

Watcher, the Infrastructure Optimization Service for OpenStack Plans for Pike and Beyond



 **evConf**

[//www.devconf.ru](http://www.devconf.ru)

OpenStack Watcher


- A flexible and scalable resource optimization service for multi-tenant OpenStack-based clouds
- Provide a pluggable framework for optimization strategies (algorithms/metrics)
 - Energy-aware optimizations
 - Workload consolidations and rebalancing optimizations
- Watcher audits the cloud against a set of optimization algorithms and builds a recommended action plan to get to the right objectives
- Integration point with external analytic systems through a pluggable scoring engine

Watcher is part of the OpenStack Big Tent

Many Contributors



b com



ZTE中兴

NEC



at&t



Key Features of Watcher

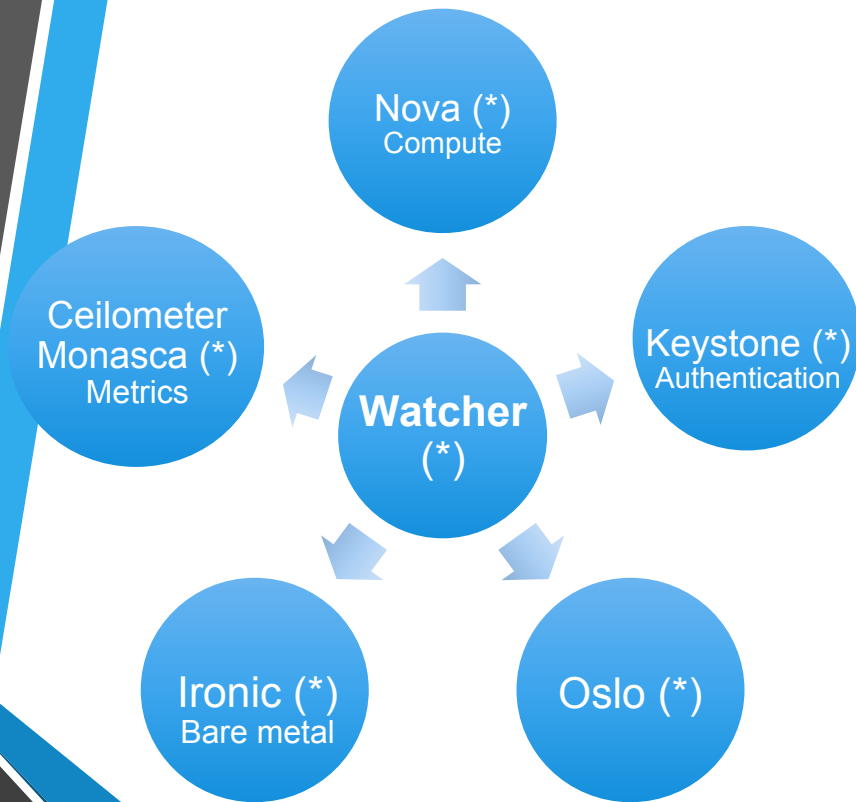
Watcher provides:

- Cloud optimization using VM live migration in case of imbalance detection
- Granularity of optimization with multiple goals setting from a set of hosts to an entire cloud
- Opportunity for evolution via its flexible plugin structure
- « On-the-shelf » optimization strategies based on CPU, memory and energy

Watcher can run in:

- "SINGLE MODE" for auditing before acting
- "CONTINUOUS MODE" for always-on optimization

Watcher in the OpenStack Ecosystem



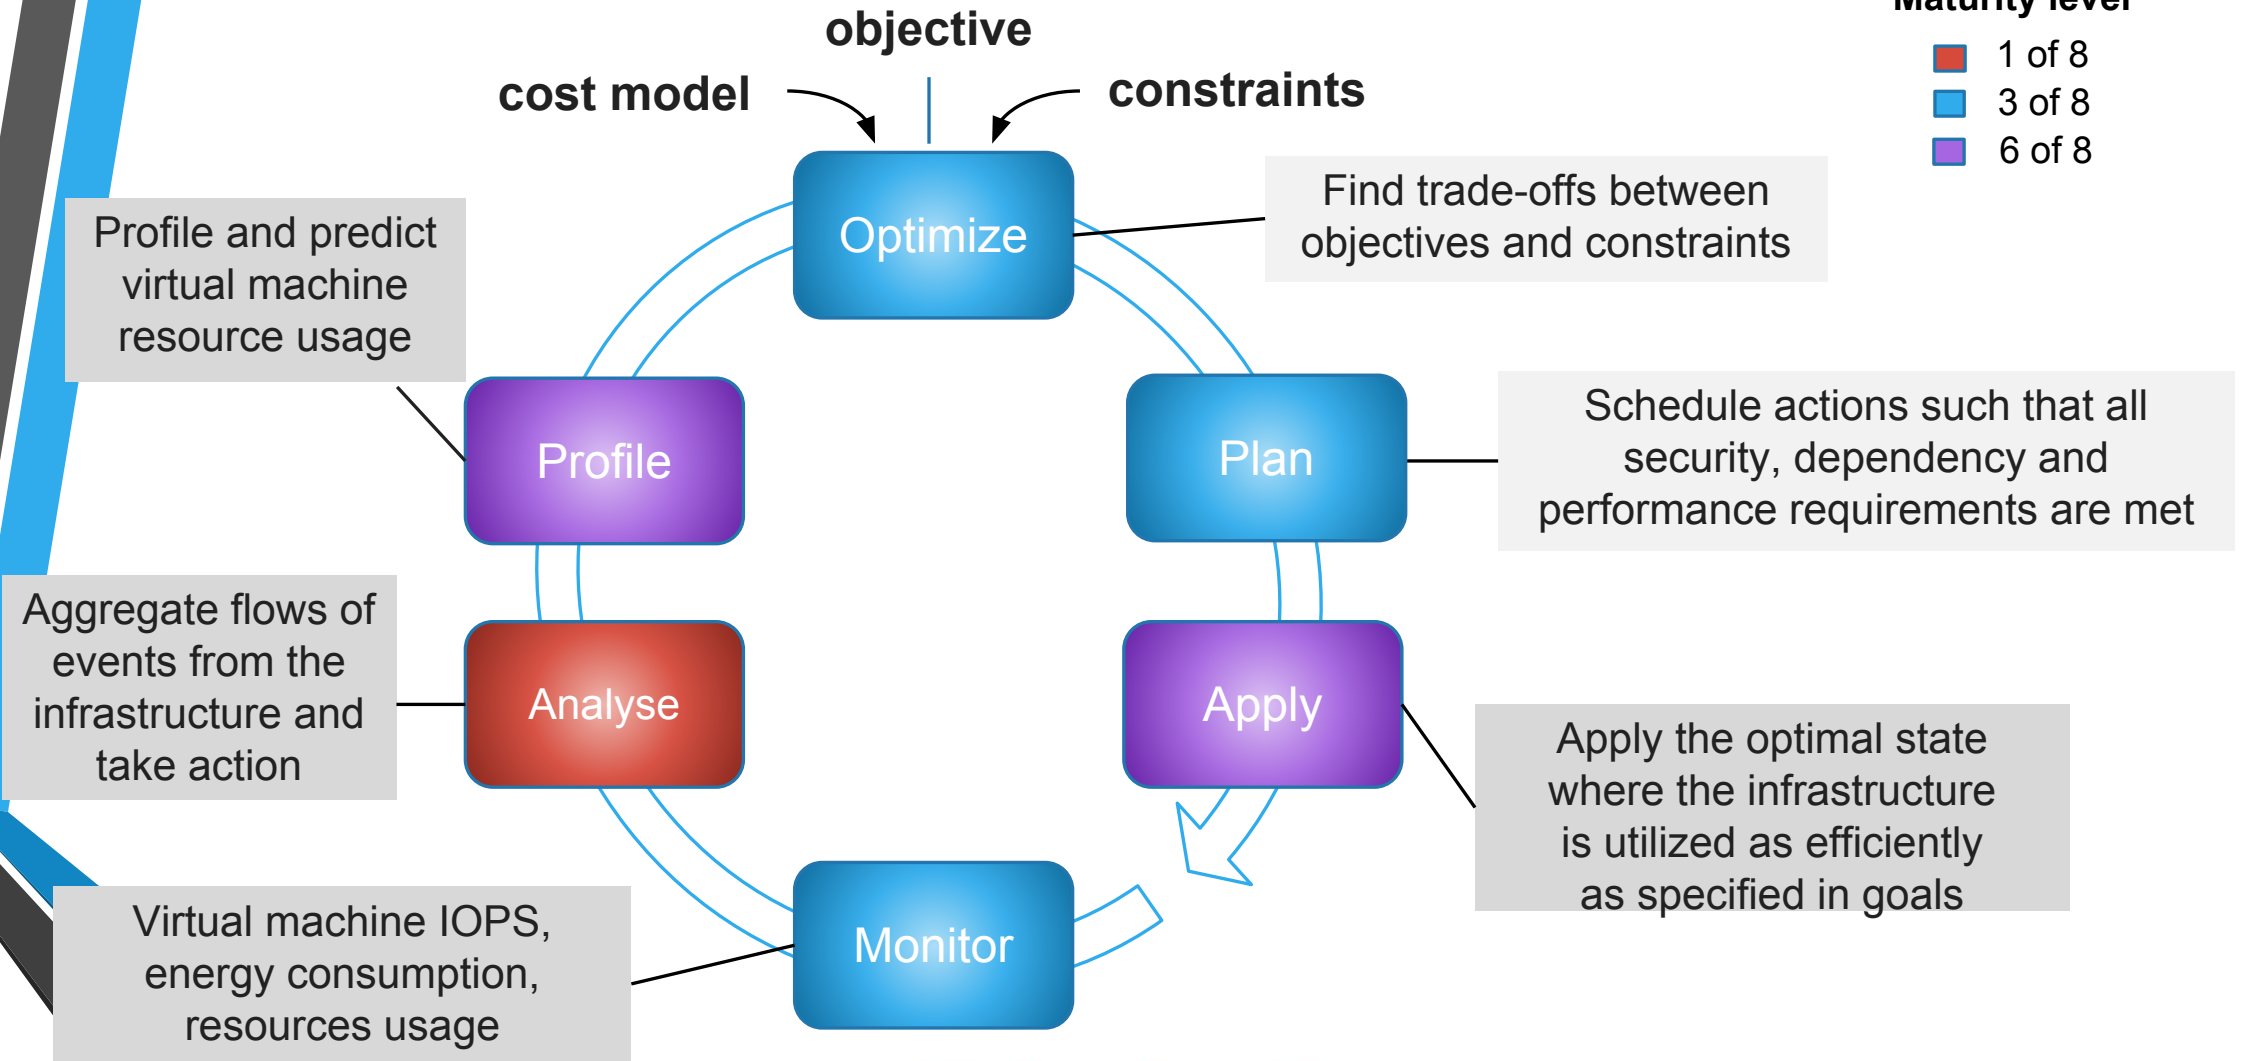
- Watcher leverages services provided by other OpenStack projects
 - VM live migration and resize
 - Metric collection
 - Power cycle bare metal nodes
- Monitors the infrastructure and performs optimizations on-demand
- Enables new ways for OpenStack administrators to reduce the cloud's TCO

*Other names and brands may be claimed as the property of others.

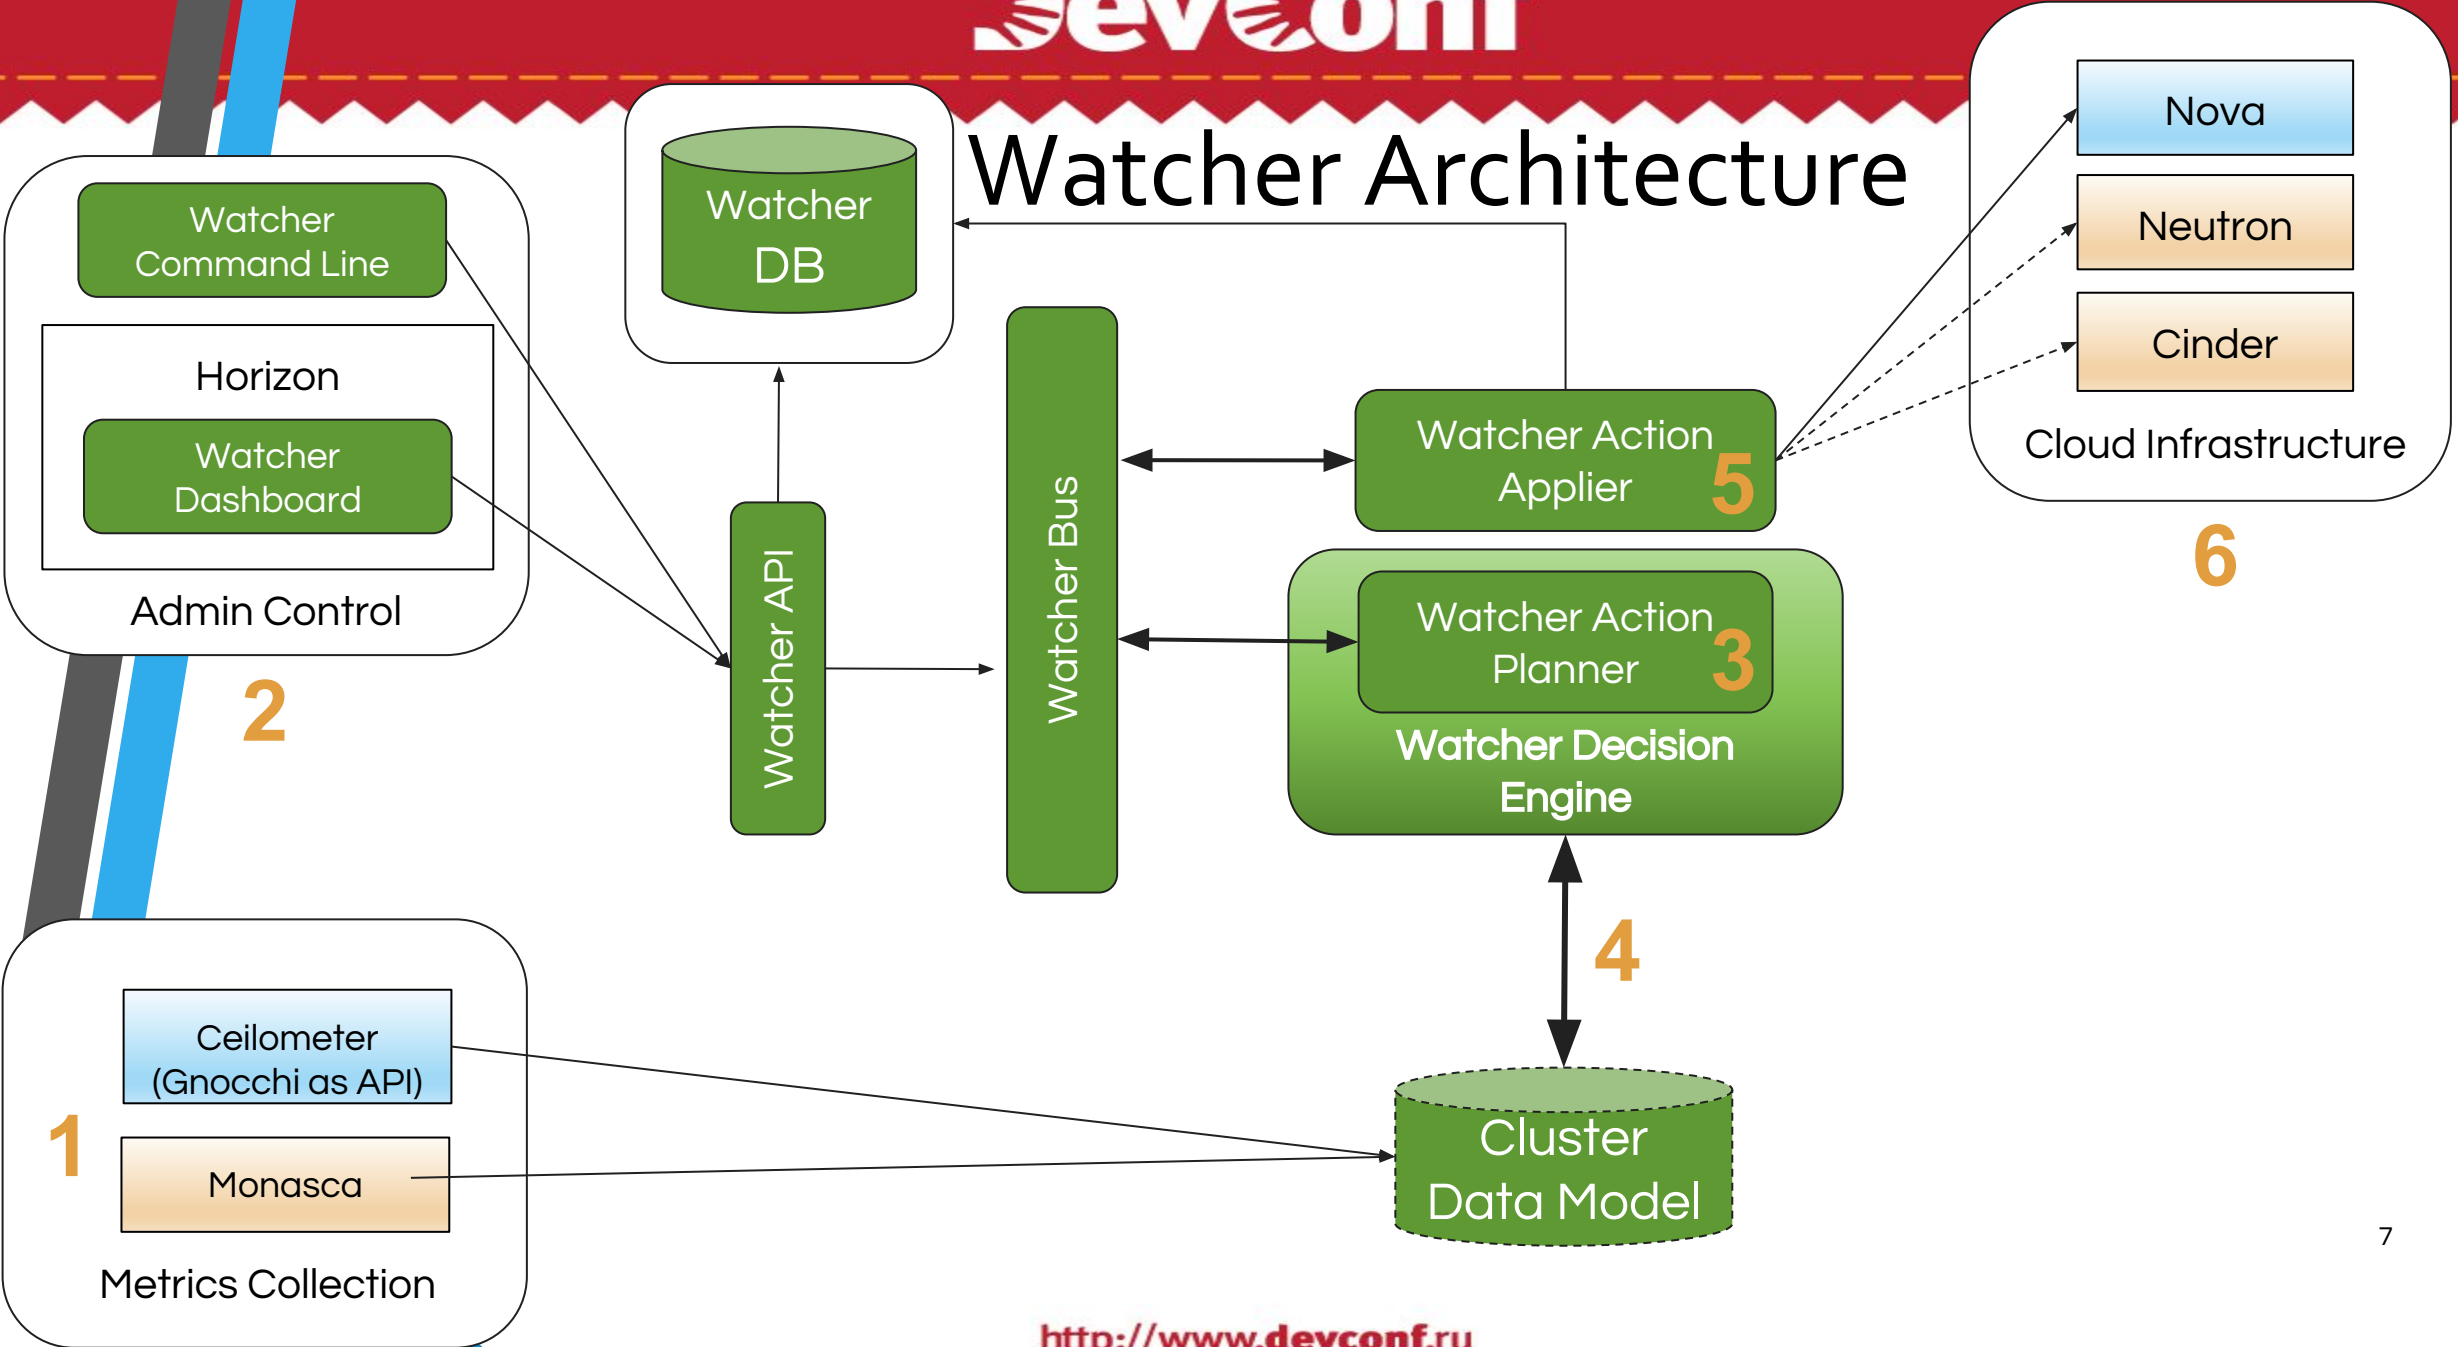
Watcher Workflow and Maturity

Maturity level

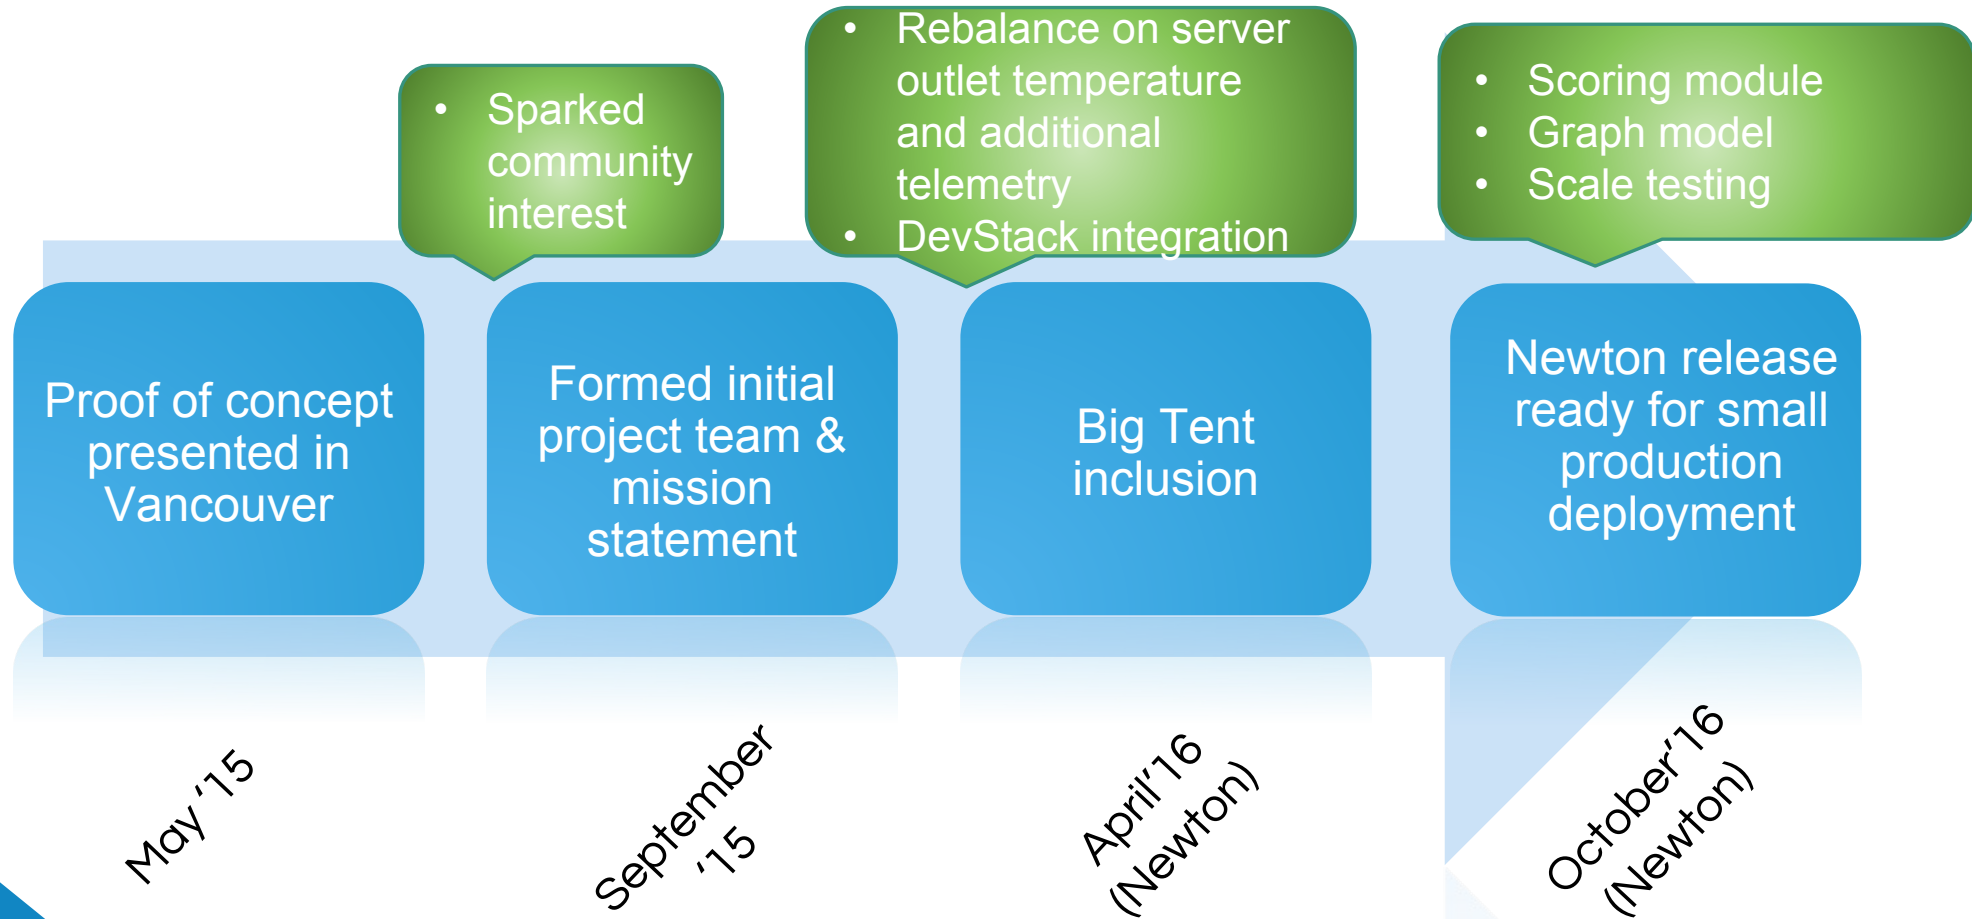
- 1 of 8
- 3 of 8
- 6 of 8



Watcher Architecture



Watcher's History



Watcher's Roadmap

- Scoring module
- Graph model
- Scale testing

OpenStack compliant and ready for large production deployment

April '17
(Ocata)

- Tagging VMs
- HA mode
- Grammar for workload characterization

Pike release extends list of supported resources

September '17
(Pike)

Ocata Release Accomplishments

- Notifications for objects
- Automatic triggering audit
- Provide alembic migrations
- Generic way to define the scope of an audit (set of resources)
- Service supervisor to monitor Watcher daemons

Strategy	Description	Telemetry used	Provider
Outlet temperature based migration strategy	Moves workload when server's outlet temperature is higher than specified threshold	Outlet temperature	Intel
Basic consolidation strategy	Implements a basic load consolidation; this is currently a heuristic algorithm which focuses on measured CPU utilization and tries to minimize hosts which have too much and too little load and achieve a target high(ish) level for all hosts	CPU, RAM, Disk	B<>com & Zurich University of Applied Sciences
Uniform airflow migration strategy	Moves workload when server's airflow is greater than specified threshold; it will also decide how to move the VMs according to the current inlet temperature and system power	airflow, inlet temperature	Intel
Workload stabilization strategy	Monitors if there is a higher load on some hosts compared to other hosts in the cluster and re-balances the work across hosts to minimize the standard deviation of the loads in a cluster	CPU, RAM	Servionica
Workload balance strategy	Makes decisions to migrate workloads to make the total VM workloads of each hypervisor balanced when the total VM workloads of hypervisor reaches threshold	CPU	Intel
VM Workload Consolidation Strategy	Leverages a modified first-fit algorithm to achieve increased server CPU and memory utilization which ultimately leads to freeing some of the hosts that can be powered down to save energy.	RAM, disk.root.size	Zurich University of Applied Sciences

Plans for the Pike Release

- Integrate Watcher Data Model with Cinder
- Add audit tag to VM metadata to let external systems know that VM is in optimization process
- Provide Gnocchi support as Data Source for strategies
- Add workload characterization to improve cloud optimization
- Use notifications in Watcher (event-driven fashion)
- Provide more “value-added” optimization strategies

Strategy	Descriptions	Telemetry used	
Noisy Neighbor Strategy	<p>L3 cache is critical and limit system level resource shared by all apps or VMs on one node. If one VM occupies most of L3 cache, other VMs on the node likely starve without enough L3 cache thus poor performance. This BP adds a new strategy to detect then migrate such cache greedy VM based on some new cache/memory metrics.</p>	<p>perf.instructions, perf.cpu.cycles, cpu_l3_cache</p>	Intel
Strategy to trigger "power on" and "power off" actions	<p>For a data center with large amount of VMs and physical hosts, the total power consumption is tremendous. When workload is not heavy, Watcher can be used to reduce power consumption by triggering a request to power off some idle hosts without VMs. And when the workload increases watcher will trigger a "power on" request to fulfill the service requirements.</p>	Not specified yet	ZTE

Any Questions?

- Want to learn more?!?
 - Wiki : <https://wiki.openstack.org/wiki/Watcher>
 - IRC : #openstack-watcher
- If you are interested, we would love for you to get involved – come and see us!

Legal Notices and Disclaimers

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.

No computer system can be absolutely secure.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit <http://www.intel.com/performance>.

Intel, the Intel logo and others are trademarks of Intel Corporation in the U.S. and/or other countries.
*Other names and brands may be claimed as the property of others.

© 2016 Intel Corporation.