

Dynamic optimization of workload on compute nodes in private, public and hybrid clouds

Keynoter:

Alexander Chadin

OpenStack Nova LoadBalancer Project

The goal is to provide dynamic distribution of resources among compute nodes in context of OpenStack distributed system

LoadBalancer
contains

Overload implementation

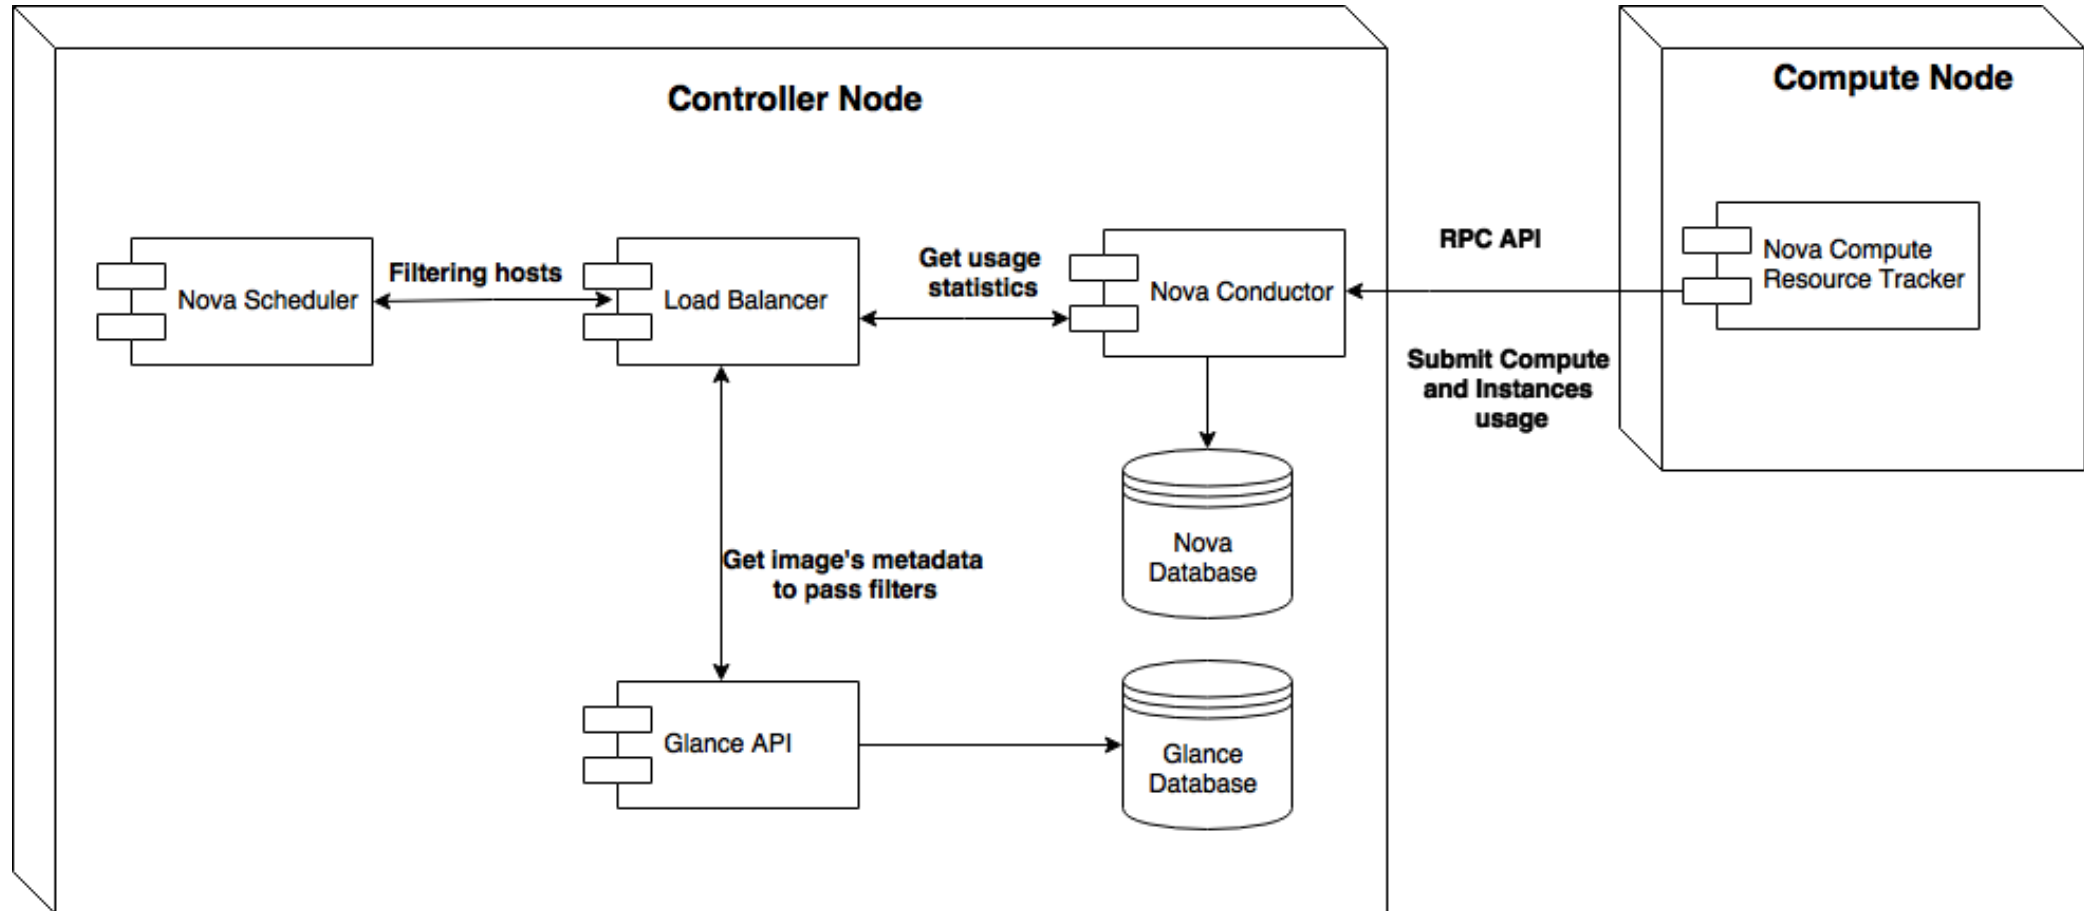
Underload implementation

API

Service files

LoadBalancer Rules

LoadBalancer Deployment Diagram



Typical Threshold Use Case

Collect CPU and RAM values among nodes

```
graph TD; A[Collect CPU and RAM values among nodes] --> B[Get arithmetic mean value]; B --> C[Calculate standard deviation]; C --> D[SD is higher than threshold value]; D --> E[LB should perform some actions to minimize SD];
```

Get arithmetic mean value

Calculate standard deviation

SD is higher than threshold value

LB should perform some actions to minimize SD

Threshold algorithm

Libvirt usage


- $\text{vmload} = k_1 \text{cpu} + k_2 \text{ram}, k_1, k_2 \in \mathbb{Q}$
- $\text{hostload} = \sum \text{vmload}$
- $\text{hostmean} = \frac{\sum \text{hostload}}{N}$
- $L = \sqrt{\frac{\sum_{i=1}^N (\text{hostload}_i - \text{hostmean})^2}{N}}$

Host usage

- $\text{hostload} = \{\text{cpu}, \text{ram}\}$
- $\text{hostmean} = \frac{\sum \text{hostload}}{N}$
- $L = \sqrt{\frac{\sum_{i=1}^N (\text{hostload}_i - \text{hostmean})^2}{N}}$

Overload algorithm

Compute 1



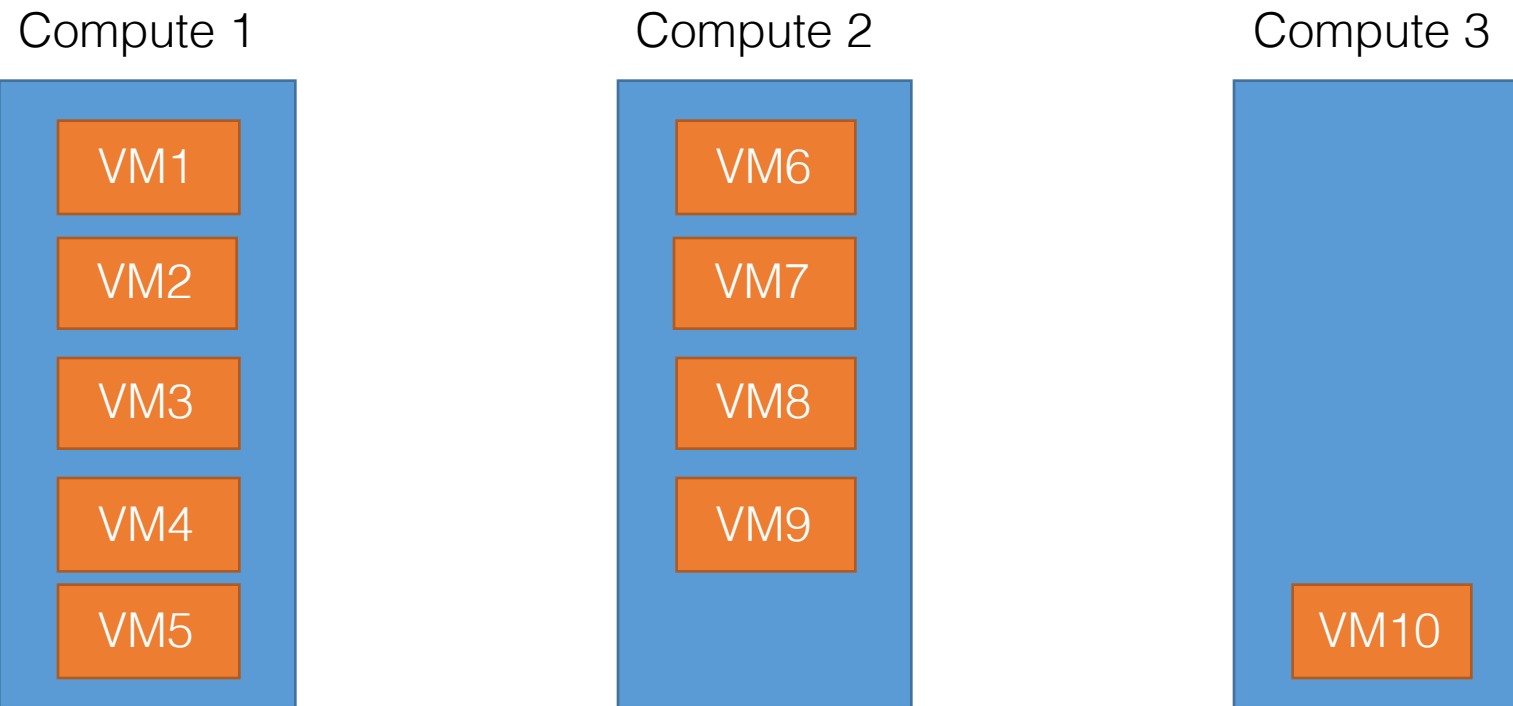
Compute 2



Overload algorithm features

- If standard deviation is greater than the threshold value, the system finds a host-vm pair, that minimizes the standard deviation the best
- If this pair hasn't passed filtering in Nova Scheduler, the balancer will choose the next one

Underload algorithm

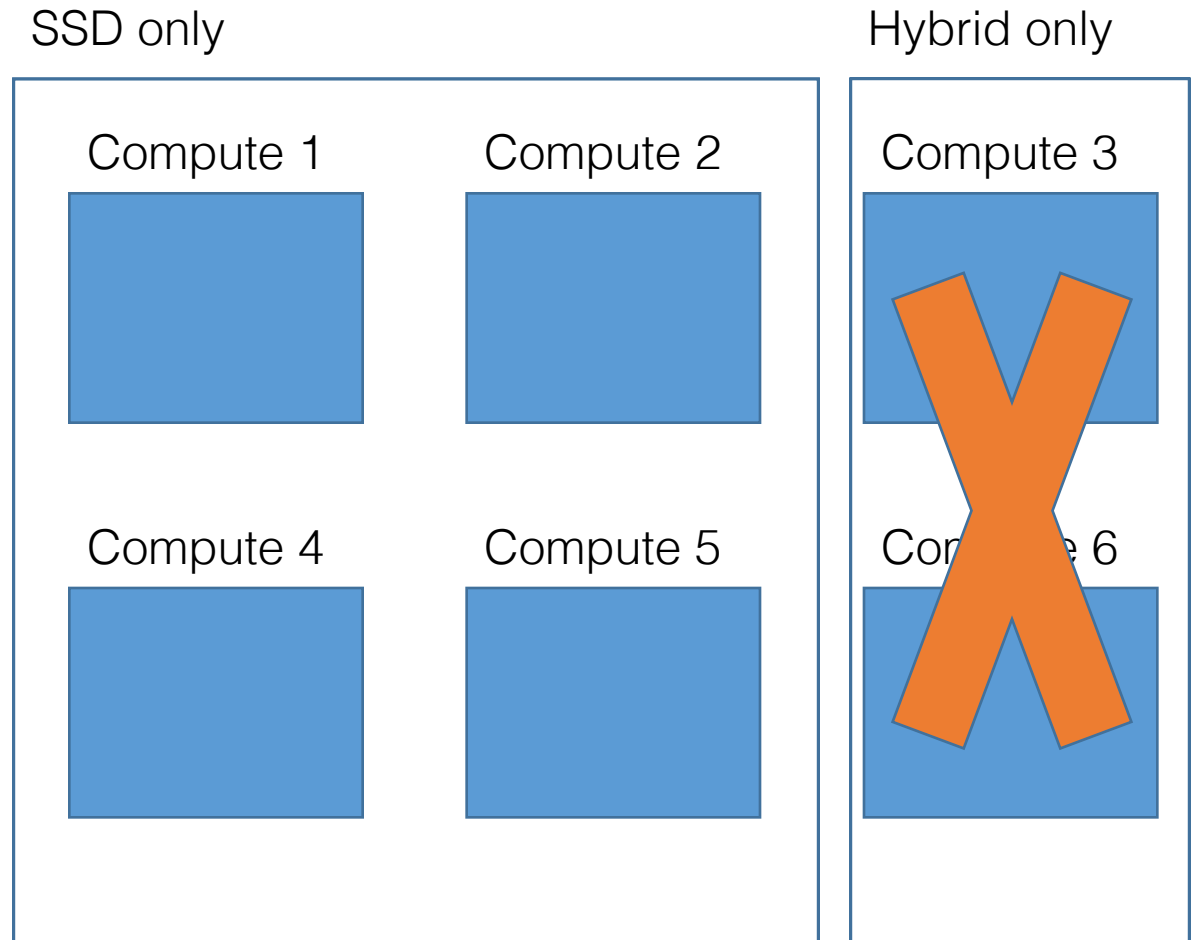


Underload algorithm features

- If one of the nodes is underloaded, the balancer will try to reduce energy consumption of the cluster by migrating all VMs from that node to other ones
- If one of VMs cannot be migrated (using Nova Scheduler), the node will continue working as normal
- ACPI S3 mode, that allows storing system state to RAM, is used in order to reduce energy consumption

LoadBalancer Rules

- Rules allow determining which nodes, host aggregates or availability zones to balance



Future plans

- Integrating Loadbalancer into OpenStack Watcher project
- Providing flexible data usage by introducing a system of plugins
- Inviting more developers to our driver team

Contact us

- Launchpad Project: <https://launchpad.net/nova-loadbalancer>
- Alexander Chadin e-mail: a.chadin@servionica.ru
- Alexander Stavitskiy e-mail: stavitskiy@servionica.ru
- GitHub project commits:
<https://github.com/joker946/nova/commits/drs>