CHINA
OpenStack Days

Workload Balance Based on Watcher Project

Canwei Li
ZTE

AGENDA

- Motivation
- Openstack Watcher
- Workload Balance Strategy

The Reason

- When the VM is deployed, no workload is considered

- During the VM running, the workload may change

The Problem

- Increased energy consumption

- influence the application of VM

Watcher

Watcher provides a flexible and scalable resource optimization service for multi-tenant OpenStack-based clouds.

Provide a framework to implement a wide range of optimization goals

Release Version: 1.0.1

https://wiki.openstack.org/wiki/Watcher

Many Contributors

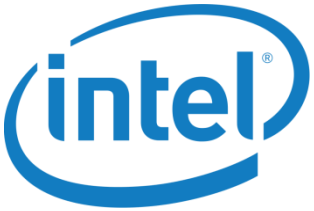

Watcher in the OpenStack Ecosystem

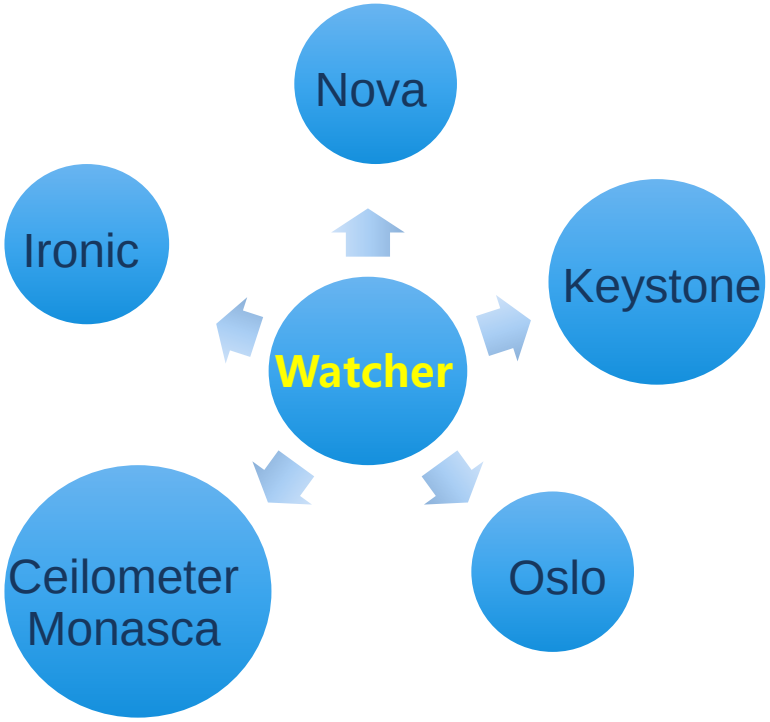

Watcher leverages services provided by other OpenStack projects

- VM live migration and resize
- Metric collection
- Power cycle bare metal nodes

Monitors the infrastructure and performs optimizations on-demand

Enables new ways for OpenStack administrators to reduce The Cloud's TCO

Glossary

Goal: A Goal is a human readable, observable and measurable end result having one objective to be achieved.

Strategy: A Strategy is an algorithm implementation which is able to find a Solution for a given Goal.

Action Plan: An Action Plan specifies a flow of Actions that should be executed in order to satisfy a given Goal. It also contains an estimated global efficacy alongside a set of efficacy indicators.

Action: An Action is what enables Watcher to transform the current state of a Cluster after an Audit.

Audit: In the Watcher system, an Audit is a request for optimizing a Cluster.

Watcher Architecture

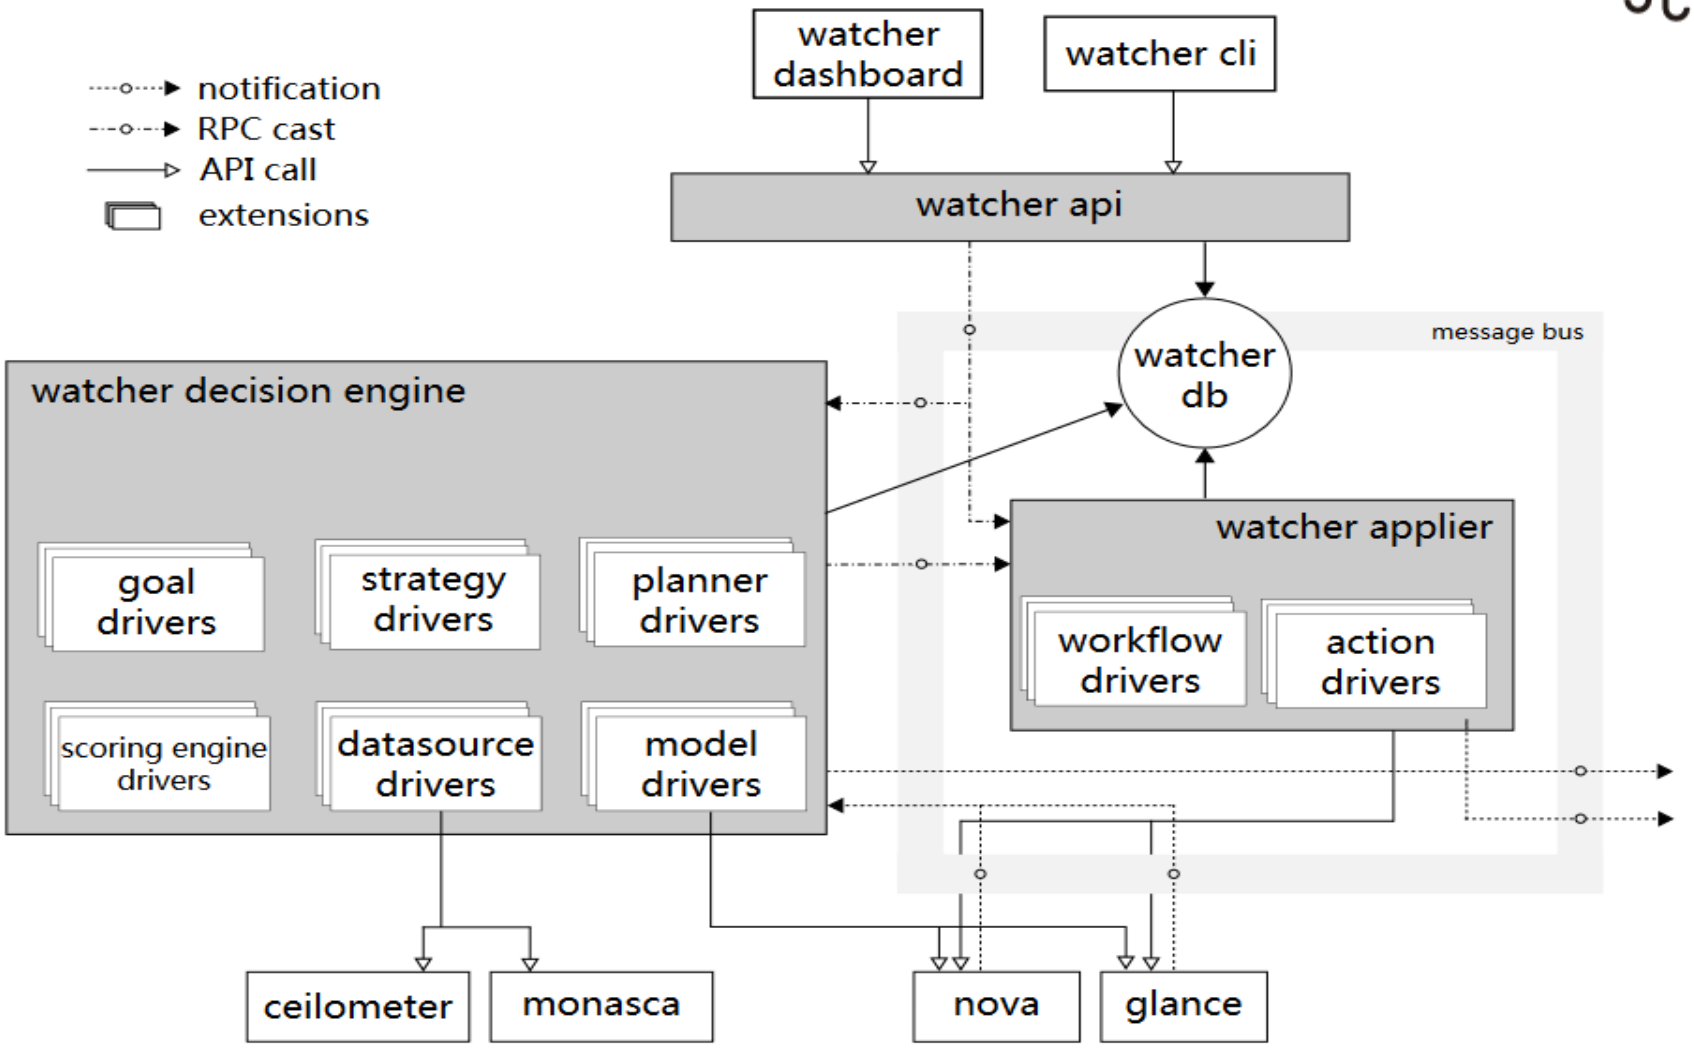

Data Model

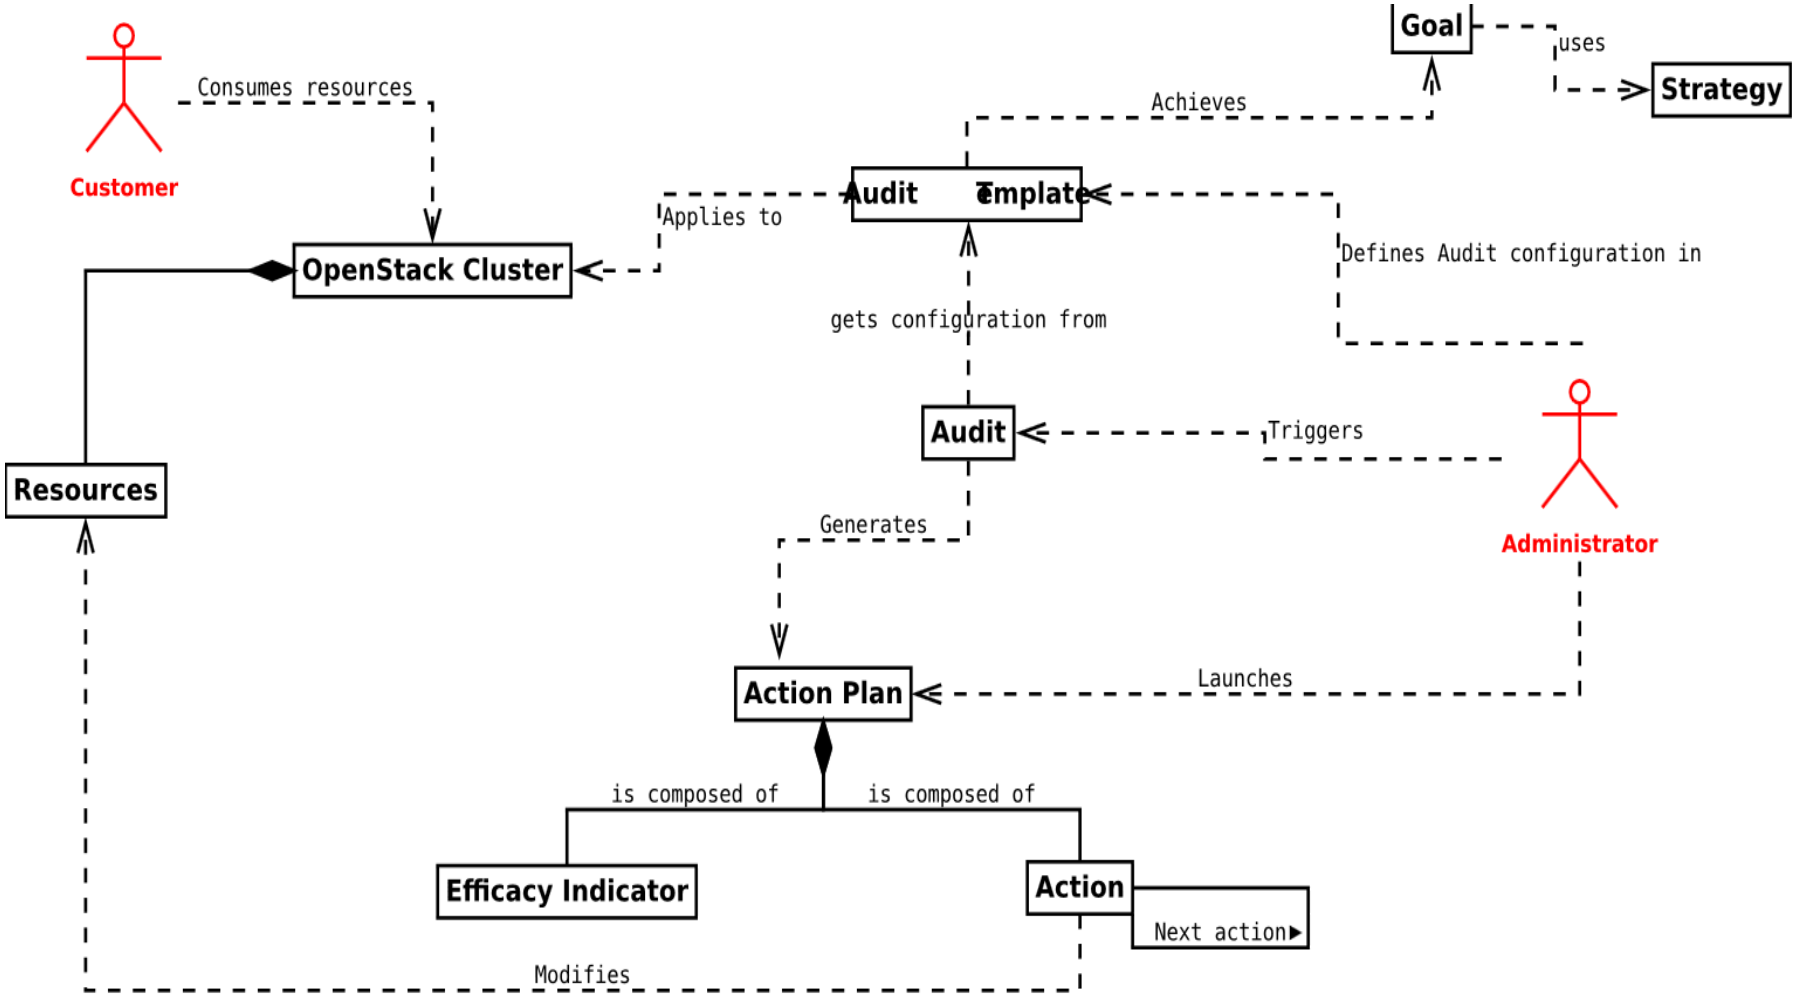
Goal
uses
Strategy
Achieves
Consumes resources
Customer
Audit Template
Applies to
OpenStack Cluster
Defines Audit configuration in
gets configuration from
Audit
Triggers
Administrator
Resources
Generates
Action Plan
Launches
is composed of
is composed of
Efficacy Indicator
Action
Next action
Modifies

Goal of Workload Balance

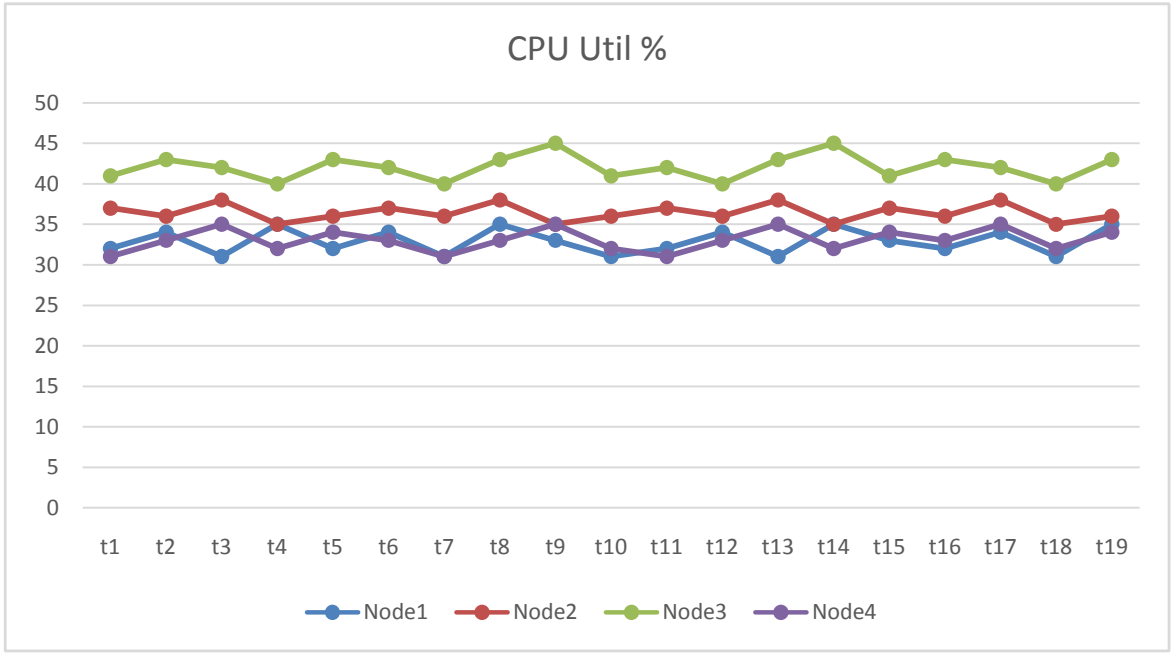

Workload Balance Strategy

Metrics:
cpu_util
memory.usage

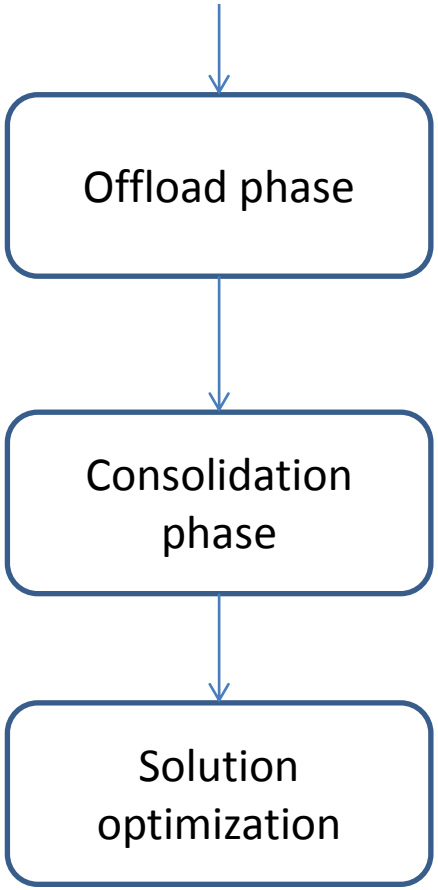

Thresholds:
High
Low

Offload phase

Thresholds High :60%
Low :30%

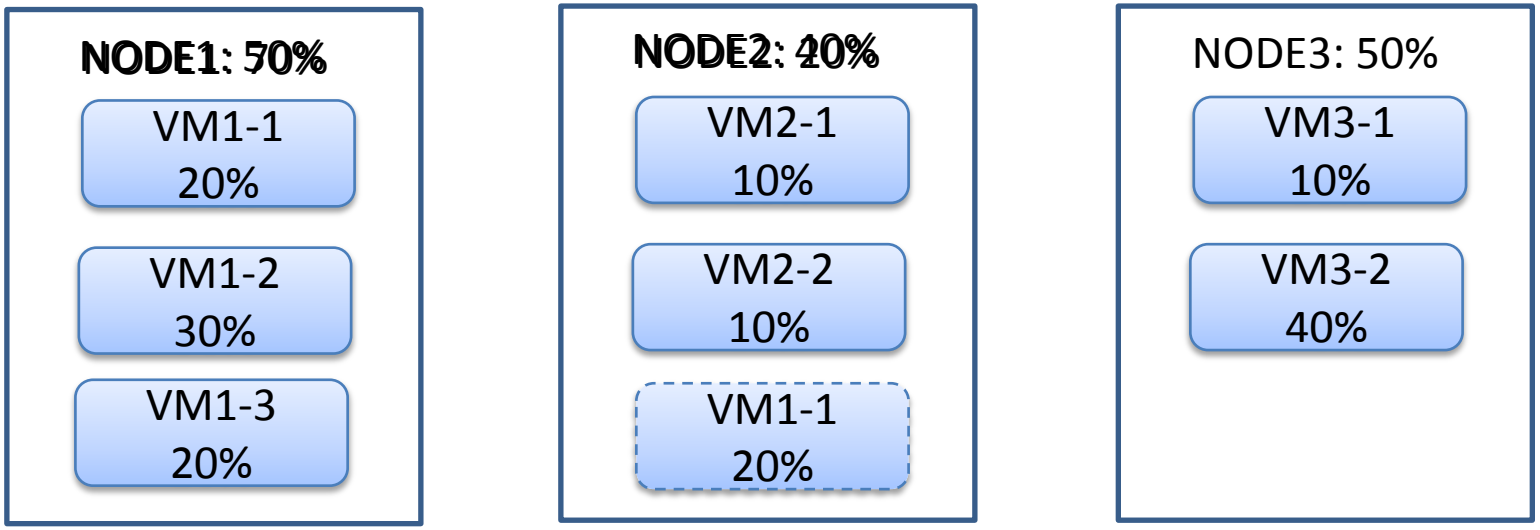

Node CPU %: High->Low Node1 Node3 Node2

Node1: 70% > 60% Node3: 50% Node2: 40%

VM CPU %: Low->High VM1-1 VM1-3 VM1-2

Destination Node: Node2 VM1-1:20% + Node2:20% = 40% < 60%

Consolidation phase

Thresholds High :60%
Low :30%

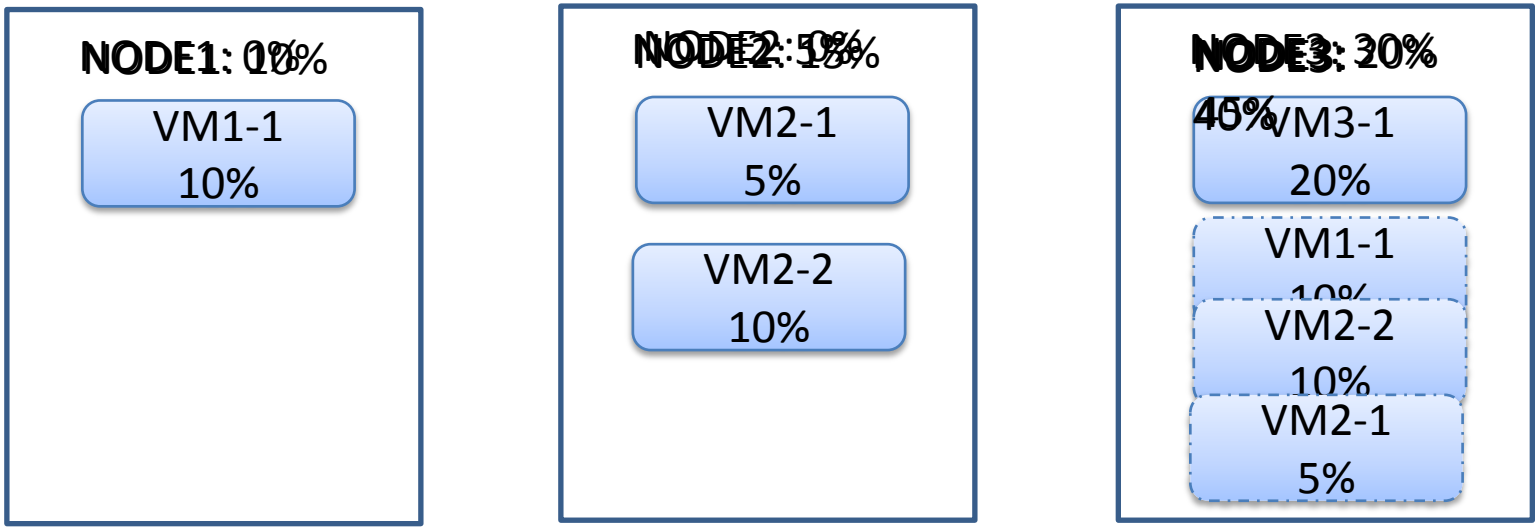

Node CPU %: Low->High

Node1: 0% Node2: 0% Node3: 45%

VM CPU %: High->Low

VM1-1 VM2-1

VM2-2 10% + Node3: 20% = 40% < 60%

Destination Node: Node3

VM2-1 5% + Node3: 40% = 45% < 60%

Node1 Node2 Node3

THANK YOU