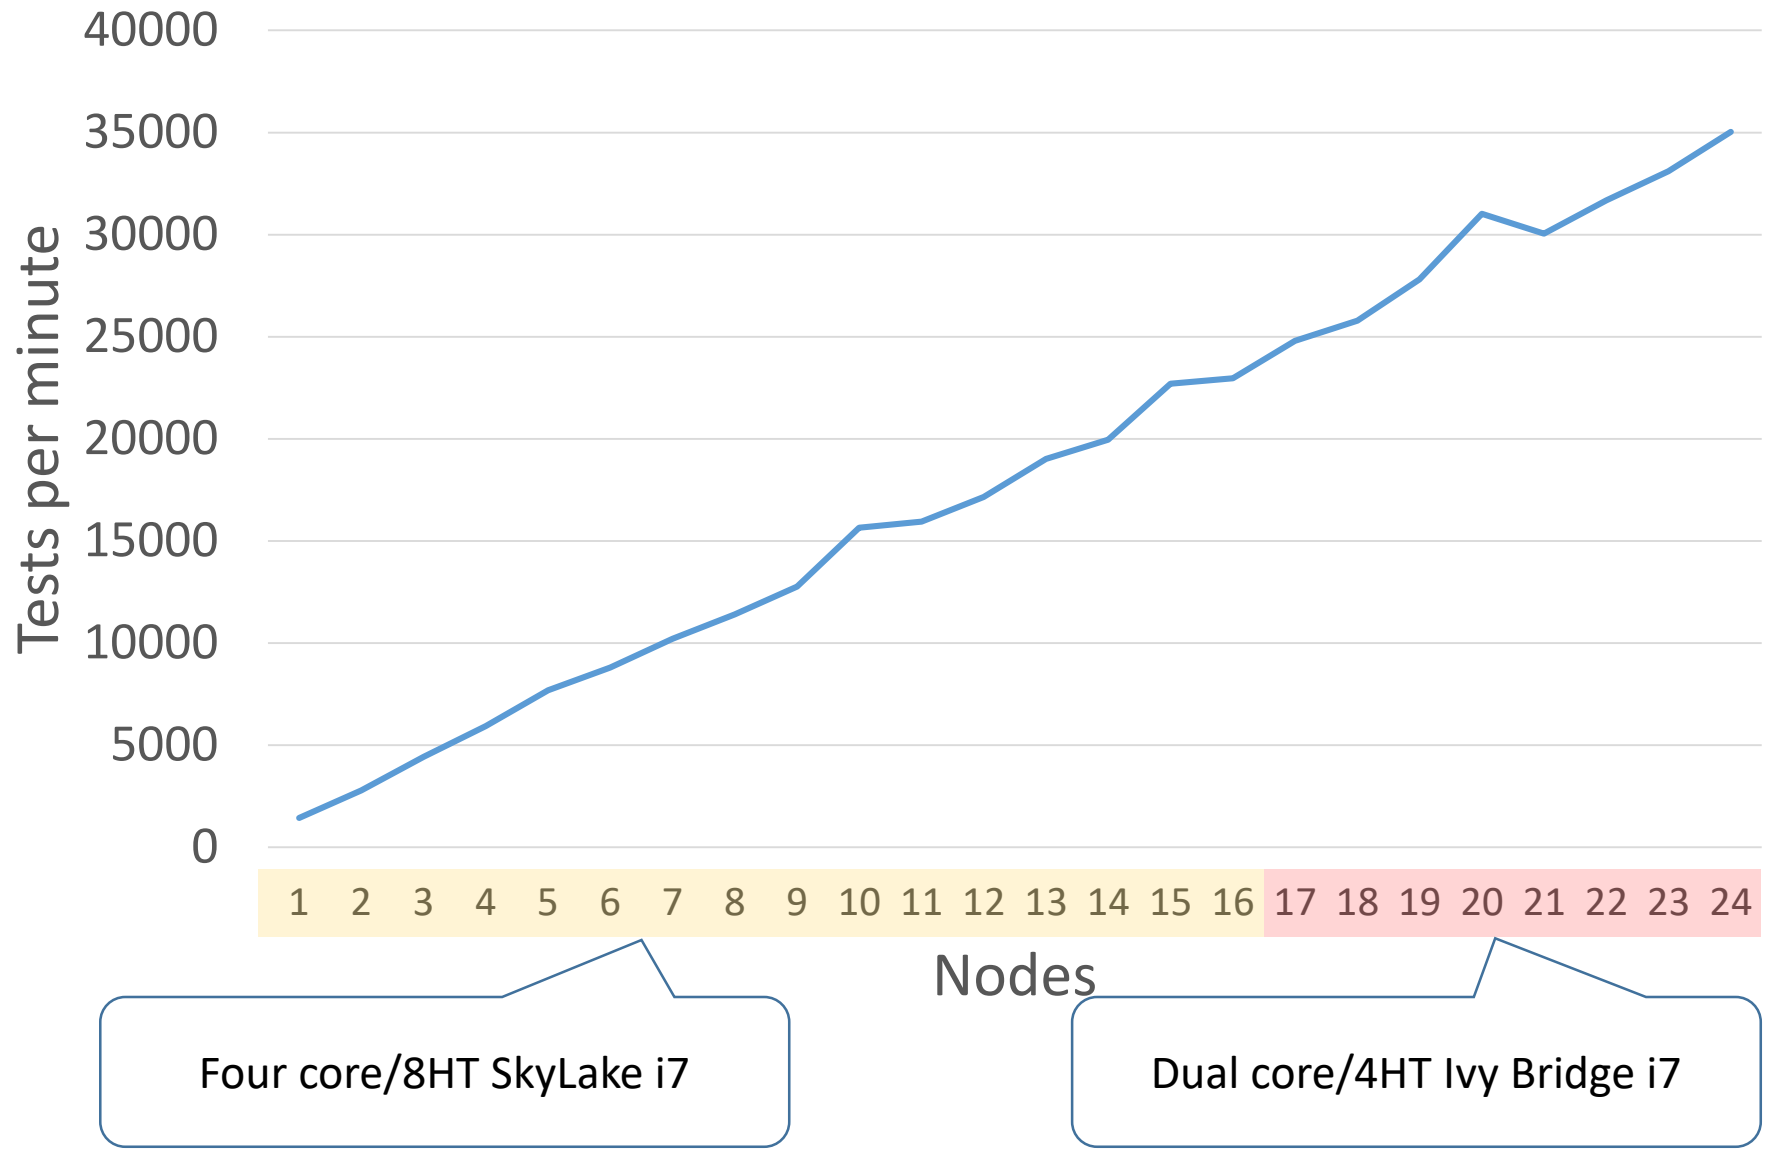


# QuickCheck

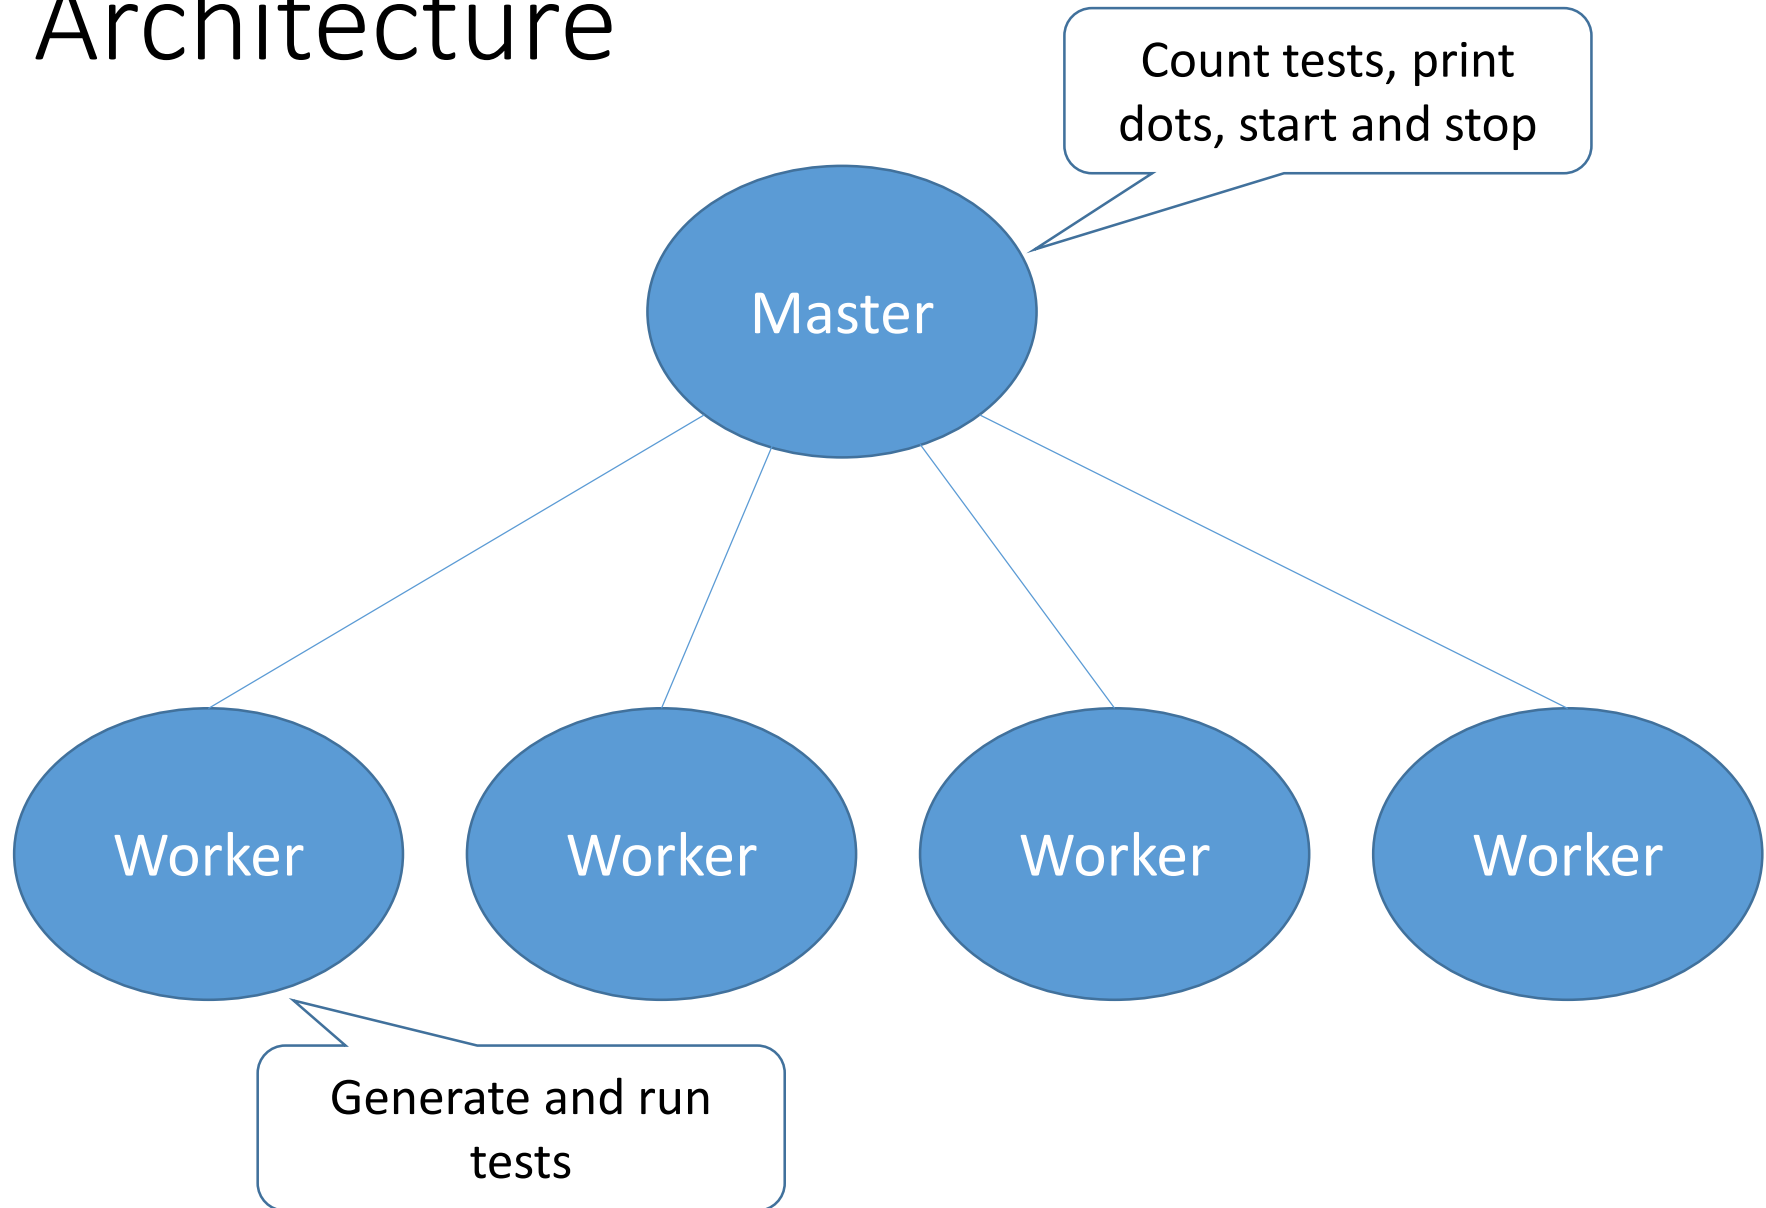
John Hughes

DEMO


# Registry tests on multiple nodes




# Architecture




# How can the Master count the tests?



# How can the Master stop the workers?









# Benchmark



1.2 million/second

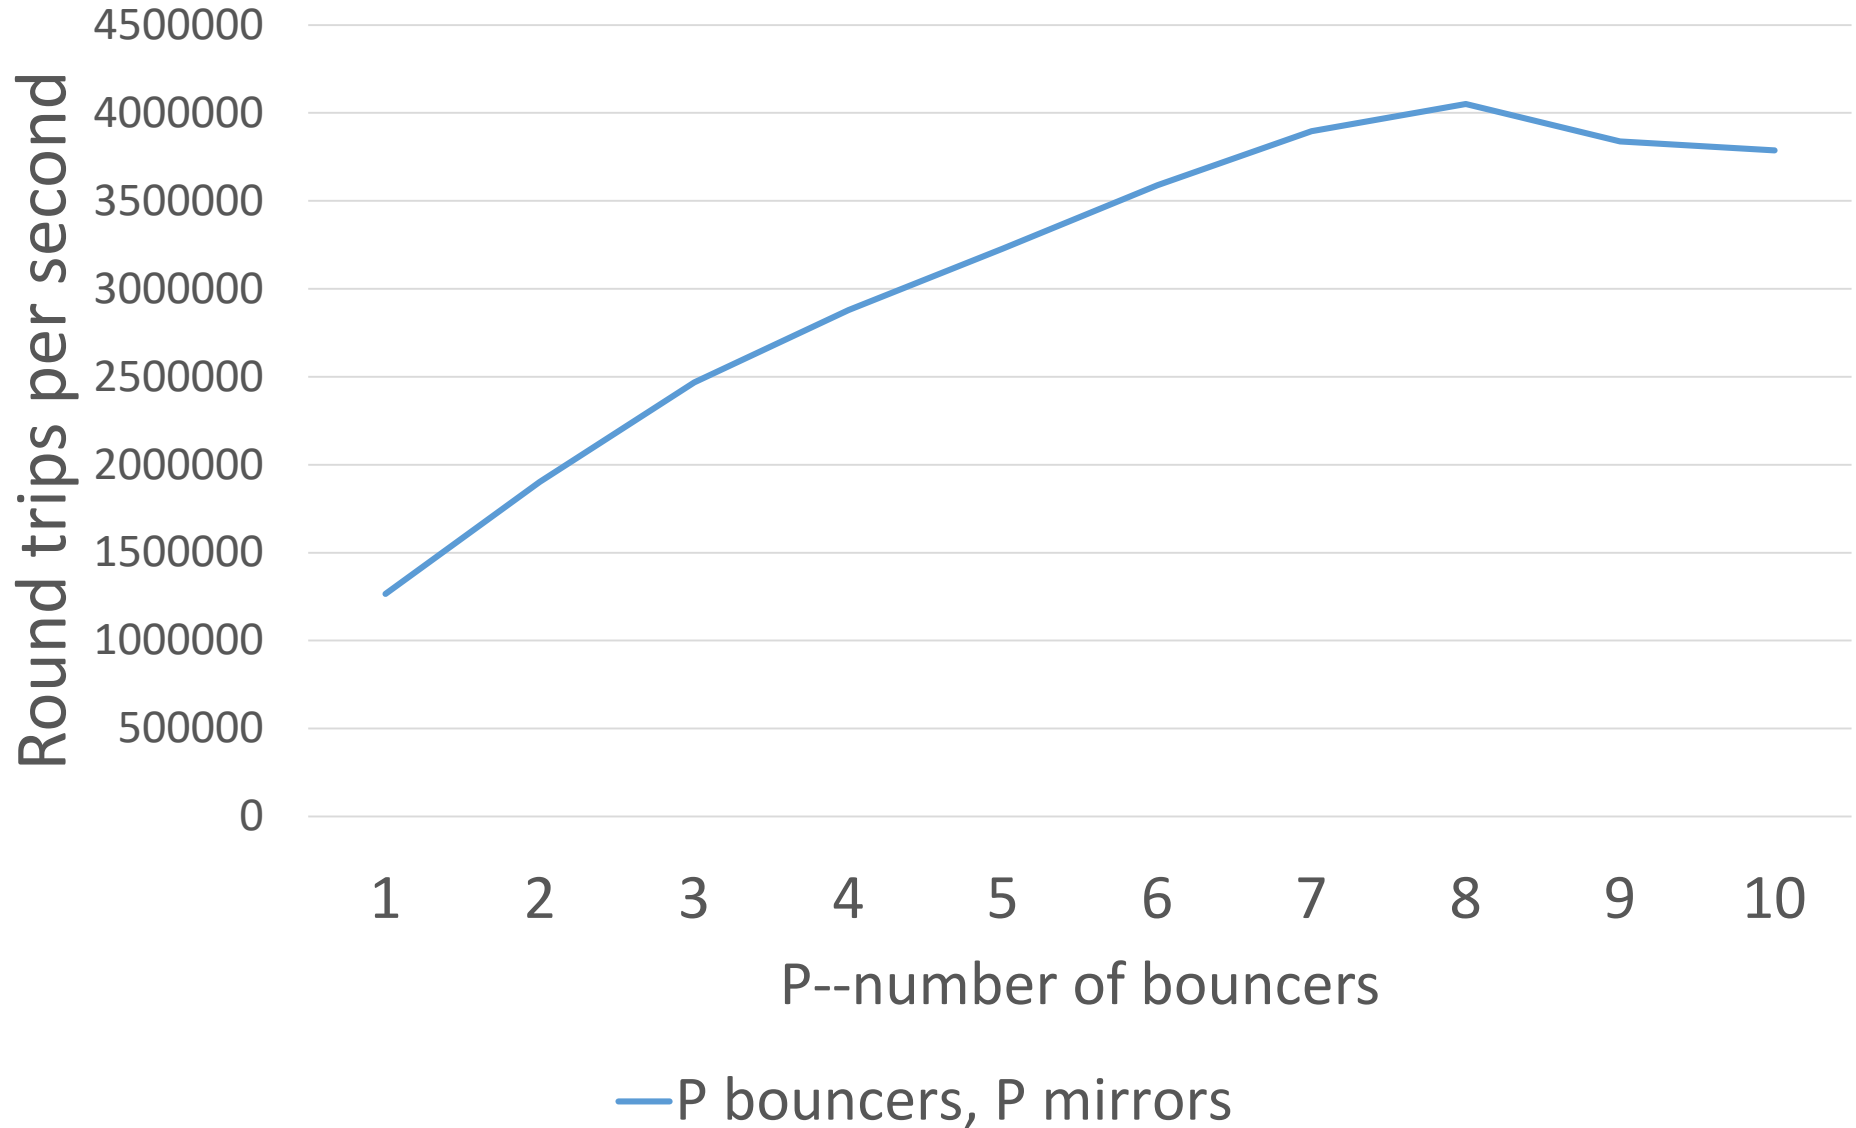
# Multiple bouncers



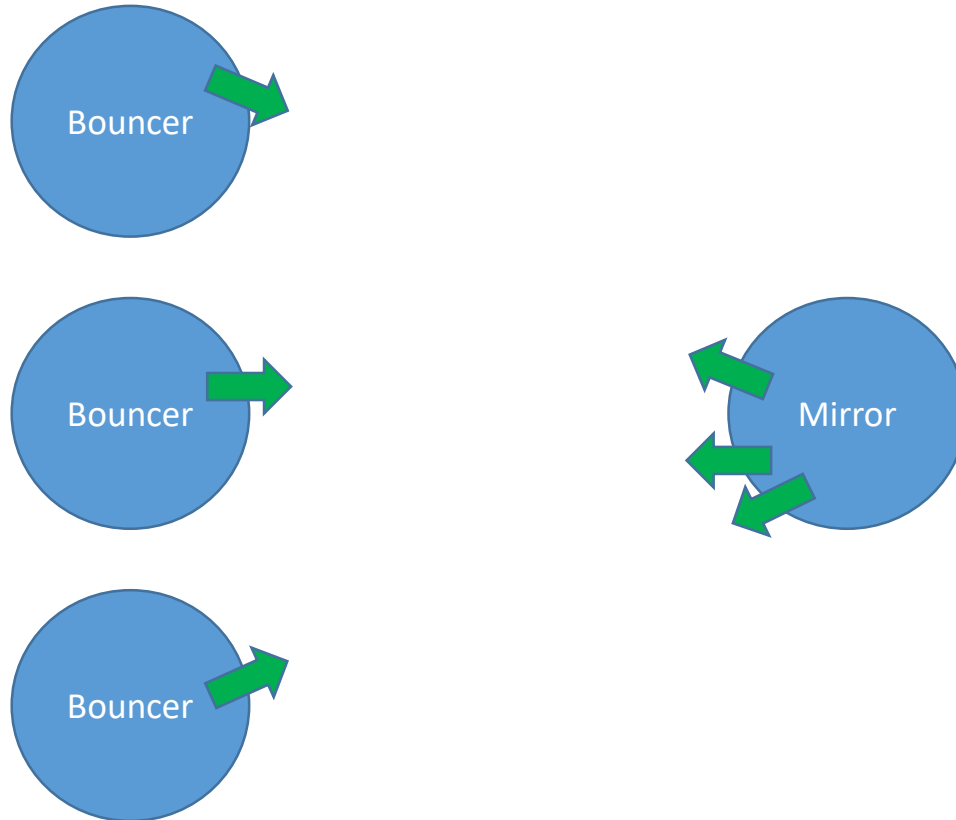


# Message round-trips per second

## Four core/8 thread i7




# Multiple bouncers, one mirror




# Message round-trips per second

## Four core/8 thread i7




# Batch permissions?



BUT this may delay termination!

# Alternative Architecture





- Every worker communicates with its own manager—scalable!




- Stopping can be slightly delayed

```
> eqc:quickcheck(examples:prop_reverse()).
.....
OK, passed 100 tests
true
> eqc:quickcheck(examples:prop_reverse()).
.....(x10).(x1).....
OK, passed 119 tests
true
> eqc:quickcheck(eqc:in_parallel(examples:prop_reverse())).
.....(x10)..(x1).....
OK, passed 126 tests
true
> eqc:quickcheck(eqc:on_nodes(examples:prop_reverse())).
.....(x10).....
.....(x1).....
OK, passed 669 tests
```


But how bad is it to run a few extra tests?

# What about node placement?

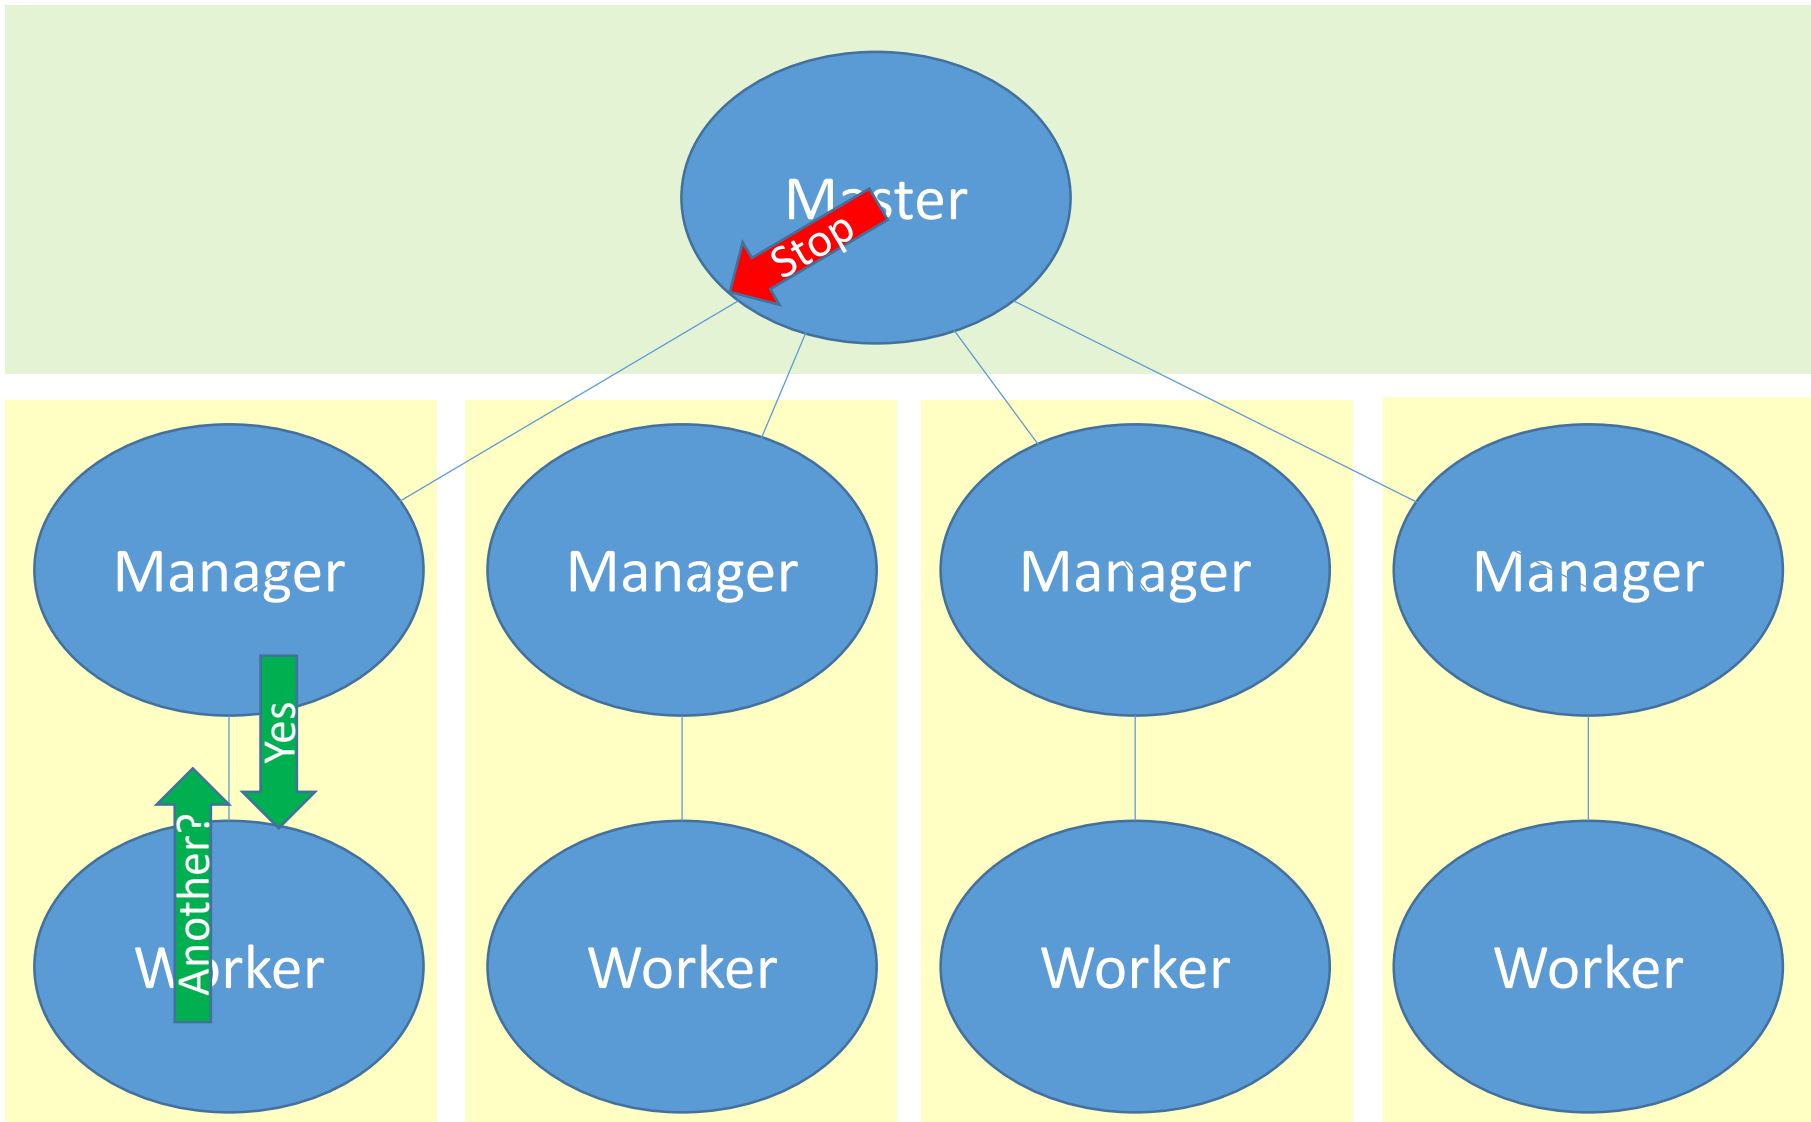




# What about node placement?




# What about node placement?



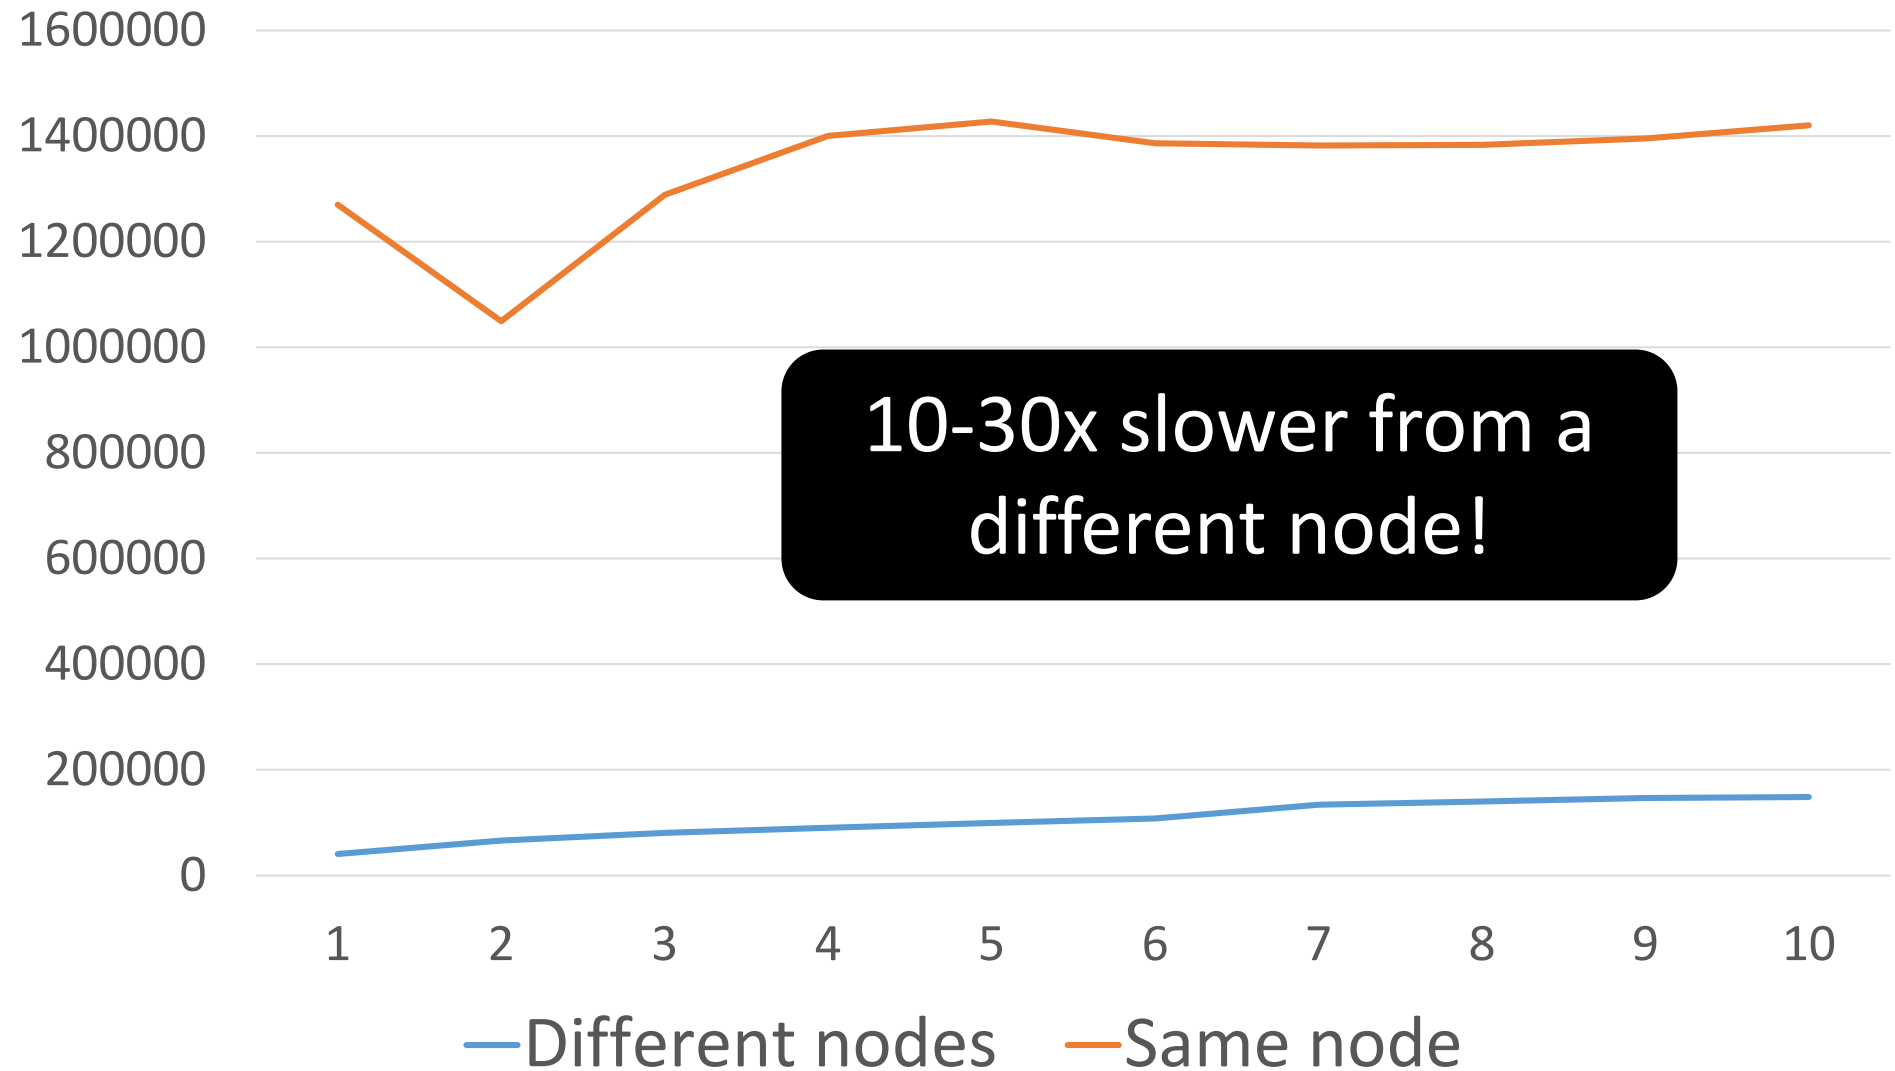
# P bouncers, 1 mirror, different nodes

## Four core/8 HT i7




# P bouncers, 1 mirror

## Four core/8HT i7

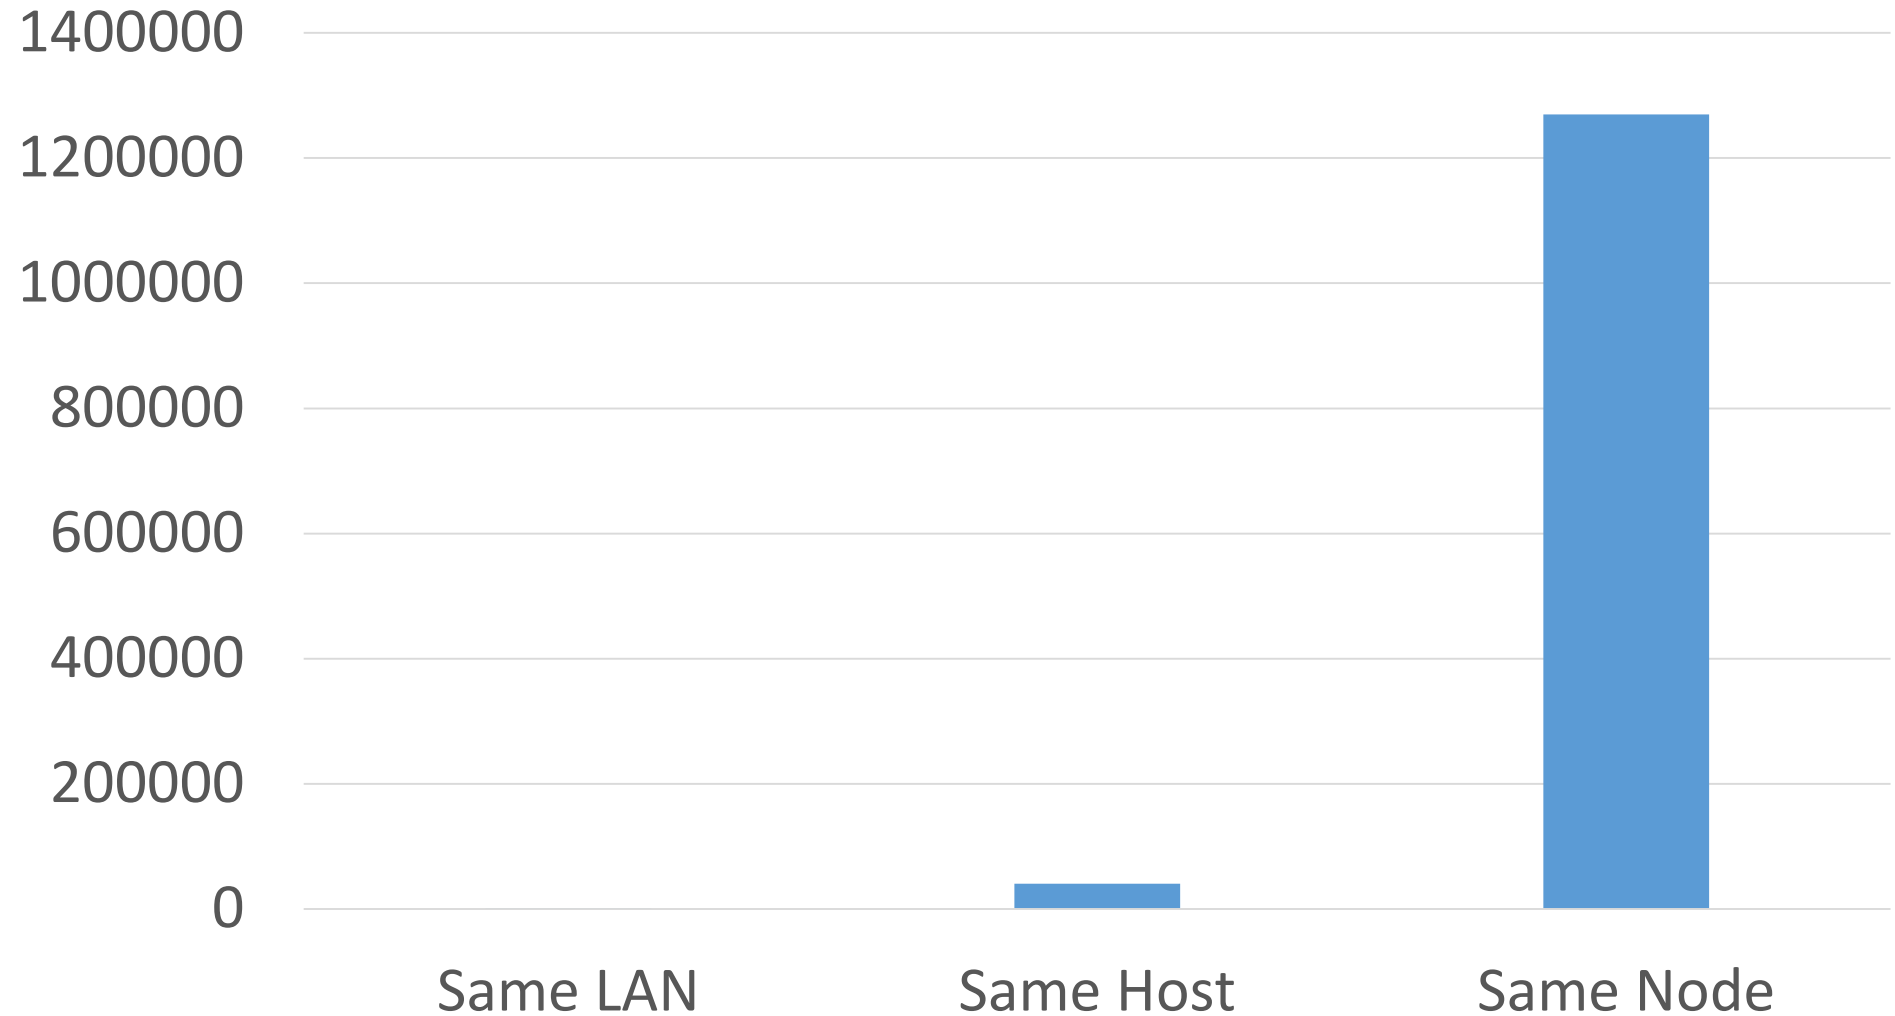


# P bouncers, 1 mirror

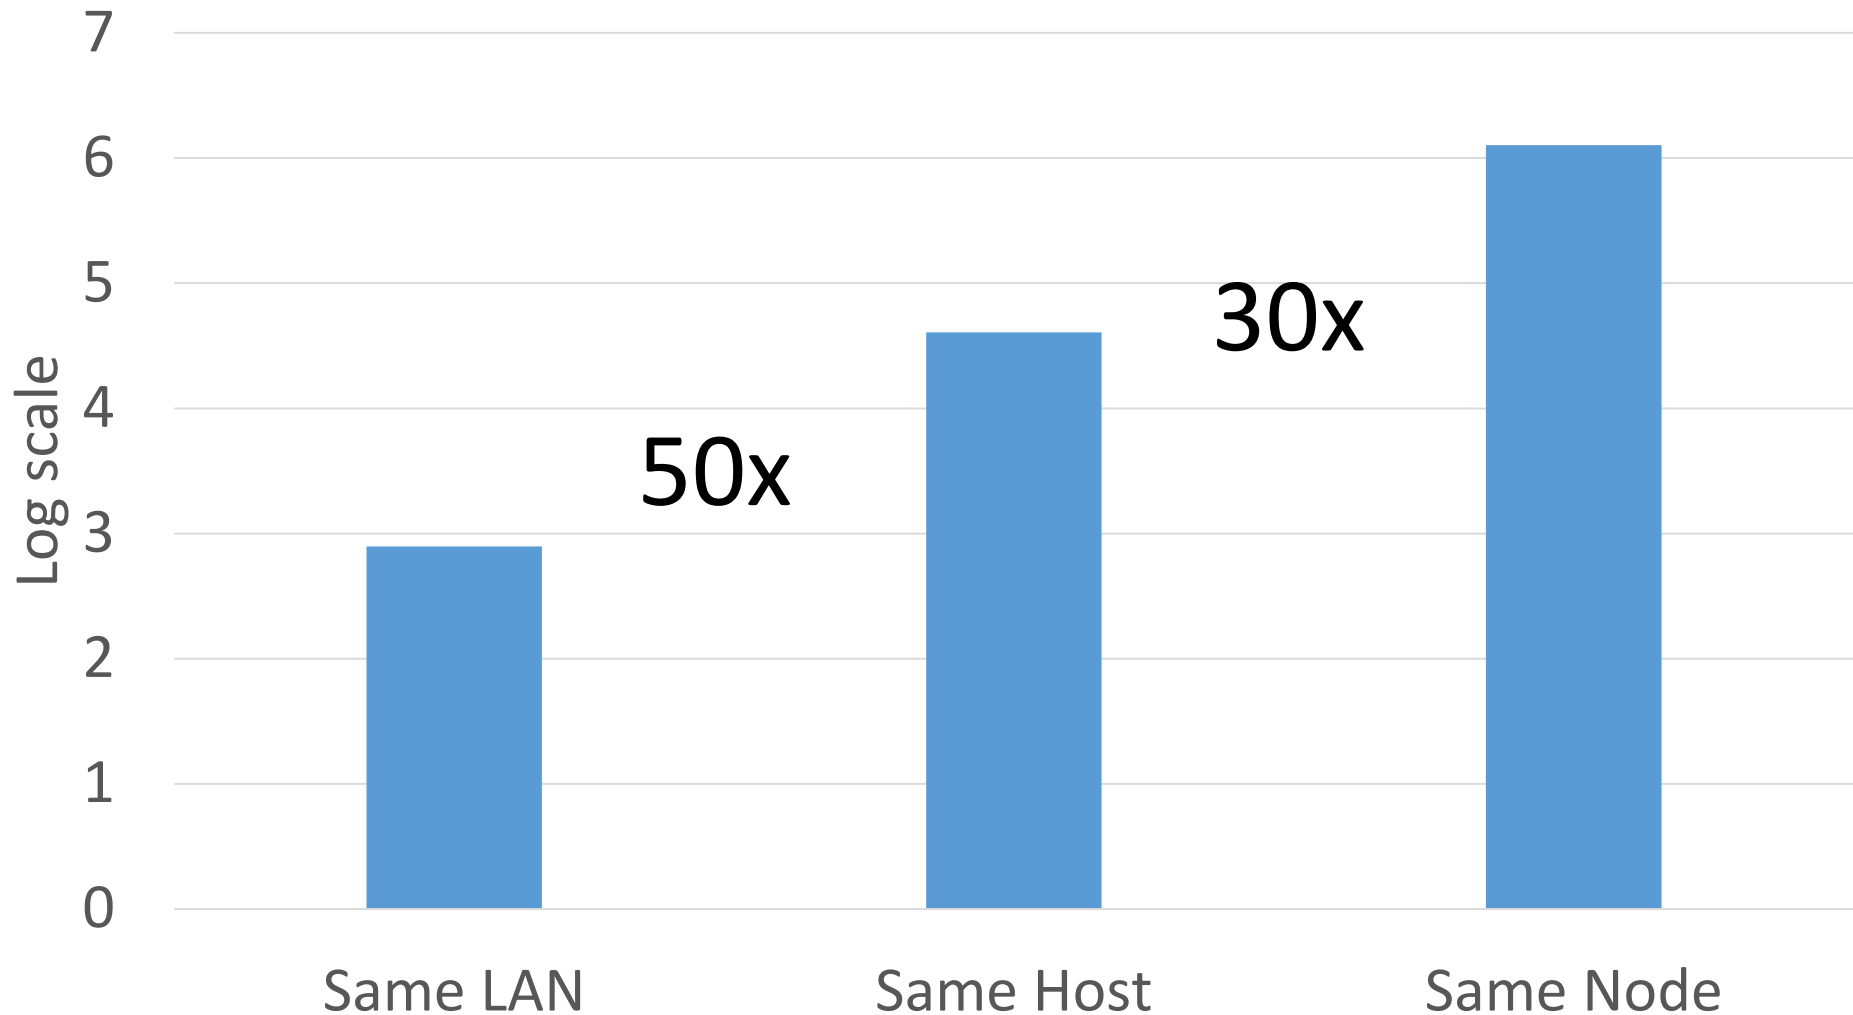
## Bouncers on dual core laptop, mirror on quad core




# Bounces per second (single bouncer)



# Bounces per second (single bouncer)




# What about success messages?





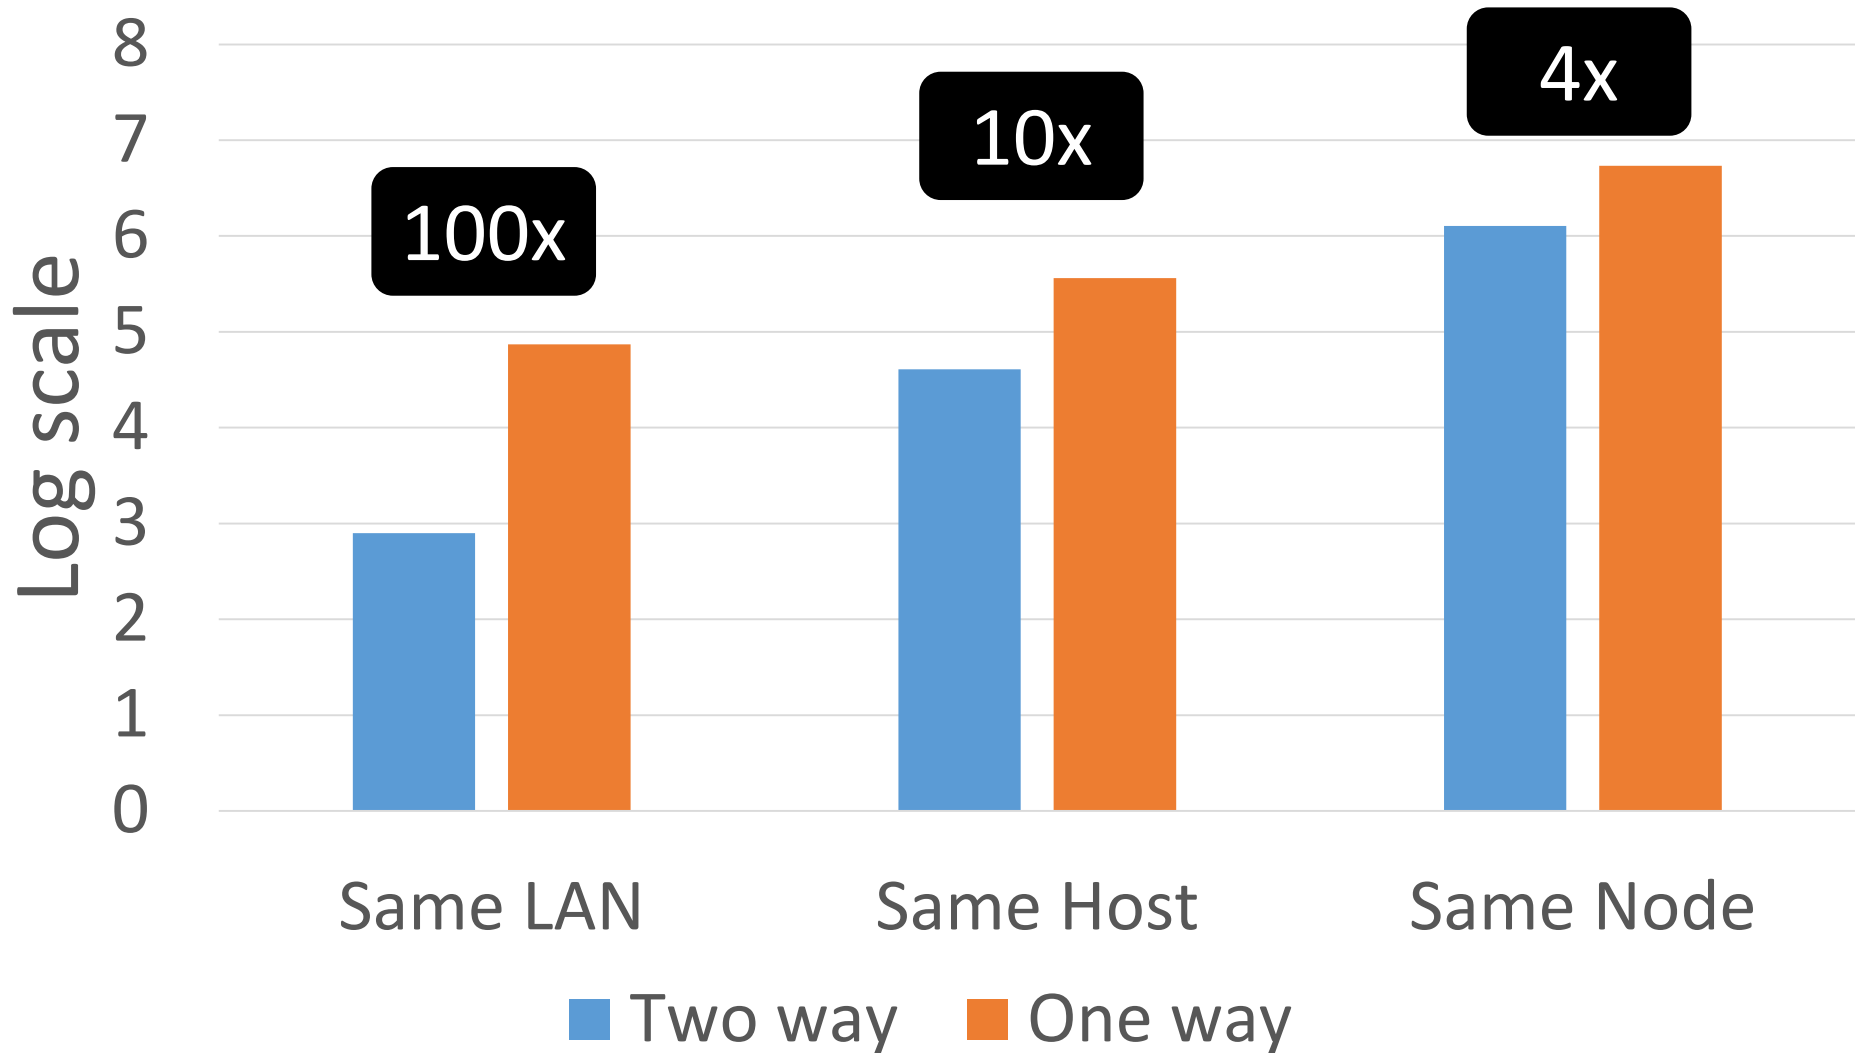
# Two-way vs one-way



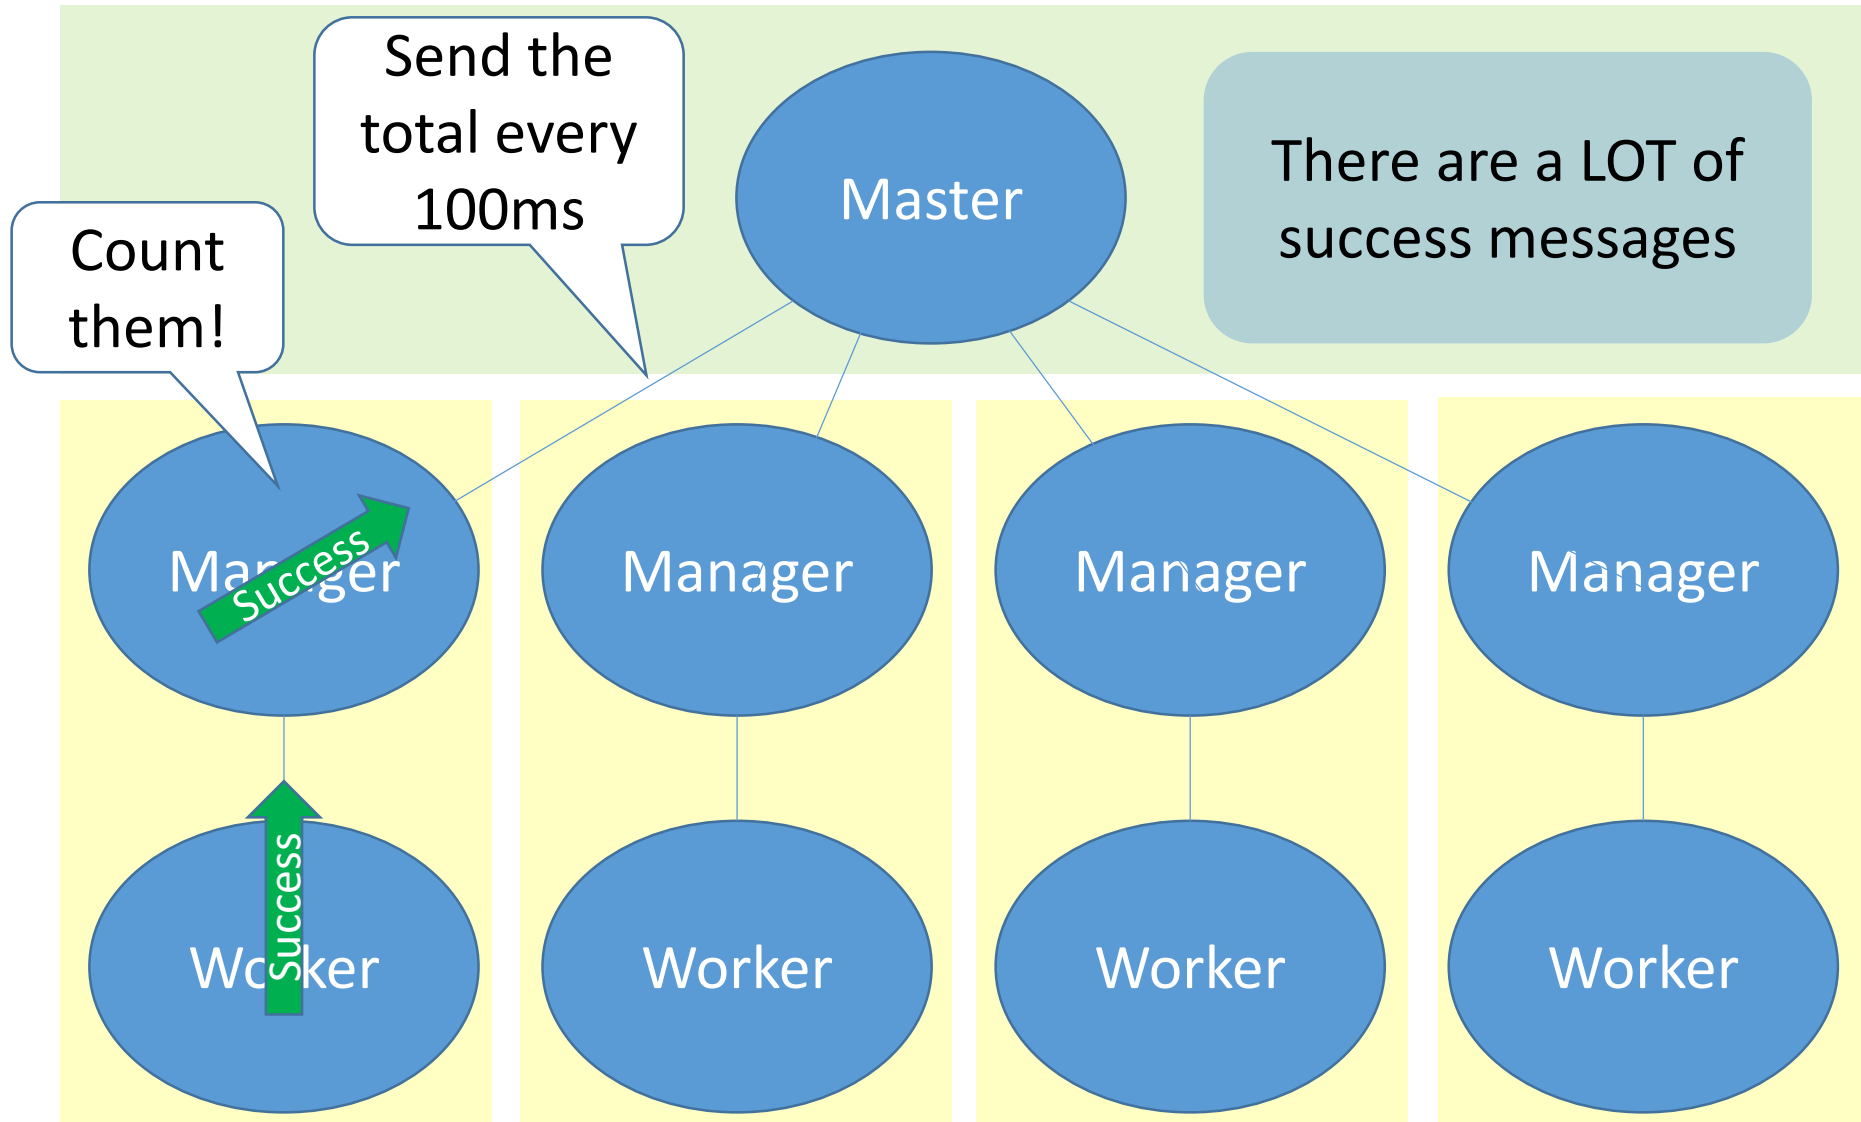
1.2 million/second

5.4 million/second

# Bounces/Messages per second



# One more optimization...





- *Tripled* the speed of quickcheck(true)!



- Stopping can be even more delayed

# Lessons

- There is at least *an order of magnitude* difference between communication costs
  - Within a node
  - Between nodes
  - Between hosts
- Latency is much worse affected than bandwidth
- This affects design for performance
  - Favours asynchronous over synchronous communication between nodes
- Optimising performance may require *changes to observable behaviour*
- ...and we didn't even consider fault tolerance