Tardigrade:
Leveraging Lightweight Virtual Machines to Easily and Efficiently Construct Fault-Tolerant Services

Jacob R. Lorch

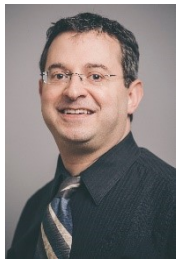

Andrew Baumann

Microsoft Research

Lisa Glendenning

Dutch T. Meyer

Andrew Warfield

Our goal:

*Turn **existing** binaries into fault-tolerant services.*

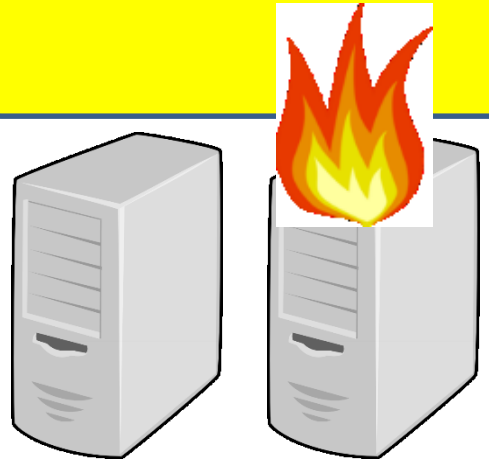

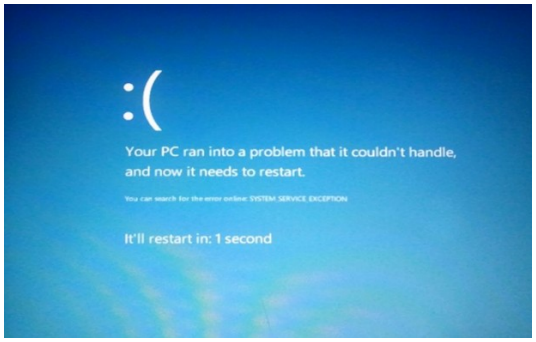

Example: FDS Metadata Service

FDS Cluster

[Nightingale et al., OSDI 2012]

Example: FDS Metadata Service

FDS Cluster

[Nightingale et al., OSDI 2012]

Techniques for making code fault-tolerant

Use state machine replication library

hav

Better:

Transparently make the binary fault-tolerant

Potential for oversight

- Non-determinism
- Failing to persist state
- Exposing non-persisted data
- Bugs in crash recovery

Outline

- Motivation
- Background: Asynchronous VM replication
- Our solution: Lightweight VM replication
- Challenges and solutions
- Evaluation

Outline

- Motivation
- Background: Asynchronous VM replication
- Our solution: Lightweight VM replication
- Challenges and solutions
- Evaluation

Asynchronous virtual machine replication - Remus

[Cully et al., NSDI 2008]

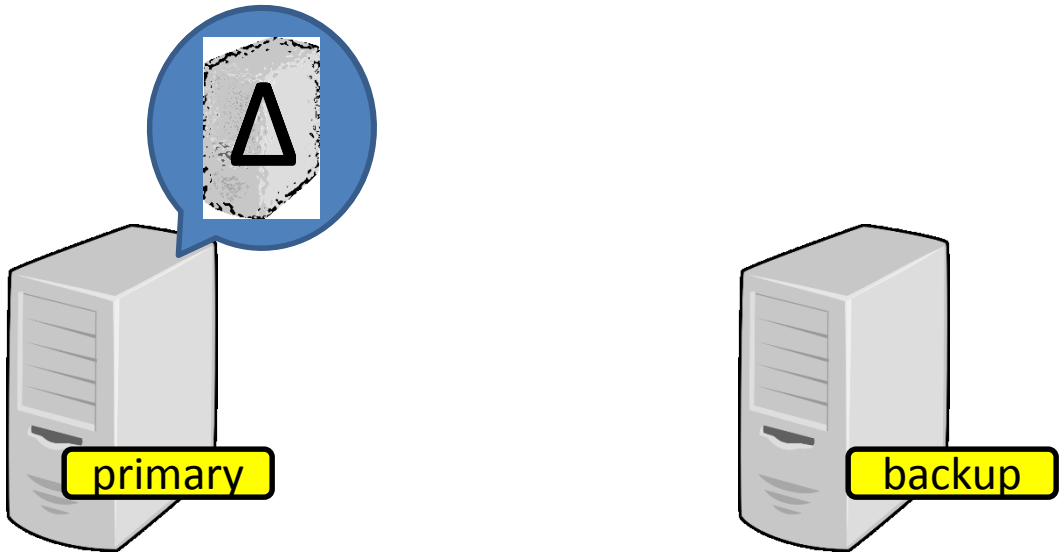

Primary can crash at any time; backup is always a bit behind.

Asynchronous virtual machine replication - Remus

[Cully et al., NSDI 2008]

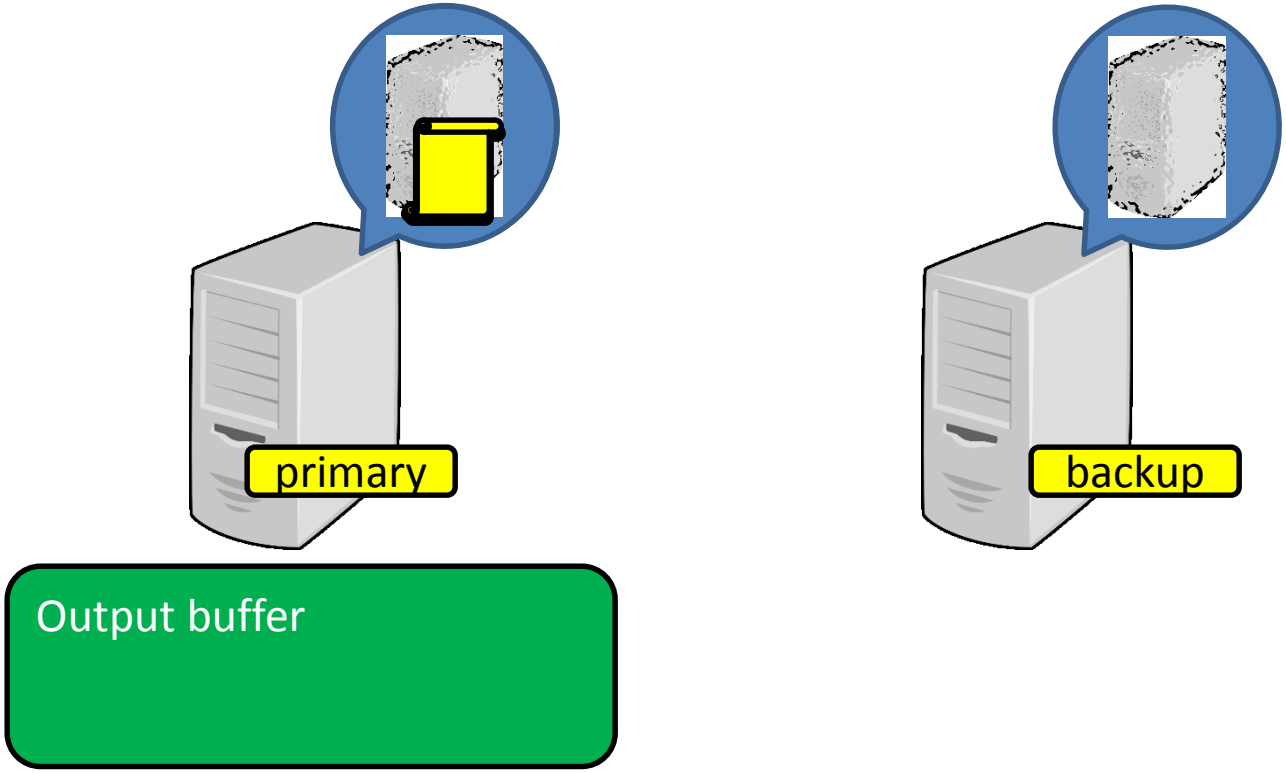

Asynchronous virtual machine replication - Remus

[Cully et al., NSDI 2008]

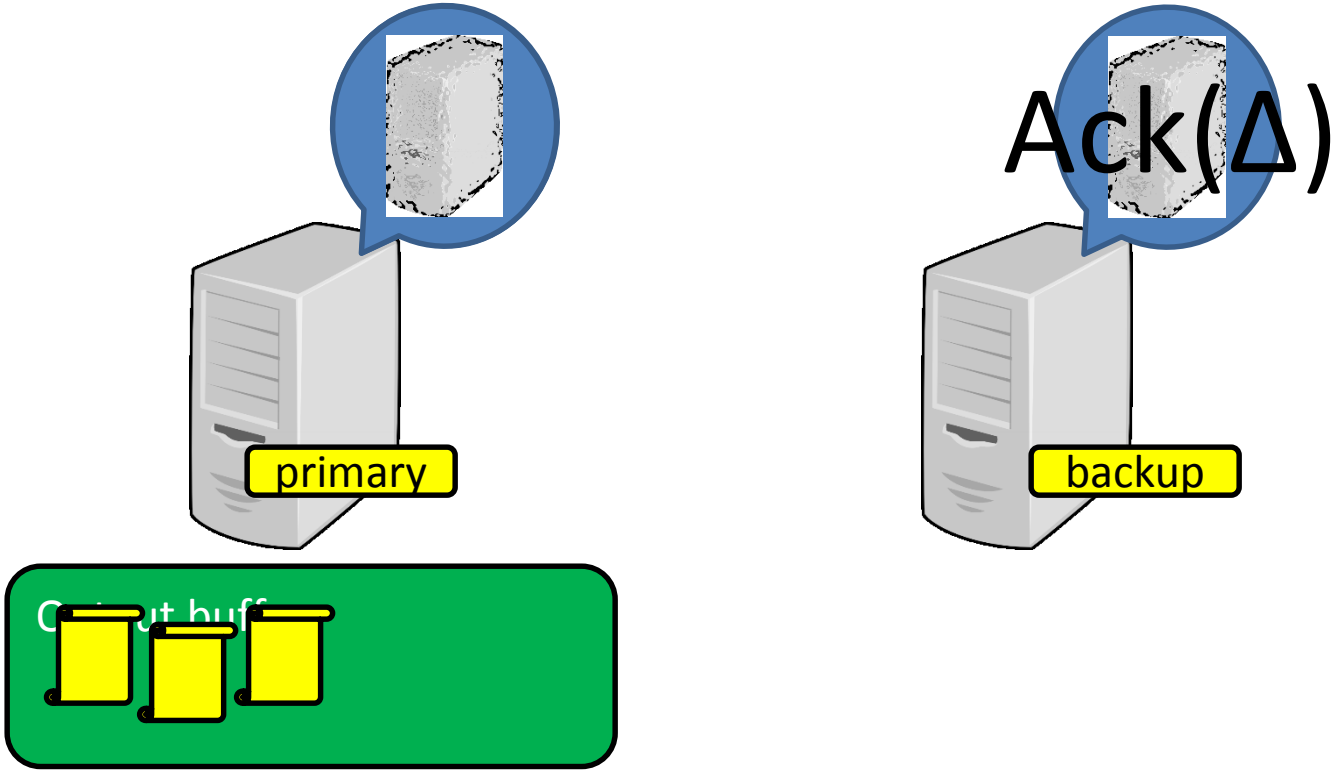

High VM activity can delay packets

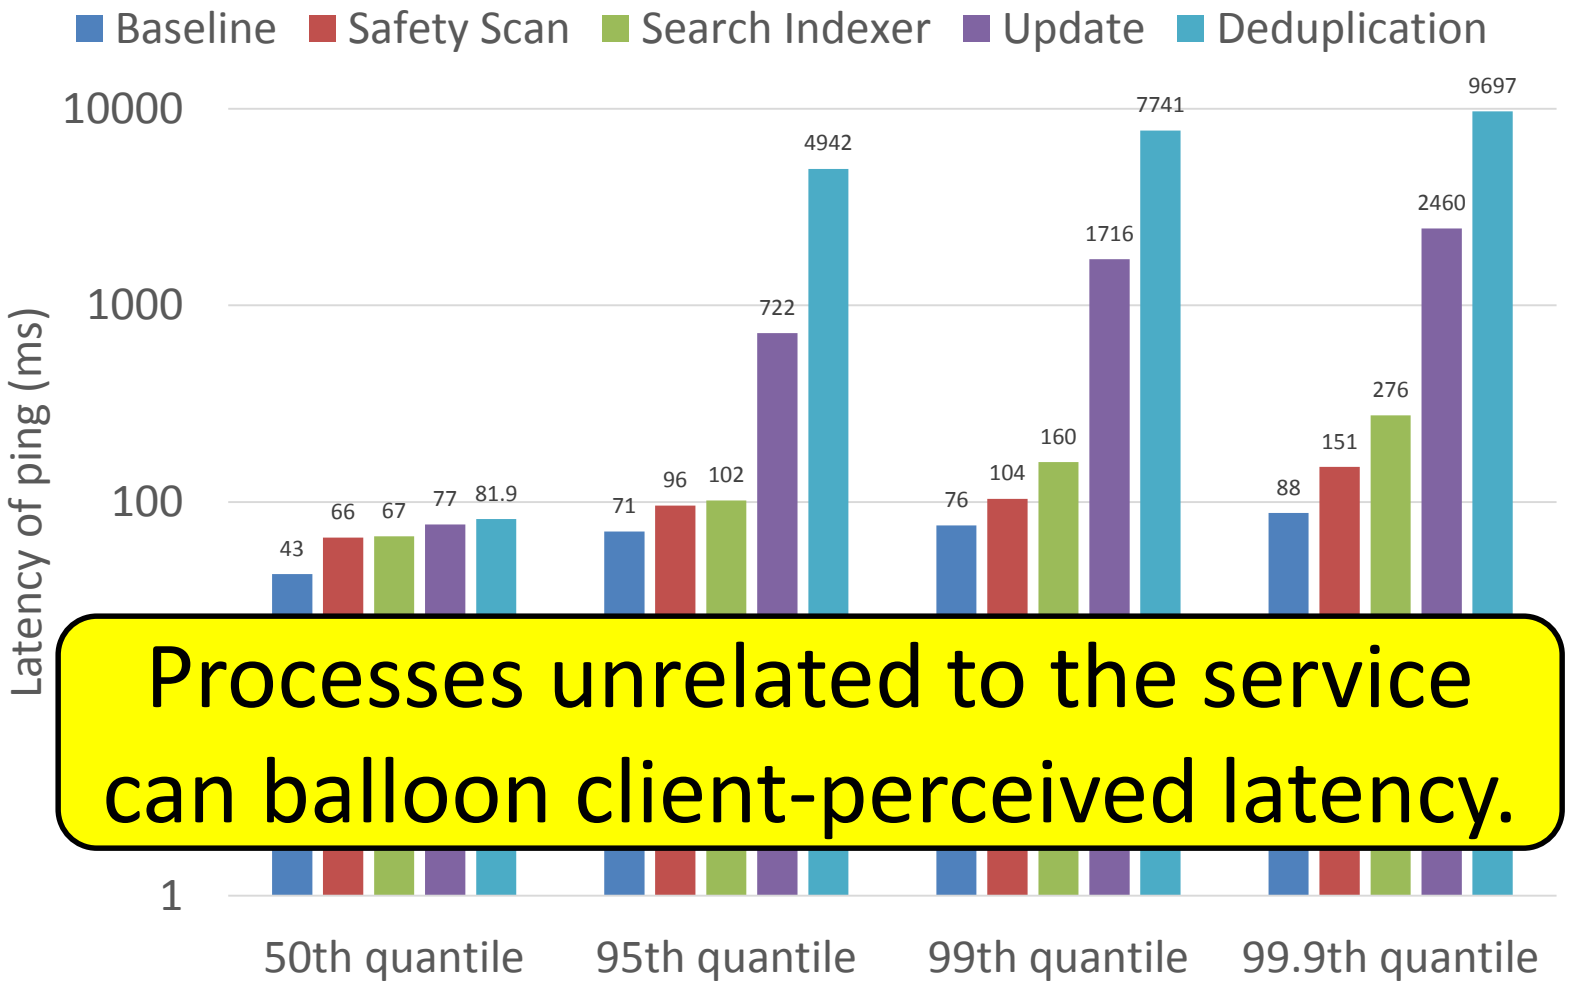

Processes unrelated to the service can balloon client-perceived latency.

Outline

- Motivation
- Background: Asynchronous VM replication
- Our solution: Lightweight VM replication
- Challenges and solutions
- Evaluation

Our solution: Use *lightweight* VMs instead

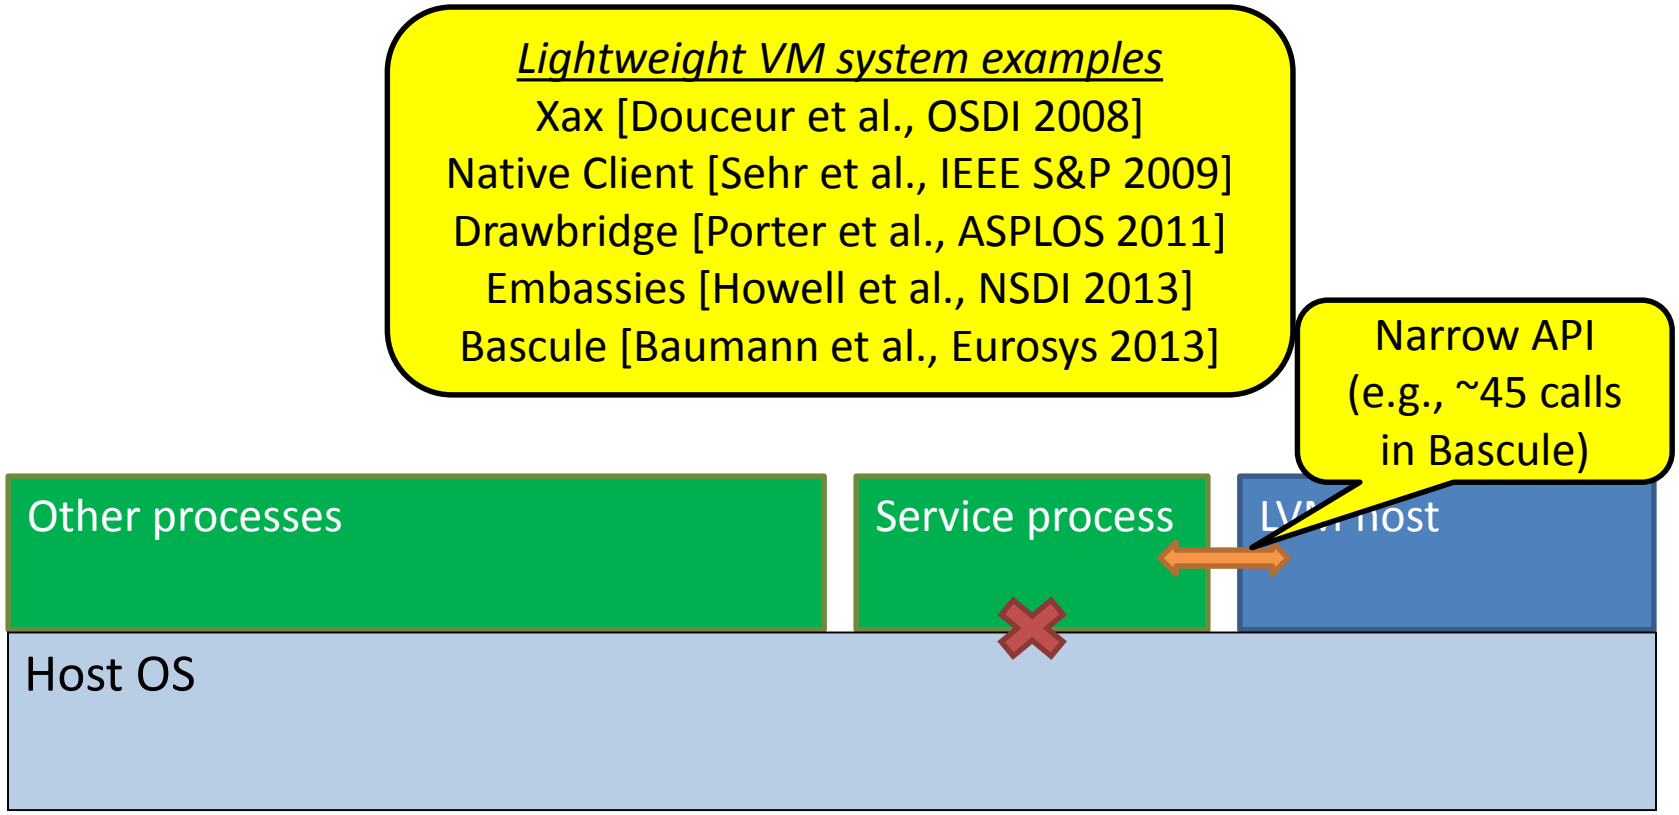

Lightweight VMs can support unmodified binaries via a library OS

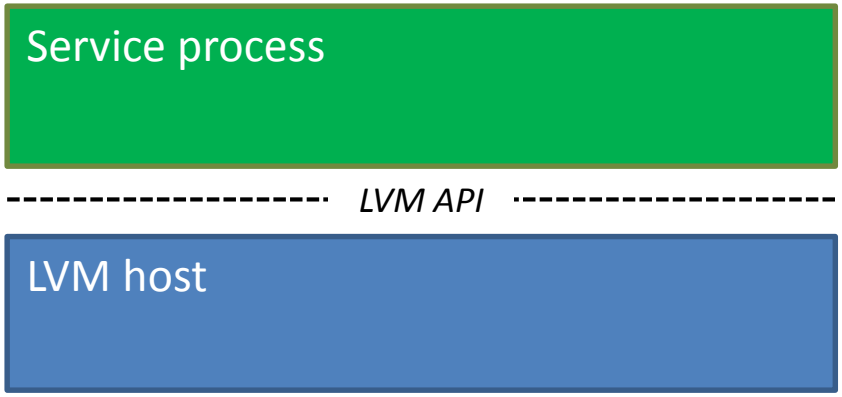

Lightweight VMs can support unmodified binaries via a library OS

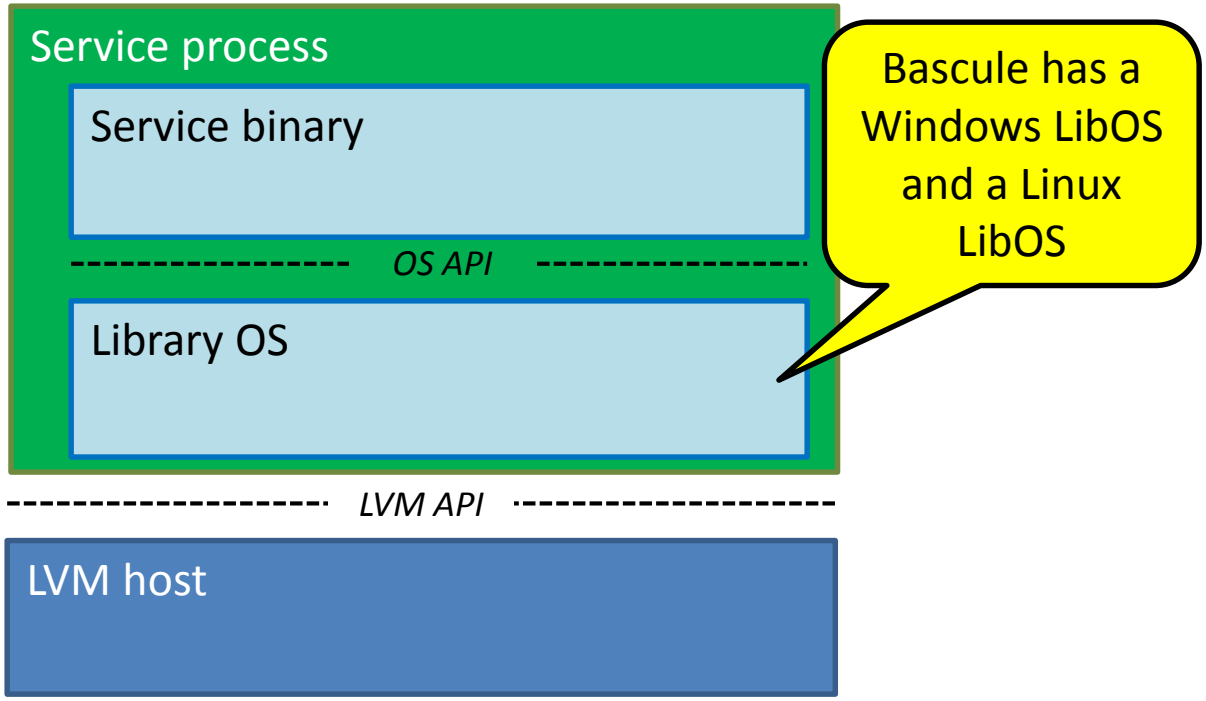

A lightweight VM is encapsulated by virtue of having a narrow interface

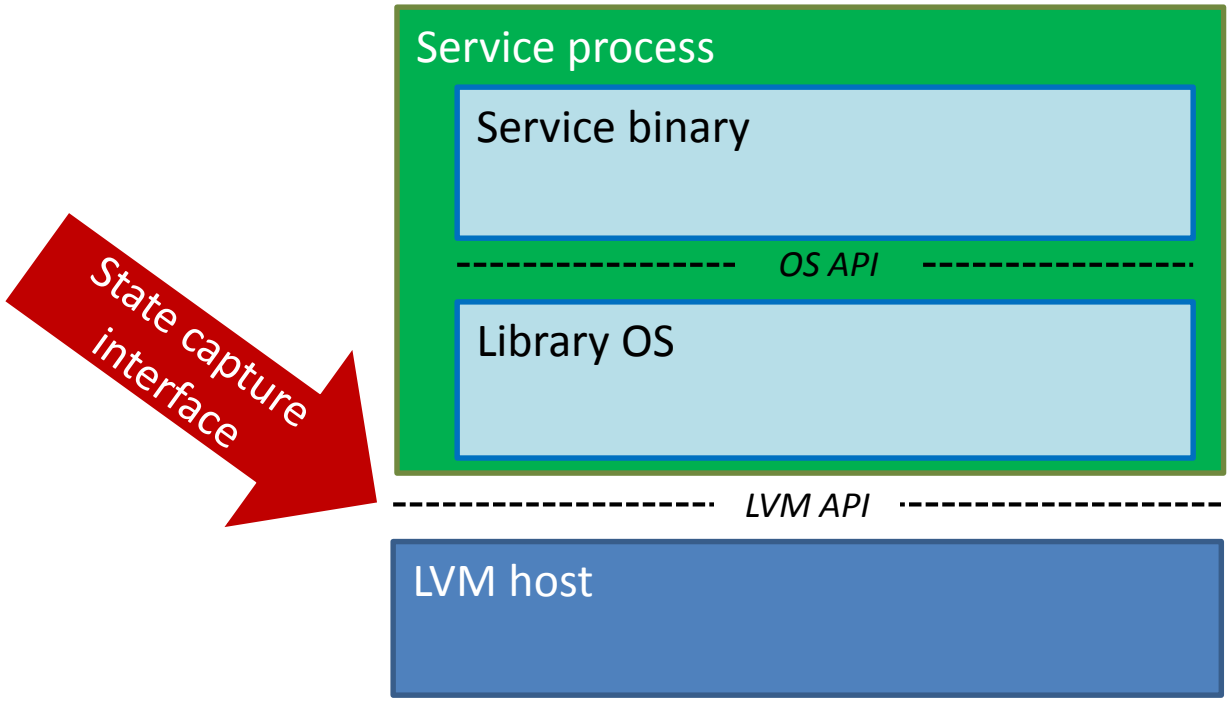

Our approach: Checkpoint by interposing on existing LVM API

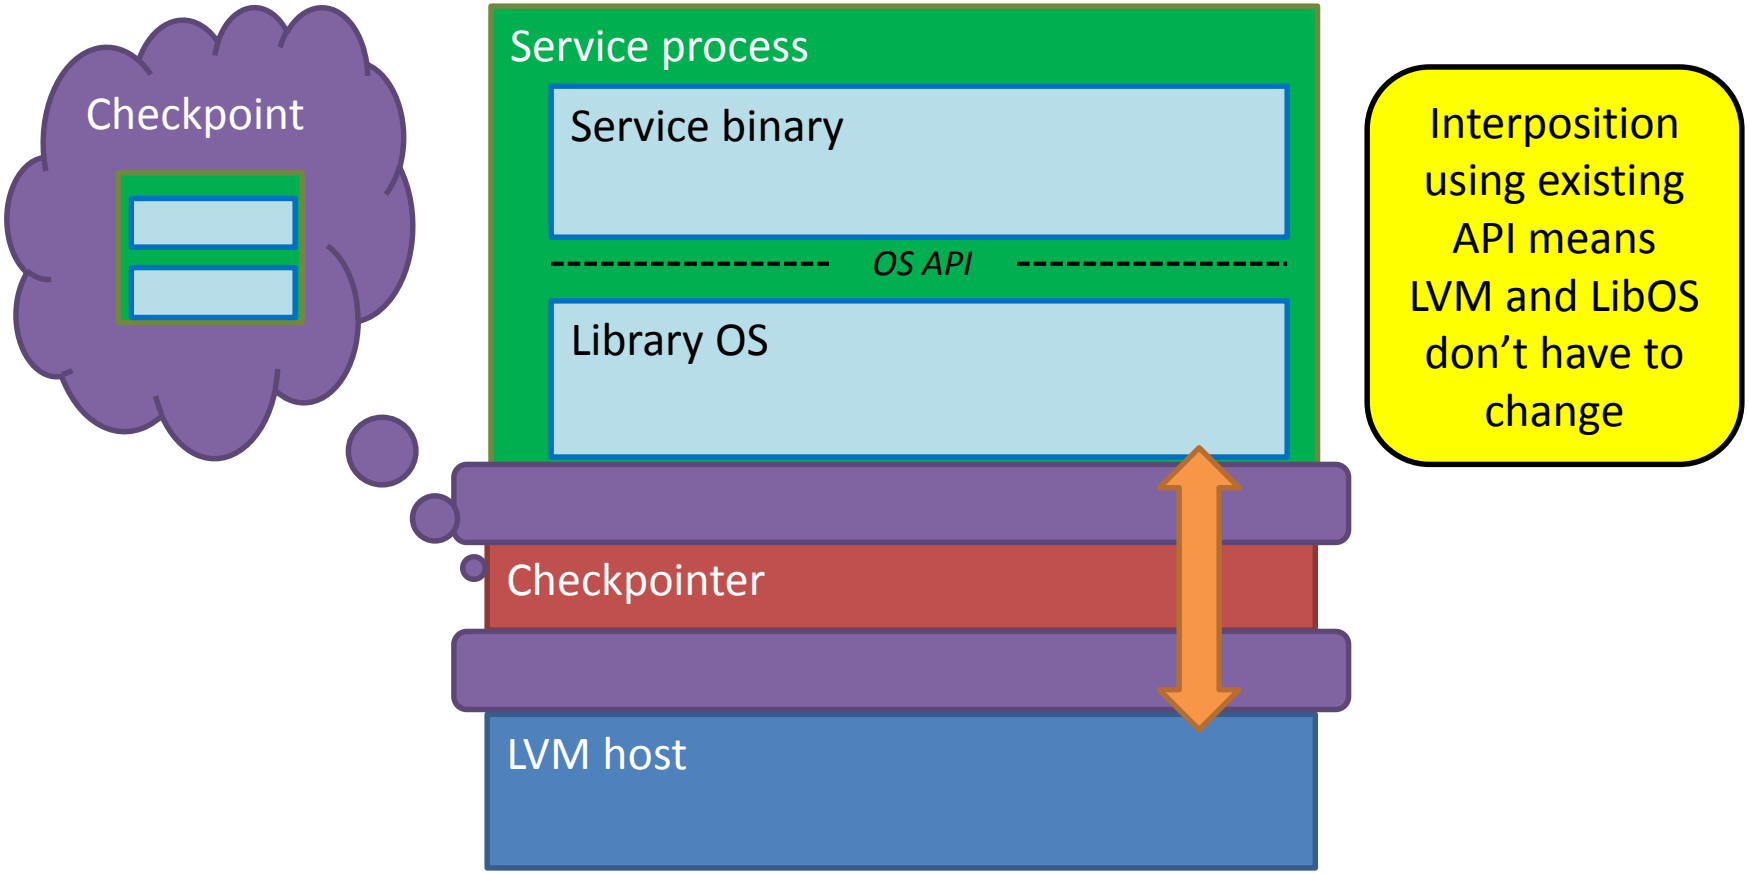

[Cully et al., NSDI 2008]

Asynchronous
Virtual
Machine
Replication

Lightweight
Virtual
Machine
Replication

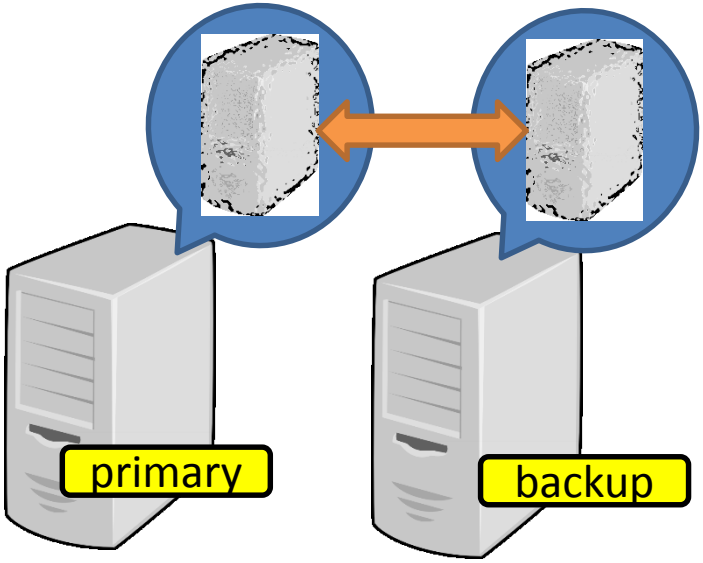

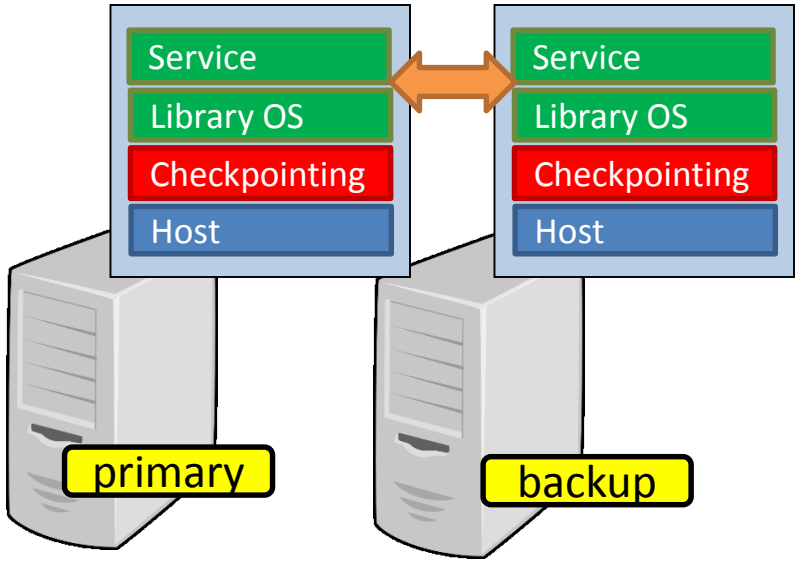

[Cully et al., NSDI 2008]

Asynchro
Virtual
Machine
Replication

Our implementation of LVMR is called Tardigrade

Lightweight
Virtual
Machine
Replication

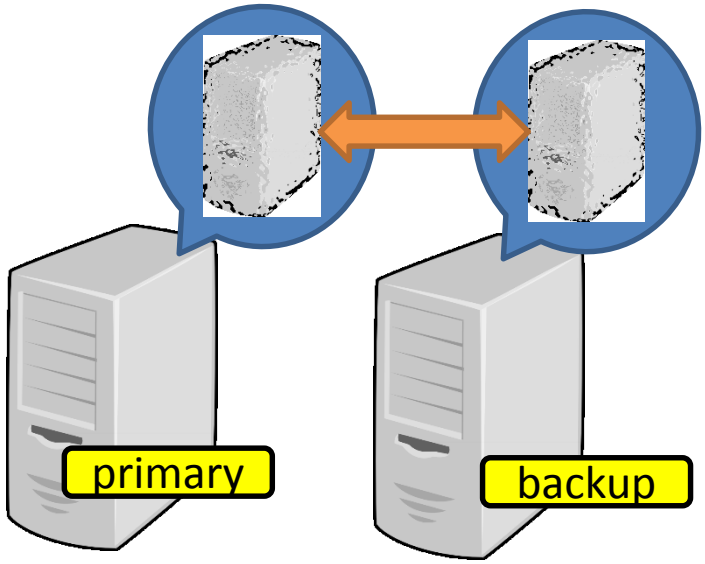

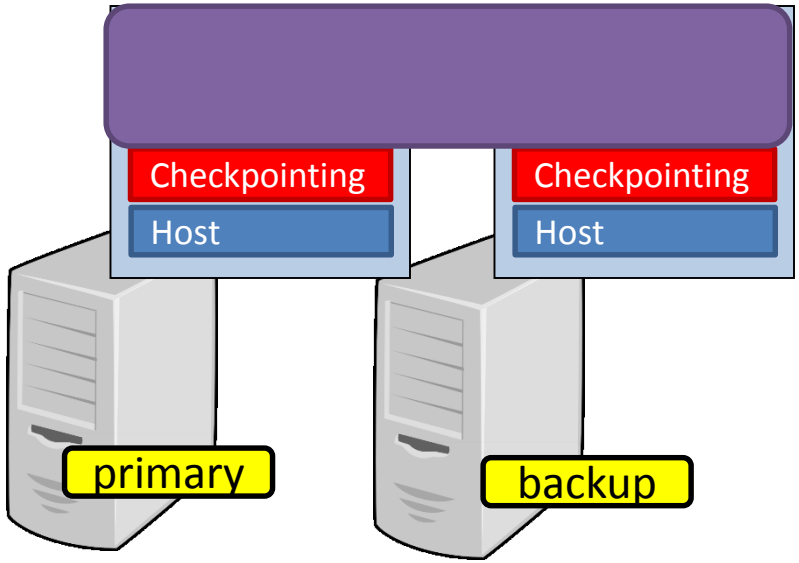

Outline

- Motivation
- Background: Asynchronous VM replication
- Our solution: Lightweight VM replication
- Challenges and solutions
- Evaluation

Practical LVMR poses challenges

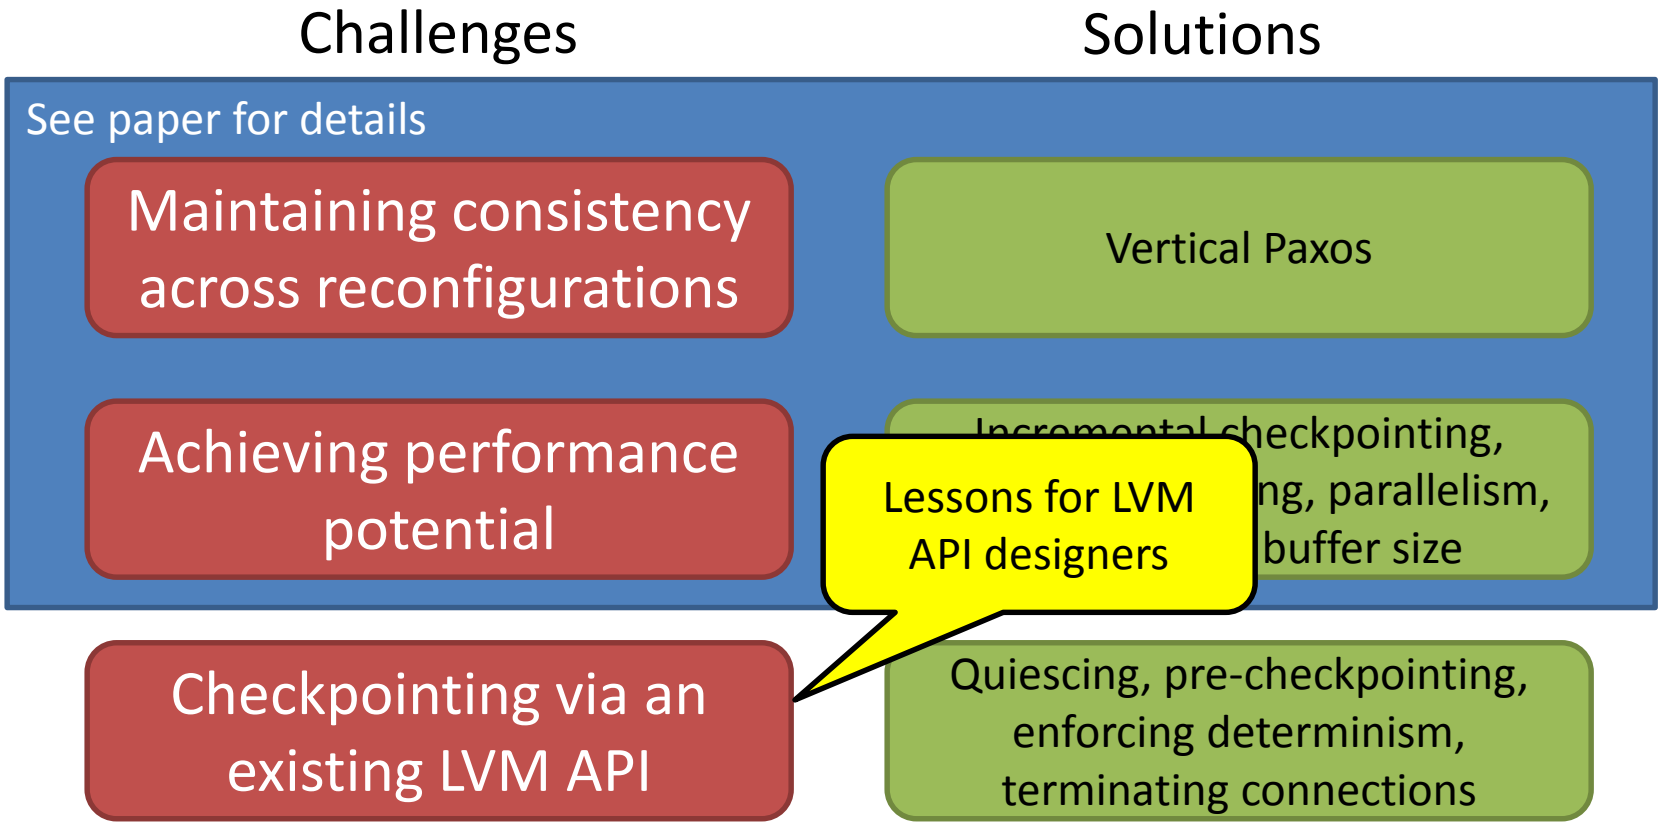

Checkpointing uses certain LVM API features

Feature	Purpose
Ability to track changed memory pages	Efficiently compute checkpoint deltas
Ability to suspend and inspect other threads	Capture consistent snapshot
Determinism when API calls are replayed	Prevent divergence on failover
Host state either replayable or regeneratable	Recreate host state on backup

Features may not always be in LVM APIs

Feature	Workaround
Ability to track changed memory pages	
Missing ability to suspend and inspect other threads	Use exceptions, pre-checkpointing
Non-determinism when API calls are replayed	Hide non-determinism
Host state **not** replayable or regeneratable	Expose divergence as error condition

To capture a checkpoint, we must quiesce and capture all threads' state.

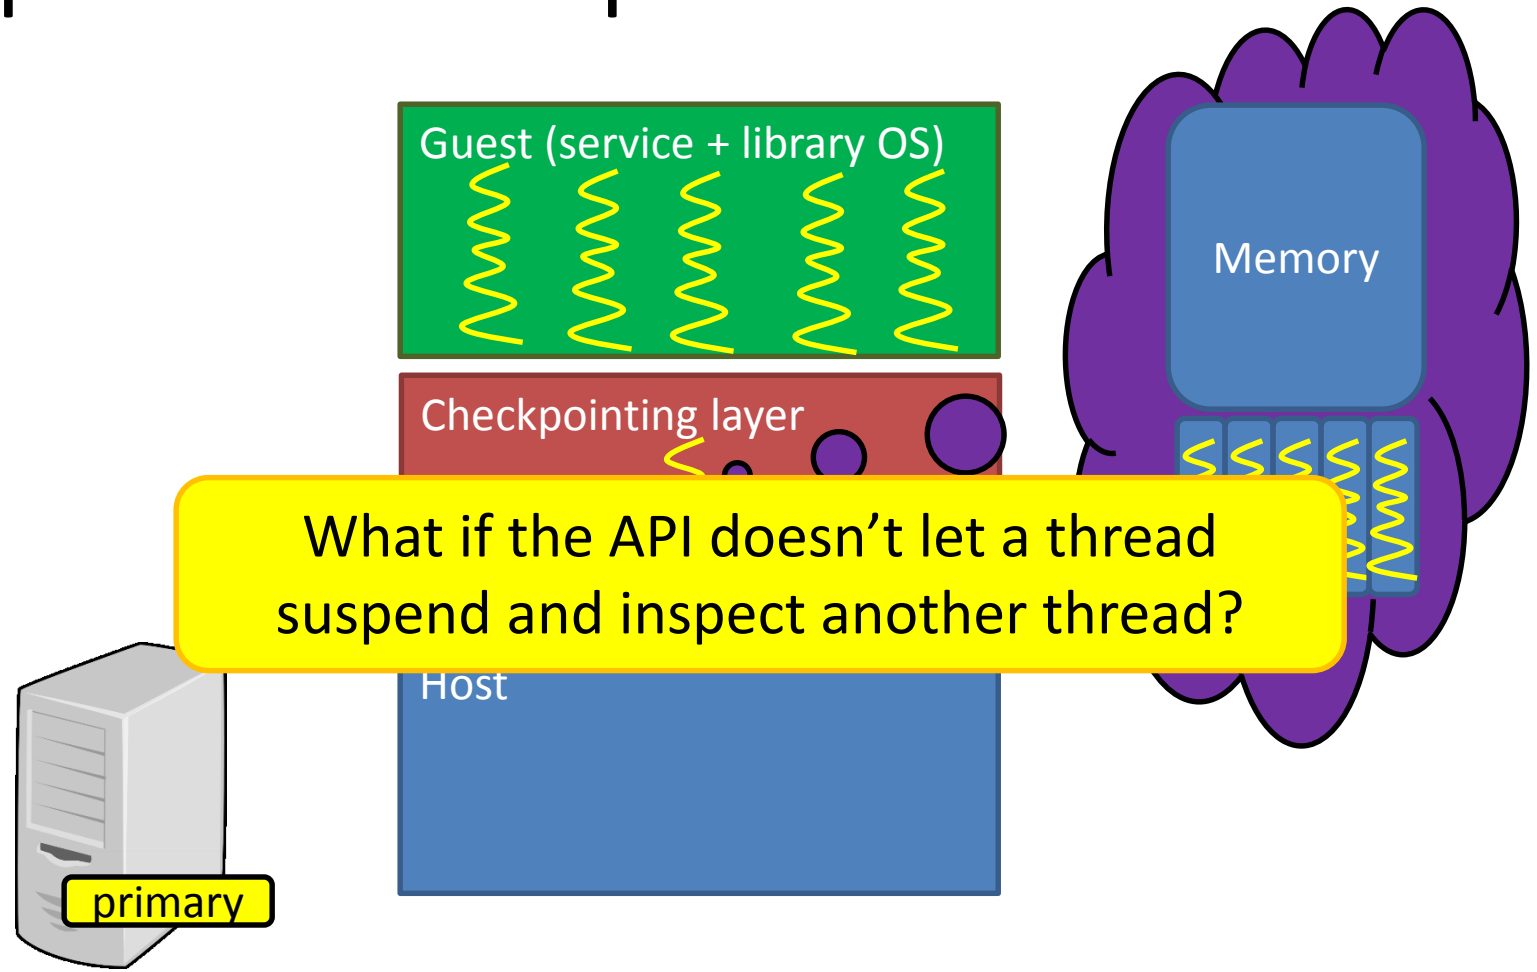

We can use exceptions to quiesce guest threads

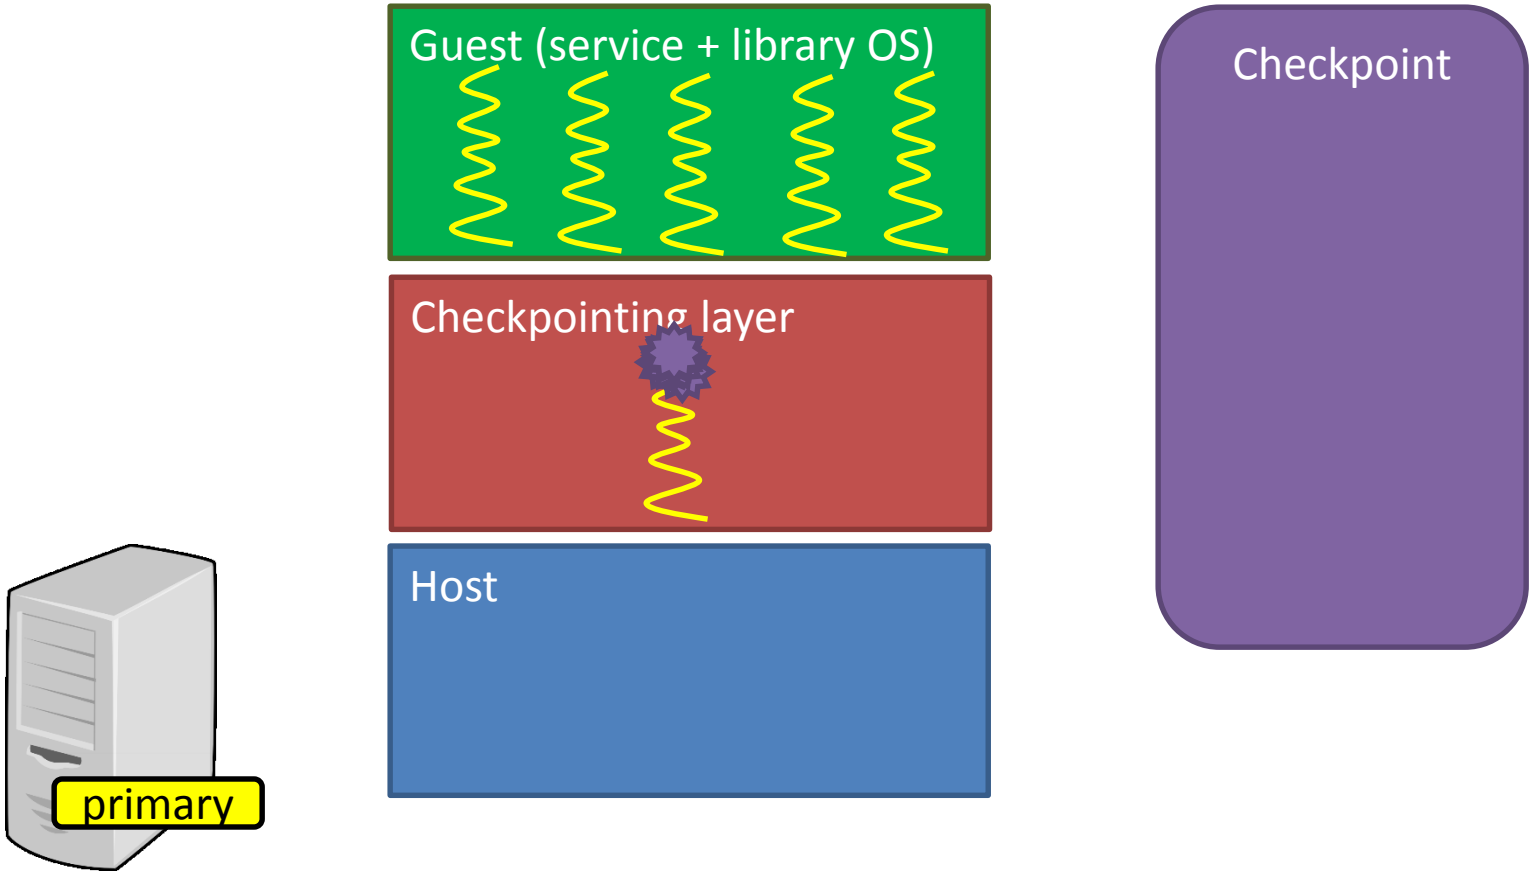

Exception handler quiesces and captures each guest thread's state

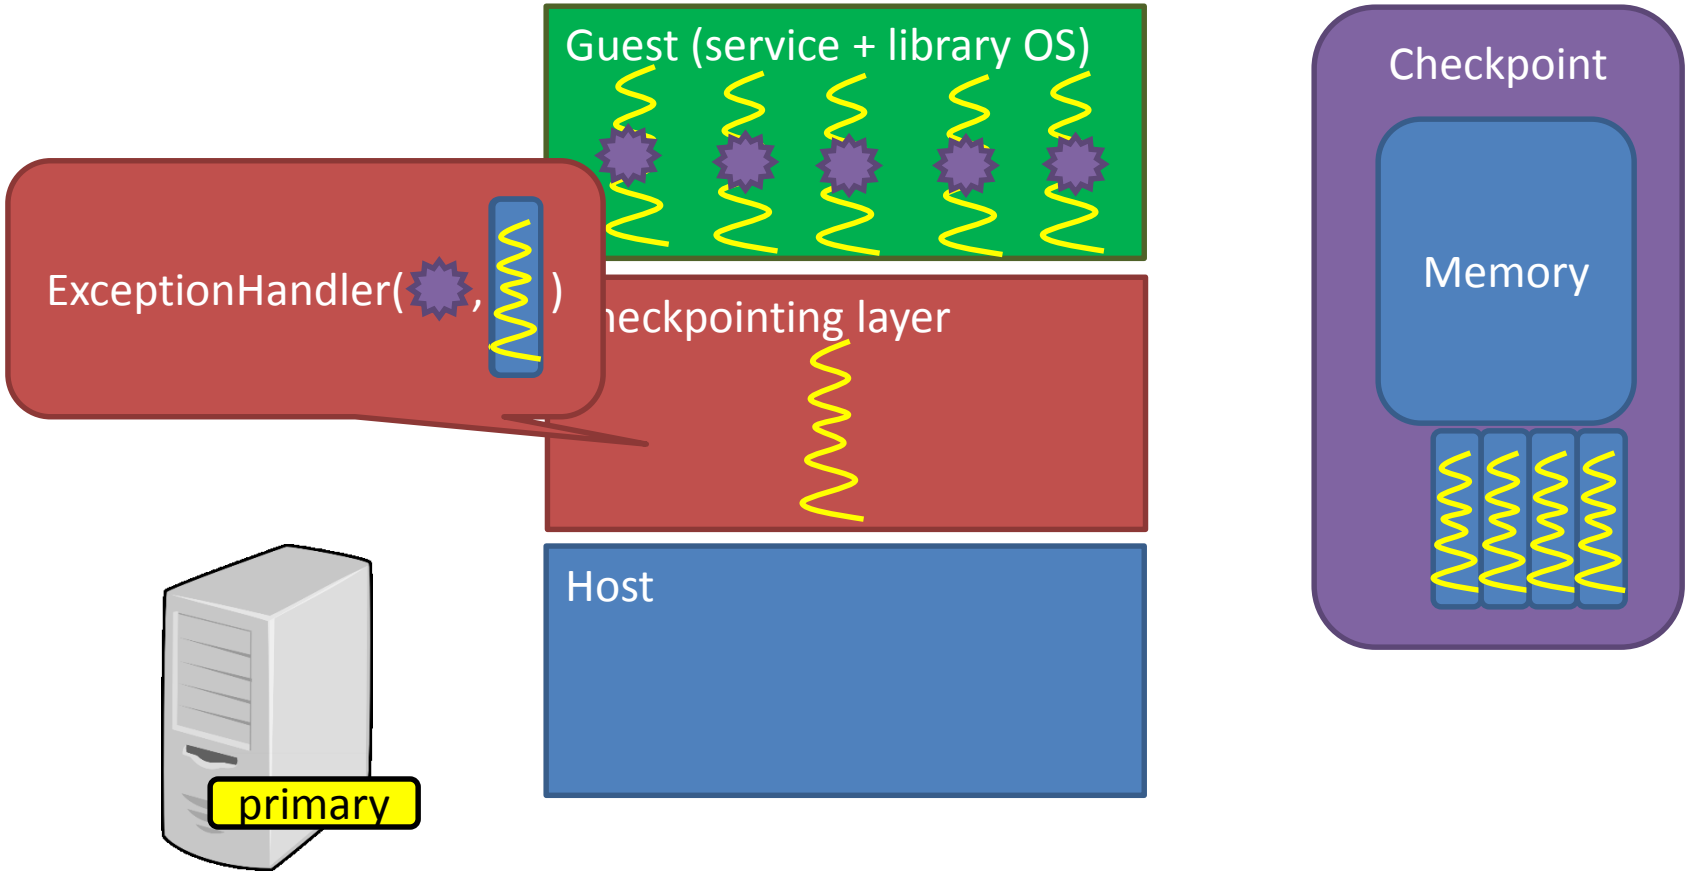

Synchronous system calls complicate quiescence

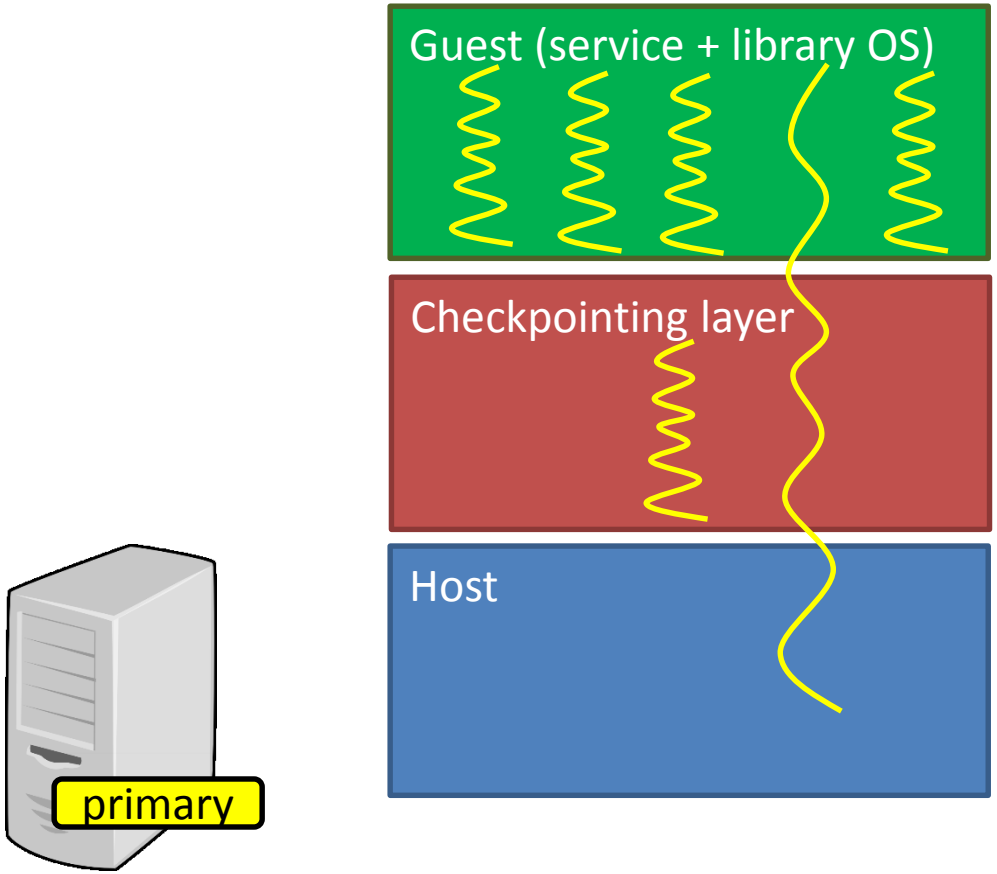

The wait system call is easy to deal with

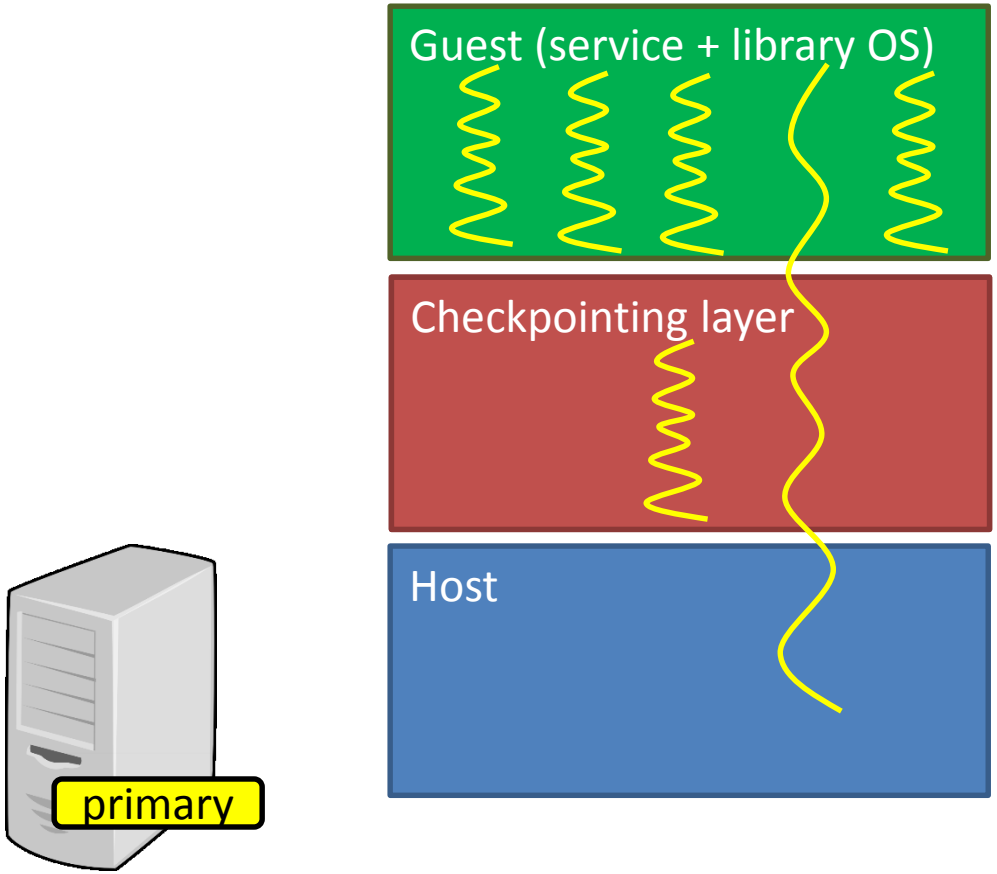

select() file descriptor list
0x1AC
0x3BB
0x907
time-to-checkpoint

General synchronous system calls require *pre-checkpointing*

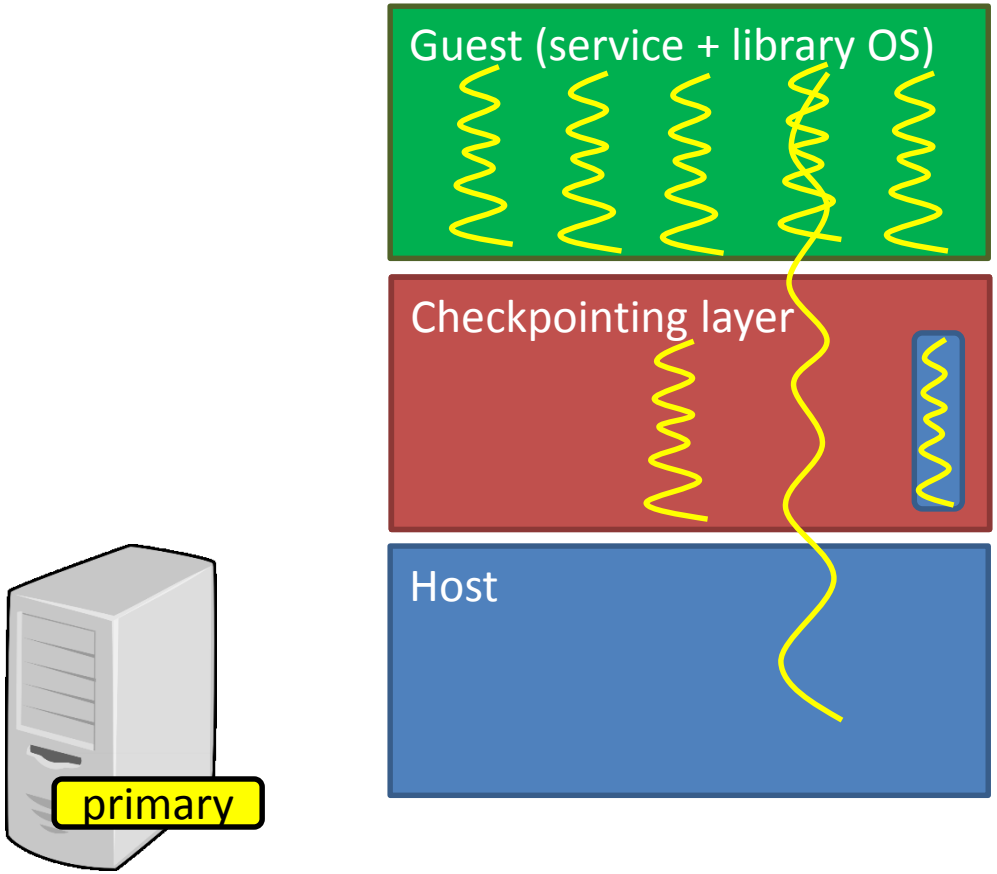

API non-determinism undermines replay

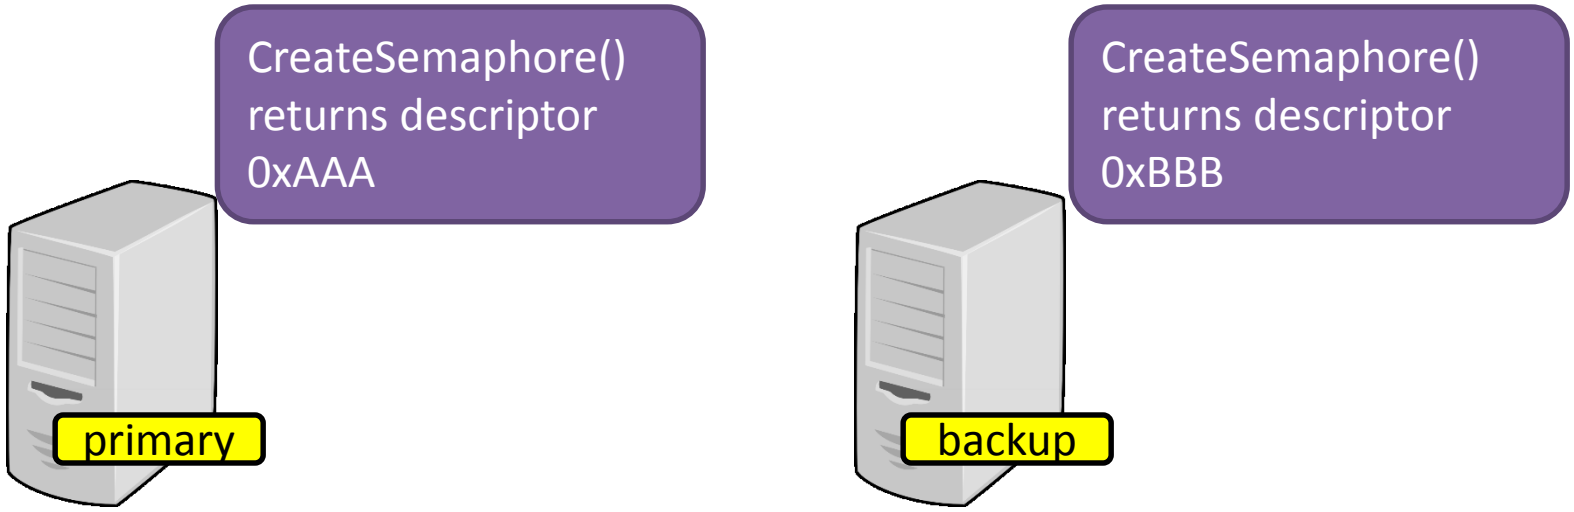

An indirection table can hide non-determinism

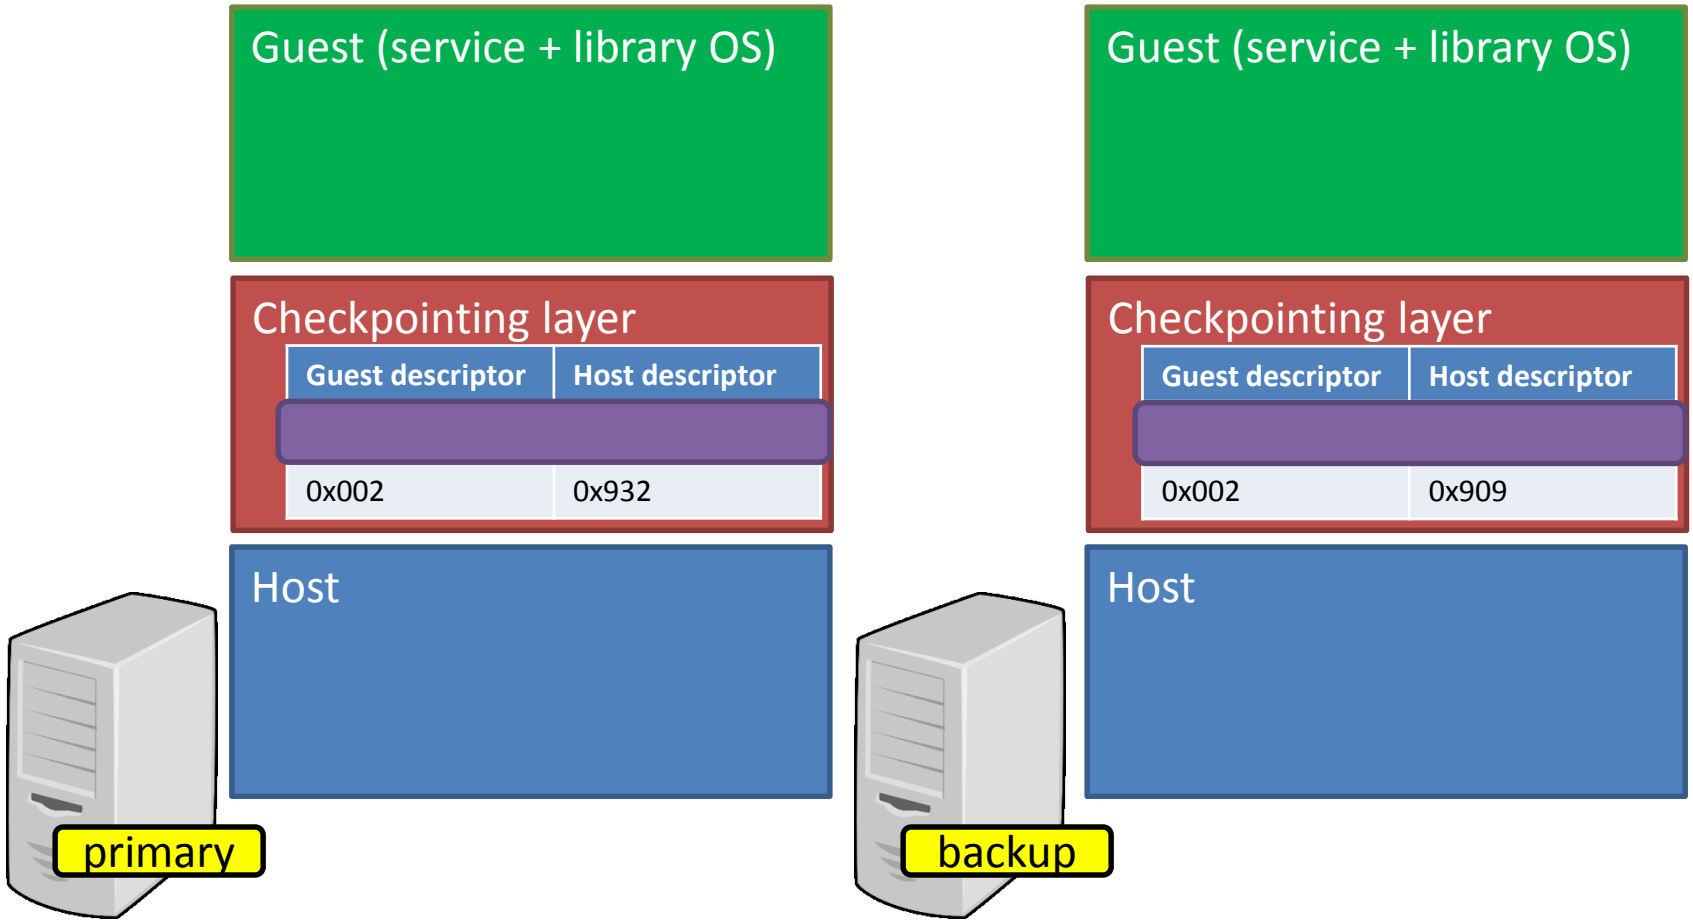

State external to guest needs to be replayable or regeneratable

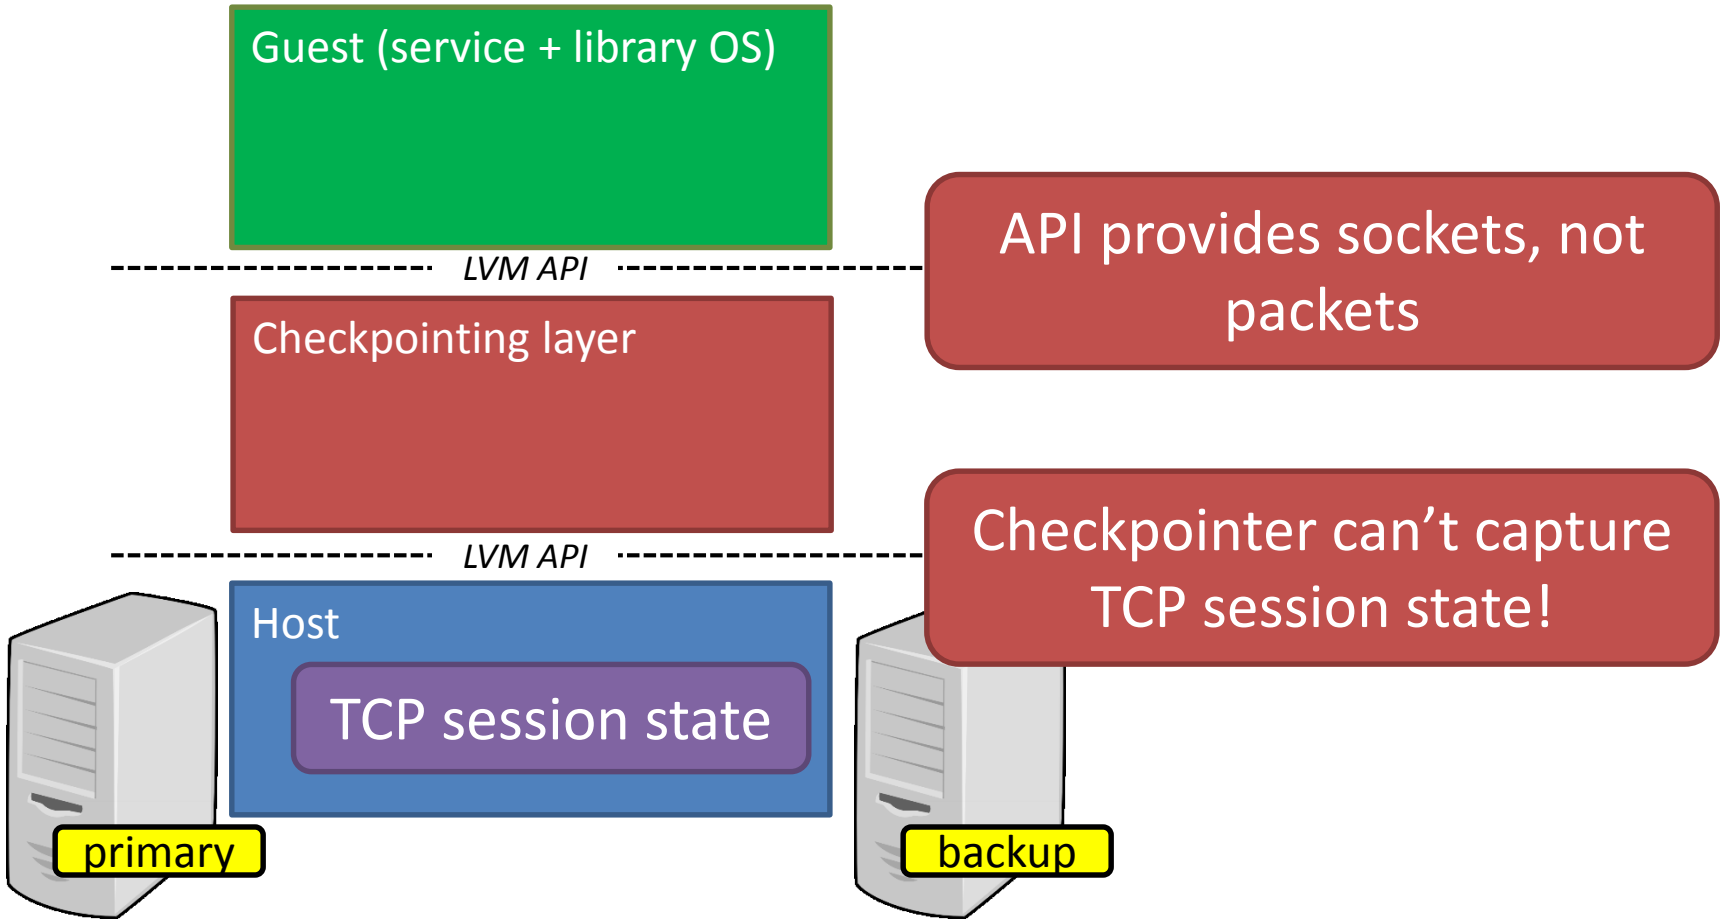

System-specific modifications may be necessary

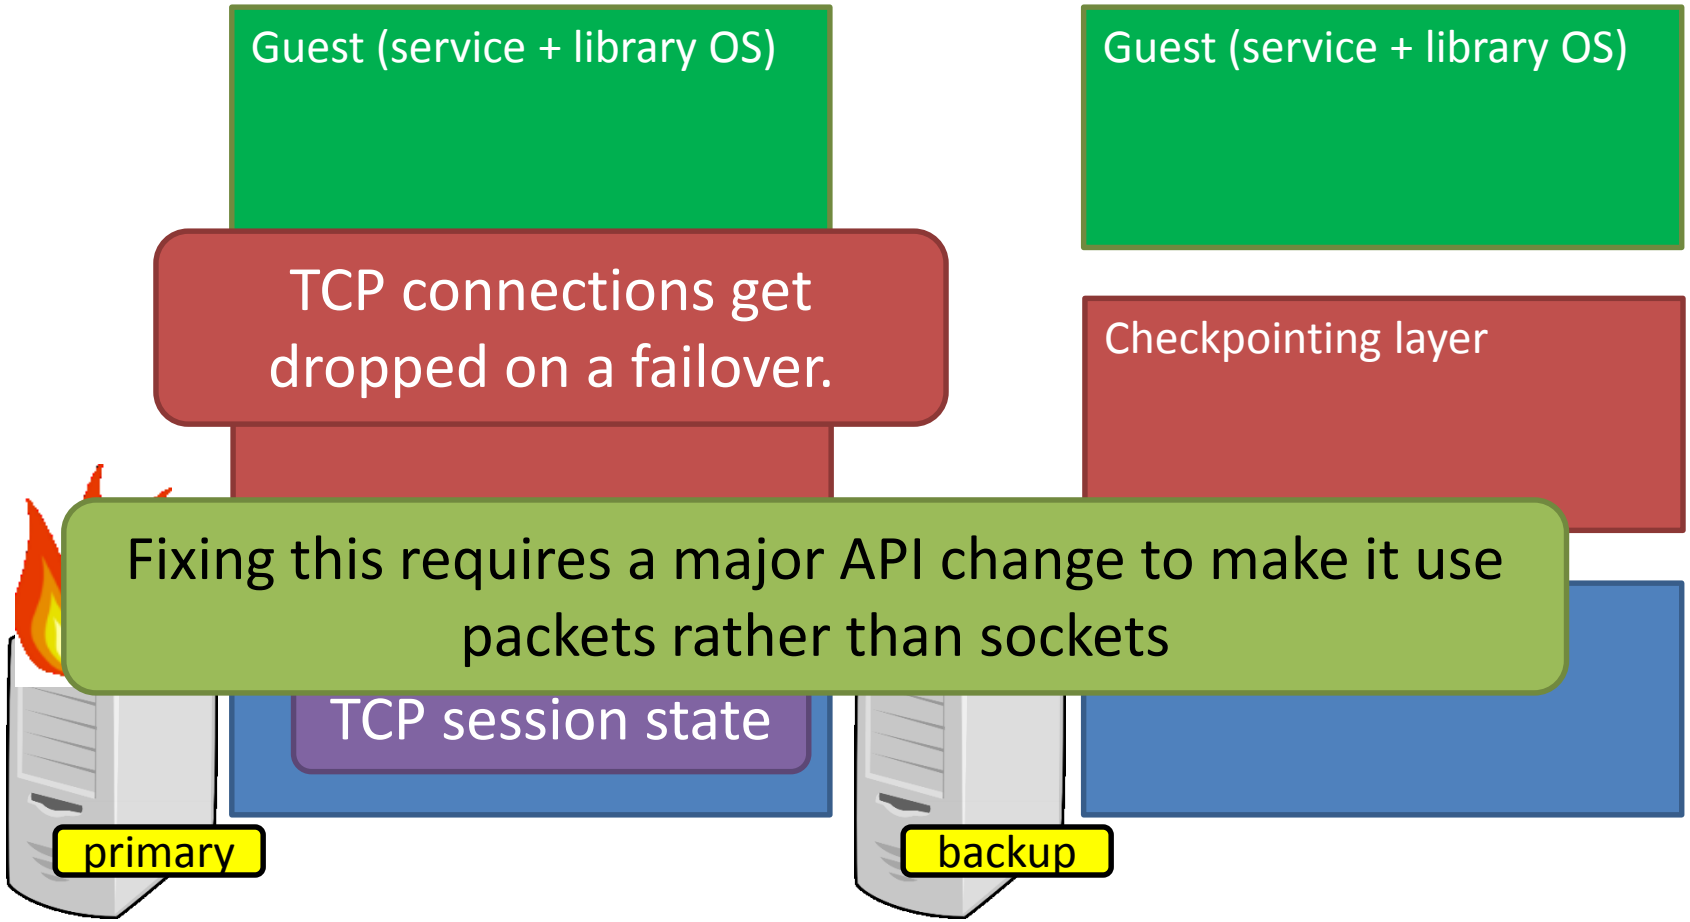

Outline

- Motivation
- Background: Asynchronous VM replication
- Our solution: Lightweight VM replication
- Challenges and solutions
- Evaluation

Effect of external processes - Remus

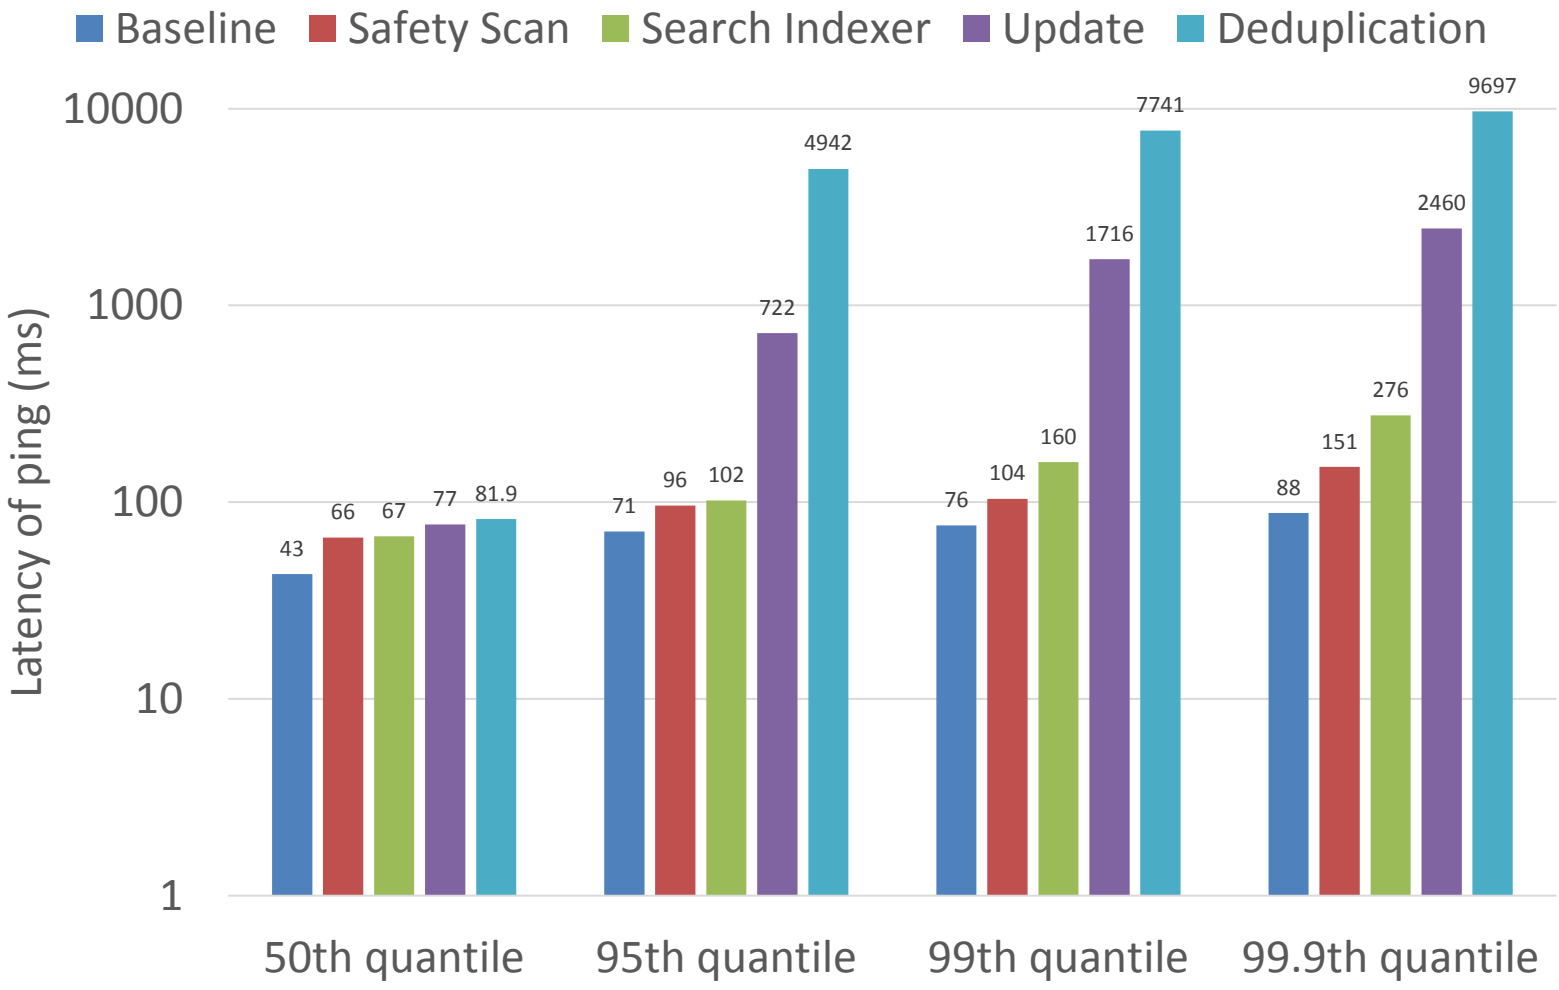

Effect of external processes - Tardigrade

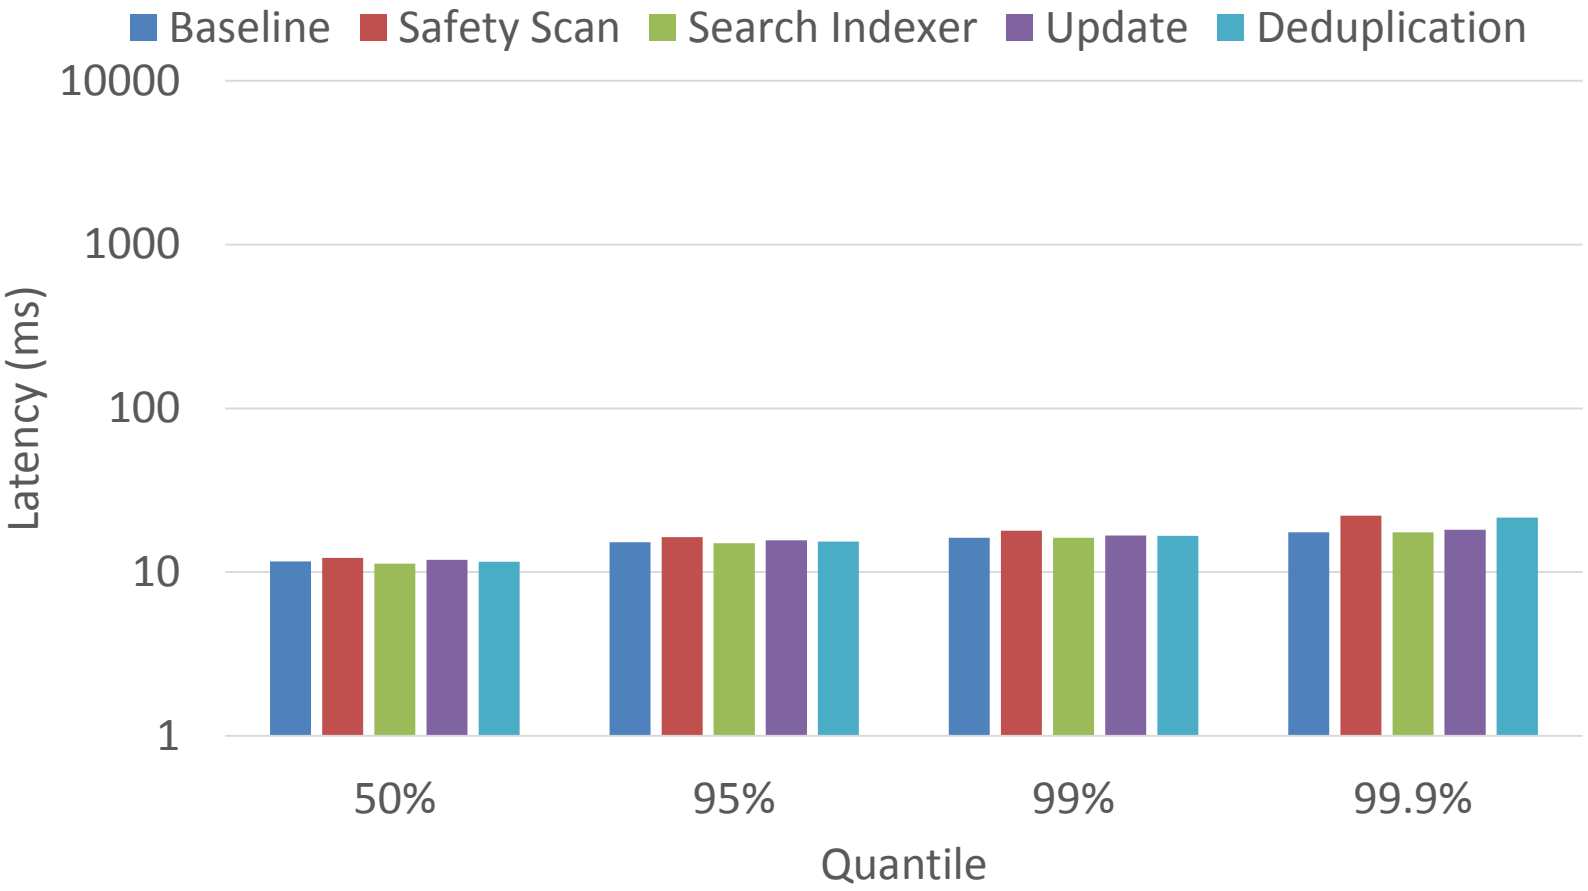

Effect of external processes - Tardigrade

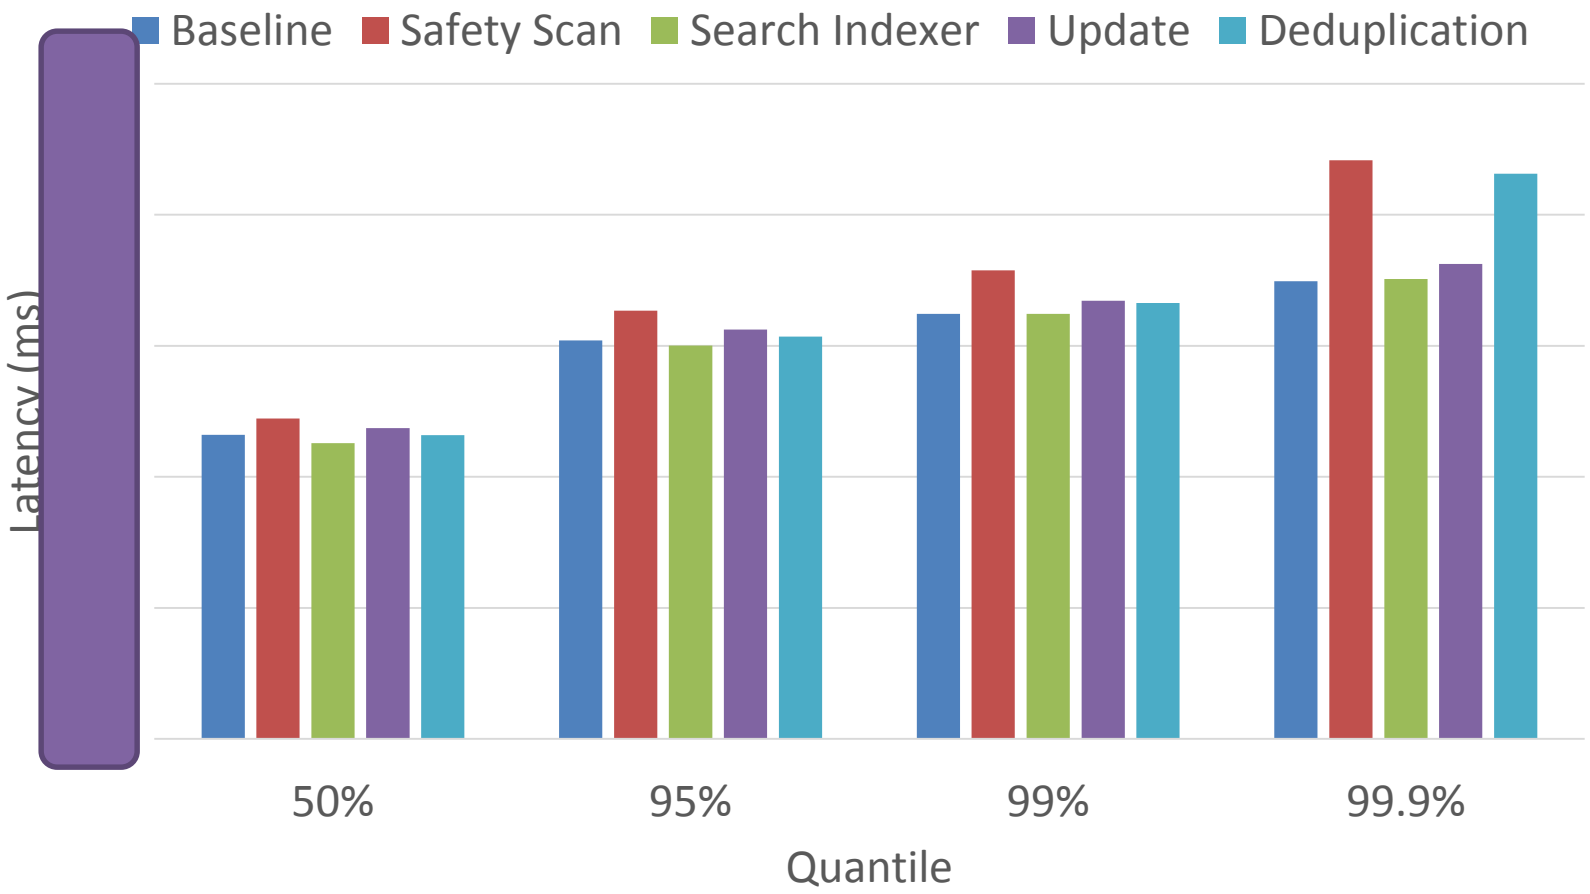

Memory dirtying affects checkpoint latency

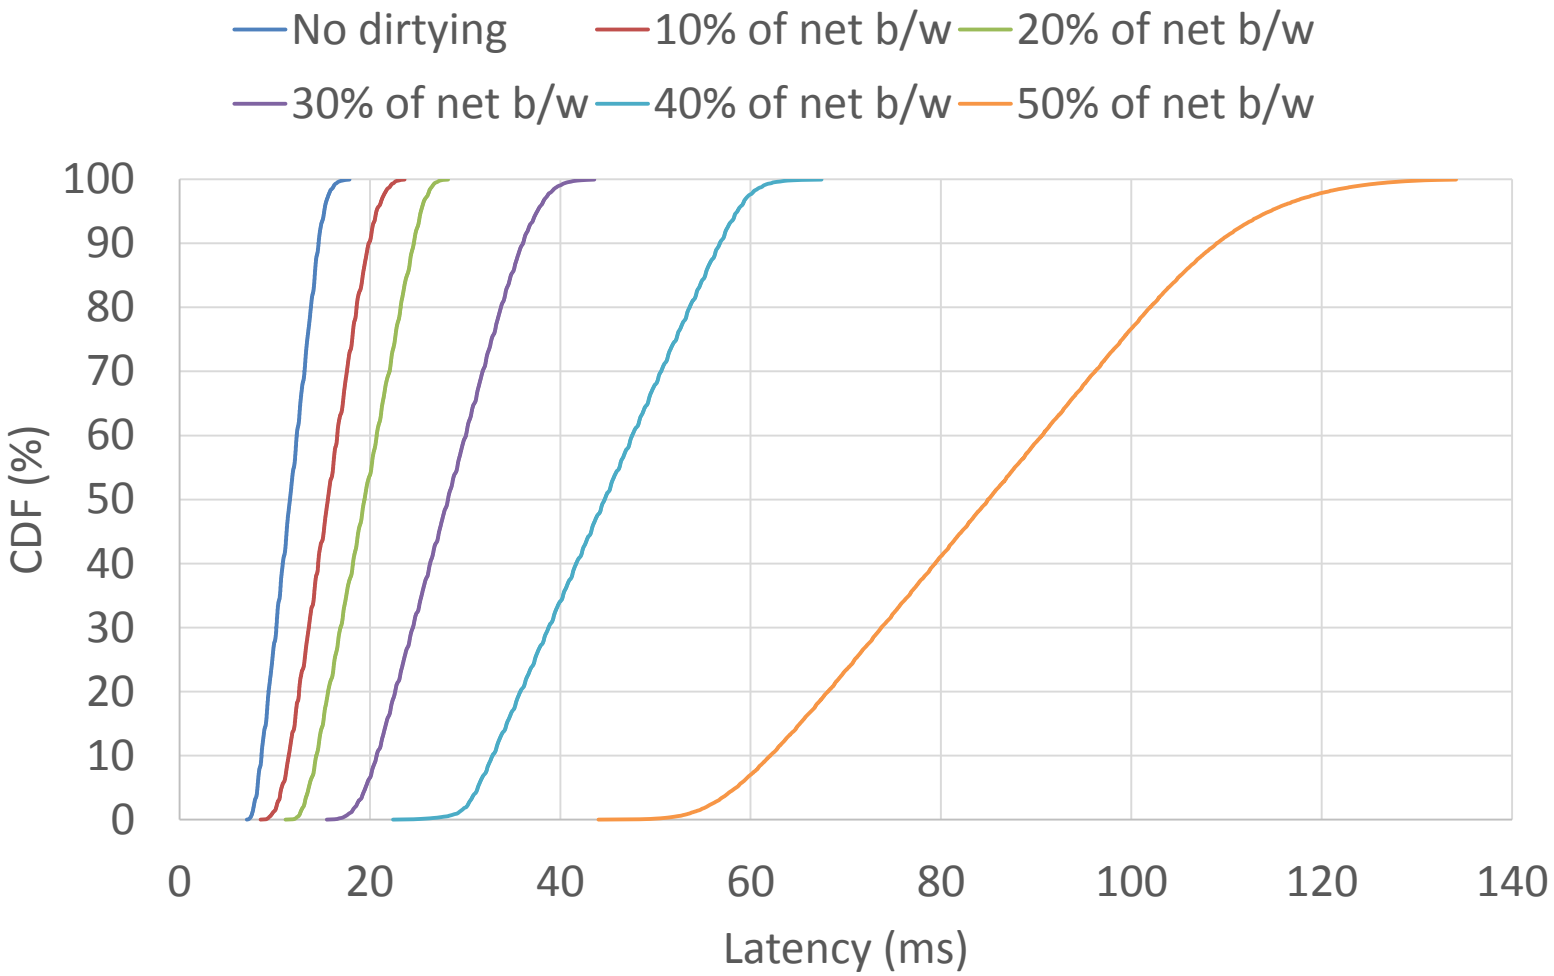

FDS metadata service

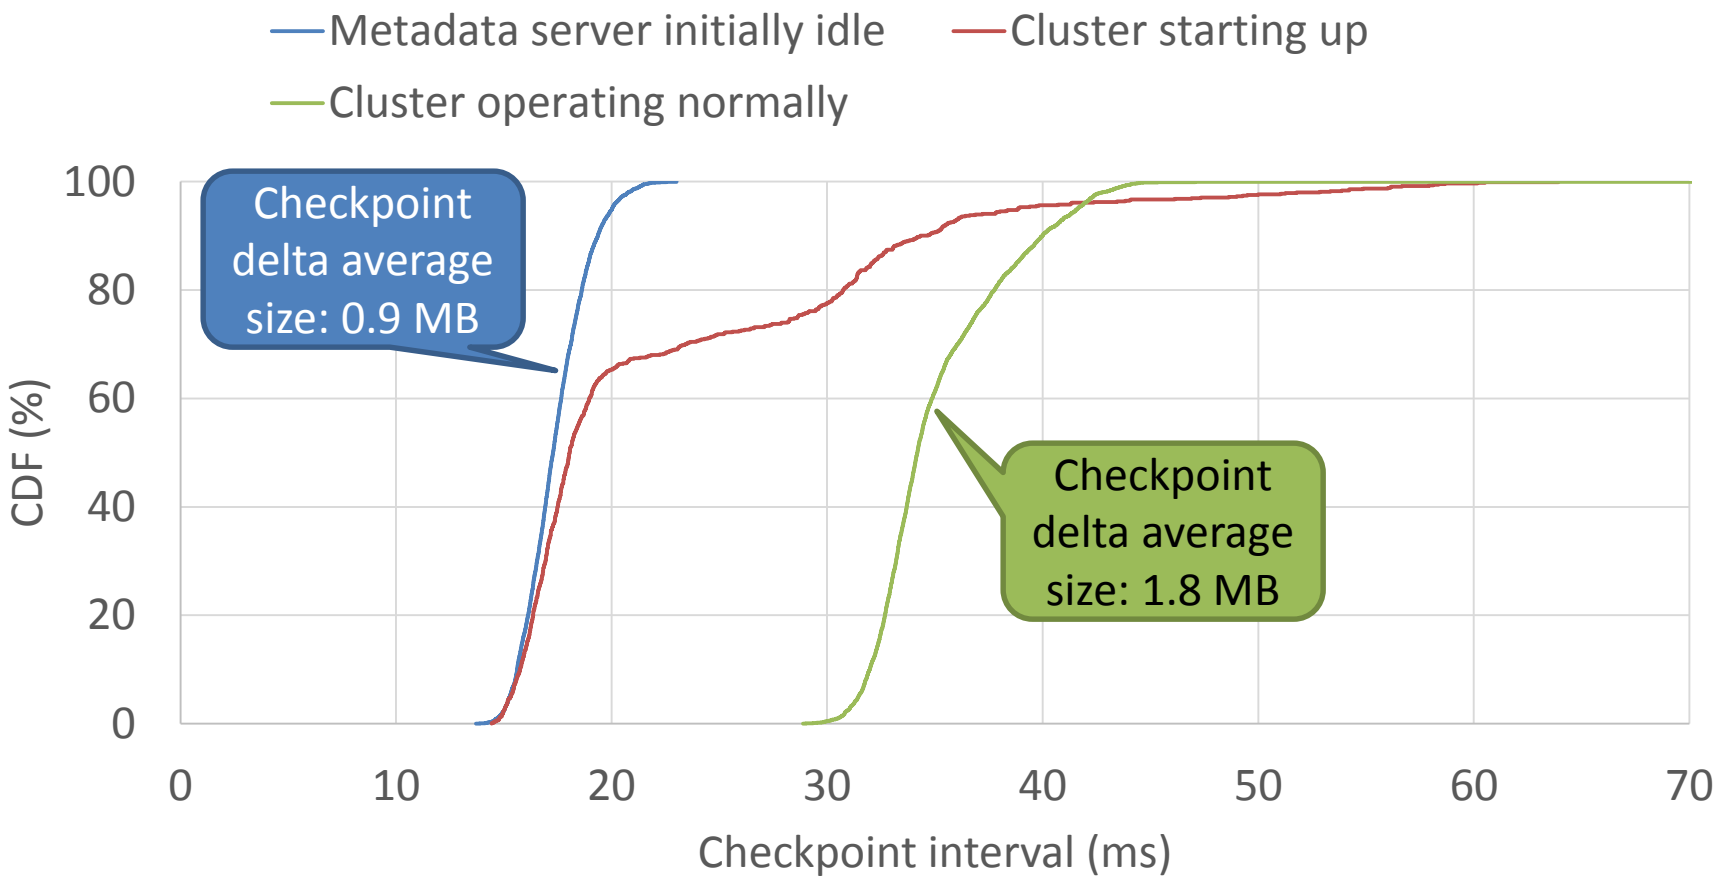

ZKLite, a simple non-fault-tolerant Java implementation of the Zookeeper API

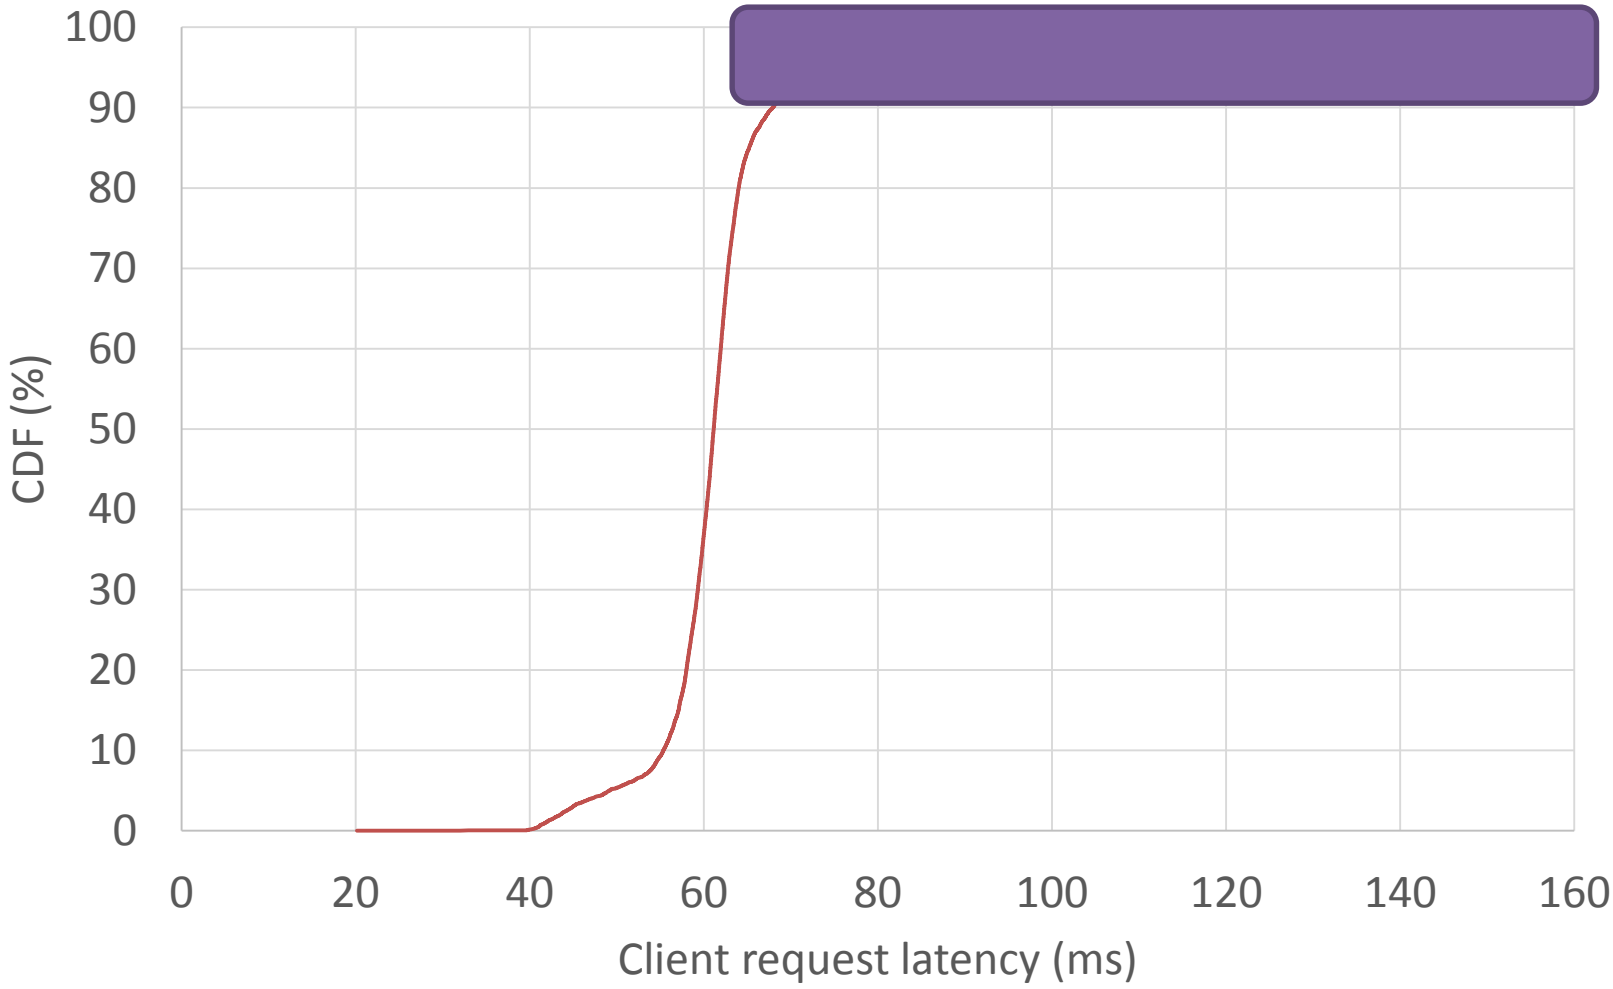

Conclusions

No changes to binaries needed, making deployment simple

Re

Li

gets:

s

es

s

Lightweight VM replication is practical for making existing service binaries fault-tolerant

Reasonable performance if memory dirtying rate and load are low