




Joint Institute for Nuclear Research

Cloud Meta-Scheduler for Dynamic VM Relocation

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

GRID'2018, Dubna, 13 September 2018

JINR Cloud

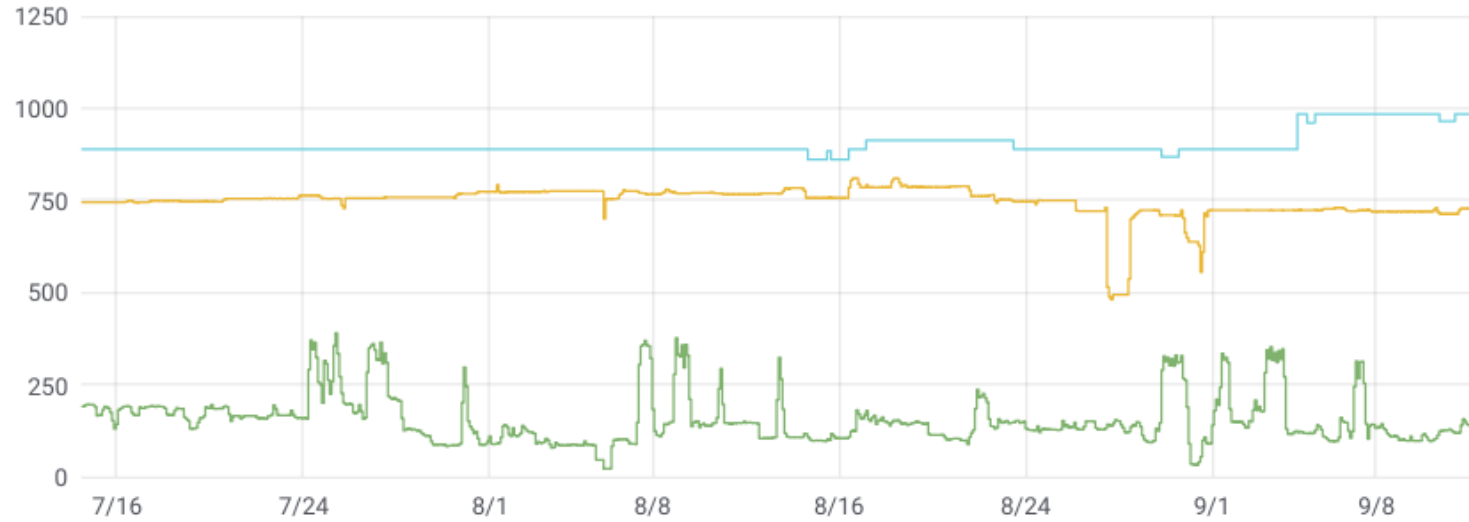


Our users

- **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**
 - Batch-schedulers (DIRAC, HTCondor)

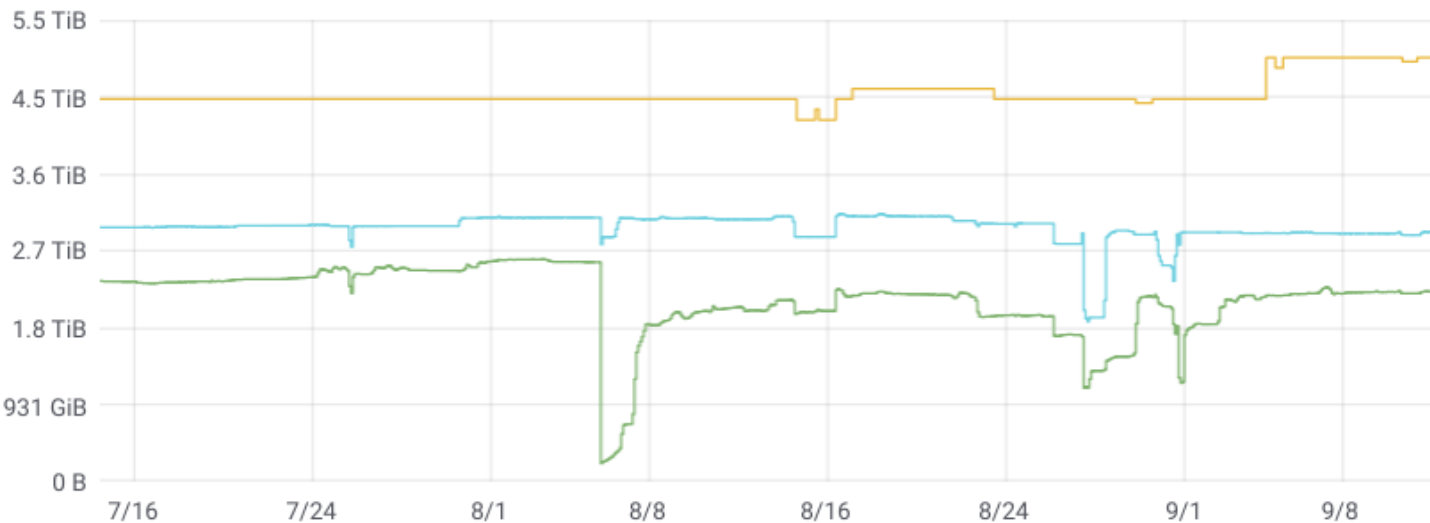
JINR cloud workloads

CPU Load



| | min | max | avg |
|-----------------|-----|-----|-----|
| Load | 21 | 390 | 155 |
| Allocated cores | 481 | 810 | 746 |
| Available cores | 860 | 984 | 902 |


Memory



| | min | max | avg |
|-----------|----------|----------|----------|
| Used | 216 GiB | 2.63 TiB | 2.15 TiB |
| Available | 4.28 TiB | 5.01 TiB | 4.59 TiB |
| Allocated | 1.89 TiB | 3.16 TiB | 3.00 TiB |

Overcommitment ratios

Capacity




Host 1:

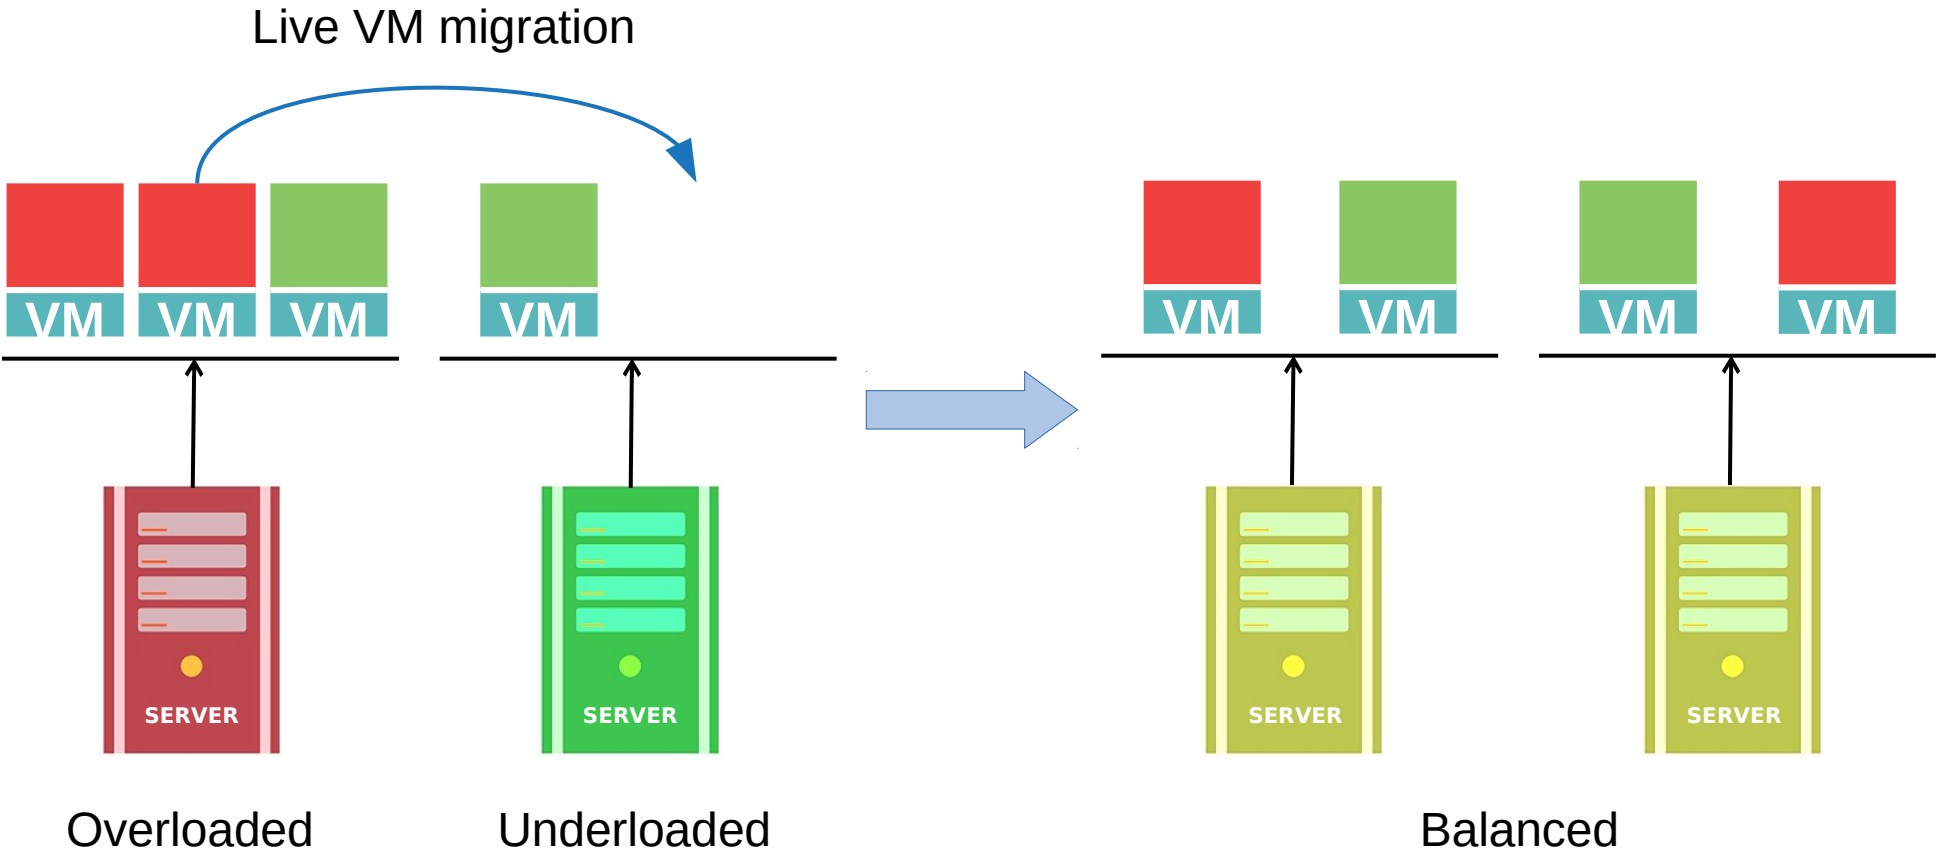
- High CPU overcommitment
- No memory overcommitment

Host 2:

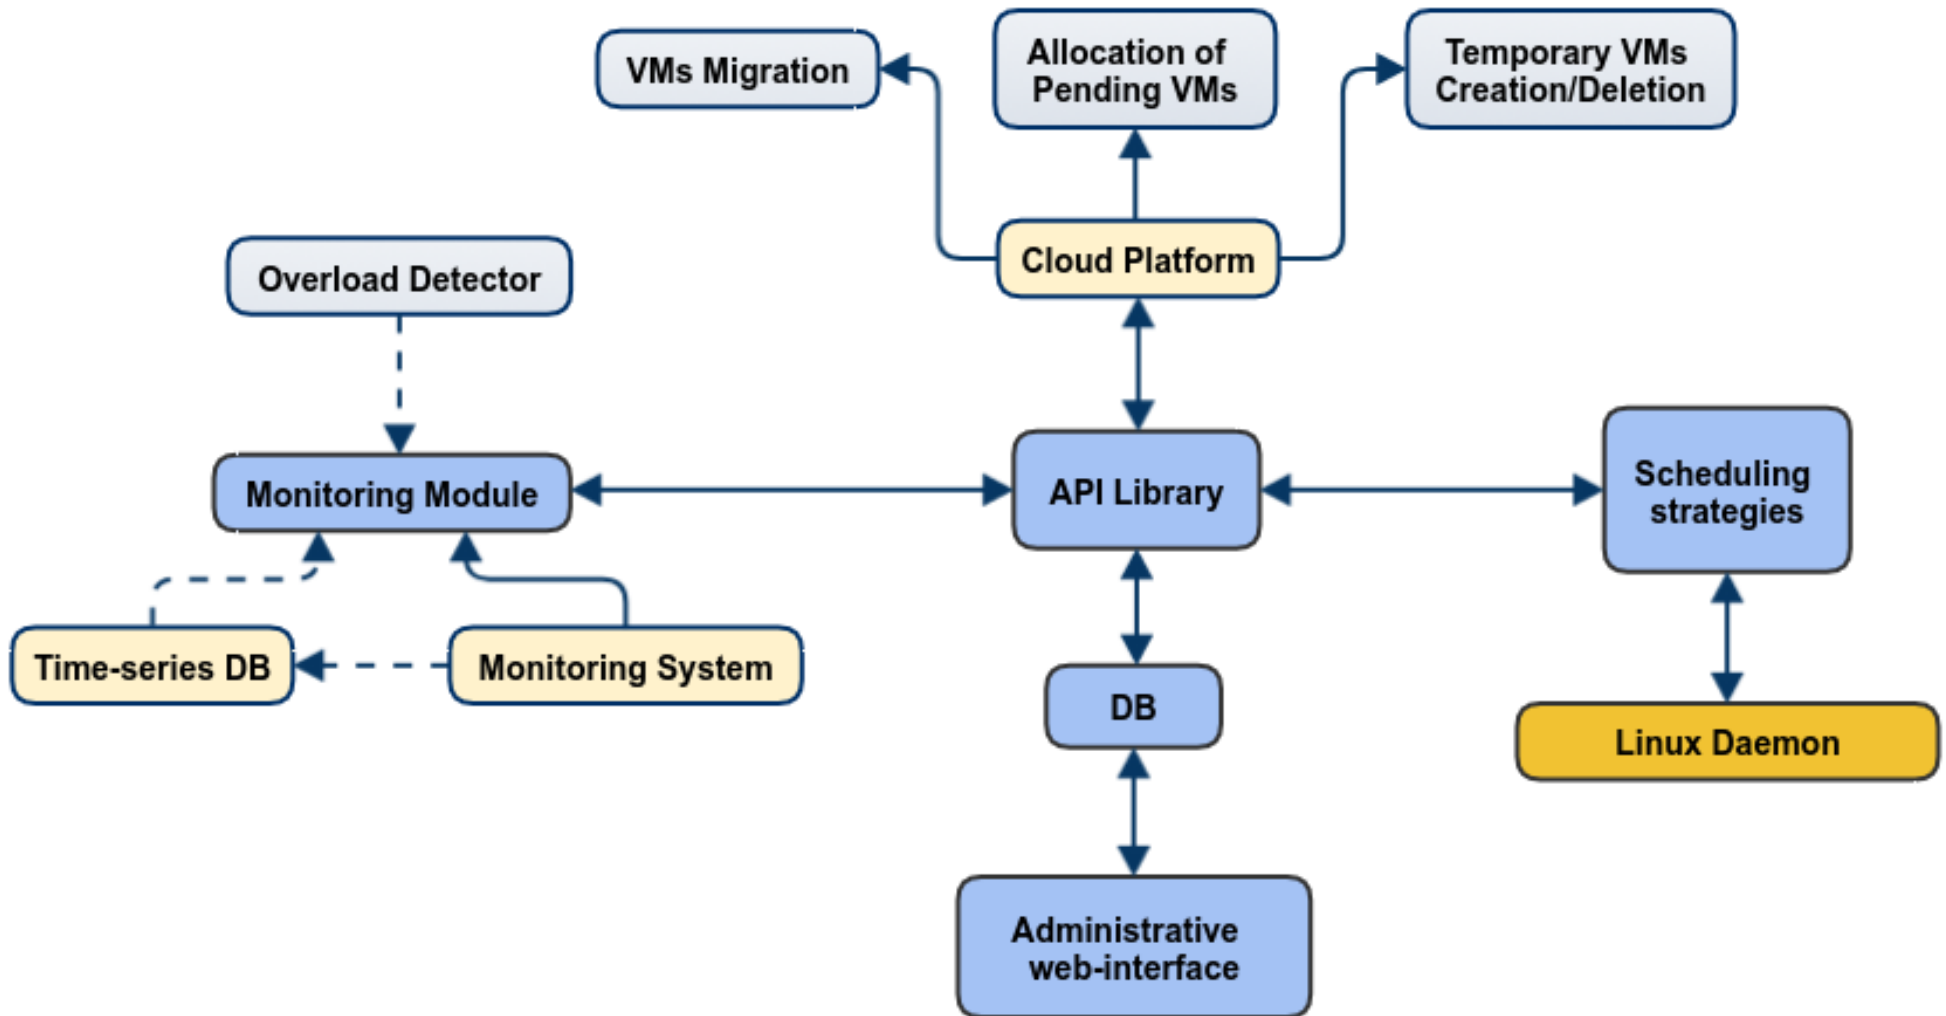
- Low CPU overcommitment
- High memory overcommitment



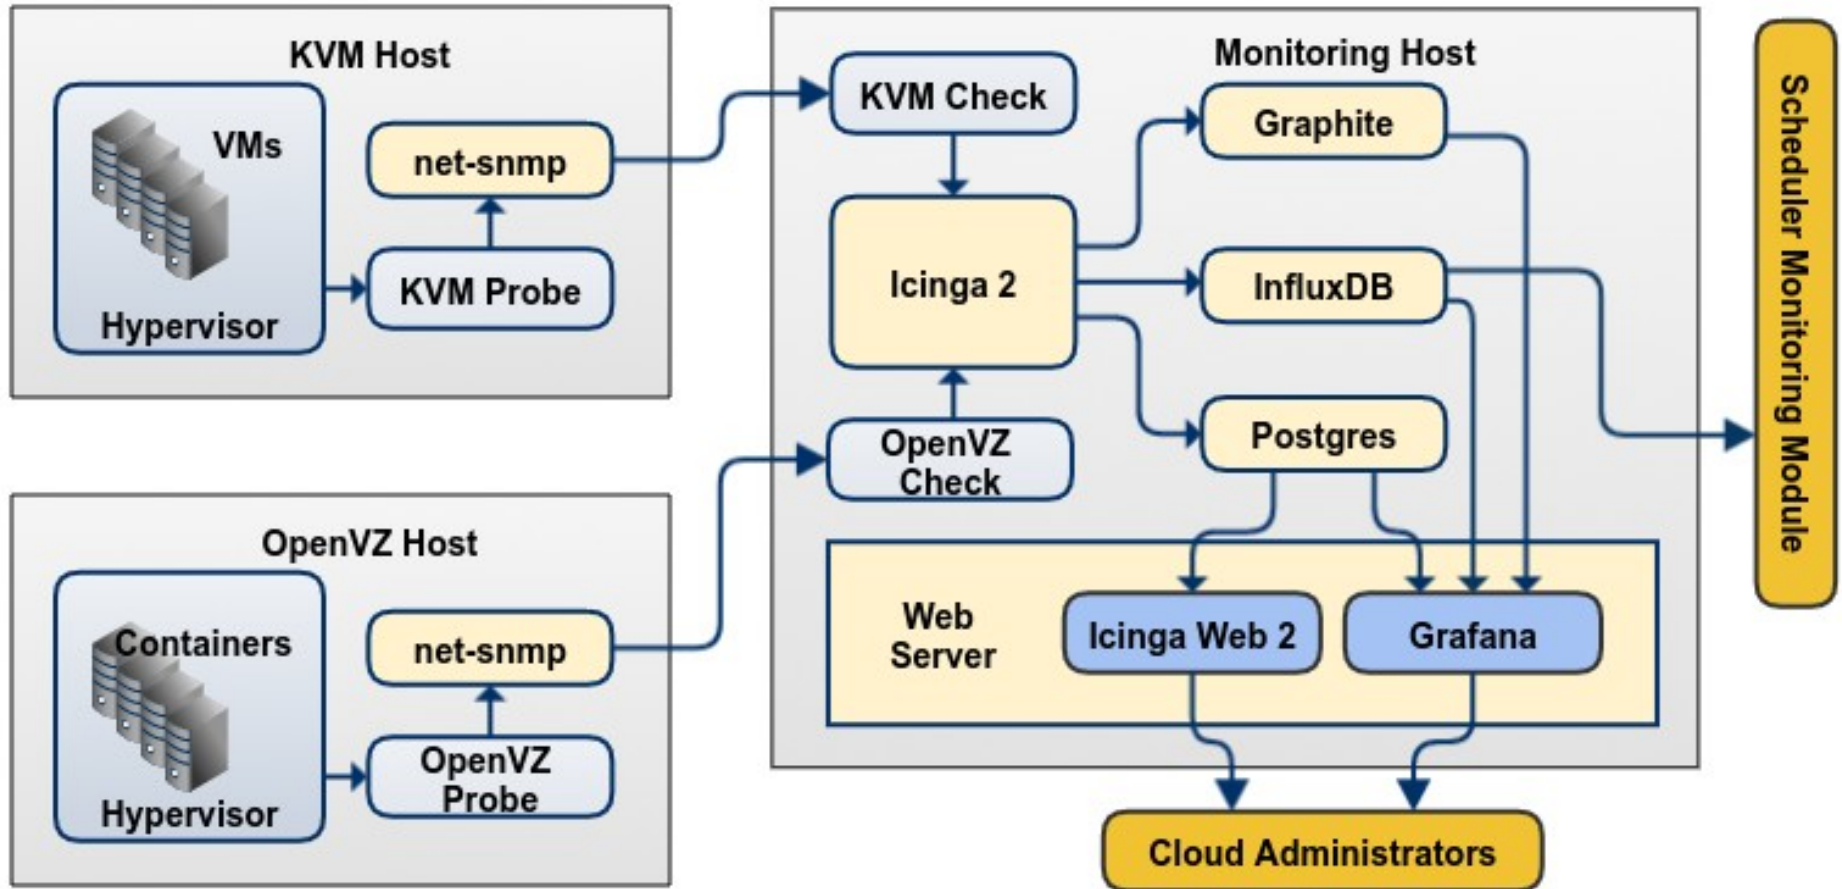
Dynamic relocation



Scheduler scheme



Monitoring



Administrative web-interface

SmartScheduler administration Welcome, root ▾ Recent Actions ▾

Home

Site administration Applications ▾

Authentication and Authorization

Groups

Users

Smartsched_Admin

Load config

Drop config

Clusters

Host groups

Hosts

Vms

Home / Load cloud configuration

Are you sure?

Your configuration is not empty

Show me the differences

Changes in clusters:
No changes.

Changes in hosts:


| ID | Old CPU | New CPU | Old RAM | New RAM | Old cluster | New cluster |
|----|---------|---------|---------|----------|-------------|-------------|
| 92 | - | 2000 | - | 98583256 | - | 106 |

Changes in VMs:


| ID | Old CPU | New CPU | Old RAM | New RAM | Old host | New host |
|-------|---------|---------|---------|---------|----------|----------|
| 33666 | - | 5.0 | - | 23552 | - | 92 |
| 33659 | - | 5.0 | - | 23552 | - | 92 |
| 33660 | - | 5.0 | - | 23552 | - | 92 |
| 33661 | - | 5.0 | - | 23552 | - | 92 |

Yes, I'm sure Take me back

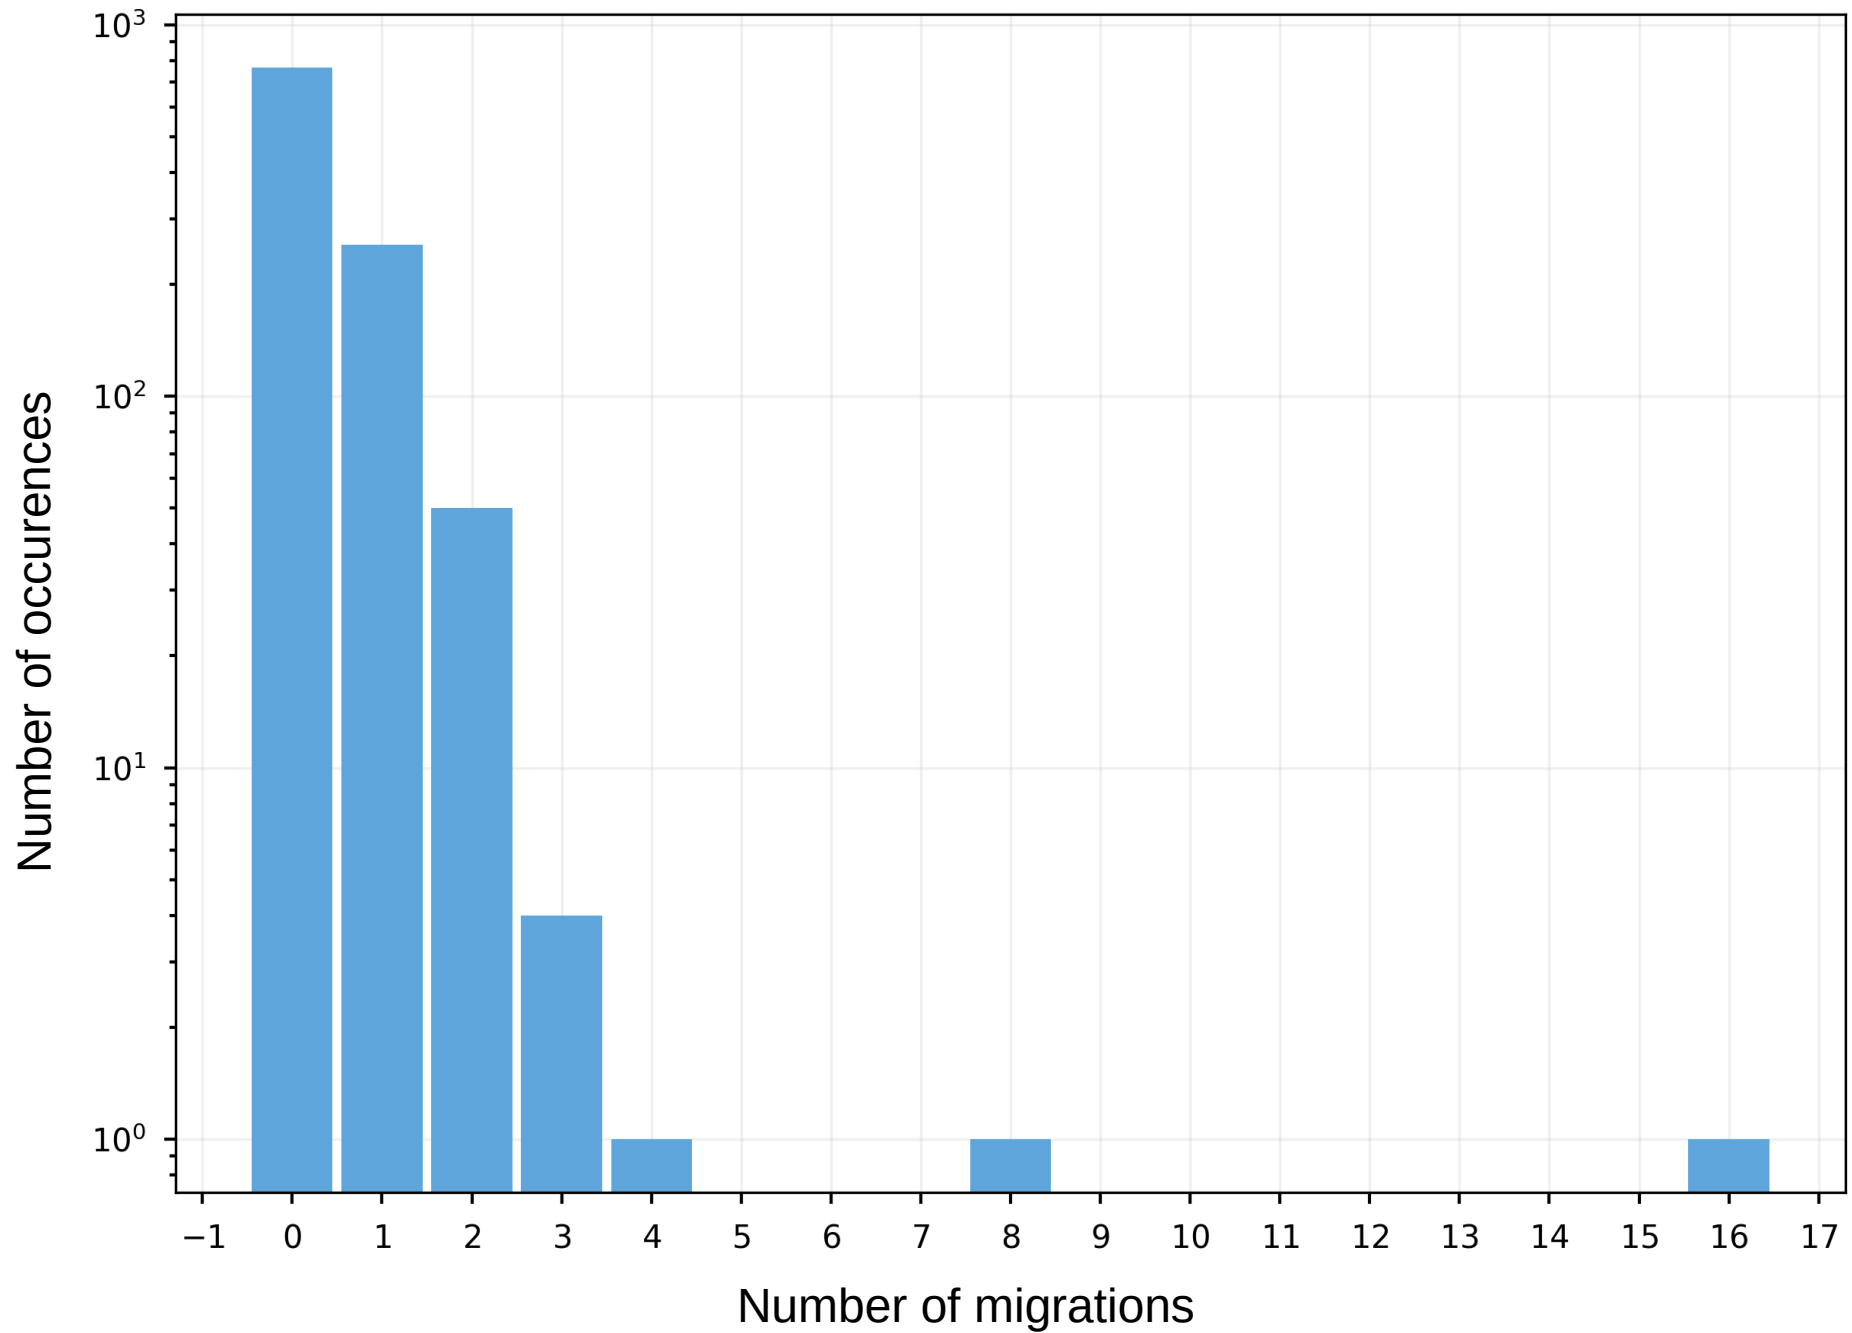
VMs classification



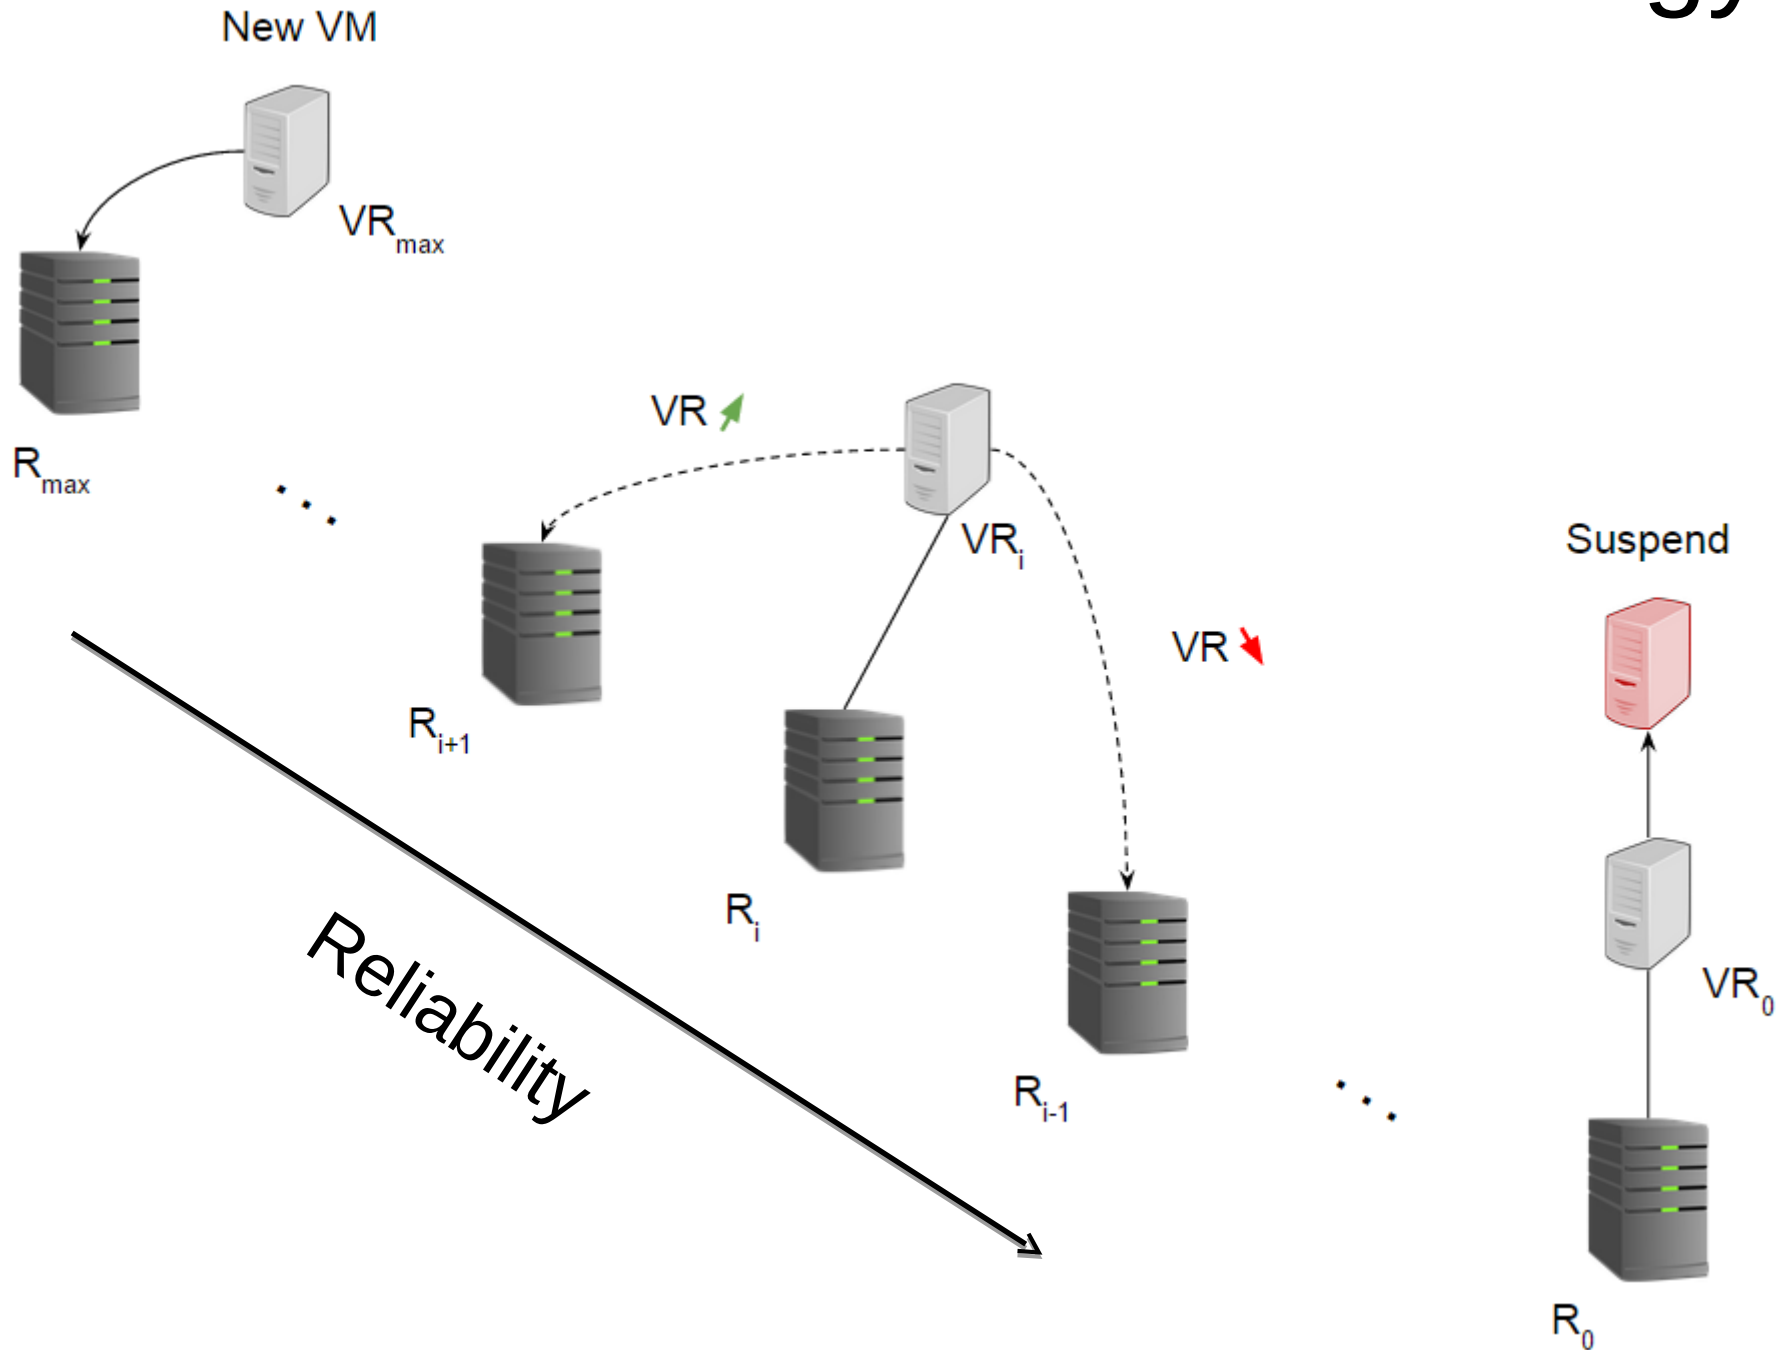
Simple VM relocation strategy



Number of VM class changes



Generalized strategy



Summary

- The goal of the project is to optimize cloud resources utilization
- A basic software framework was designed:
github.com/jinr-lit

Plans

- Implement interfaces for reutilizing idle cloud resources with preemptable batch jobs

Thanks!