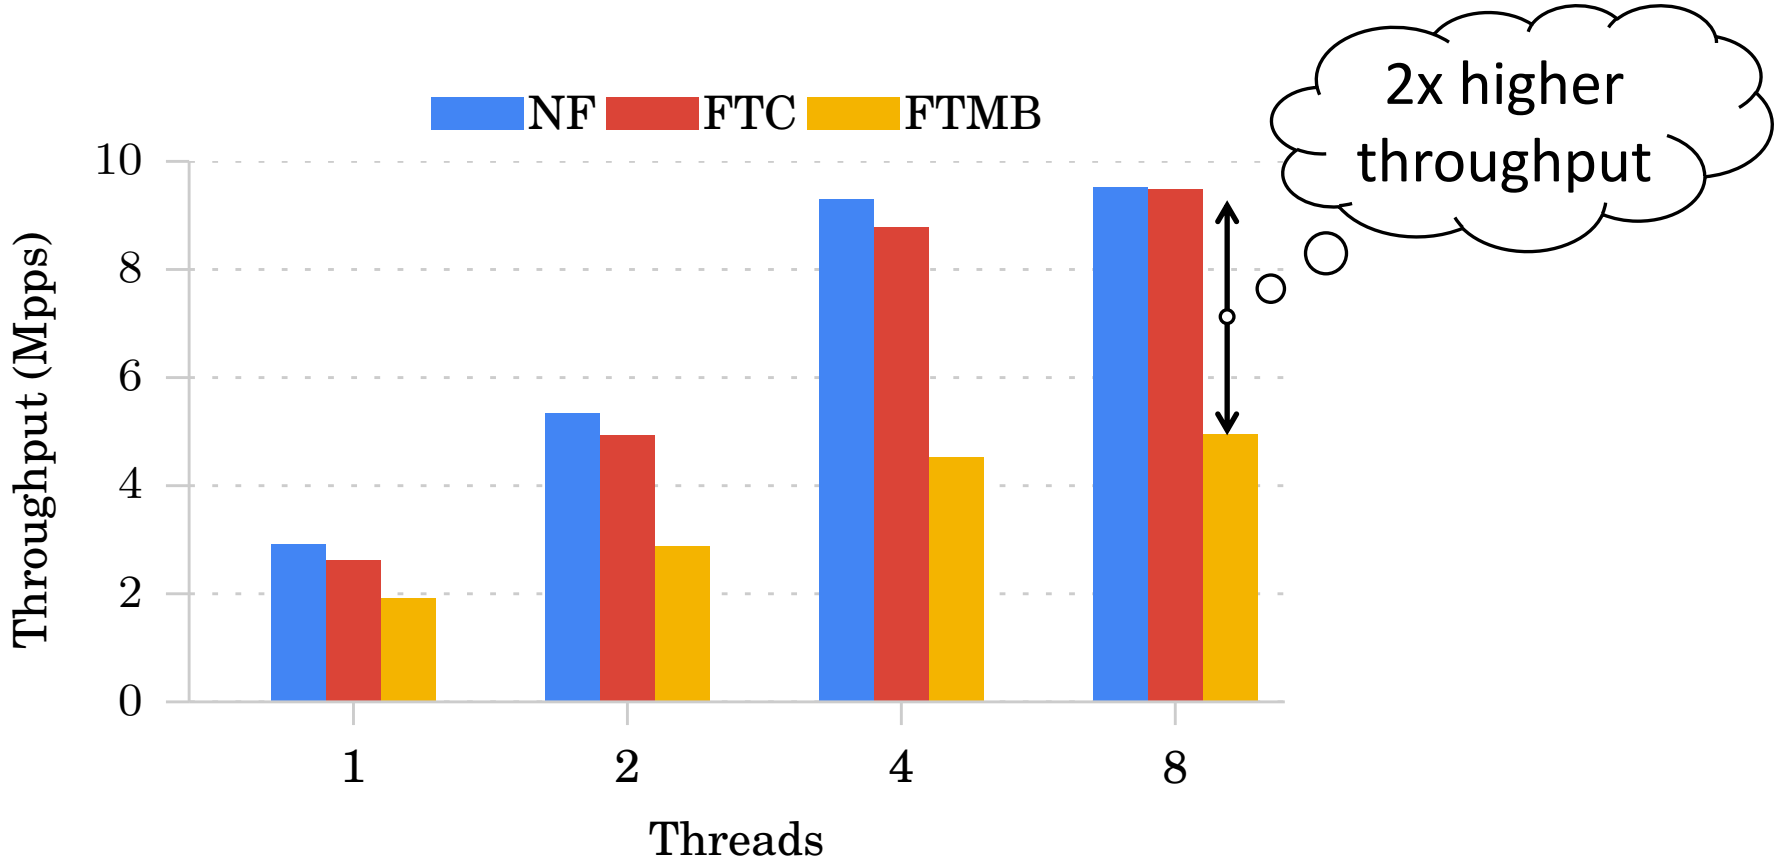

Fault Tolerant Chains – Throughput

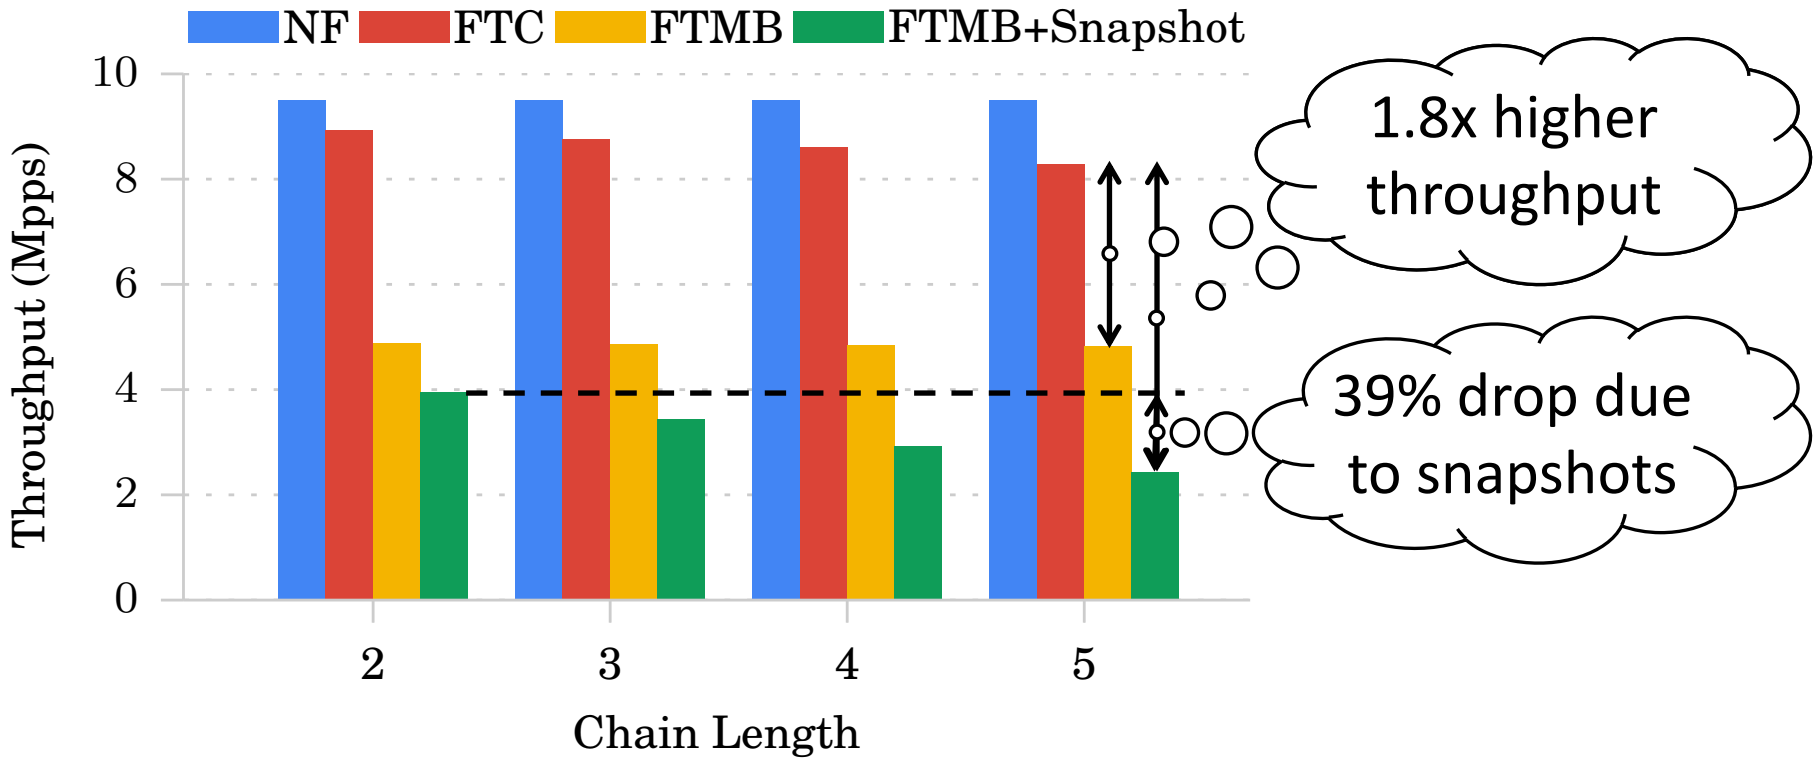

Fault Tolerant Chains – Latency

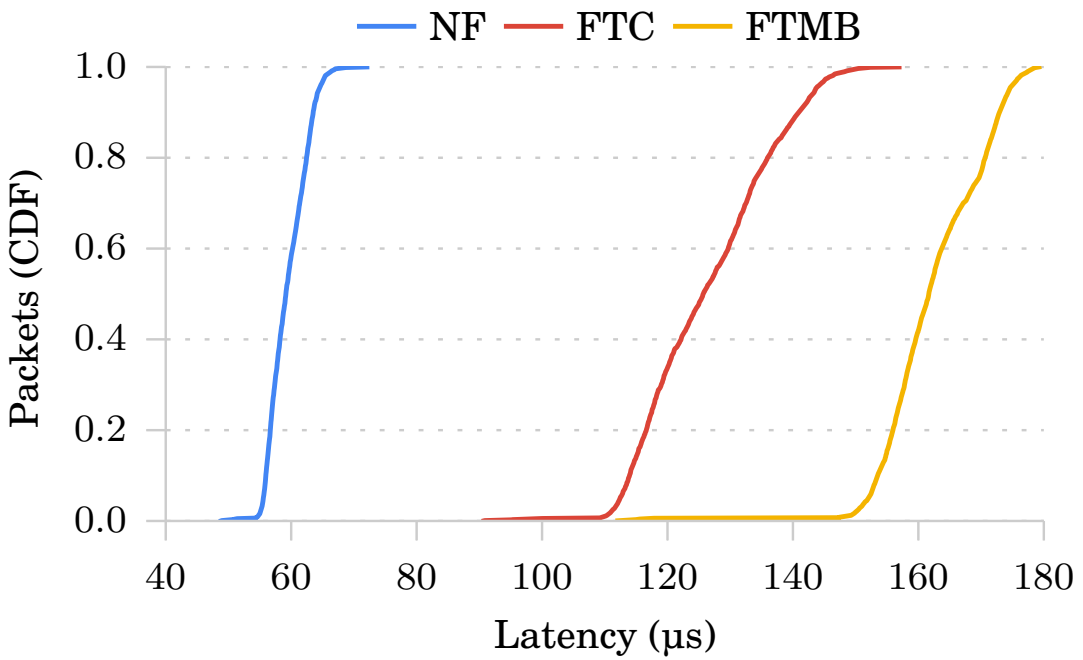

Conclusion

Keep operation of a chain of middleboxes online after f middleboxes fail

- In-chain replication
- Transactional packet processing
- Data dependency vectors