




# Keeping a Hypervisor Fleet Up to Date with Minimal Customer Disruption

**Atalay Kutlay**

#SRECON26  
March 26, 2026



# \$whoami



**Atalay Kutlay**

Senior Software Engineer  
Akamai

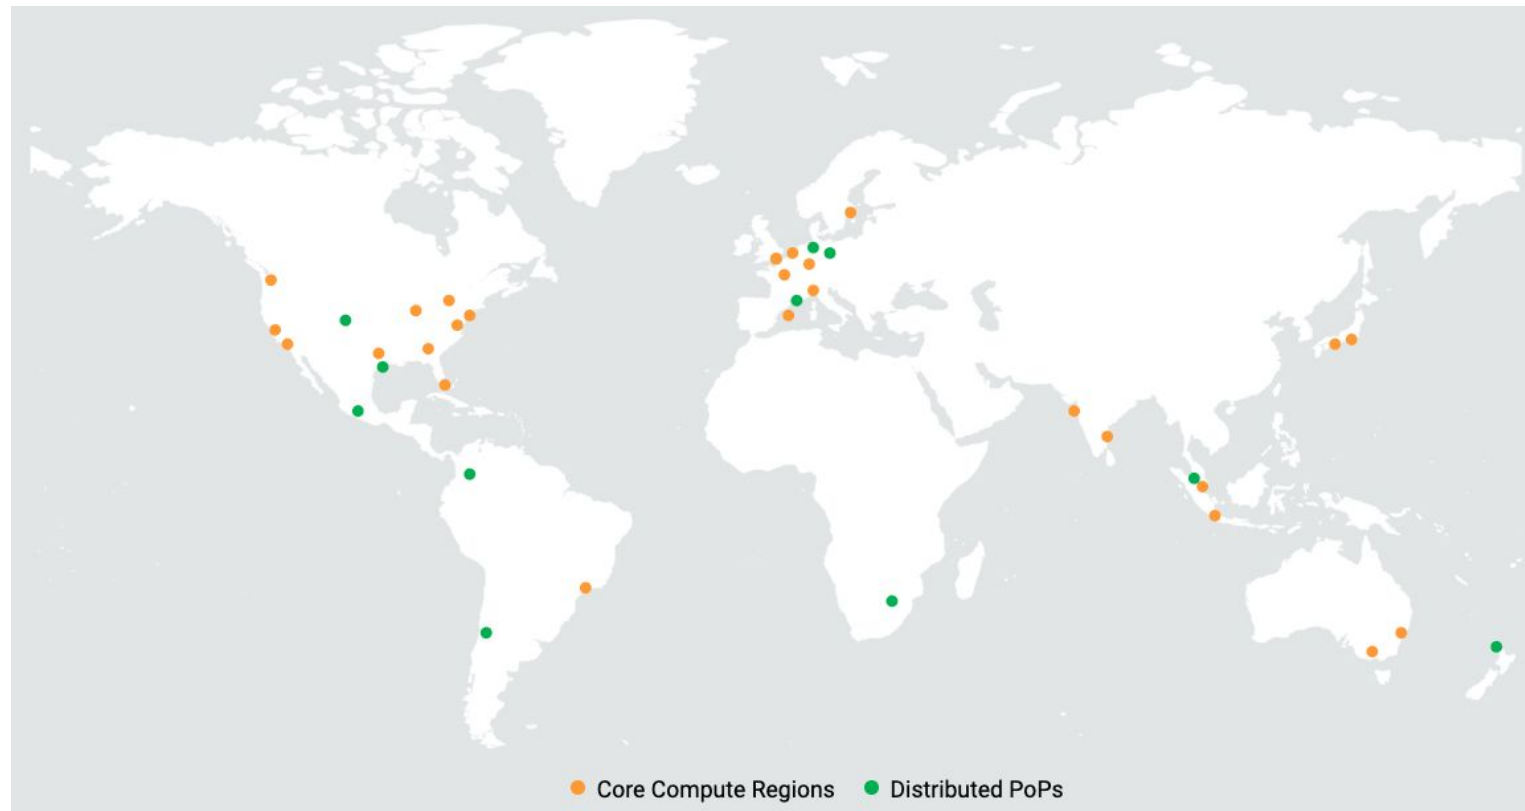
Senior Software Engineer at Akamai

Network Optimization team

I've worked on:

- CDN capacity planning / content allocation
- Virtual machine allocation
- Large scale scheduling problems

# Akamai's Cloud Infrastructure






# Hypervisor Maintenance


- Kernel, OS, firmware updates
- Preventative hardware maintenance
- Most maintenance activities require reboot

# Hypervisor Maintenance

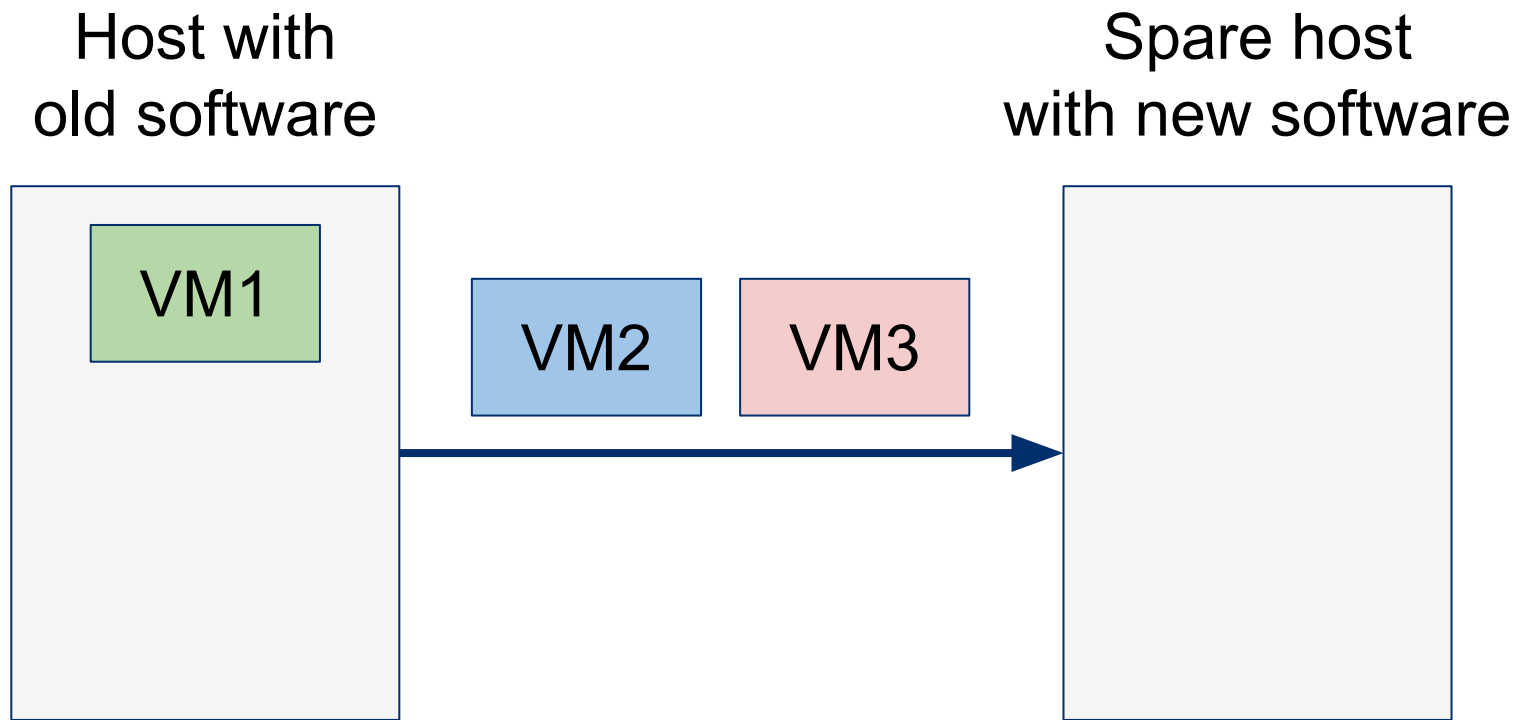
Host with  
old software



Spare host  
with new software




# Hypervisor Maintenance




# Hypervisor Maintenance

Host with  
old software



Host with new  
software



# Hypervisor Maintenance

~~Host with  
old software~~

Updating  
to new  
software

Host with new  
software


VM1

VM2


VM3

# Hypervisor Maintenance

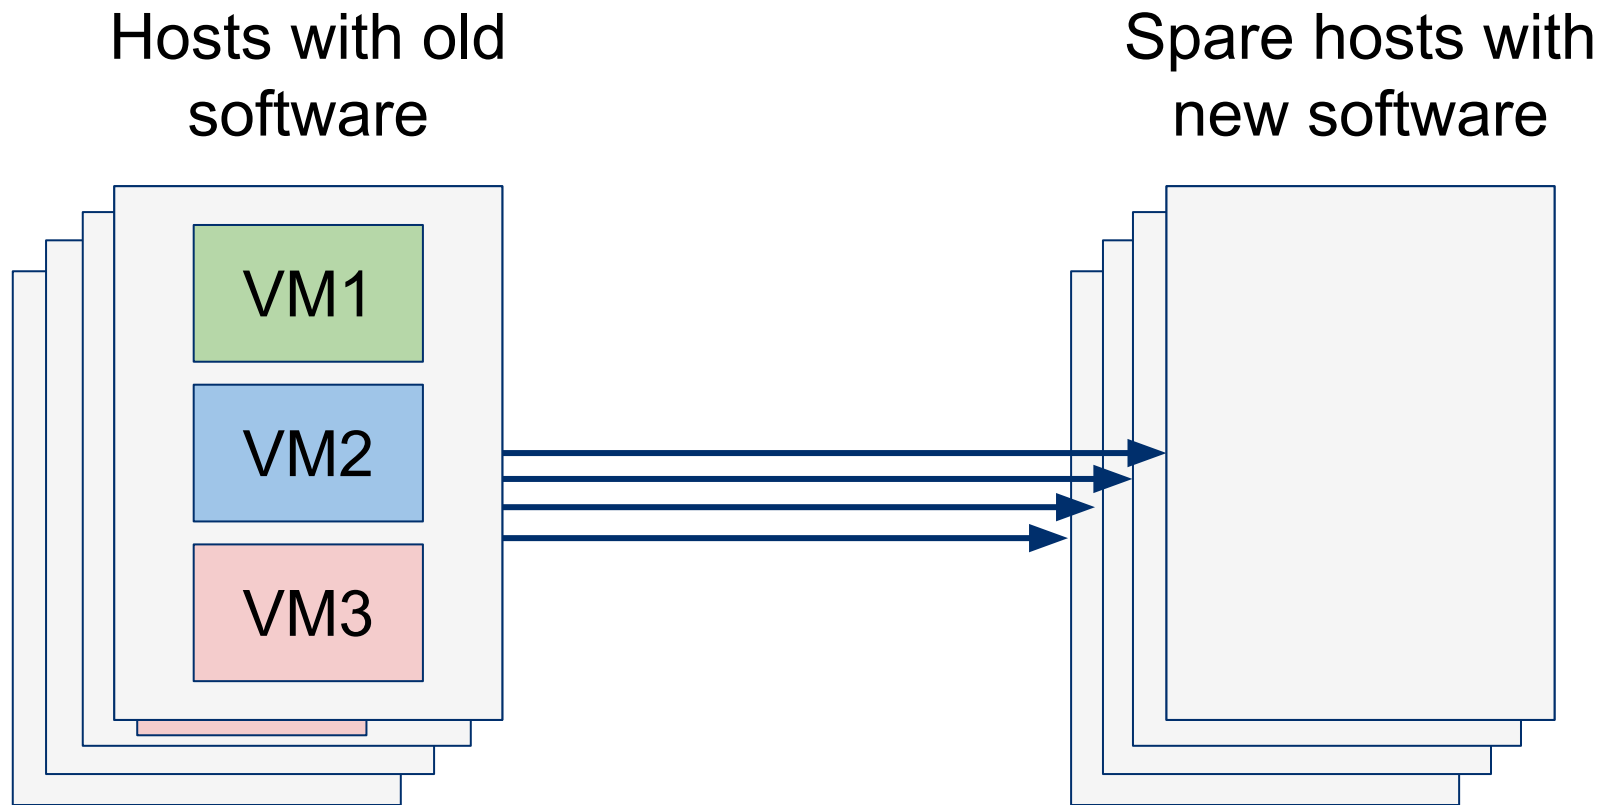
Spare host  
with new software



Host with new  
software



# Hypervisor Maintenance at Scale



# Scenario

~100 hosts day



- You're asked to update **10,000+ hypervisor hosts**
- Hosts need reboot, VMs need to be migrated
- Customers need to be notified in advance
- You have **90 days**



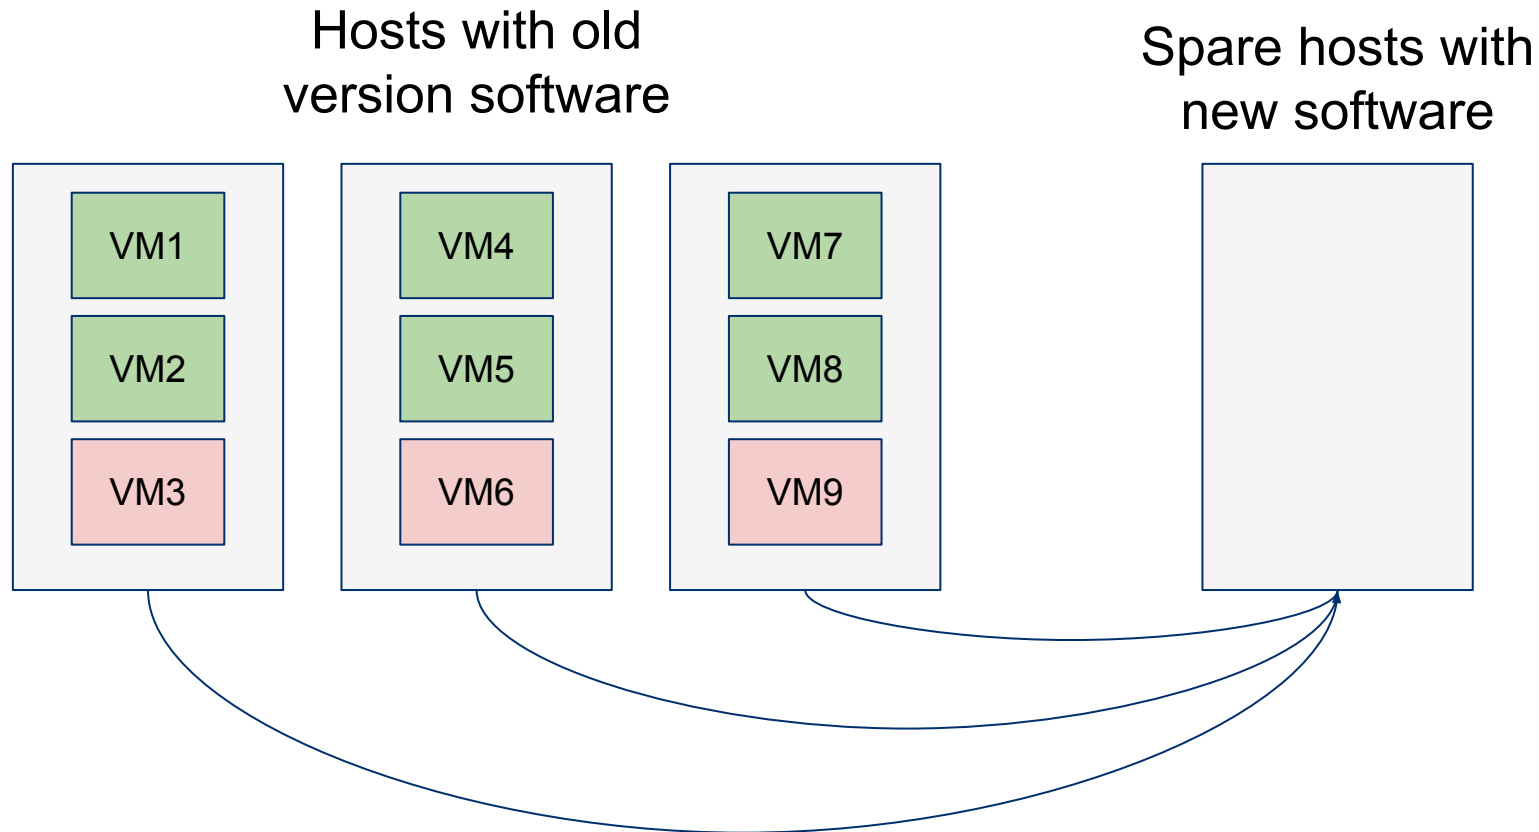
# Capacity, Conflict, Concurrency

## Problem #1: Capacity

- Finding **up-to-date** capacity is crucial
- Keeping availability in the datacenter stable is important

# Capacity, Conflict, Concurrency

## Problem #1: Capacity






# Capacity, Conflict, Concurrency

## Problem #2: Conflict

- Customers tolerate disruption to a certain limit
- Host selection may impact customers unfairly / asymmetrically

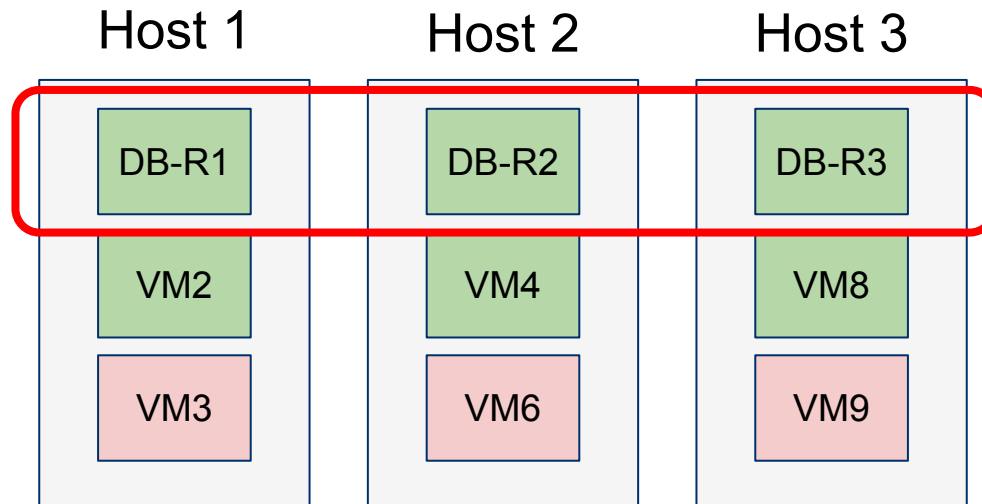
# Capacity, Conflict, Concurrency

## Problem #2: Conflict



# Capacity, Conflict, Concurrency

## Problem #2: Conflict





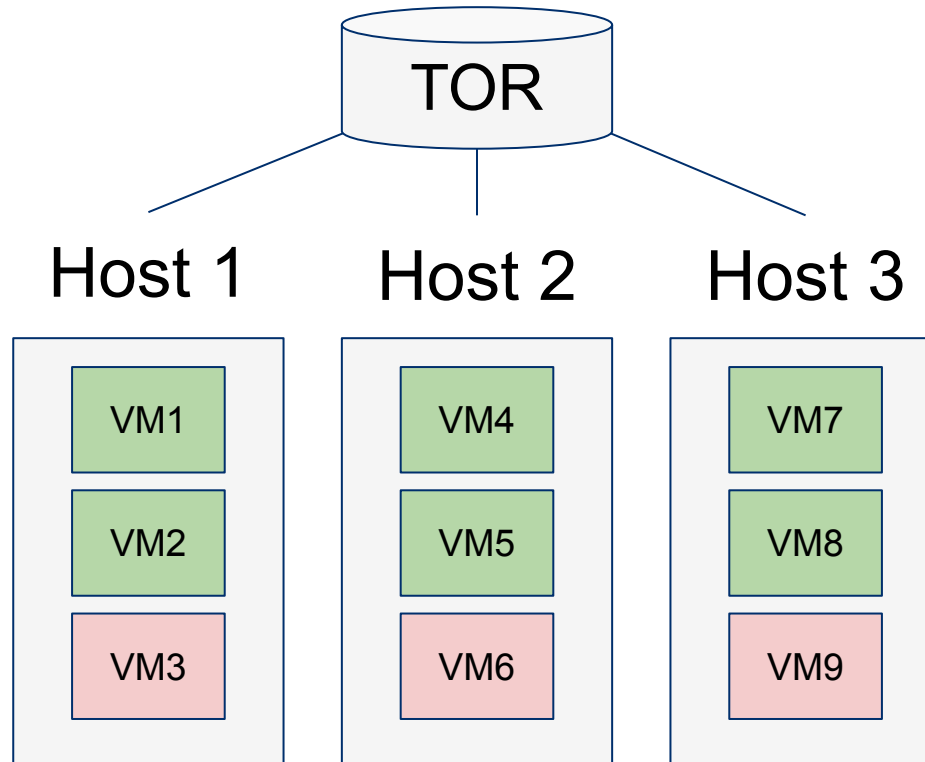
# Capacity, Conflict, Concurrency

## Problem #3: Concurrency

- Migration takes resources from the host (CPU, Disk IO, NIC)
- Top-of-rack router bandwidth is limited


# Capacity, Conflict, Concurrency

## Problem #3: Concurrency



# Capacity, Conflict, Concurrency

## Problem #3: Concurrency





# Capacity, Conflict, Concurrency

## Solution: spreadsheets?

- Takes hours of work
- Rinse and repeat in 3 months



## Solution: write an algorithm?

- Easy to say, hard to get right:
  - Hard to measure the quality of solution
  - Not flexible enough for changing constraints / business priorities




# Operations Research

- Operations research provides ways to model and solve complex problems to improve decision making
- **Mathematical optimization** is a fundamental tool in OR



# What Is Mathematical Optimization?

# Mathematical Optimization



# Mathematical Optimization



Give me ice cream

Chocolate or oreo  
At most 3 scoops

Make me happy

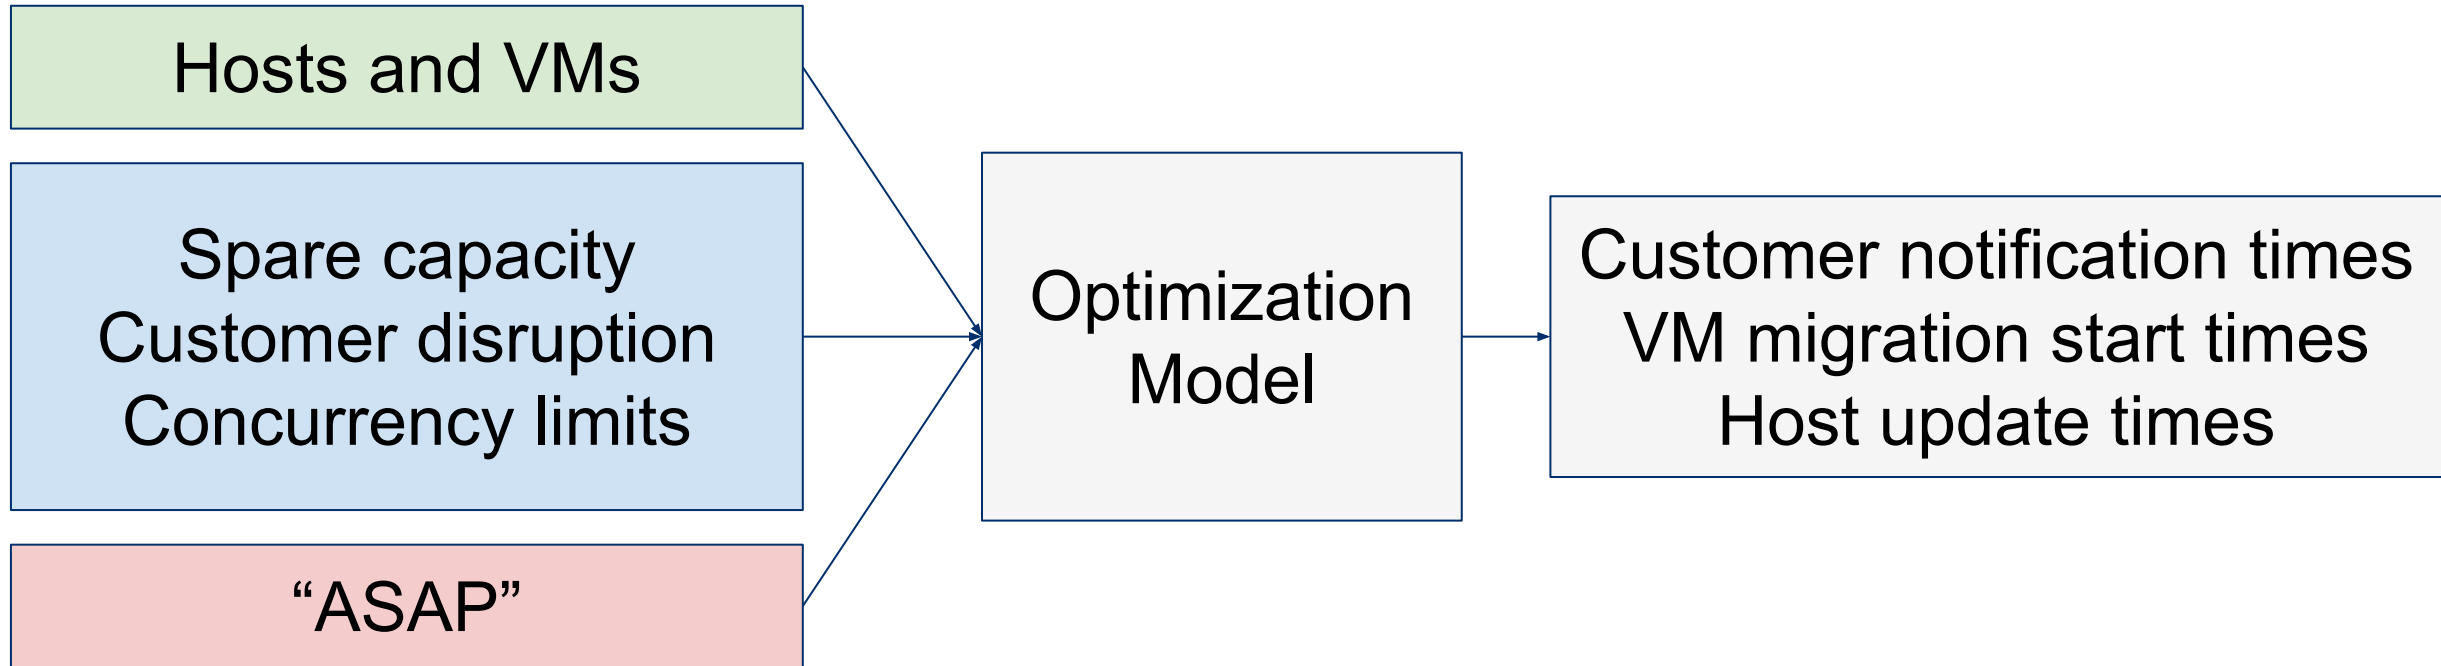


“Garlic flavored ice cream” wasn’t a joke, unfortunately  
[https://en.wikipedia.org/wiki/Garlic\\_ice\\_cream](https://en.wikipedia.org/wiki/Garlic_ice_cream)

# Mathematical Optimization



# Maintenance Scheduling Mathematical Model



# OR-Tools CP-SAT

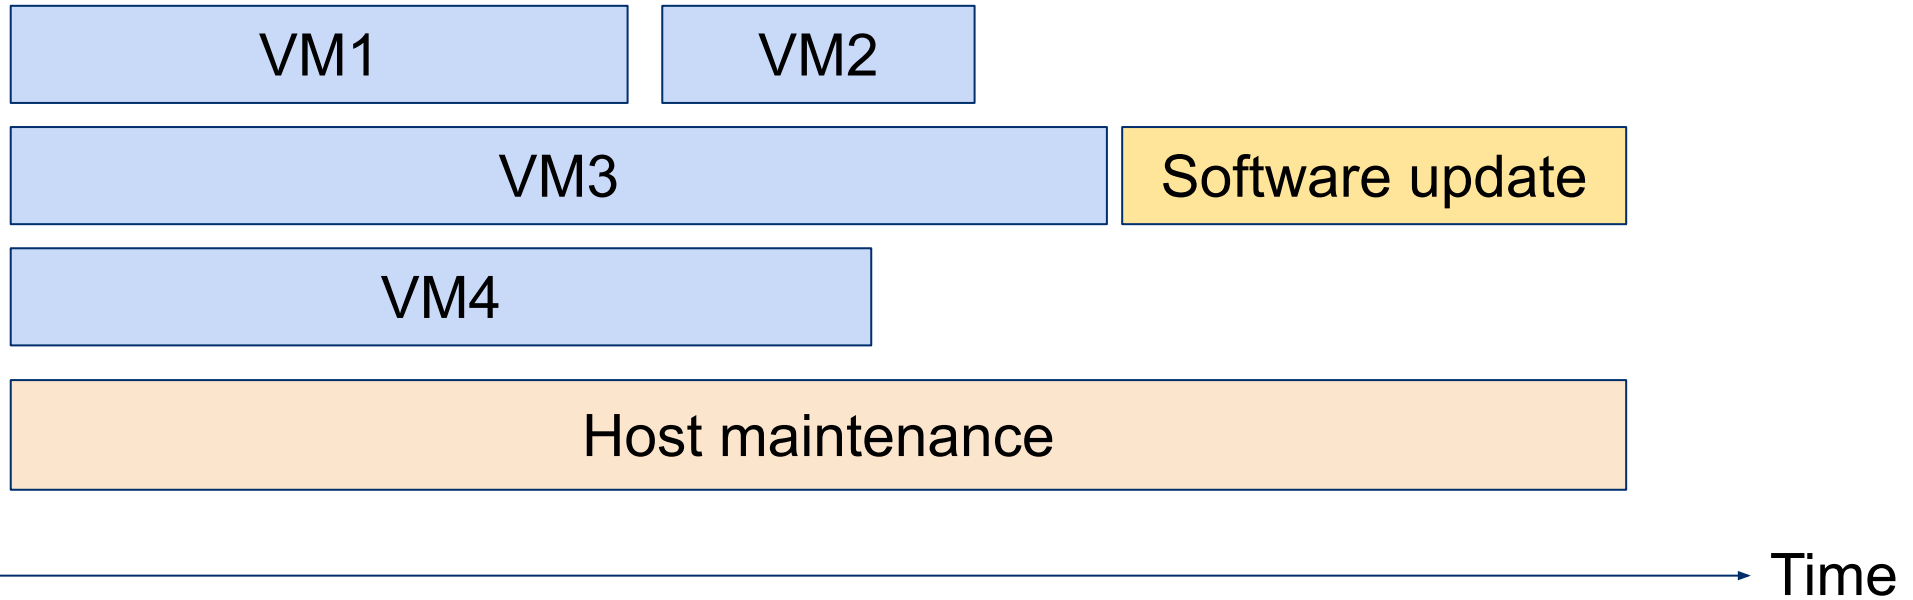
- Open source “portfolio solver” developed by Google ([developers.google.com/optimization](https://developers.google.com/optimization))



Google OR-Tools

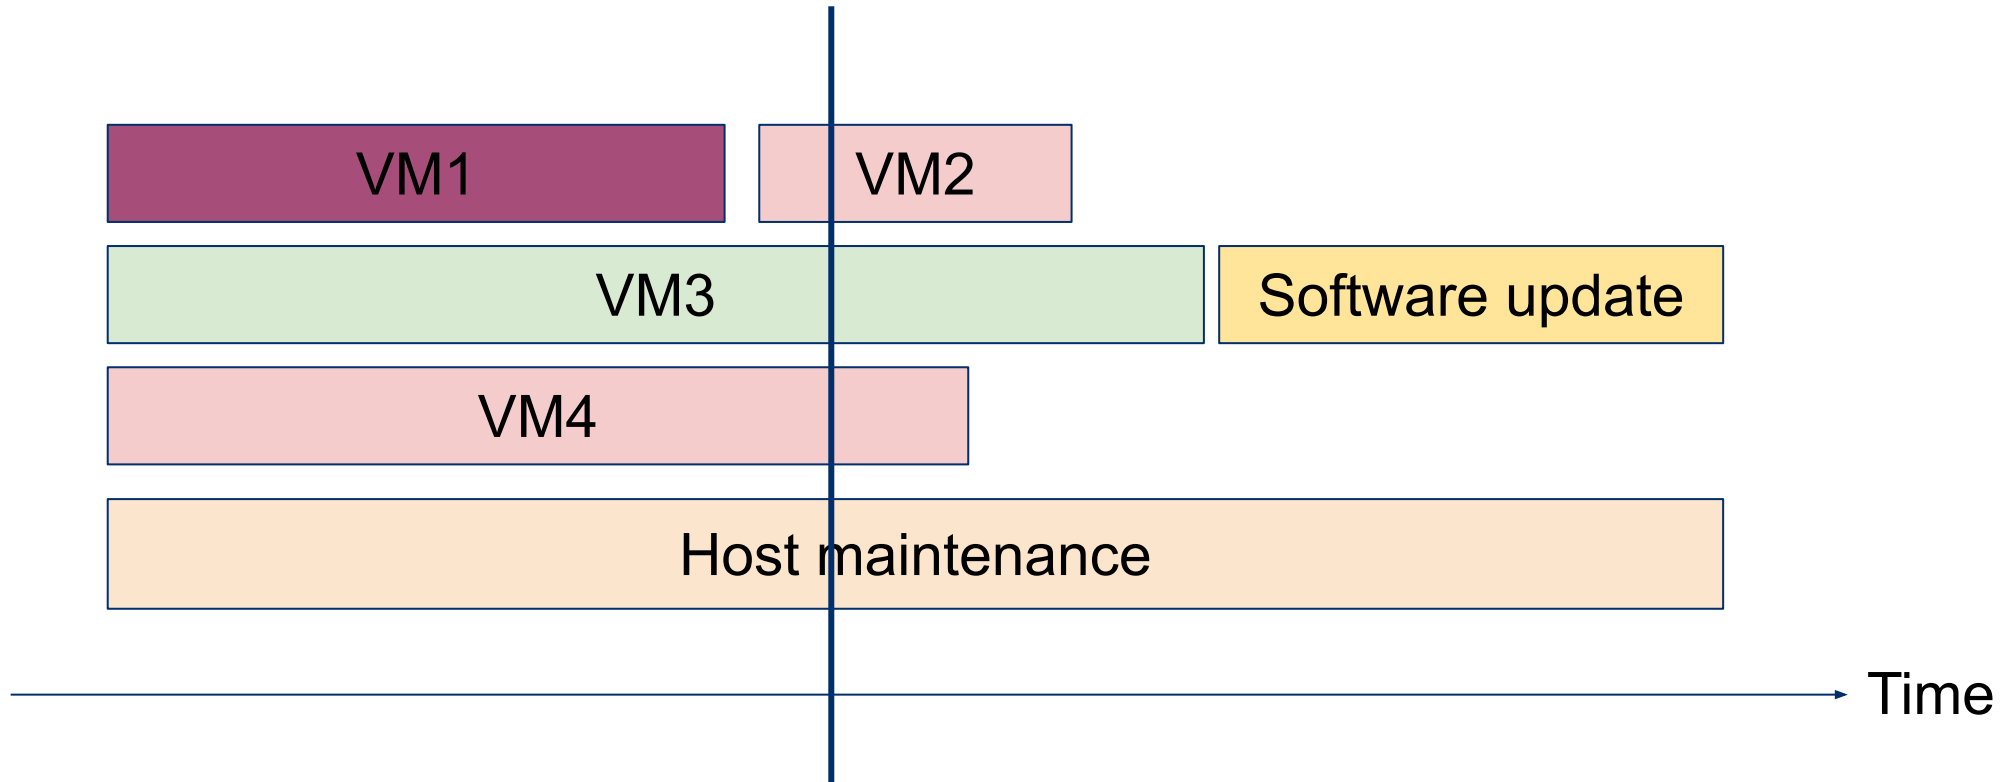
- Easy to use interface for scheduling problems:
  - Interval Variables
  - Specialized constraints:
    - AddNoOverlap
    - AddCumulative

# Example: AddNoOverlap

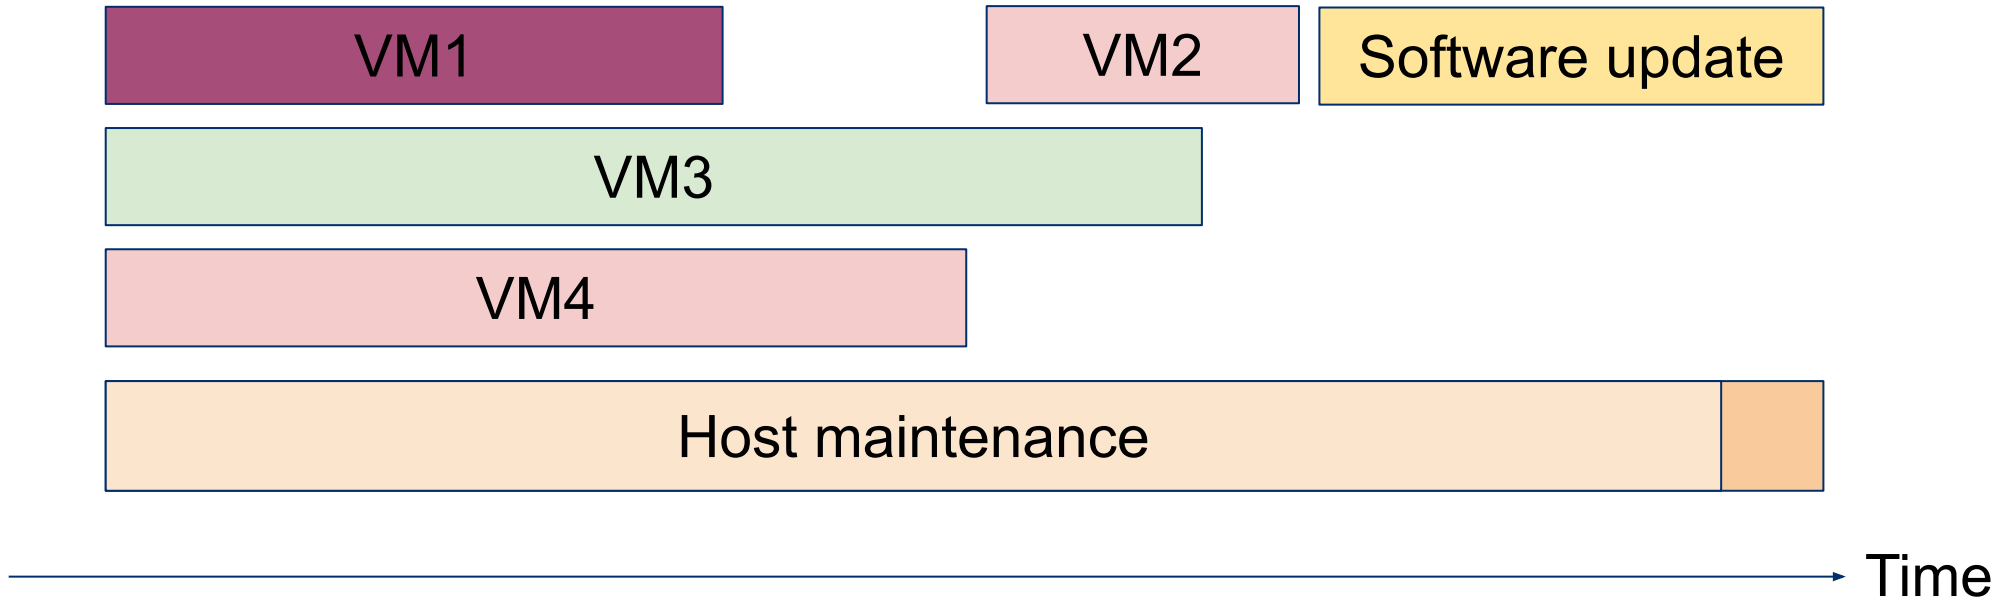


# Example: AddNoOverlap

Customer disruption at most 1



# Example: AddNoOverlap

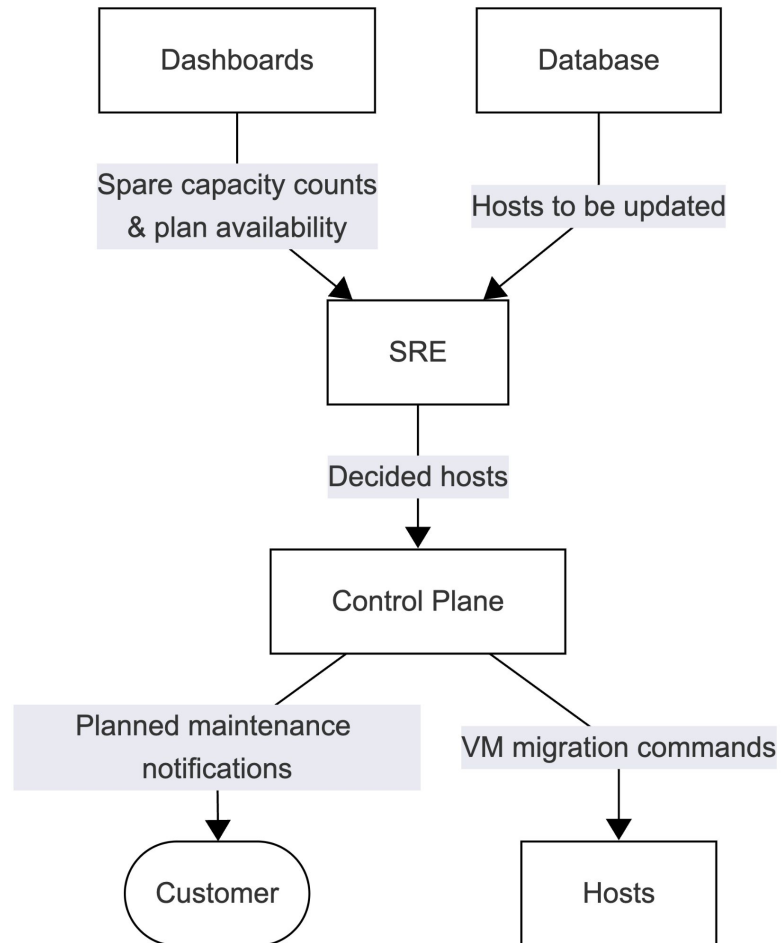




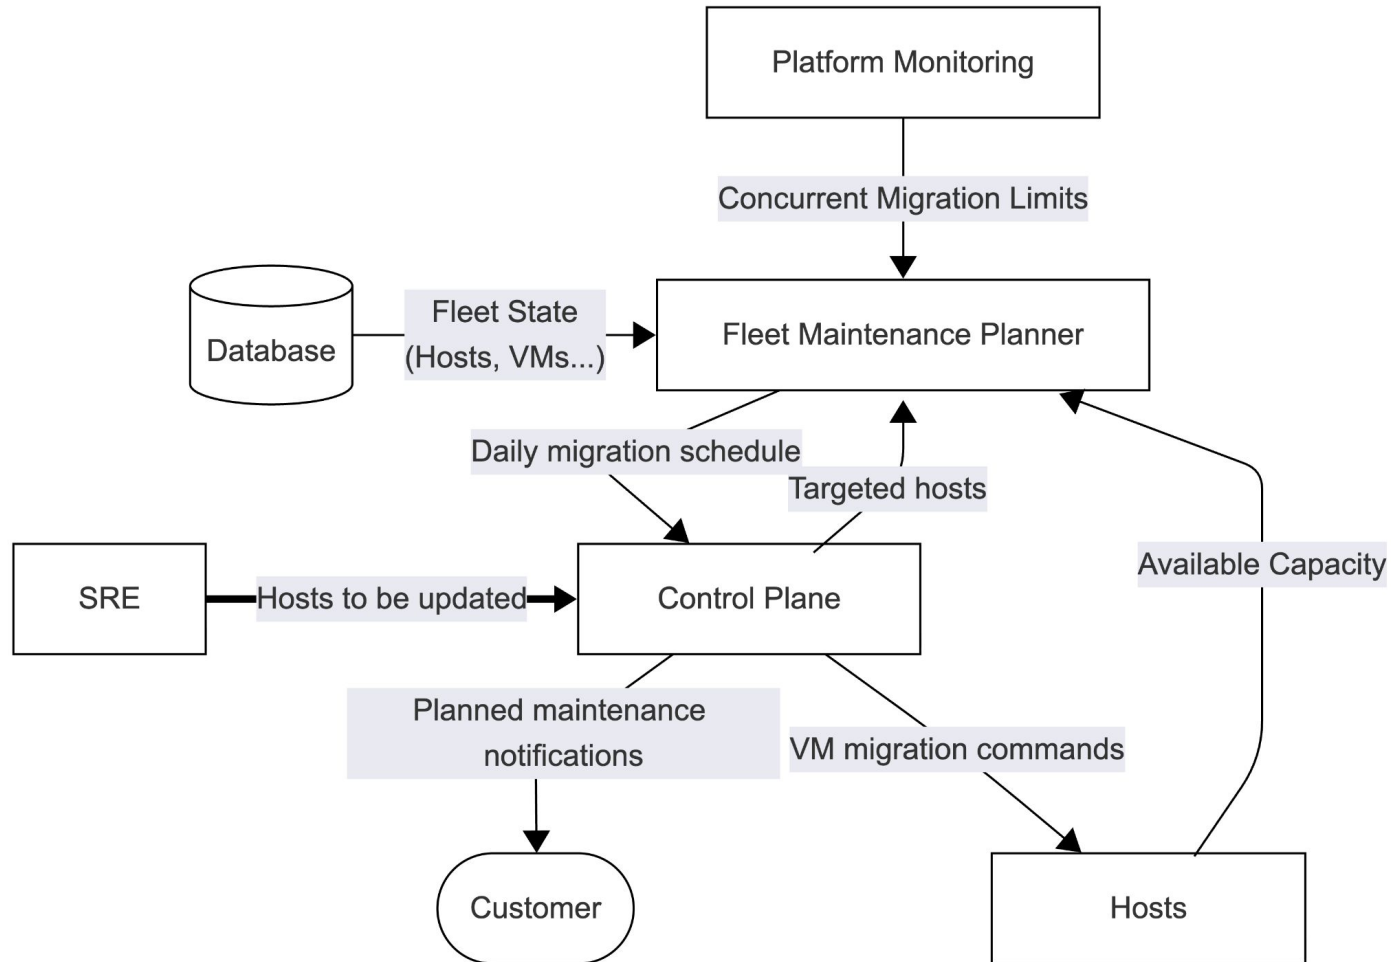
# Fleet Maintenance Planner

- Centralizes decisions of “who” and ”when”
- Produces all the data needed to schedule maintenance
- Enables optimal maintenance scheduling at scale

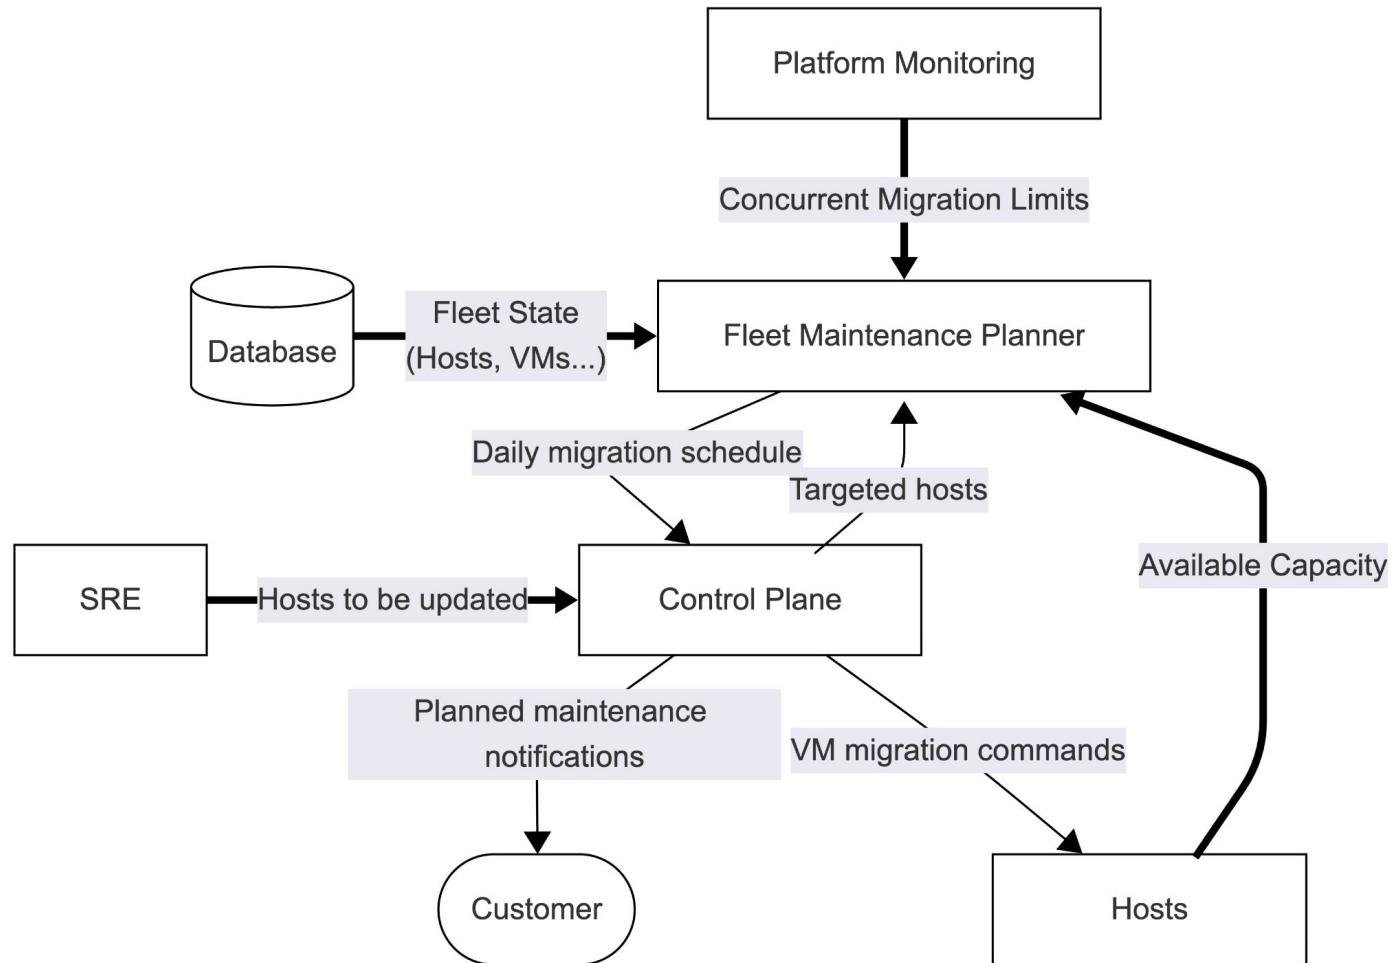
# Maintenance Scheduling (before)



# Maintenance Scheduling (after)



# Maintenance Scheduling (after)





# Feedback Loops

Fleet state is dynamic:

- Spare capacity depends on successful updates
- Migration limits change with infrastructure health
- VMs are short-lived

# Optimization Model Observability

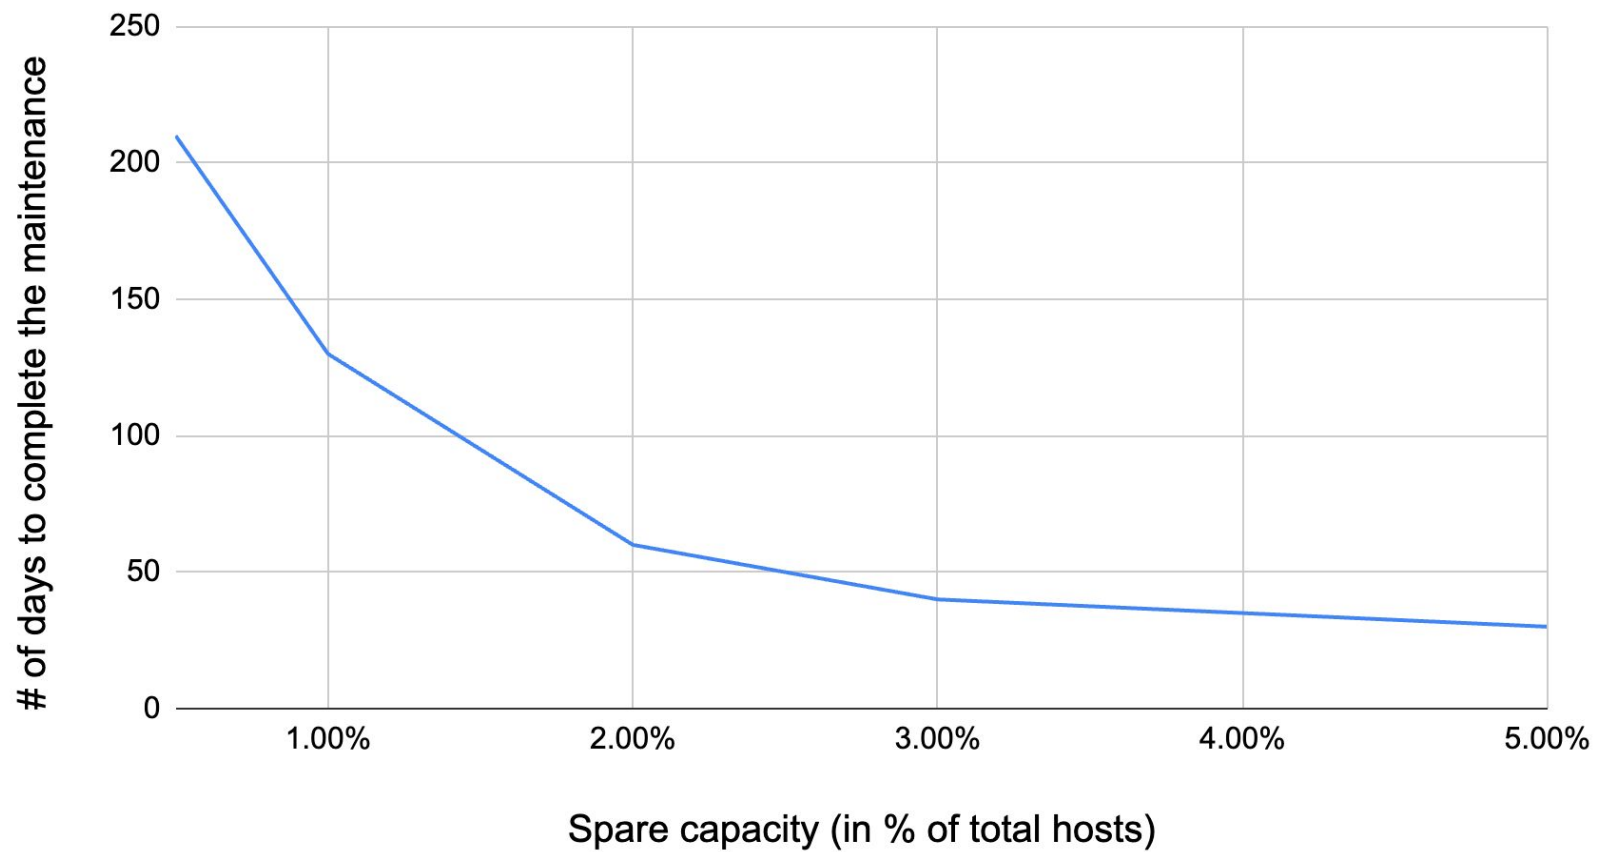
Tracking “optimality”



## Side Benefit: Scenario Analysis

- “What if we need to update the whole fleet tomorrow?”
- “What if we scale the data center 2x in a year?”
- “What if customer X scales 10x?”

Fleet update time and spare capacity





# Summary

## Mathematical optimization:

- Solve problems at scale
- Flexible & allows for incremental development
- Gives you tools to answer business questions


**Thank you!**  
**Questions?**



**Atalay Kutlay**



[/in/atalay-kutlay](https://www.linkedin.com/in/atalay-kutlay)





# Extra Slides

## Why not MILP?

- Performance: even commercial MILP solvers are not competitive for RCPSP problem (very similar to cloud maintenance scheduling)
- Ease of development: OR-tools has specialized variables for scheduling
- MILP models are hard to get right for scheduling problems:
  - Time-index based models are too expensive, doesn't scale well
  - Compact models are more complex

# CP-SAT Interval & AddNoOverlap Example

```
planning_horizon = 24 * 120 # 120 days

start_var = model.new_int_var(0, planning_horizon, f"{vm}_start")
end_var = model.new_int_var(0, planning_horizon, f"{vm}_end")
migration_duration = 10

vm_interval_var = model.new_interval_var(
    start=start_var, size=migration_duration, end=end_var, name=f"
{vm}_interval")
```

```
vm_intervals = list()

model.AddNoOverlap(vm_intervals)
```



## Complete Problem Statement for Hypervisor Maintenance in Cloud

We need to schedule host evacuations so that:

- Customer impact is below the “disruption budget”
- Only the capacity carved out for the maintenance is used
- Migrations do not cause performance issues / network bottlenecks

Our goal is to:

- Minimize the total duration of software updates

# Maintenance Scheduling Problem

## Variables:

- Each VM's notification / migration start time
- Each host's status change time
- Each host's software update time

## Objective:

Finishing the updates as fast as possible.

## Constraints:

- Number of hosts “under maintenance” can not exceed spare capacity
- No “maintenance group” has more VMs migrating than disruption budget
- Each VM gets the notification period according to its maintenance policy
- Concurrent VM migrations are limited in host, rack, and datacenter levels

