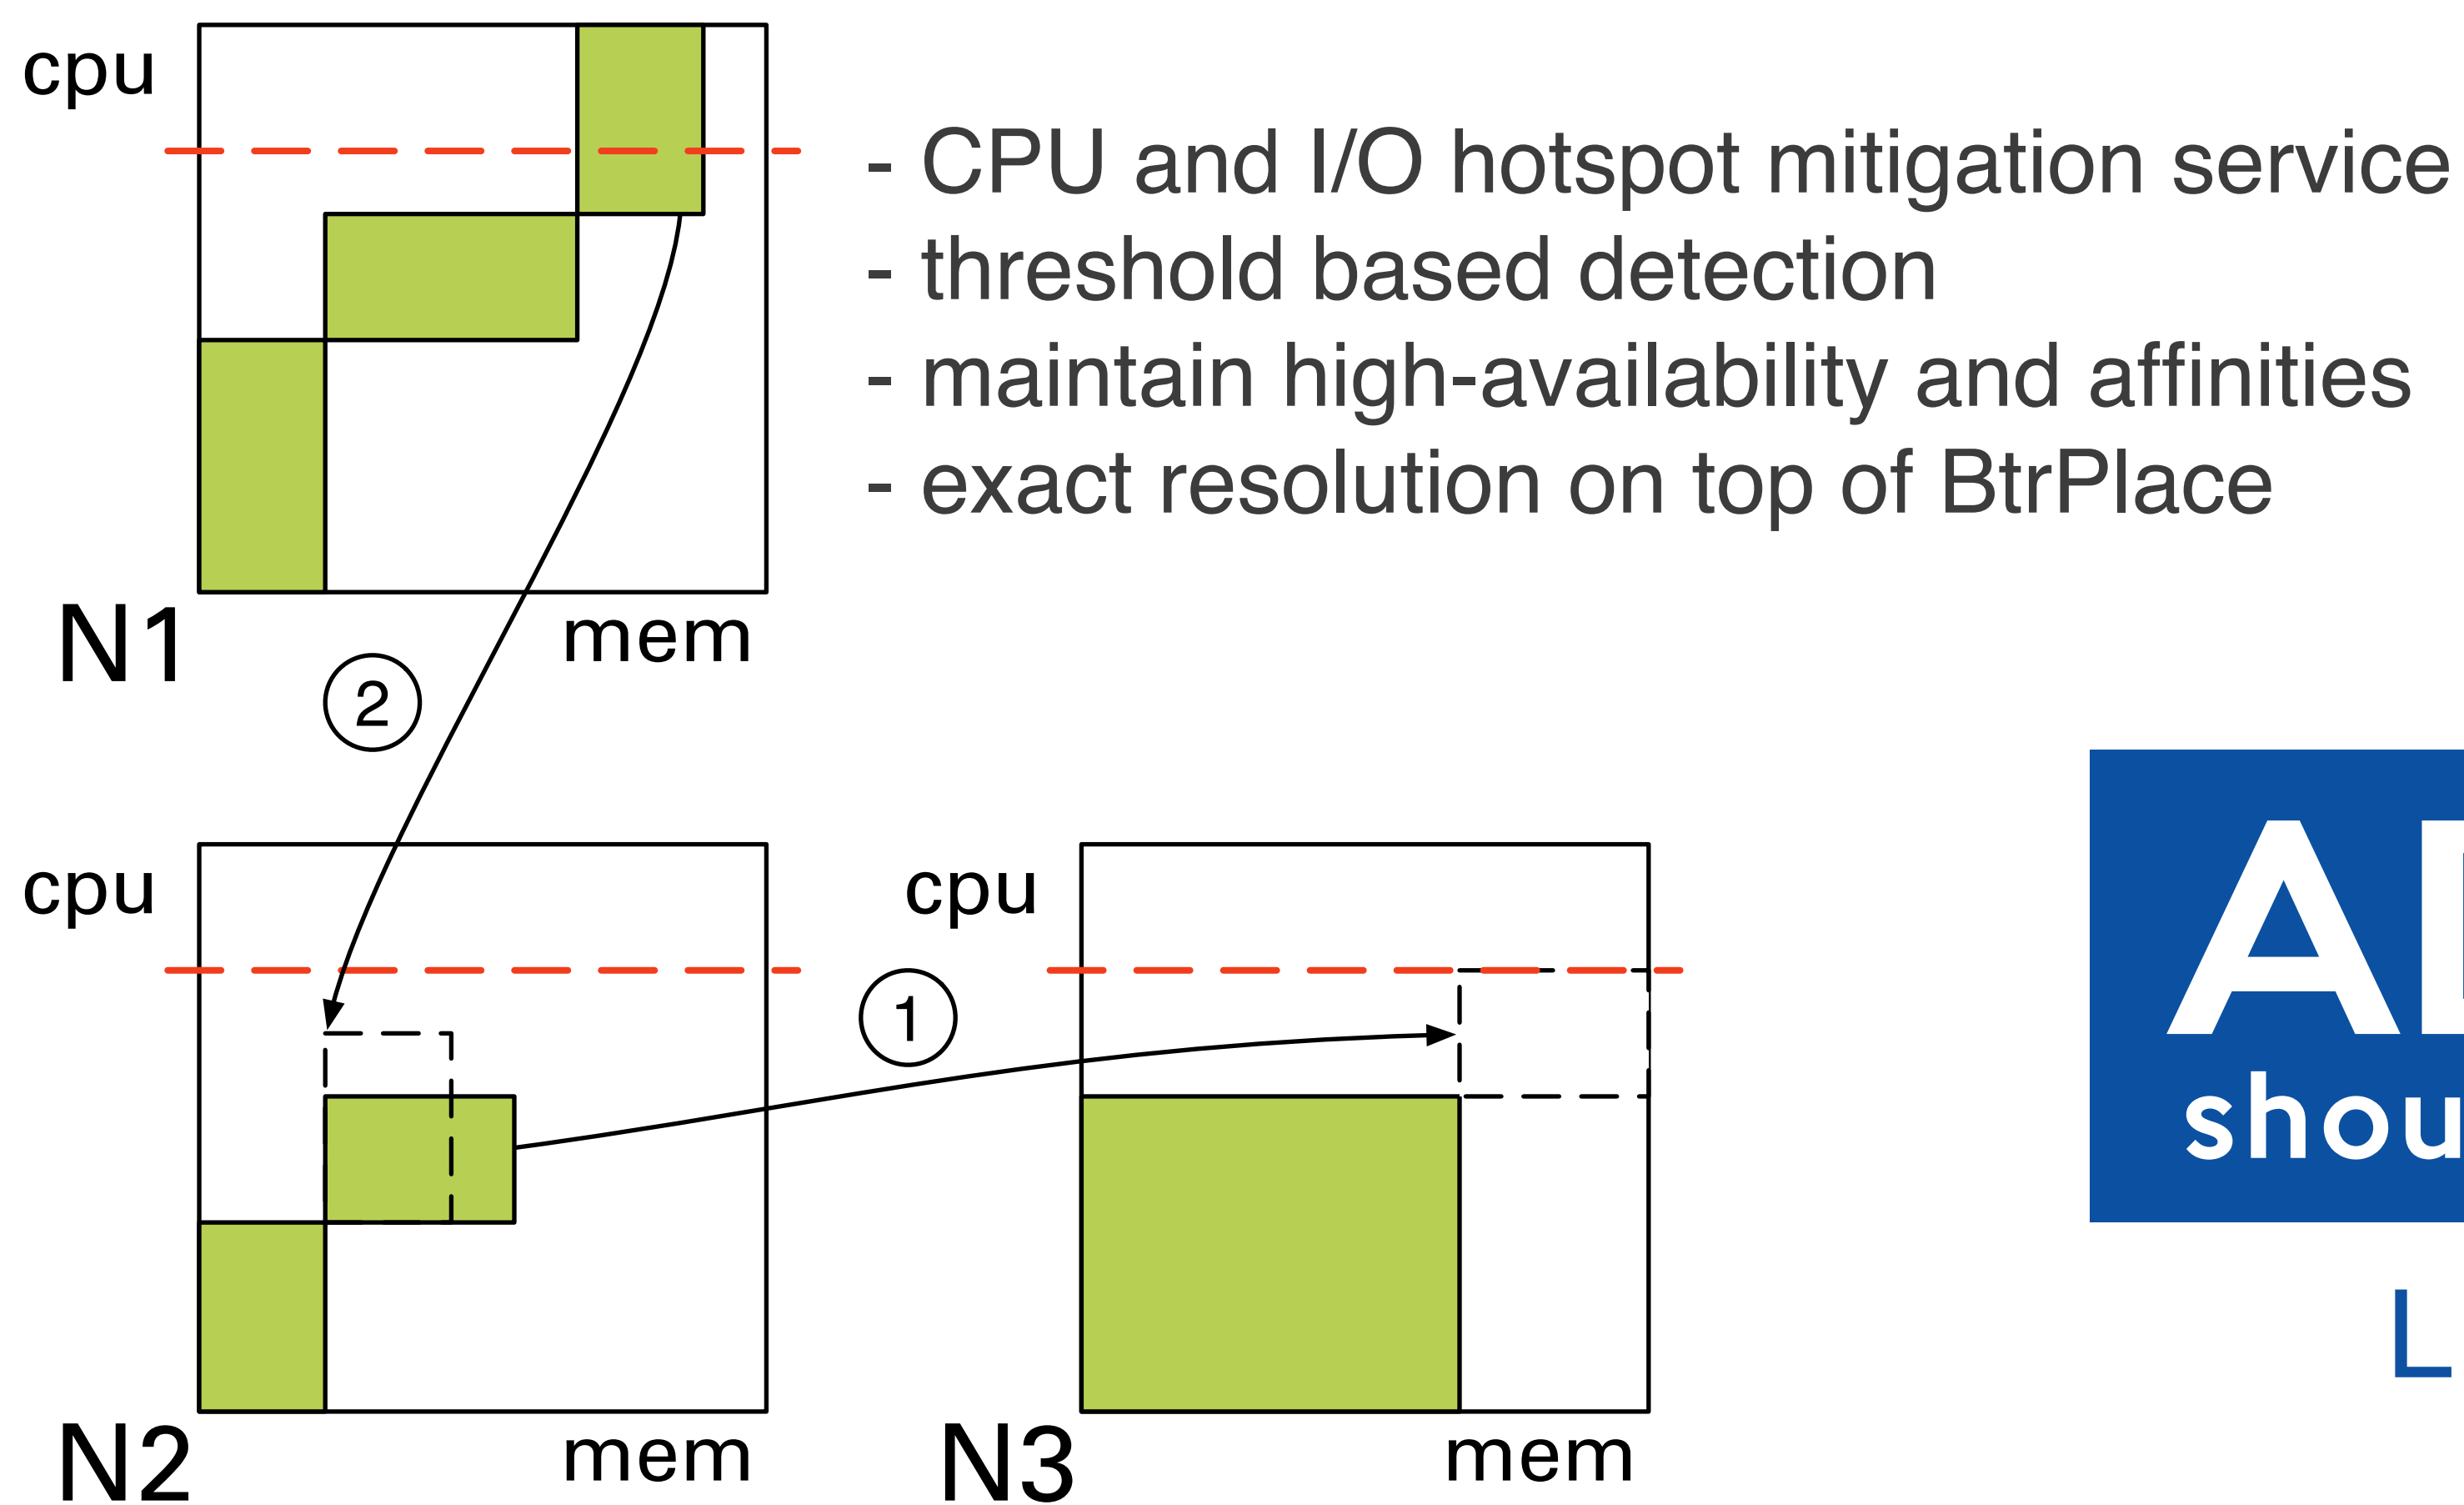


Acropolis Dynamic Scheduler (ADS)



Workload diversity in private clouds


- small clusters
- beefy nodes
- overcommitted CPUs
- moderate, non-uniform loads
- non-uniform VM resource usage

ADS perfect for someone should be **good for everyone**

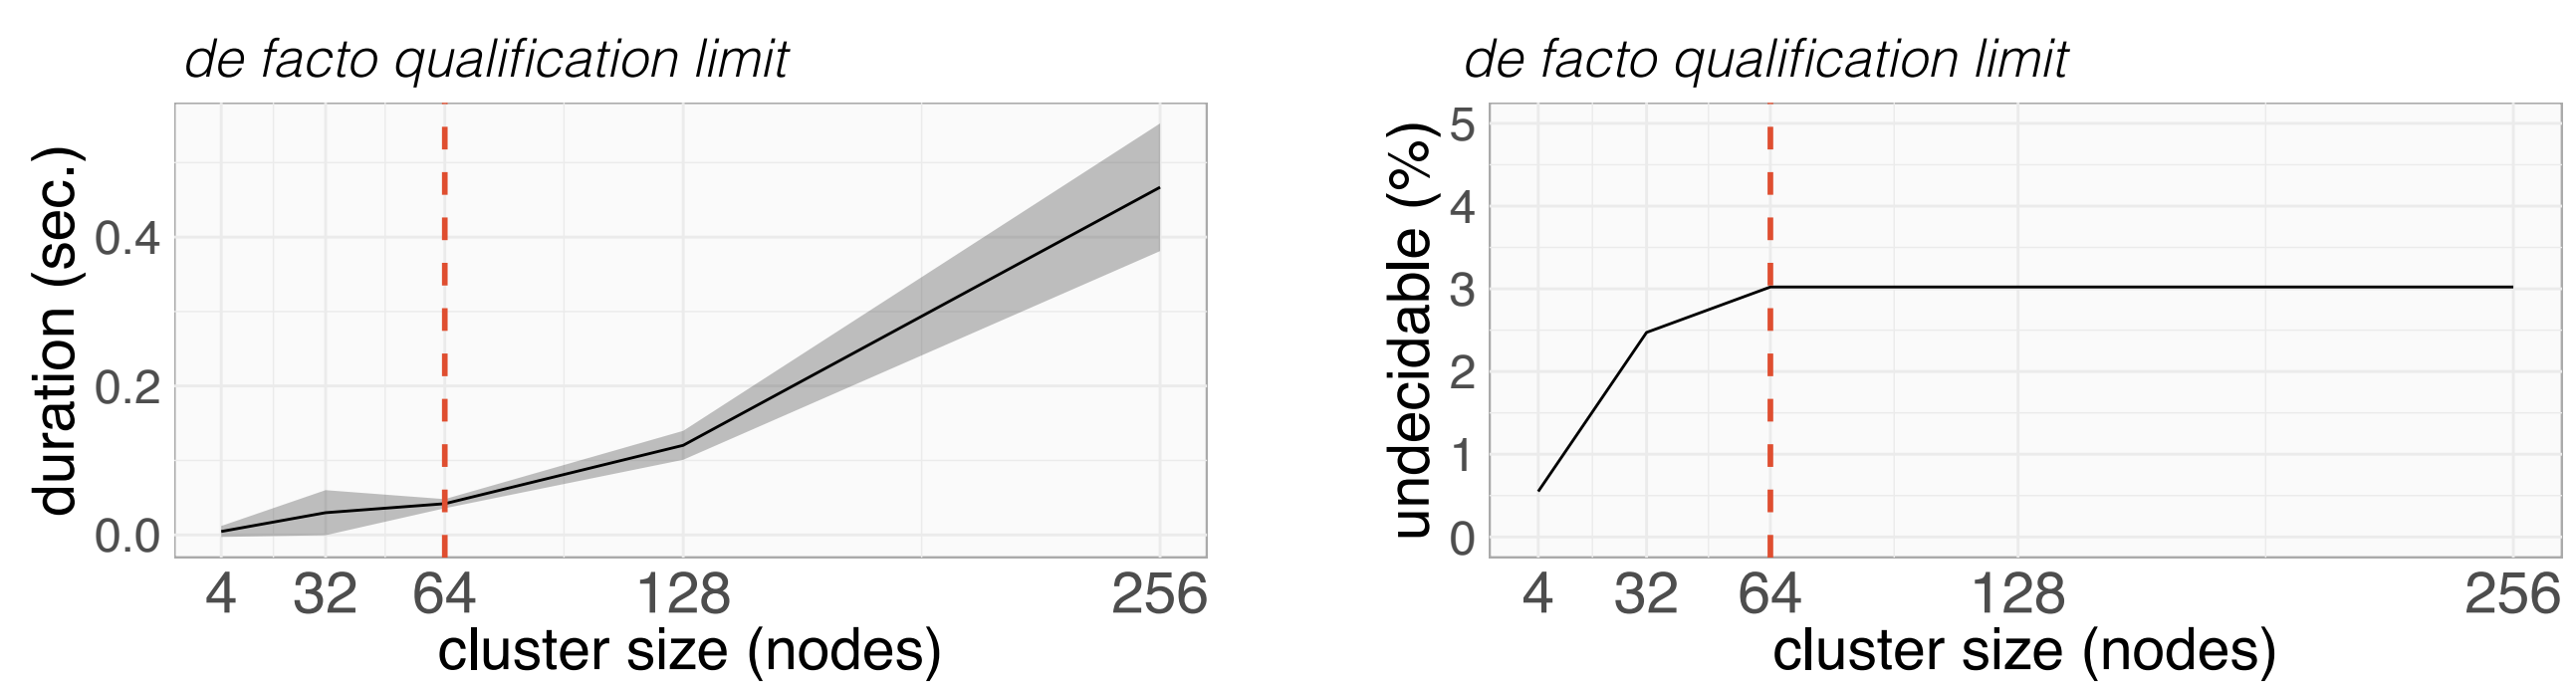
Lessons learnt facing 10000s private clouds

On the use of an exact approach

- continuous search helps yield better mitigation plans




- no scalability issue



- engineering bootstrapping overhead
- quality self-assessment

The right local search method ?

- reduce the problem size
- wise with a moderate load and local hotspots
- motivated by tests during the prototyping phase on estimated workloads
- revised over time from customer workload analysis
- double edged



Practical effectiveness

- good in theory vs. good in practice
- complex to analyse without *a/b testing*
- 6 times fewer hotspots observed after ADS issued a mitigation plan
- success rate is a consequence of subjective modelling choices

Engineering lessons

- All about resource modelling
- validating model accuracy is hard
 - the love-hate relationship with thresholds


- Observations over assumptions
- learn and improve from practical data
 - do not neglect biases

- Scheduler ecosystem
- < 50% of the code for the scheduler
 - dedicated debugging and testing tools

Placement or scheduling problem ?

- action scheduling brings more mitigation possibilities
- complex to implement, theoretically costly

+2.41% solving rate, but lower quality for 1% of the workloads



Decision capabilities and overheads

- smarter decisions may be expensive
- supporting new workloads must not lead to regressions

e.g. knapsack filtering:
-0.56% undecidable problems
but lower quality for the outliers

