



Infrastructure Optimization for OpenStack

CERN Openlab Lightning Talk

Author: **Mohammed Henni**

Supervisor: **Jose Castro Leon**

August 14th, 2018



Introduction



+300 k cores



~40 k VMs

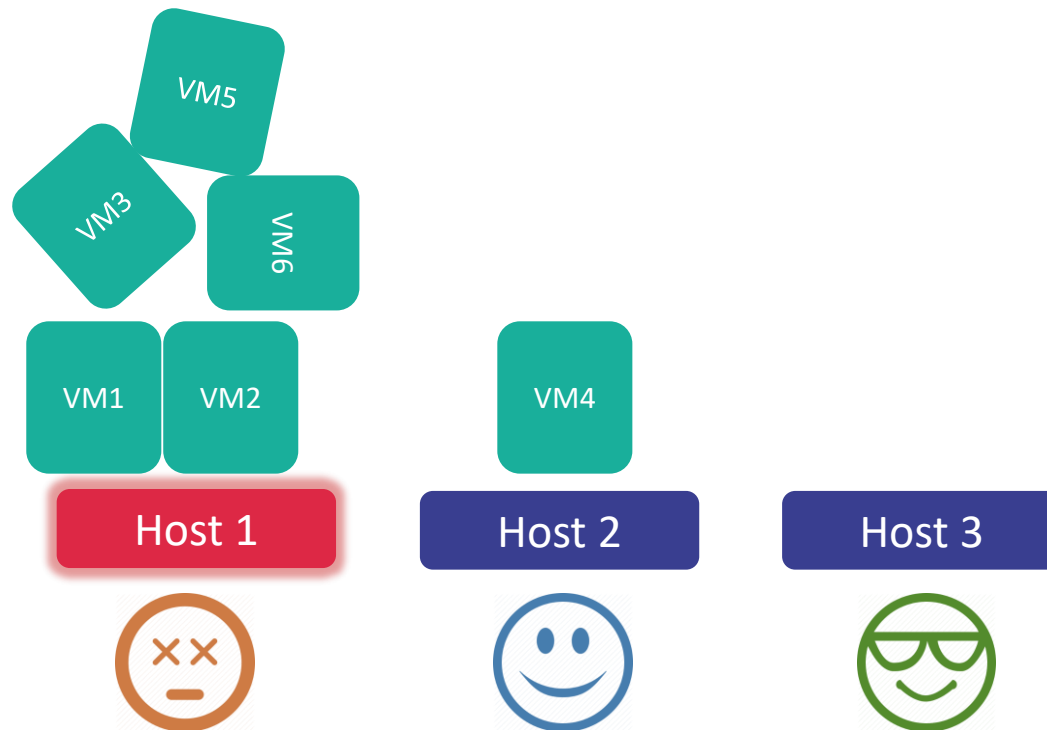


+9 k hosts



VM creation/deletion every 10s

Sometimes, things can look like this

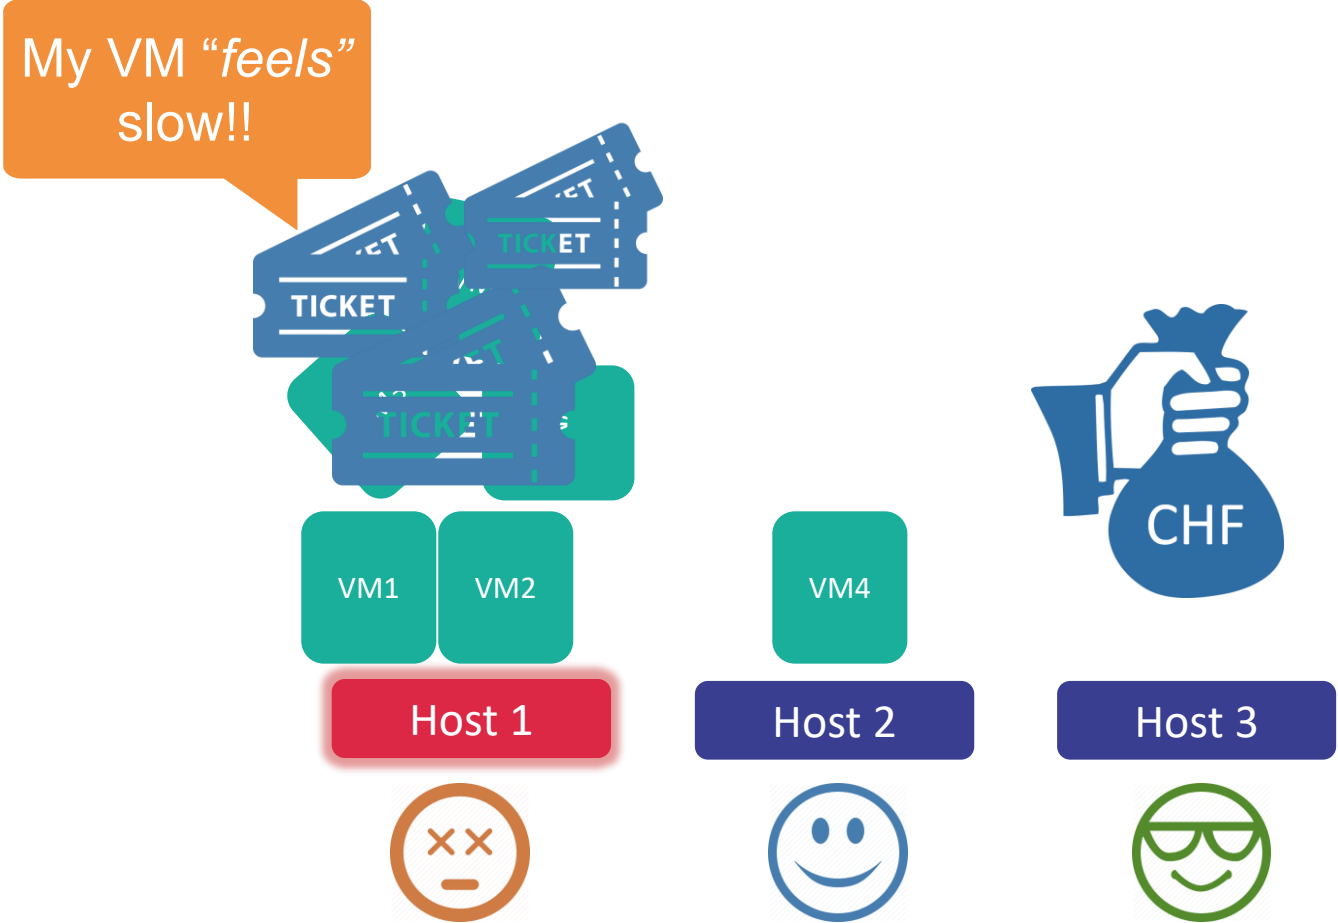


This results in...

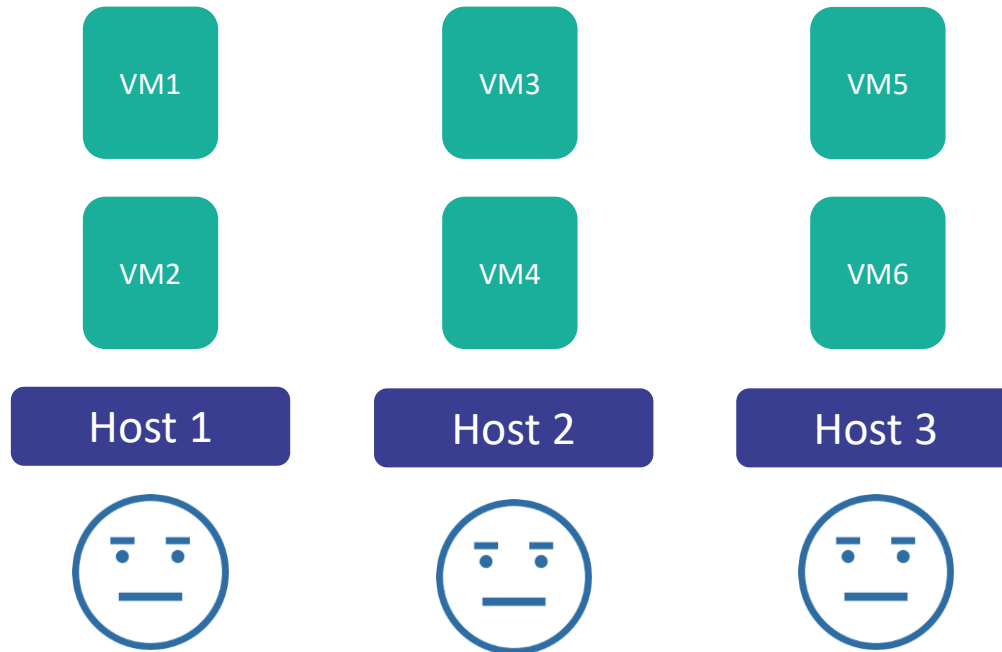
My VM "feels" slow!!



This results in...




So, we need someone to keep watching






OpenStack Watcher



Watcher workflow




<https://drive.google.com/file/d/0B6U3RY726GT9aElfaDNFbnJPVFE/view>



My work

Deploying on Devstack

Implementing an optimization algorithm



Exploring Watcher

Trying on pre-production

Testbed

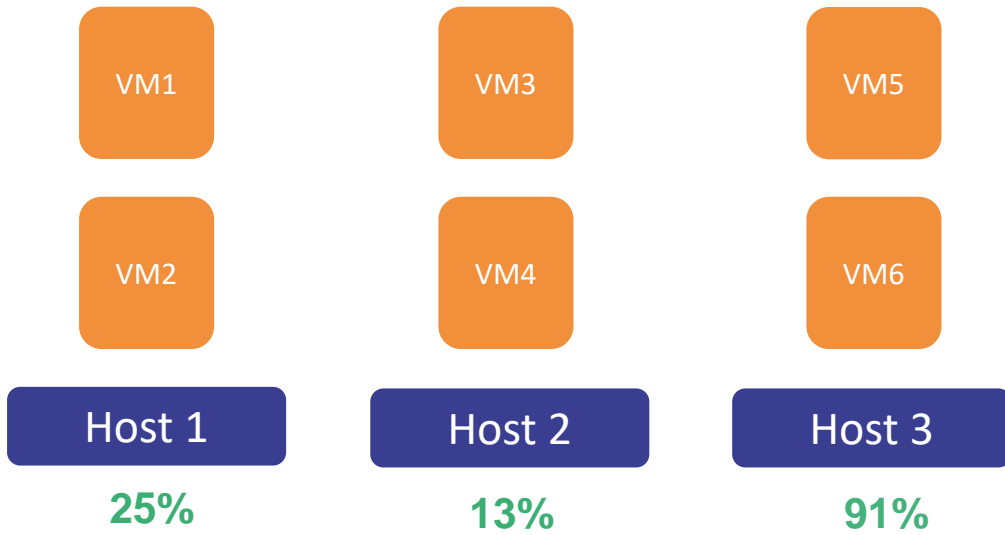
Hyperconverged servers = compute + storage

 9 servers

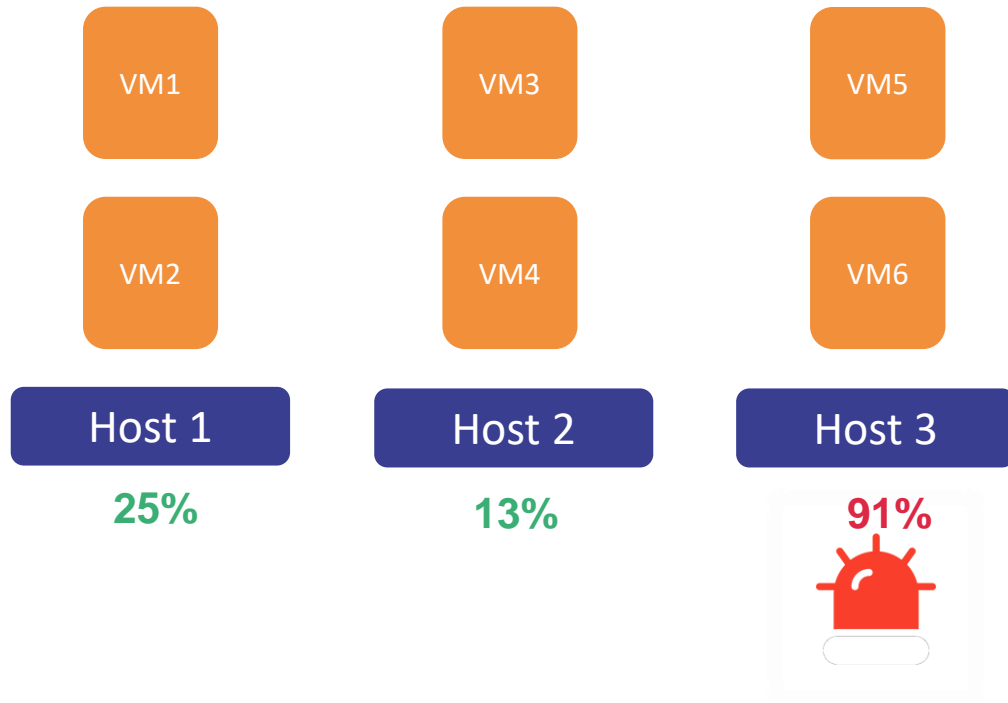
 24 virtual machines

Goal: optimize the workload balance

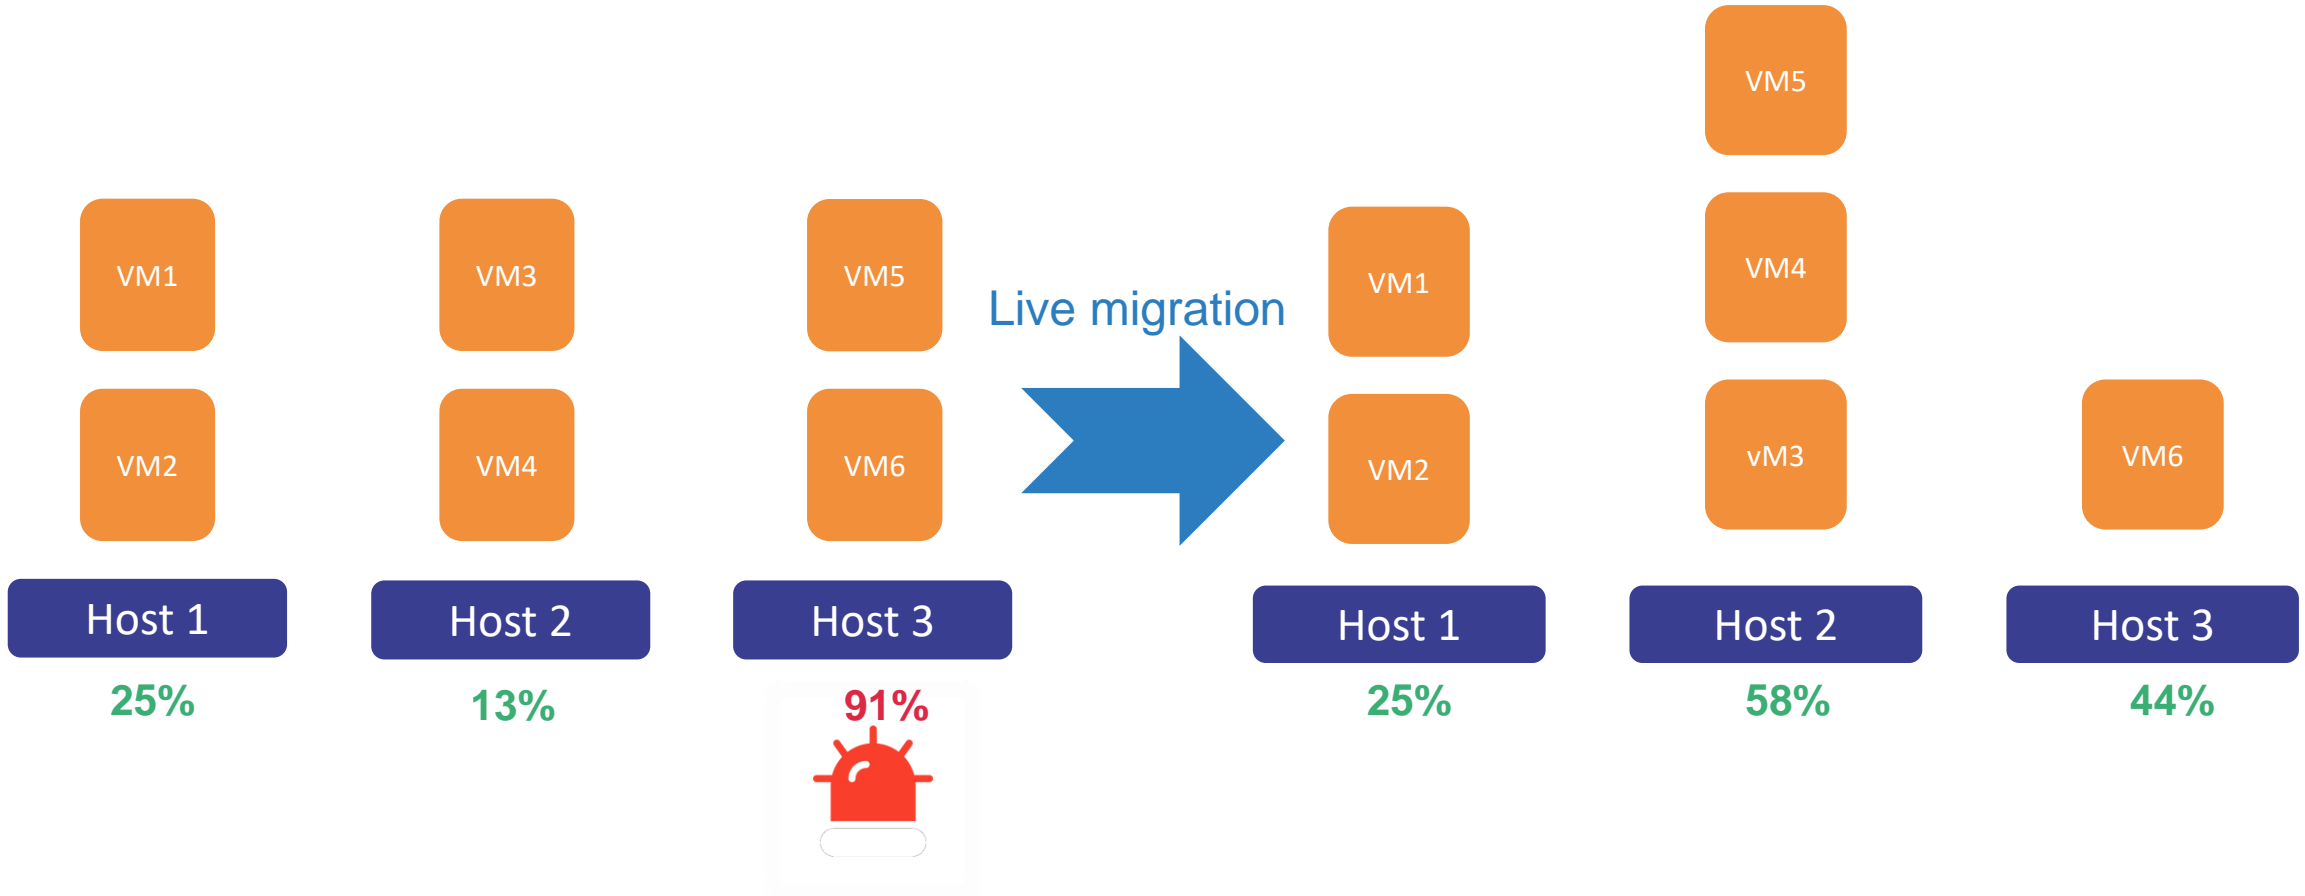
Watcher in action



Watcher in action



Watcher in action





Results

- ✓ Watcher deployed on pre-production
- ✓ Extended with custom optimization
- ✓ Tested on hyperconverged servers



Future work

- ✓ Watcher deployed on pre-production
- ✓ Extended with custom optimization
- ✓ Tested on hyperconverged servers
- ⦿ Further integration with CERN monitoring
- ⦿ Find more use cases
- ⦿ Implement more optimization algorithms

Thank you



*So what have you been
doing at CERN?*

Watching clouds..

Contact

 em_henni@esi.dz

 [linkedin.com/in/hennimohammed](https://www.linkedin.com/in/hennimohammed)