Optimization of the JINR Cloud's Efficiency

N. Balashov, A. Baranov, S. Belov, I. Kadochnikov,
V. Korenkov, N. Kutovskiy, A. Nechaevskiy, I. Pelevanyuk

NEC'2017, Budva, 28 September 2017
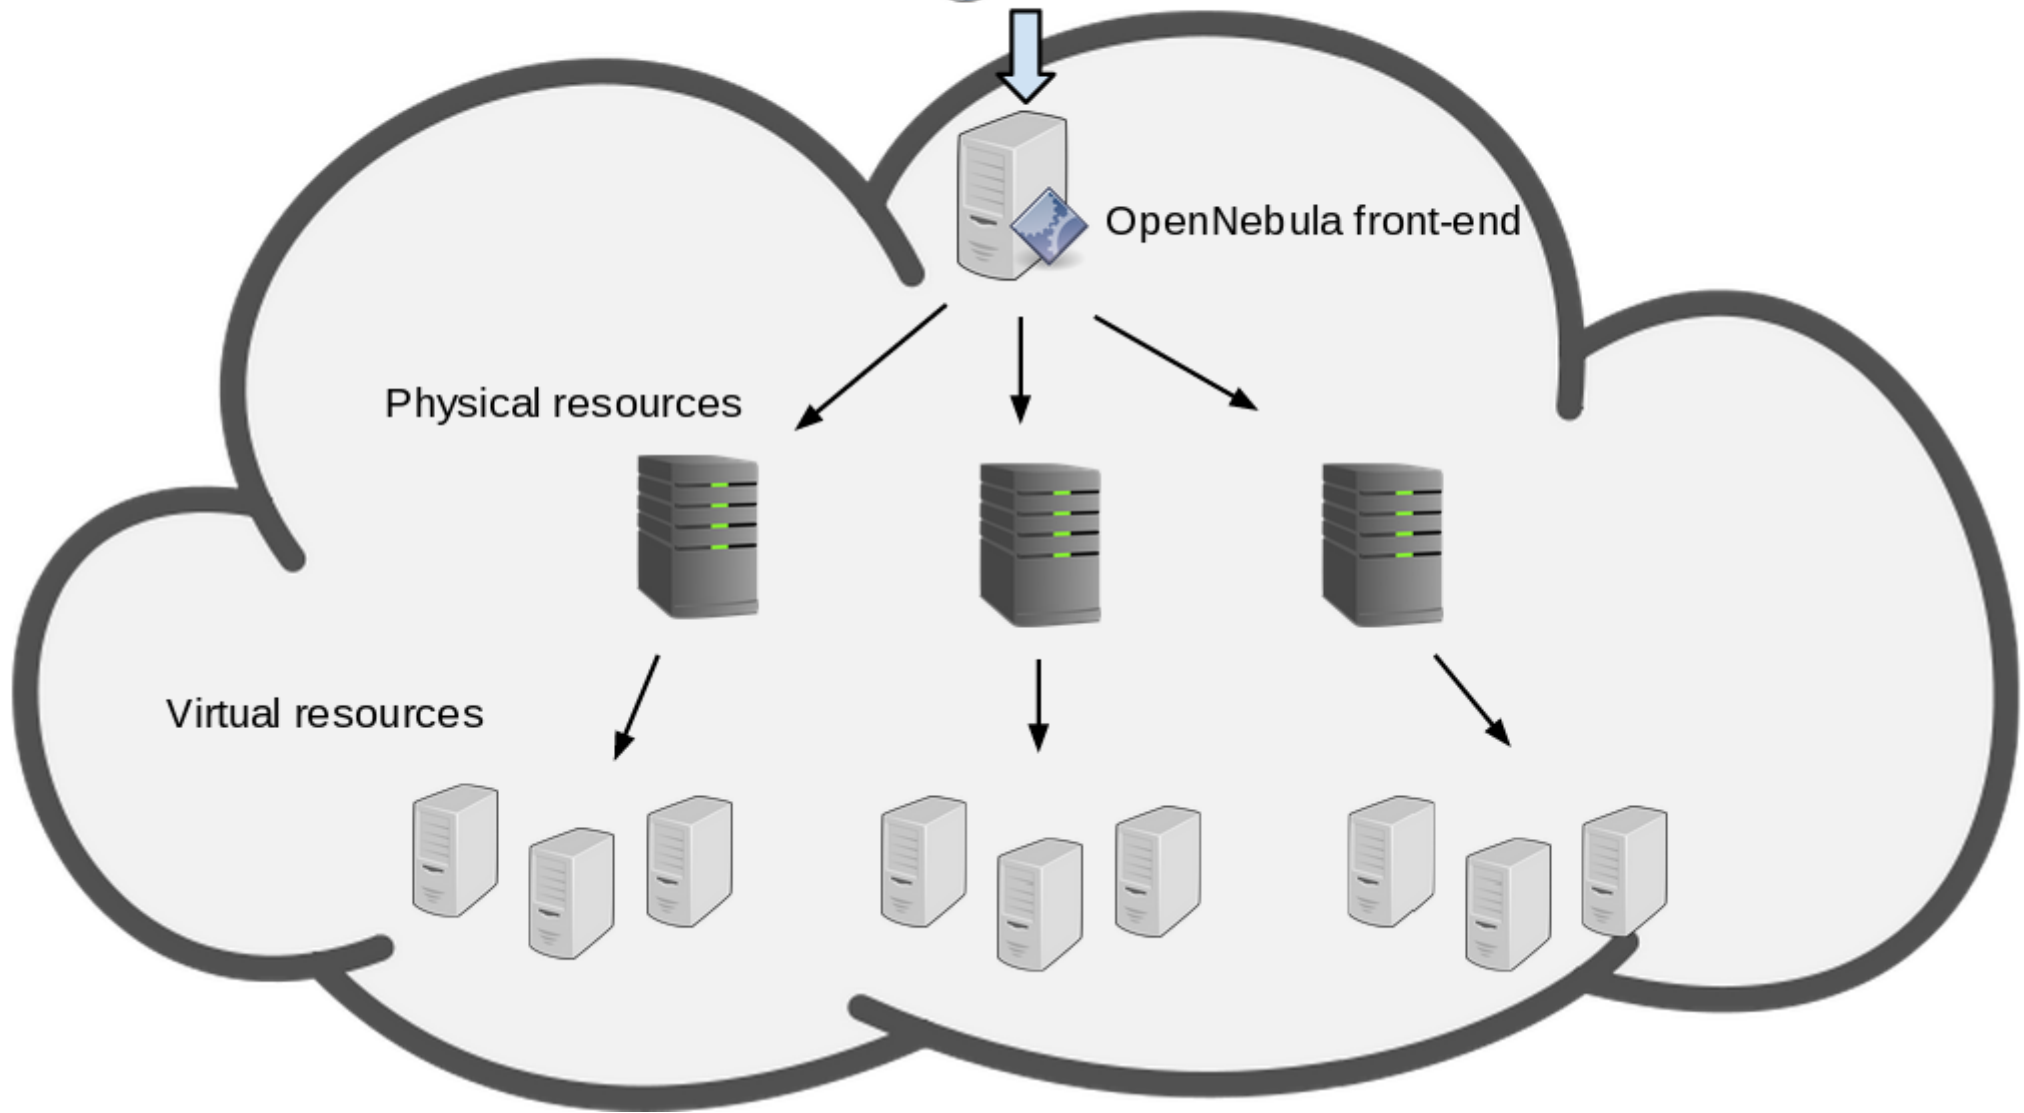
JINR Cloud

Sunstone GUI

OpenNebula front-end
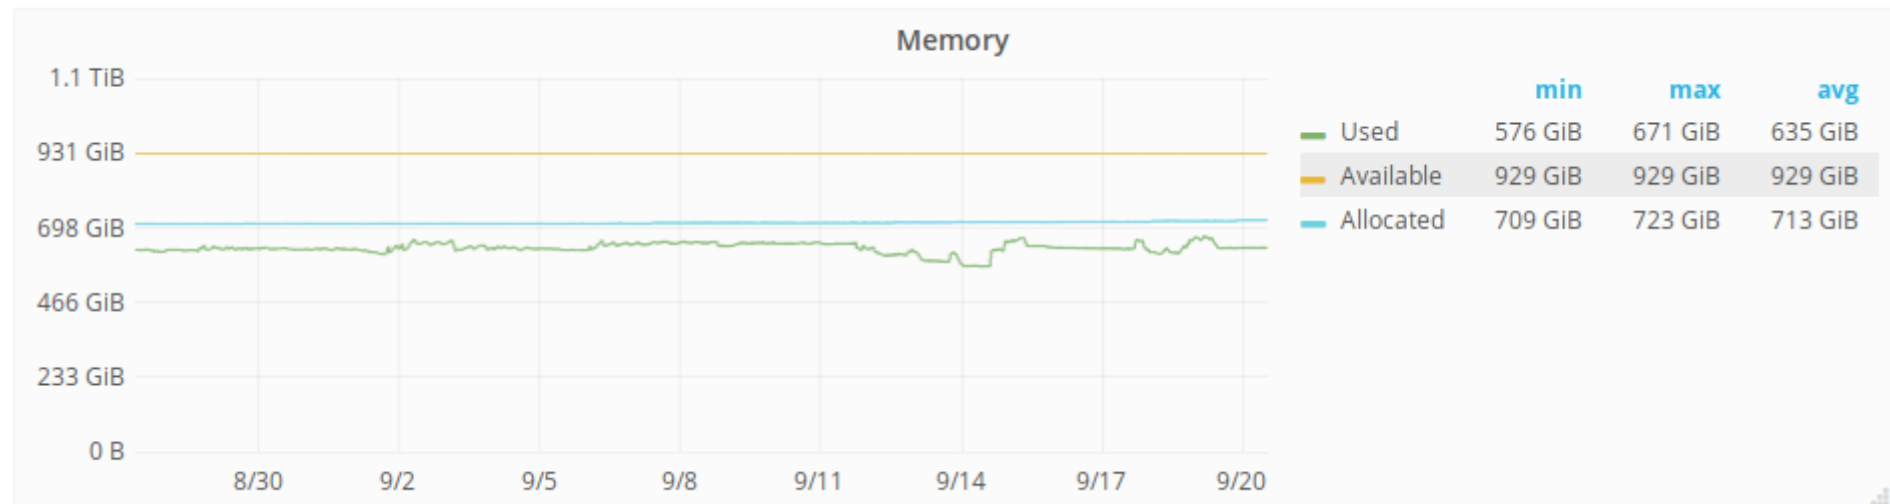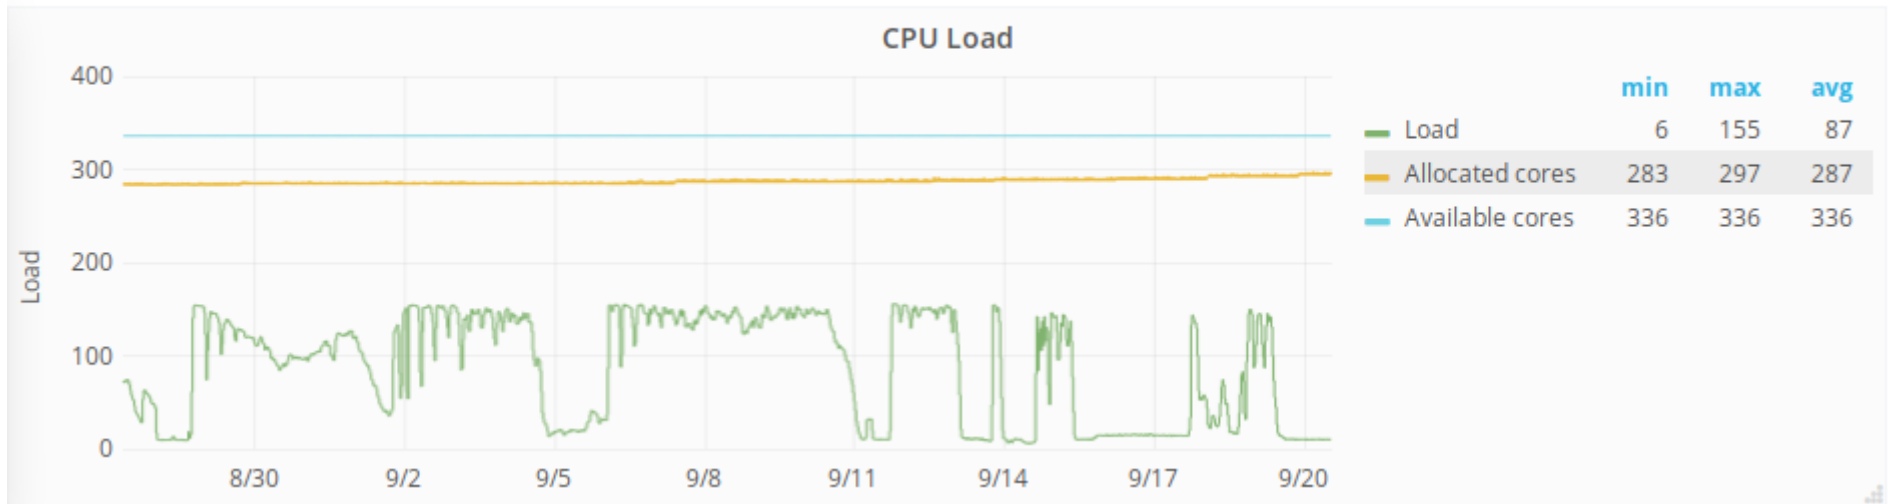
Physical resources

Virtual resources

Our users

- ## Developers
  - development, testing and debugging various apps in various environments.

- ## System administrators
  - testing and studying specifics of installation and operation of new apps or testing updates

- ## PC-style users
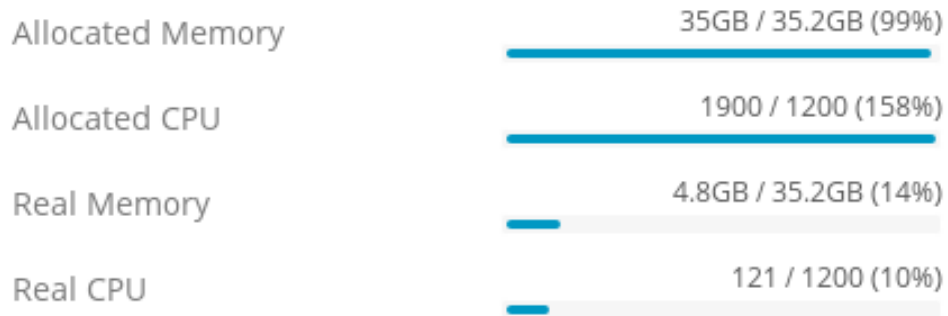 - Physicists

- ## Automated systems
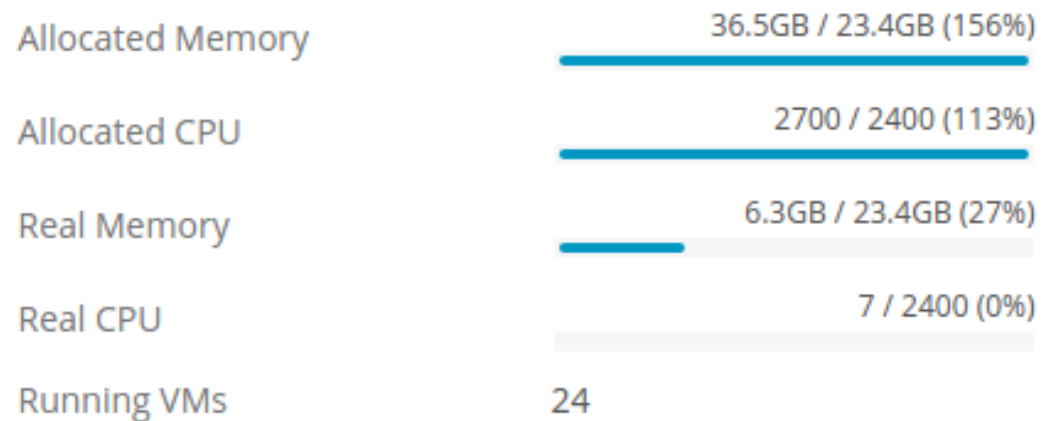 - BES-III, NOvA

Workloads

Overcommitment ratios

Capacity

Allocated Memory	35GB / 35.2GB (99%)
Allocated CPU	1900 / 1200 (158%)
Real Memory	4.8GB / 35.2GB (14%)
Real CPU	121 / 1200 (10%)

Host 1:

- High CPU overcommitment
- No memory overcommitment

Host 2:

- Low CPU overcommitment
- High memory overcommitment

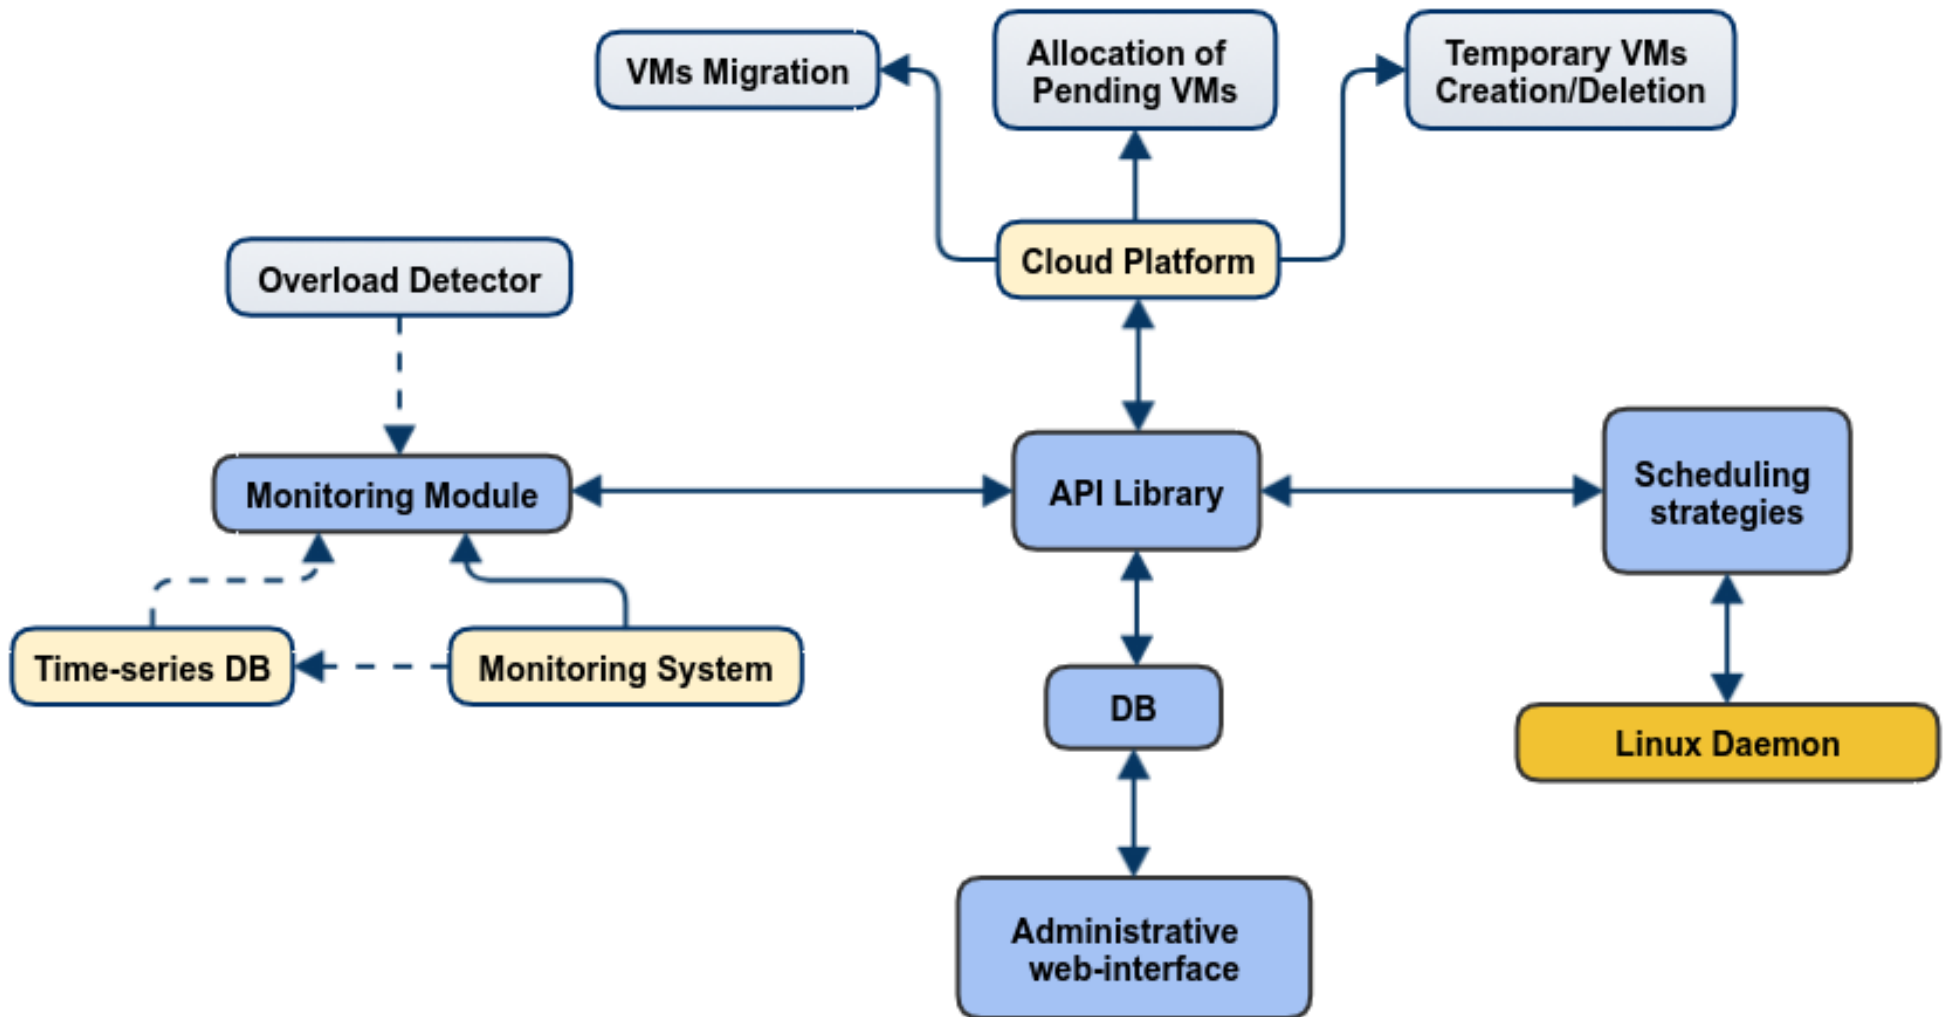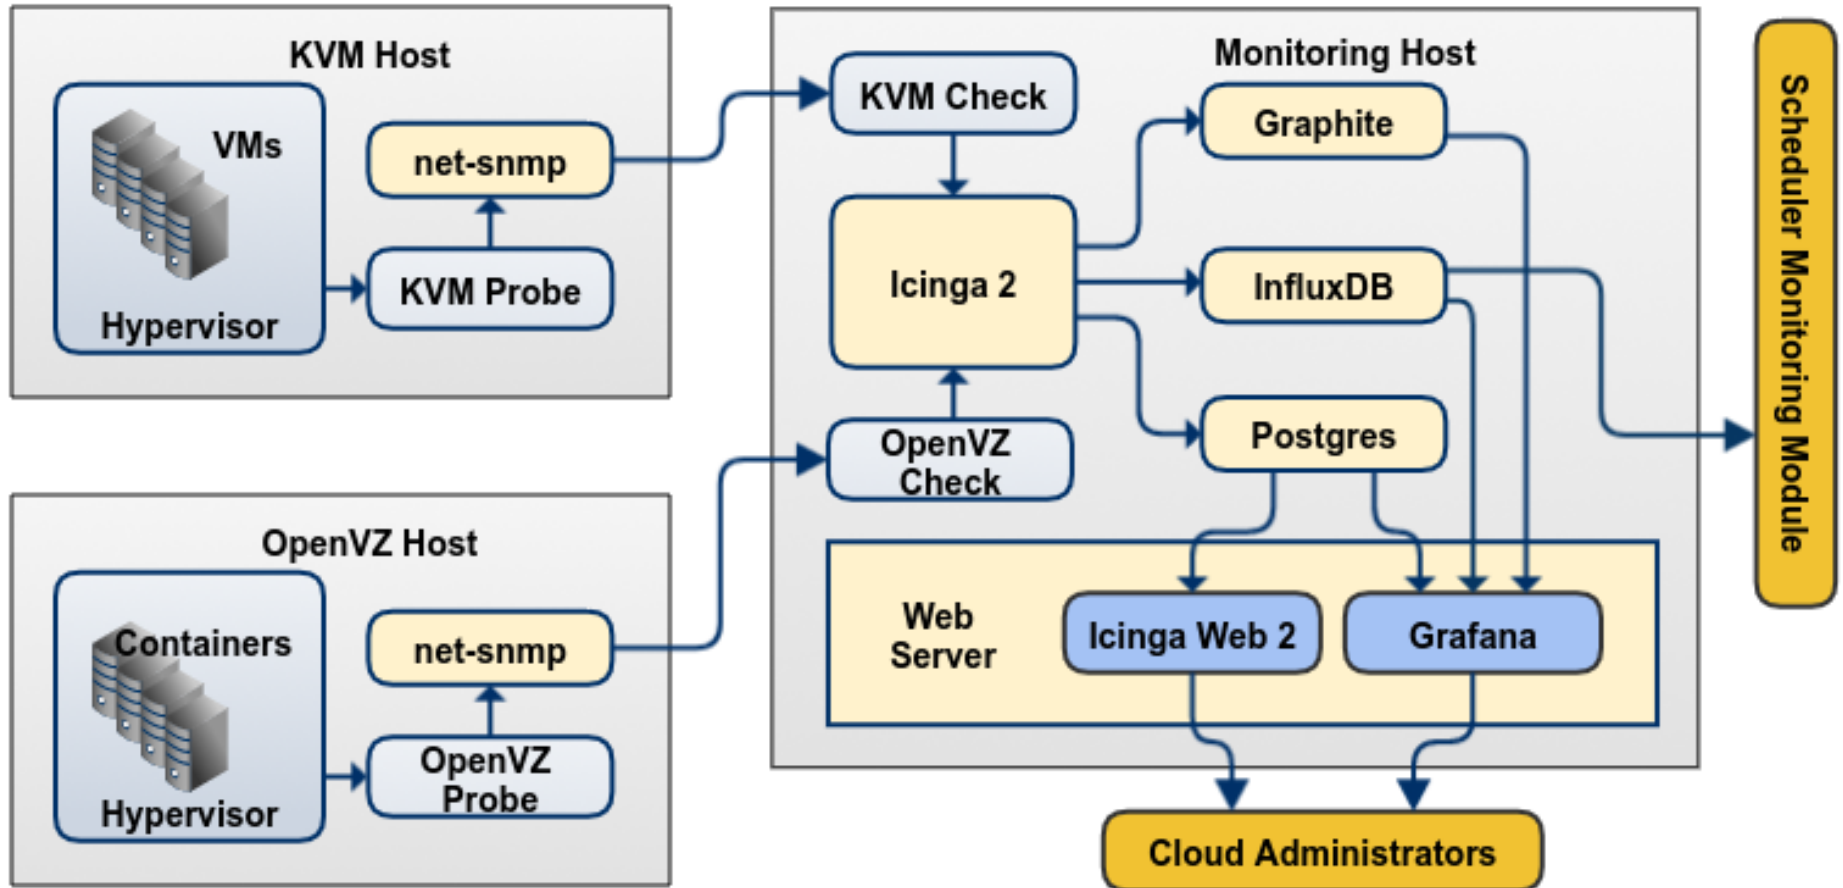|	
Allocated Memory	36.5GB / 23.4GB (156%)
Allocated CPU	2700 / 2400 (113%)
Real Memory	6.3GB / 23.4GB (27%)
Real CPU	7 / 2400 (0%)
Running VMs	24

Possible solutions

- OpenStack Neat
  - Supports OpenStack only
  - Development has stopped
- Green Cloud Scheduler for OpenNebula
  - Outdated

The goal: minimize energy consumption

Scheduler scheme

Monitoring

Administrative web-interface

VMs classification

Dynamic VM re-allocation

New VM

No overcommitment   Low overcommitment   High overcommitment

Dynamic VM re-allocation

New VM

VR_{max}

R_{max}

. . .

VR ⬈

VR_i

VR ⬊

R_{i+1}

R_i

R_{i-1}

Reliability

. . .

Suspend

VR_0

R_0
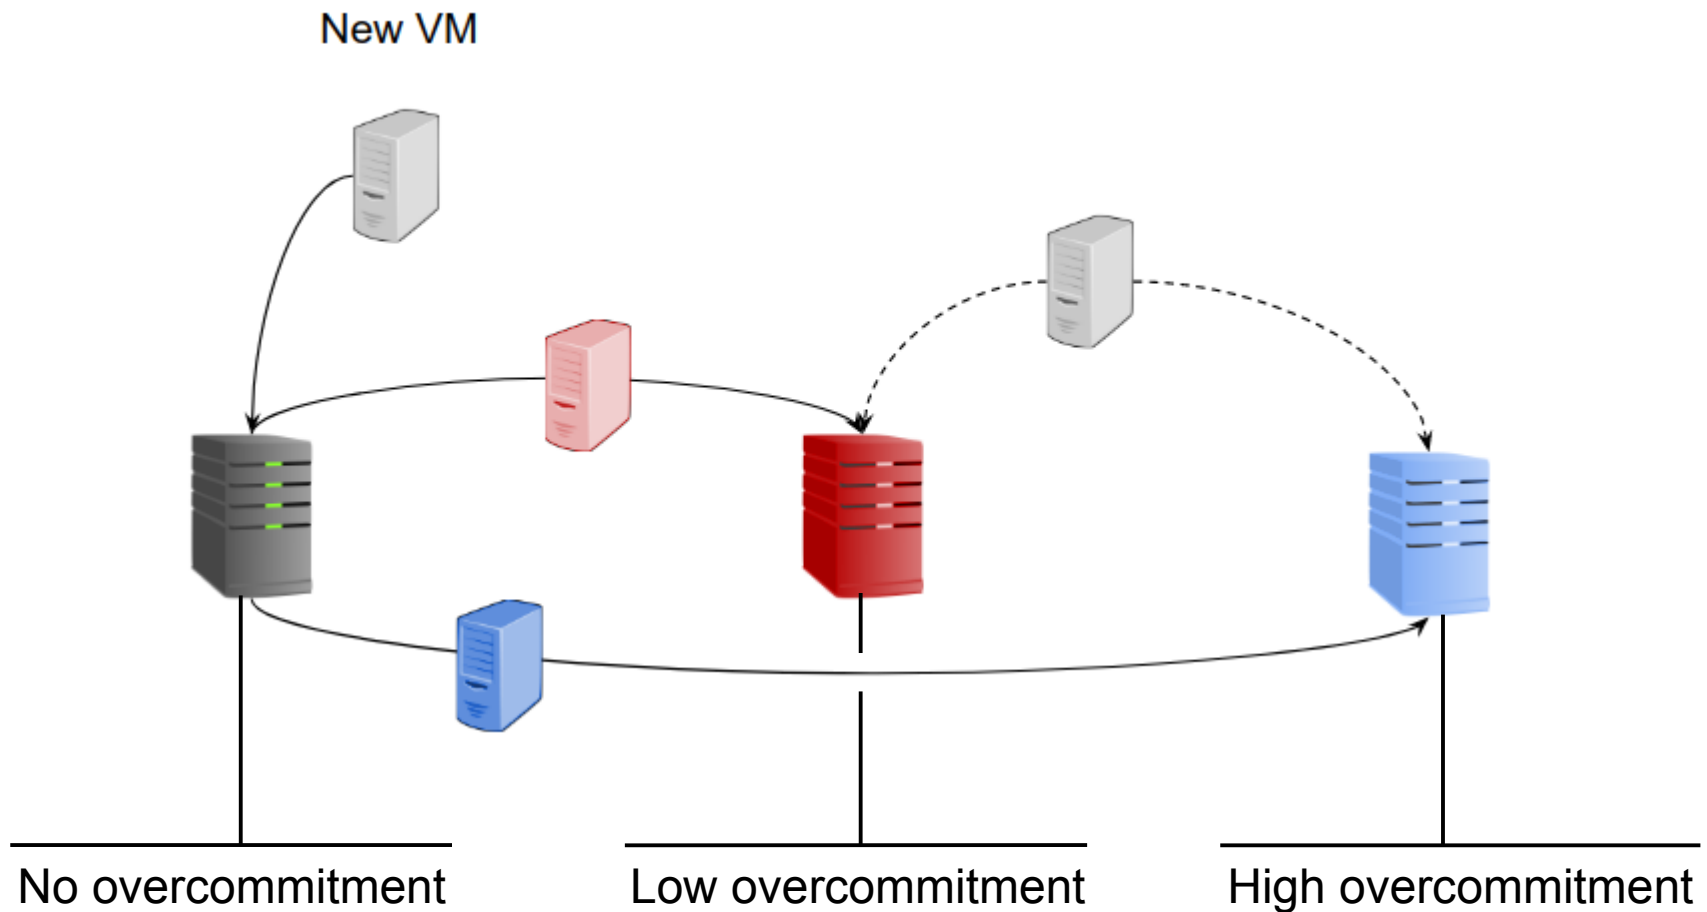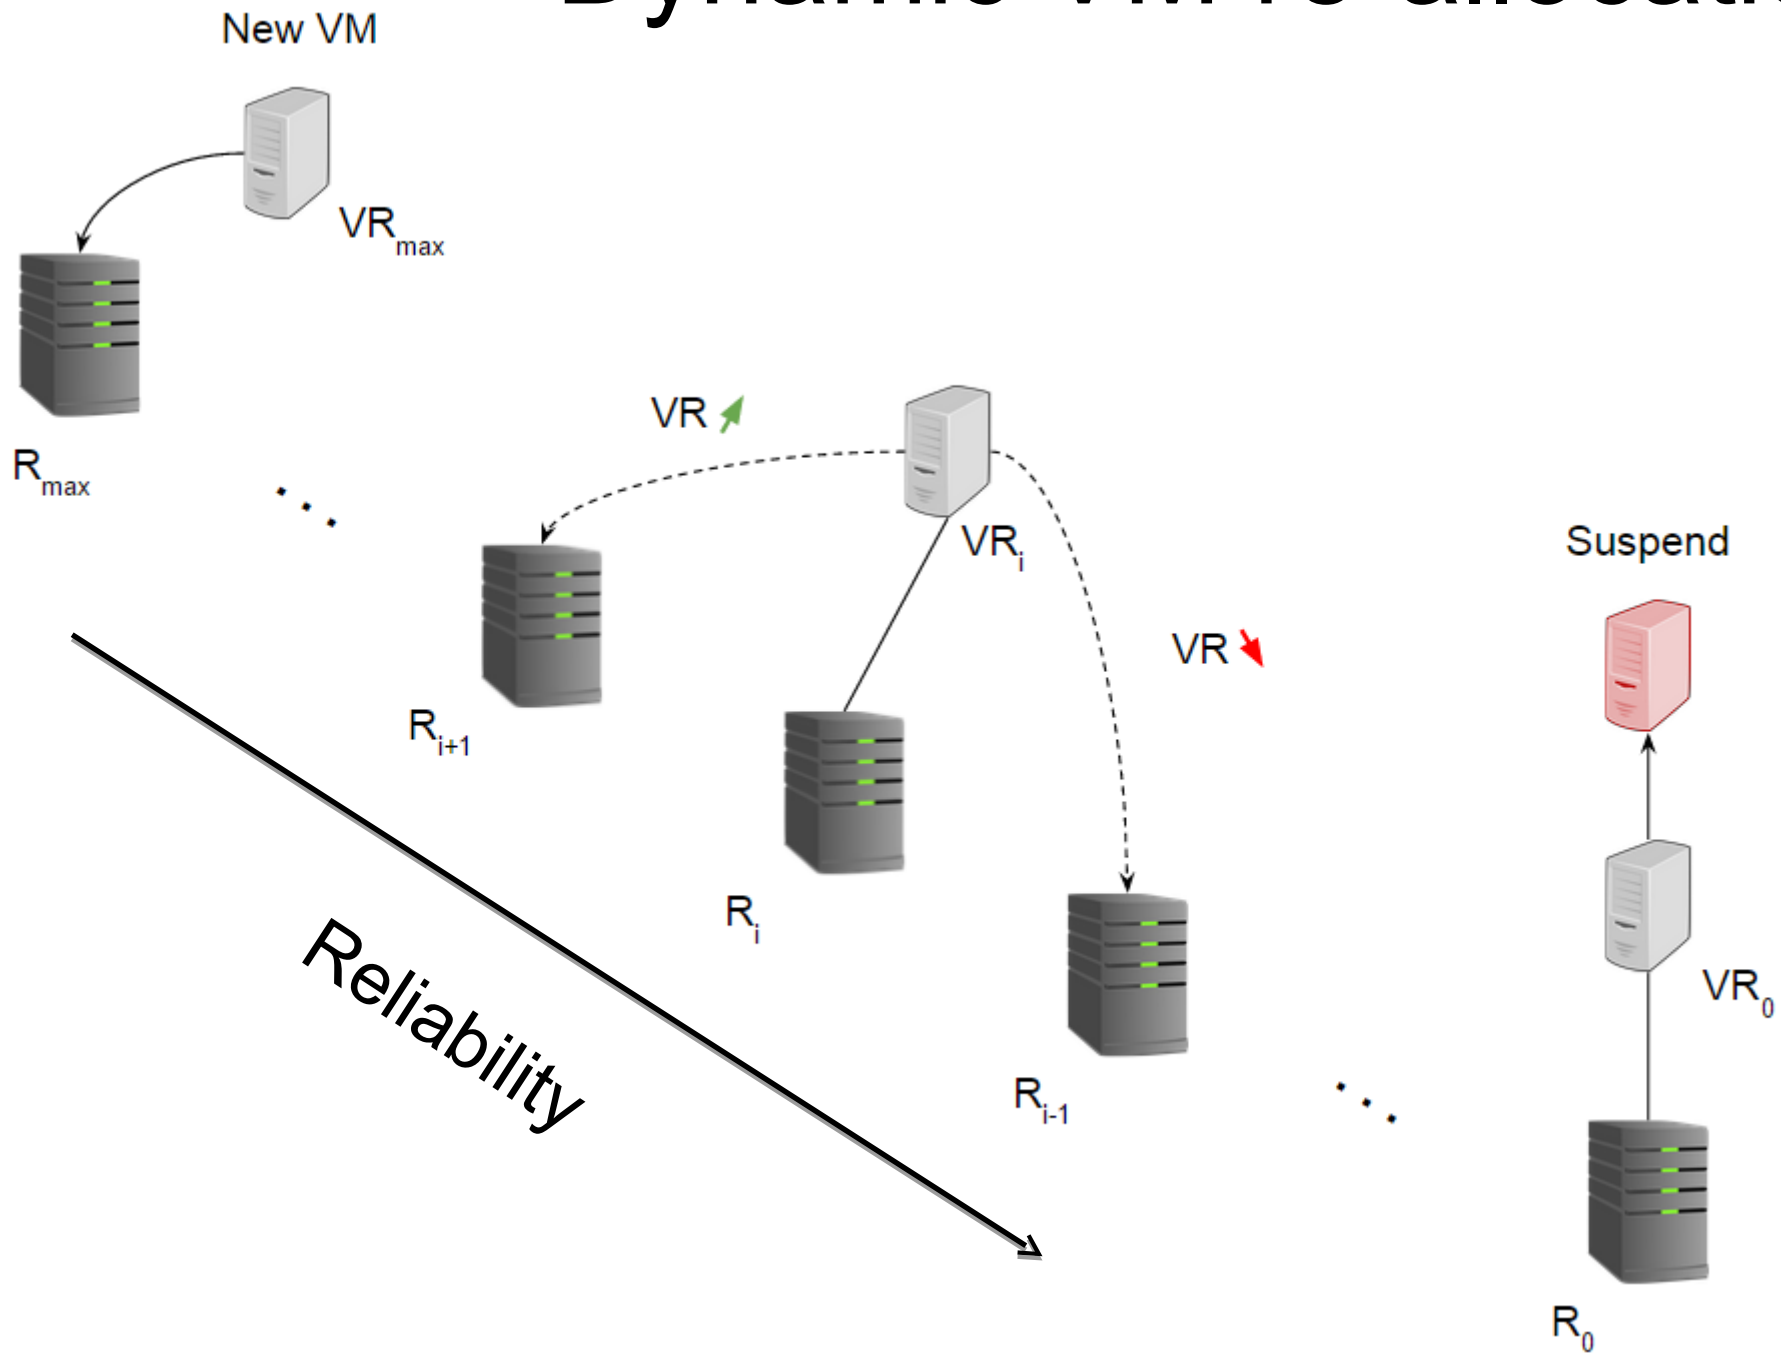
Summary

- The goal of the project is to optimize cloud resources utilization

- A basic software framework was designed
  - Open-source: github.com/jinr-lit

Thanks!