

Mistral performance in numbers

Will it blend?

- Andras Kovi akovi@nokia.com
- 07-06-2017



<http://tinyurl.com/osmistralbp>

Keep it open!
Instant feedback to me!




<http://tinyurl.com/osmistralbp>

Instant feedback to me!

Play with us! Keep this
page open during
The presentation.



Mistral




<http://tinyurl.com/osmistralbp>



2

Workflow

- “orchestrated and repeatable pattern of business activity” – Wikipedia



Main properties of a workflow system

Language to describe workflows

Tasks are **distributed** across enterprise/cloud environment

Support **parallelism** for running tasks

Synchronization of parallel branches

Data flows from one task to another

APIs to **create/update/delete workflows**

APIs and active component to **execute workflows**

APIs to **monitor state** of running workflows and tasks

Persistent state of running workflows and tasks

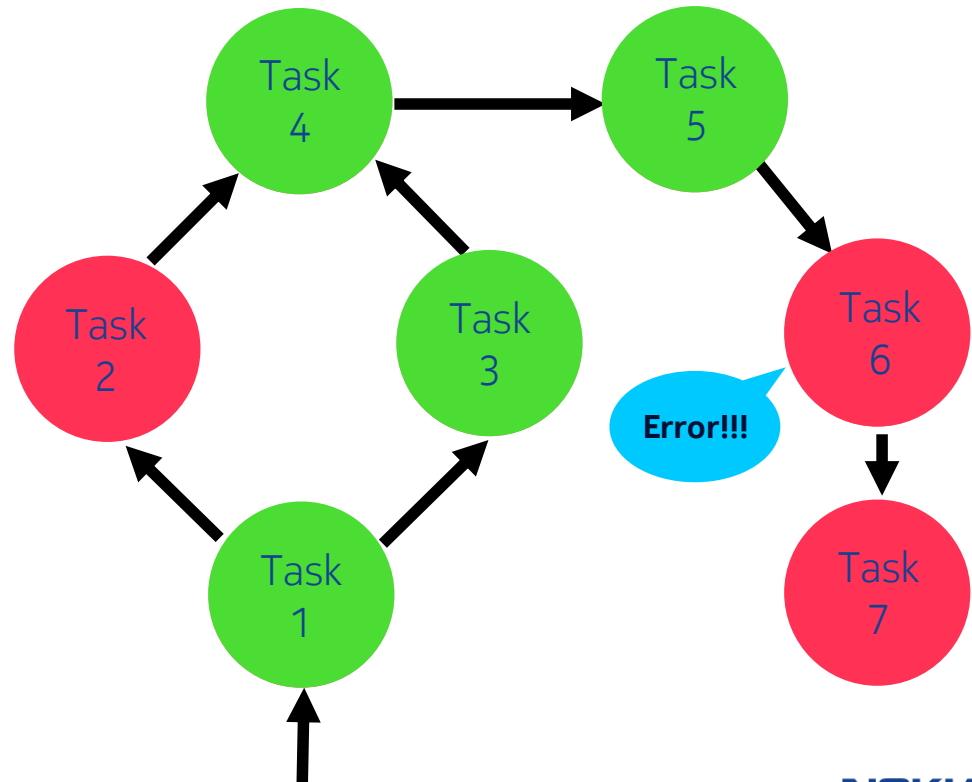
Persistent state is a key!

<http://tinyurl.com/osmistralbp>



3

- Progress
- Observability
- Execution history
- Persistent result
- Asynchrony
- Effective error handling
- Recovery



Right tool for the right task!

Scripts

Error handling is challenging

- All is OK while a script is working
- On an error we're left with a mess

Not scalable

Lack of monitoring

No parallelism & synchronization

Not suitable for large scale orchestration!



Mistral Workflow language

<http://tinyurl.com/osmistralbp>



4

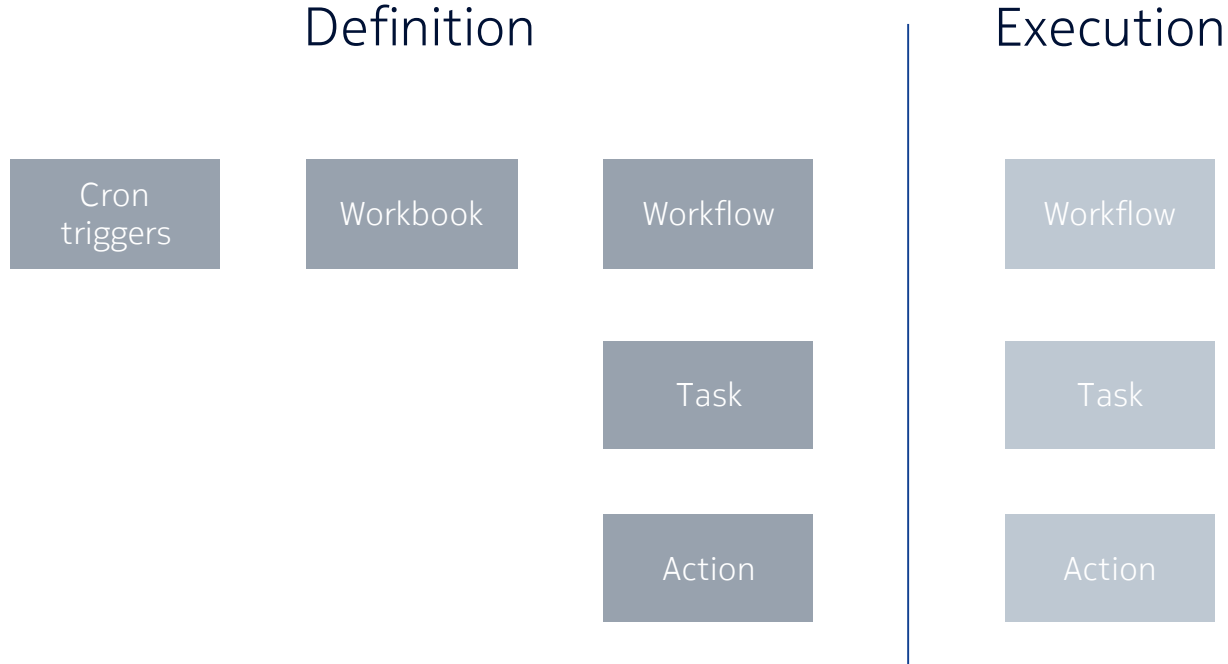
YAML-based

Relies on expression languages (YAQL, JINJA)

Main features

- Define workflow tasks
- Conditional task transitions
- Publish persistent variables
- Fork & Join
- Loop (to process data collections)
- Engine commands (fail, succeed, pause, noop)
- Task policies (retry, timeout, wait-before, wait-after, pause-before)
- Nested workflows

Building blocks



Example: deleting VMs by criteria

<http://tinyurl.com/osmistralbp>



5

```
version: '2.0'
```

tenant_cleanup:

```
description: Deletes virtual machines containing "mistral" in their names.
```

```
tasks:
```

get_vms:

```
action: nova.servers_list
```

```
publish:
```

```
vms: <% $.get_vms['mistral' in $.name]
      .select(dict(id=>$.id,name= >$.name)) %>
```

```
on-success:
```

```
- delete_vms
```

```
on-error:
```

```
- fail
```

delete_vms:

```
with-items: vm in <% $.vms %>
```

```
action: nova.servers_delete server=<% $.vm.id %>
```

Available actions

Common system actions:

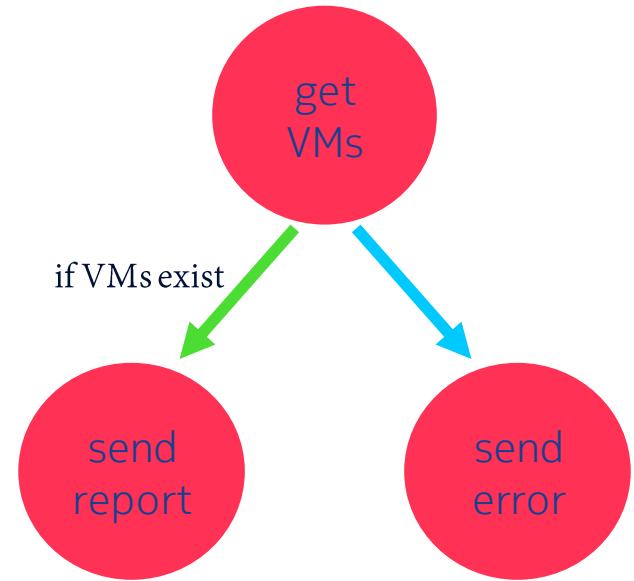
- std.echo
- std.email
- std.fail
- std.http
- std.mistral_http
- std.java_script
- std.ssh
- std.ssh_proxied
- std.wait_ssh

OpenStack services actions available for:

- Ceilometer
- Cinder
- Glance
- Heat
- Keystone
- Neutron
- Nova

Conditional transitions

```
report_vms:  
  input:  
    - email_info  
  
  tasks:  
    get_vms:  
      action: nova.servers_list  
      on-success:  
        - send_report: <% $.get_vms.len() > 0 %>  
      on-error:  
        - send_error  
  
    send_report:  
      action: std.email  
      ...  
  
    send_error:  
      action: std.email  
      ...
```



Publishing persistent variables

```
version: '2.0'

tenant_cleanup:
  tasks:
    get_vms:
      action: nova.servers_list
      publish:
        vms: <% $.get_vms['mistral' in $.name]
              .select(dict(id=>$.id,name=>$.name)) %>
      on-success:
        - delete_vms

    delete_vms:
      with-items: vm in <% $.vms %>
      action: nova.servers_delete server=<% $.vm.id %>
```

Fork & Join

```
version: '2.0'
```

```
deploy_app:
```

```
  tasks:
```

```
    install_db:
```

```
      action: install_mysql
```

```
      on-success:
```

```
        - install_app
```

```
    install_web_server:
```

```
      action: install_web_server
```

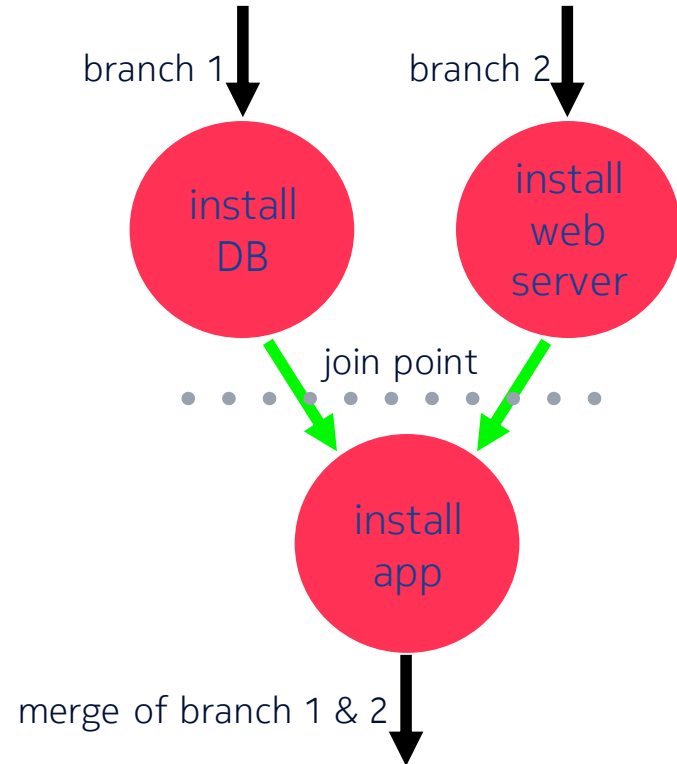
```
      on-success:
```

```
        - install_app
```

```
    install_app:
```

```
      join: all 1, 2, 3...
```

```
      action: install_my_app
```



Loops

```
version: '2.0'
```

```
tenant_cleanup:
```

```
  description: Deletes virtual machines containing "mistral" in their names.
```

```
  tasks:
```

```
    get_vms:
```

```
      action: nova.servers_list
```

```
      publish:
```

```
        vms: <% $.get_vms['mistral' in $.name]  
              .select(dict(id=>$.id,name=>$.name)) %>
```

```
      on-success:
```

```
        - delete_vms
```

```
      on-error:
```

```
        - fail
```

```
delete_vms:
```

```
  with-items: vm in <% $.vms %>
```

```
  action: nova.servers_delete server=<% $.vm.id %>
```

Engine commands

```
version: '2.0'
```

```
deploy_app:  
  tasks:  
    get_vms:  
      action: nova.servers_list  
      on-success:  
        - fail: <% $.get_vms.len() = 0 %>  
        - install_app  
  
    install_app:  
      ...
```

<http://tinyurl.com/osmistralbp>



6

- **fail** - ends entire workflow with a failure
- **succeed** - ends entire workflow with success
- **pause** - puts entire workflow on hold (can be resumed manually)

Task policies

```
cinder_backup:
  tasks:
    start_backup:
      action: cinder.backups_create ...
    publish:
      backup_id: <% $.start_backup.id %>
    on-success:
      - wait_for_backup

    wait_for_backup:
      action: fail_if_not_ready id=<% $.backup_id %>
      wait-before: 30
      retry: count=3 delay=5
    on-success:
      - report_backup_is_ready

    report_backup_is_ready:
      ...
```

- **retry** - repeats a task on failure
- **wait-before** - waits a number of seconds before starting a task
- **wait-after** - waits a number of seconds after task completion
- **pause-before** - puts workflow on hold before running a task
- **timeout** - fails a task is running longer than configured value

Nested workflows

```
version: '2.0'
```

```
deploy_app:
```

```
  input:
```

```
    - vm_ids
```

```
  tasks:
```

```
    install_db:
```

```
      workflow: install_db vm_id=<% $.vm_ids[0] %>
```

```
      on-success:
```

```
        - install_app
```

```
    install_app:
```

```
      with-items: id in <% $.vm_ids %>
```

```
      action: std.ssh host=<% $.id %>
```

```
      input:
```

```
        cmd: 'sudo apt-get install my_app'
```

Cron trigger

Cloud crontab

HA

Runs workflow by configured cron pattern

Can limit number of executions

Example using CLI:

\$ mistral cron-trigger-create trig **tenant_cleanup** --pattern '* / 5 * * * *



More at

https://docs.openstack.org/developer/mistral/dsl/dsl_v2.html




Numbers

Measurement and Predictions

<http://tinyurl.com/osmistralbp>



- Memory
- CPU
- IO
- KPIs



Numbers

Performance, Reliability, Scalability

Number of tasks

Types of tasks

- Sleep

- CPU intensive

- External communication

- Data processing (filter, collect data from input)

Data flow between tasks

- No data

- Small documents

- Large documents

- Increasing size documents

Task dependencies

- #of join

- #of fork

- Loops

- Recursion

- Type

 - spider web (entagled, radial (directed inward, directed outward))

 - Tree

Sequence length

- Short sequence

- Long sequence

Parallelism

- Tasks

 - Sequential

 - Parallel

- Workflows

 - Many at once

 - Few at once

Workflows

- Flat

- Hierarchical

Trigger types

- on-demand

- cron

Size of database

#API requests

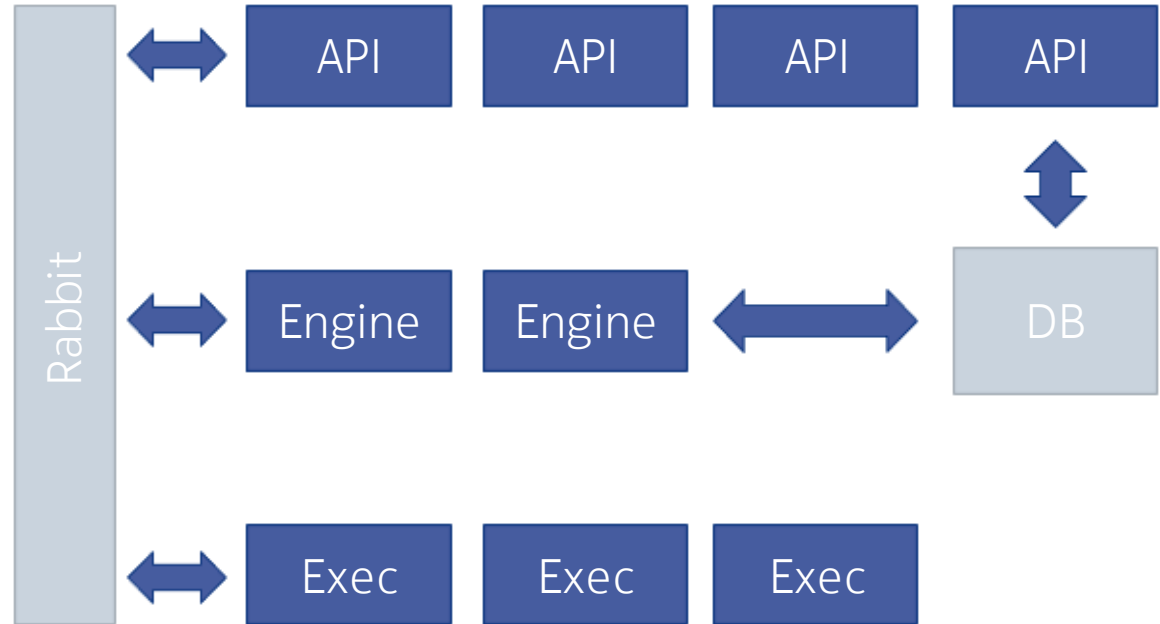
- create/delete

- read

Mistral services

Responsibilities

- API (any number)
- Engine (any number 🚚)
- Executors (any number)
- Transport - RabbitMQ
- DB - MySQL, Postgres



Mistral services - API

<http://tinyurl.com/osmistralbp>



8

Functions

- REST Model query
- Work initiation

Stats

- Memory per instance (~70MB initial)
- Load distribution (socket accept based)
 1. `fork()`
 2. `accept()`

Jmeter ☺

Tests

- Create executions
- List all
- List limited
- API traversal...

Responsiveness

Load distribution

- Forked API processes

API traversal

- Flat workflow
- Sub-sub-sub

Responsiveness

Workflow CRUD

#workflows	Create	List	limit=10
100	6.5/sec	30/sec	
1000	7/sec	3/sec	50/sec
7000	7/sec	3/sec 913Kb	34/sec

Mistral services - Engine

Functions

- Workflow processing
- Data flow management

Input-output calculation

YAQL <% \$.vm_ip %>

JINJA {{ _.data }}

Data size

keep result ?

Redundant engines

- Ongoing stability improvements

Move actions to engine?

```
install_app:
  with-items: id in <% $.vm_ids %>
  action: std.ssh host=<% $.id %>
  input:
    cmd: `sudo apt-get install my_app`
```

Mistral services - Executor

<http://tinyurl.com/osmistralbp>



9

Actions run here

Operation mode

1. RPC request
2. Work
3. Respond

Long test

1 week

50 cron triggers (Mistral) (H/60 * * * *)

50 runs per minute

3000 runs per hour

72000 per day

~500000 runs total

Max executions policy

< 700

Age < 1 day

Run



Create/Start dedicated container

Touch files

Cleanup files

Stop container

Long test - results



War stories

Size

Rally job

Expressions

YAQL -> Jinja

Big data

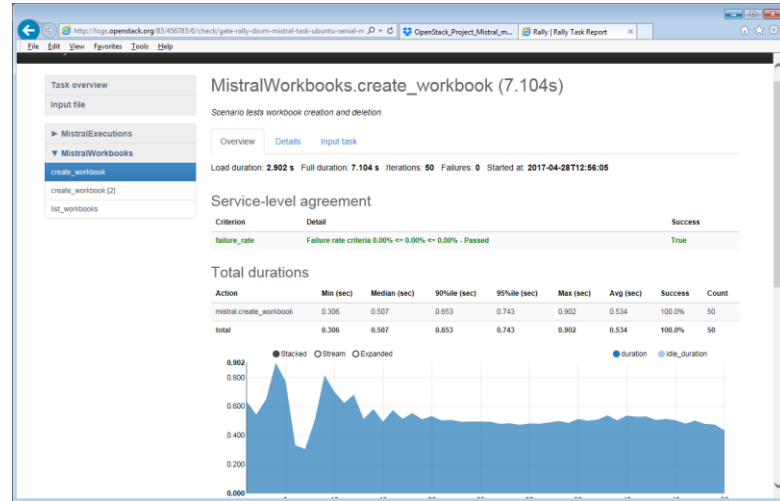
execution_field_size_limit_kb

mysqld -> max_allowed_packet value

Database trashing

execution_expiration_policy

max_executions



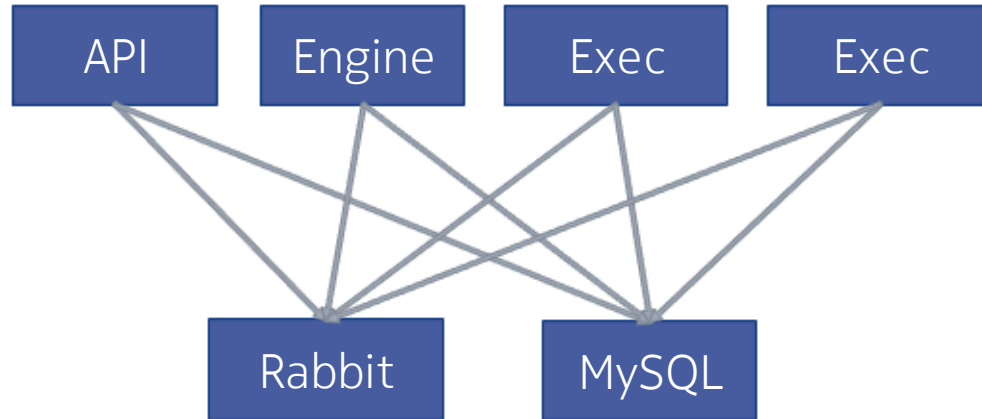
Evaluation setup

Demo time

.../mistral/tools/docker/

build.sh

start_....sh



Quotes from the (Boston) Summit

„Can I use Mistral to implement a VM expiration policy?”

„Absolutely”

„Cron jobs are one of the killer features we still don't use in TripleO”

„Go ahead and do so”

„Masakari could call Mistral workflows...” - Openstack HA discussion

NOKIA