

DevOps deployment tools - Behind the scene

Advanced distributed systems

Hélène Coullon



Associate professor at IMT Atlantique, France



Inria researcher, France



Adjunct professor at UiT, Tromsø, Norway

March 1st, 2021

Outline

1. Introduction
2. Overview of deployment tools
3. Parallelism within deployments
4. The MADEUS deployment model
5. Evaluation
6. Conclusion

Let's start with a few questions on Wooclap



WEB

- 1 Connect to www.wooclap.com/NGOACZ
- 2 You can participate



SMS

- 1 Not yet connected? Send **@NGOACZ** to **06 44 60 96 62**
- 2 You can participate


Table of Contents

1. Introduction
2. Overview of deployment tools
3. Parallelism within deployments
4. The MADEUS deployment model
5. Evaluation
6. Conclusion

Distributed software systems

General definition

- Non monolithic code,
- modular units of code - **components**,
 - black-box of code,
 - with well-defined provided and required interfaces,
- software system = **architectural assembly** of component instances,
- interactions between components through **communications**.



- Master/workers,
- peer-to-peer,
- dataflow/stream,
- service-oriented,
- layered.


Distributed software systems everywhere

Examples of distributed software systems

- Smartphone applications (e.g., Waze),
- health, bank, tax information systems,
- Netflix micro-services infrastructure (i.e., Netflix OSS),
- operating system of the Cloud (e.g., OPENSTACK),
- 5G networks (e.g., network function virtualization).

Distributed infrastructures

Computing, storage and network resources everywhere!



- Cloud computing
- Fog and Edge computing
- Internet-of-Things (IoT)
- Cyber-physical systems

Deployment

What is a deployment?

Install, configure, start, test a distributed software system.

Questions raised by deployment

- What do I need to deploy?
- Where do I need to deploy?
- How do I deploy?
- When do I deploy?

[**WHAT** + **HOW** + **WHERE** + **WHEN**] = DevOps deployment tools

Example - deploying LAMP

[WHAT] LAMP

- **L**inux operating system
- **A**pache web server
- **M**ariaDB database
- **P**HP language

[WHERE]

on **2 nodes**: machine1, machine2

Example - deploying LAMP

[HOW]

- Let's take a look at a basic **documentation**
- What if I want to configure **Apache** and **MariaDB**?
- What if I deploy on **CentOS**?

[WHEN]

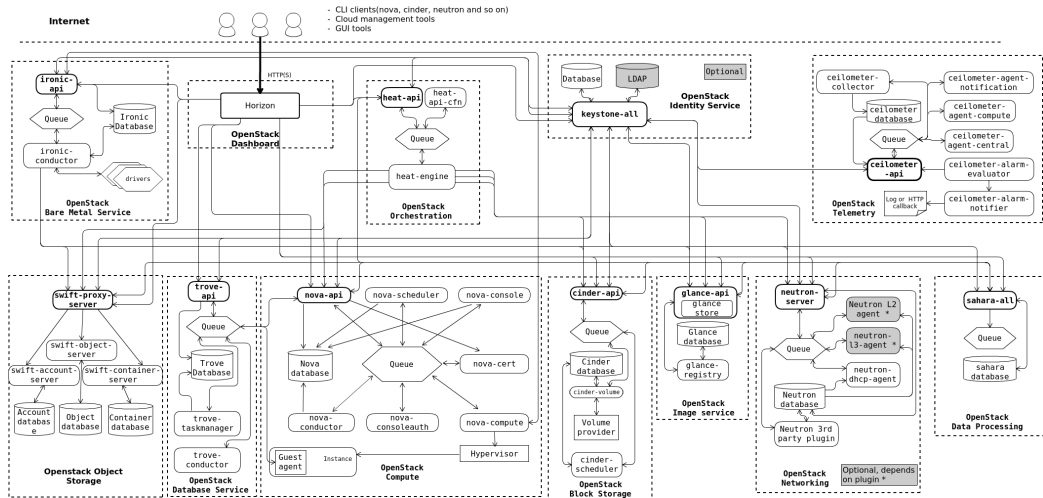
linux → apache → mariadb → php

linux → mariadb → apache → php


linux → apache → php → mariadb

Not so simple?

What about this?



Or this?



amazon.com®



NETFLIX

On this?



Deployment tools and DevOps community

documentation or README.md → ad-hoc scripts → deployment tools

Software engineering practices applied to deployment

- Automating deployments at scale,
- structuring deployments (languages, models),
- reusing deployment procedures,
- avoiding errors.

Table of Contents

1. Introduction
2. Overview of deployment tools
3. Parallelism within deployments
4. The MADEUS deployment model
5. Evaluation
6. Conclusion

Overview of deployment tools

	Config Management	Docker ecosystem	Provisioning	Orchestration
ANSIBLE (2012)	✓		(✓)	
PUPPET (2005)	✓		(✓)	
DOCKER COMPOSE (2014)		✓		
DOCKER SWARM (2014)		✓		✓
KUBERNETES (2014)		✓		✓
NOMAD (v1.0)				✓
TERRAFORM (2014)	(✓)		✓	
JUJU (2011)	✓		✓	
CLOUD FORMATION (2006)			✓	
HOT and HEAT (2010)			✓	✓
TOSCA ecosystem (2014)	✓	(✓)	✓	✓

Architecture of these tools

Agentless, master/workers and heavier software stack (i.e., bootstrap problem).

In practice - combination of tools

4 categories

Configuration management tools

Initially designed to install, configure, manage software on existing servers.

Provisioning tools

Initially designed to provision the servers, network, platforms etc.

DOCKER ecosystem

Solve portability problem and reduce configuration issues through containers.

Orchestration tools

Automated coordination, and management of a set of components that form distributed software systems, on a set of resources (virtual or physical).

Decoupling WHAT / HOW / WHERE / WHEN

Machine 1 **[WHERE]**

Database (DB) **[WHAT]**

[HOW]

1. install prerequisite 1
2. install prerequisite 2
3. install MySQL
4. configure parameters
5. start the service
6. setup the root user
7. add a user
8. create table

Machine 2 **[WHERE]**

Web-server (WS) **[WHAT]**

[HOW]

1. install prerequisite 1
2. install prerequisite 2
3. install Apache
4. configure the firewall
5. restart the firewall
6. download the website content
7. untar the website content
8. configure parameters
9. start the service

[WHEN]: DB → WS (components granularity)

Academic contributions

Enhancing **[WHEN]**

- enhancing **[WHEN]** by decoupling
 - **[LIFECYCLE]**
 - **[DEPENDENCIES]**


Contributions

TOSCA, DEPLOYWARE, SMARTFROG, ENGAGE, AEOLUS, MADEUS etc.

SMARTFROG (1996-2003-2009) → ENGAGE (2012) → AEOLUS (2013-2016) → MADEUS (2018-2021)

- comparable but complementary to configuration management tools
- generic to any kind of resource and action (configuration, provisioning, management)

Decoupling WHAT / HOW / WHERE / WHEN



Machine 1 **[WHERE]**

Database (DB) **[WHAT]**

[HOW] [LIFECYCLE]

1. Install
2. Configure
3. Start the service
4. Prepare the service

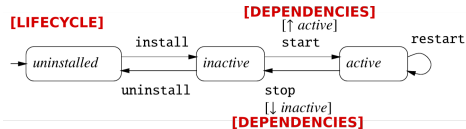
Machine 2 **[WHERE]**

Web-server (WS) **[WHAT]**

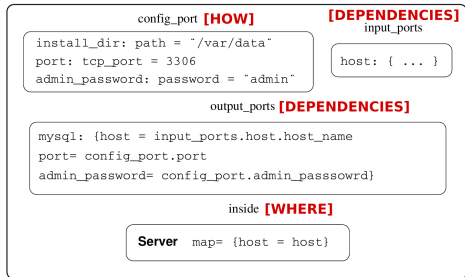
[HOW] [LIFECYCLE]

1. Install
2. Configure firewall
3. Download
4. Configure parameters
5. Start the service

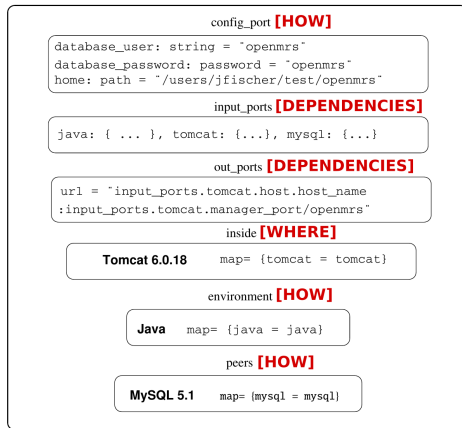
[DEPENDENCIES]: WS(4) → DB(3), WS(5) → DB(4) (lifecycle granularity)



MySQL 5.1 [WHAT]



OpenMRS 1.8 [WHAT]



Particularities of AEOLUS

- programmable lifecycle
- finer grain to model dependencies
- inspired from state machines
- inspired from component models

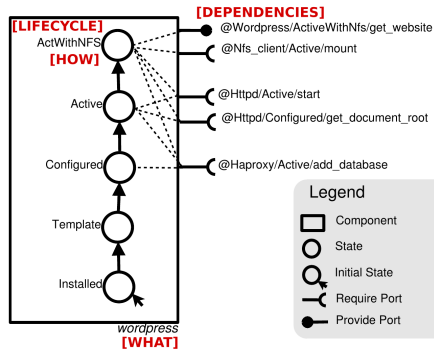


Table of Contents

1. Introduction
2. Overview of deployment tools
3. Parallelism within deployments
4. The MADEUS deployment model
5. Evaluation
6. Conclusion

Our goals

Performance related to [WHEN] \equiv [LIFECYCLE, DEPENDENCIES]

- structured parallelism
- reach quickly a configuration
- avoid disruption time


Safety

- formally-defined semantics
- tools to assist during design
- verification of properties

Performance through parallelism and dependencies

level1: multiple nodes, same action

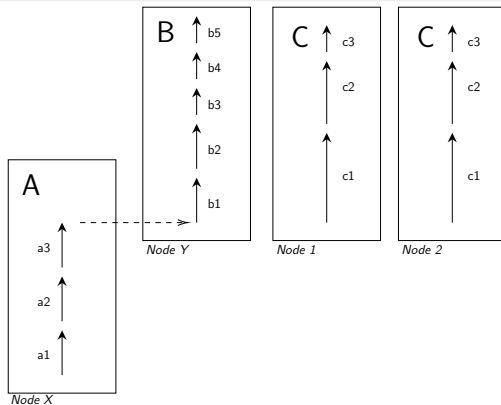
- no dependencies declared
- procedural execution order
- ANSIBLE



Performance through parallelism and dependencies

level2: level1+non-dependent components


- dependencies at the component level
- DEPLOYWARE, (basic) TOSCA, ENGAGE



Performance through parallelism and dependencies

level3: level1 + level2 + inter-component


- dependencies at the task level
- (advanced) TOSCA, AEOLUS



Performance through parallelism and dependencies

level 4: level1 + level2 + level3 + intra-component

- internal task dependencies
- MADEUS



The finer the dependencies granularity is, the better is the efficiency (related to **[WHEN]**)


Table of Contents

1. Introduction
2. Overview of deployment tools
3. Parallelism within deployments
- 4. The MADEUS deployment model**
5. Evaluation
6. Conclusion

Control components



Written by the **component developers**



Internal net [LIFECYCLE]

- places = milestones
- transitions = actions to perform
 - concretely: scripts are attached to transitions
 - in the model: exact nature/effects of actions not represented, only coordination

Interfaces [DEPENDENCIES]

- use ports = requirements
- provide ports = provisions
- during execution: active/inactive

Control components in practice



Written by the **component developers**

```
1 class Apache(Component):
2     def create(self):
3         self.places = ['uninstalled', 'installed', 'configured', 'running']
4
5         self.initial_place = 'uninstalled'
6
7         self.transitions = {
8             'install1': ('uninstalled', 'installed', self.install1),
9             'install2': ('uninstalled', 'configured', self.install2),
10            'configure': ('installed', 'configured', self.configure),
11            'start': ('configured', 'running', self.start)
12        }
```

Control components in practice



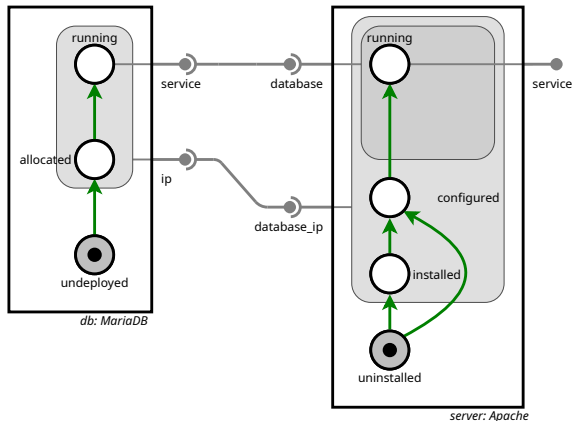
Written by the **component developers**

```
1 class Apache(Component):
2     def create(self):
3         ...
4
5         self.dependencies = {
6             'database_ip': (DepType.USE, ['installed', 'configured', 'running']),
7             'database': (DepType.USE, ['running']),
8             'service': (DepType.PROVIDE, ['running'])
9         }
10
11 # Definition of the actions
12 def install1(self):
13     remote = SSHClient()
14     remote.connect(host, user, pwd)
15     remote.exec_command(cmd)
16     ...
```


Assembly of components

Assembly of components [DEPENDENCIES]

instantiation of component types and connections of ports



Assembly of components in practice




Written by the **assembly developers**, architect, DevOps

```
1 from components.mariadb import MariaDB
2 from components.apache import Apache
3
4 class ApacheWithDB (MadeusAssembly):
5     def create():
6         self.components = {
7             'server': Apache(),
8             'db': MariaDB()
9         }
10        self.dependencies = [
11            ('server', 'database_ip', 'db', 'ip'),
12            ('server', 'database', 'db', 'serv')
13        ]
14
15 if __name__ == '__main__':
16     assembly = ApacheWithDB()
17     assembly.run()
```

Execution example

Execution semantics


Initial places



Execution example

Execution semantics


Firing transitions, parallel transitions



Execution example

Execution semantics


Entering places, inter-coordination through connections



Execution example

Execution semantics


Reaching places, inter-coordination through connections



Execution example

Execution semantics


Reaching places, inter-coordination through connections



Execution example

Execution semantics


Reaching places, intra-coordination



Execution example

Execution semantics


Reaching places, intra-coordination



Execution example

Execution semantics


Reaching places, intra-coordination



Execution example

Execution semantics


Reaching places, inter-coordination through connections



Execution example

Execution semantics

Reaching places, inter-coordination through connections



Execution example

Execution semantics

Reaching places, inter-coordination through connections

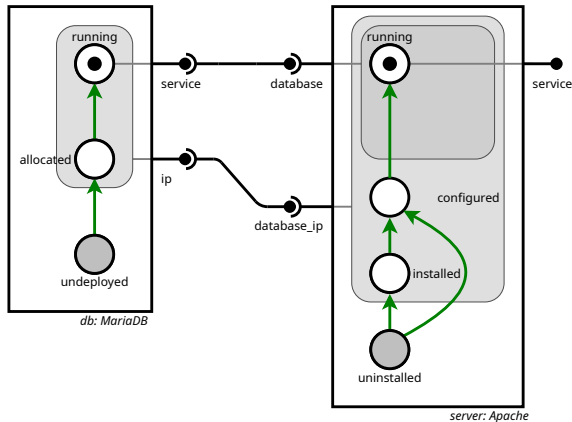


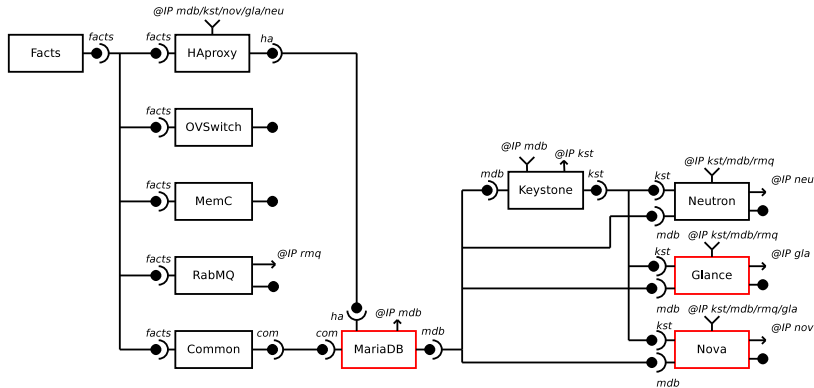
Table of Contents

1. Introduction
2. Overview of deployment tools
3. Parallelism within deployments
4. The MADEUS deployment model
- 5. Evaluation**
6. Conclusion

Use-case - deployment of OpenStack

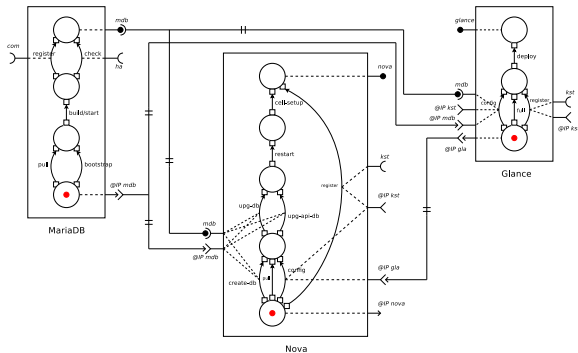
Basic production OpenStack as defined in KOLLA-ANSIBLE:

- 11 components, 36 services in total, deployed on three nodes,
- DOCKER container-based deployment,
- combination of DOCKER images, ANSIBLE,
- more than 20 minutes to deploy.



Evaluation setup


- Comparison to KOLLA-ANSIBLE (production tool), and AEOLUS (literature),
- emulation of AEOLUS (no longer maintained) with MADEUS,
- three different versions of the deployment remote, local, cached,
- **Reproducible experiments on Grid'5000.**



Evaluation on the deployment of OpenStack


Results on three nodes Ecotype (Nantes) of Grid'5000

Cluster	CPU	Memory	Network
Nantes Ecotype	2× Intel Xeon E5-2630L v4, 10 cores/CPU	128GB	2× 10Gbps




Evaluation on the deployment of OpenStack


MADEUS



ANSIBLE



AEOLUS



Evaluation on the deployment of OpenStack

- Traces of the OpenStack **continuous Integration platform**
- February 19 to February 27 2020
- Exactly 2963 deployments of OpenStack have been recorded (329 runs per day)
- Projection of the gain with deployment times of our experiments in remote mode

	Kolla	MADEUS	gain
<i>reference time(s)</i>	529	150	71%
<i>projection on 9 days(h)</i>	435	123	71%
<i>projection on av./day(h)</i>	48	14	71%

Table of Contents

1. Introduction
2. Overview of deployment tools
3. Parallelism within deployments
4. The MADEUS deployment model
5. Evaluation
6. Conclusion

Conclusion

- Deployment problem and its complexity
- Need for software engineering practices
- Overview of deployment tools and academic contributions
- Presentation of MADEUS
- Evaluation of MADEUS

Questions?

Configuration management

Focus of configuration management tools

Initially designed to install, configure, manage software on existing servers.

- **ANSIBLE**: Procedural approach, agentless on top of SSH
 - sequential order of roles and tasks
 - similar to a well structured script
 - abstraction on top of SSH and system commands
- **PUPPET**: Declarative approach with a master/worker architecture
 - **[HOW]** is mostly hidden for the user
 - specify what you want not how to get it
 - interesting for management

ANSIBLE **[WHERE]**

Inventory file

ANSIBLE **[WHAT]**

Playbook, roles, vars

ANSIBLE **[HOW]**

Tasks, templates, vars, handlers

Example of **ANSIBLE** Apache role on **GITHUB**.

Provisioning

Focus of provisioning tools

Initially designed to provision the servers, network, platforms etc.

- CLOUD FORMATION and HEAT: Specific to a given Cloud provider
 - resp. AWS, OPENSTACK
- TERRAFORM, JUJU, and TOSCA: Generic to any provider
 - write your own providers

TERRAFORM [WHAT]

resources, variables

TERRAFORM [WHERE]

resources (provisioning)

TERRAFORM [HOW]

variables, user_data, local_exec,
write custom providers

Example of an AWS EC2 instance provisioning with a webserver with TERRAFORM

Writing custom TERRAFORM providers

Combining ANSIBLE and TERRAFORM

Doing provisioning with ANSIBLE

DOCKER ecosystem

Focus of configuration management tools

Solve portability problem and reduce configuration issues through containers.

DOCKER [HOW]

DOCKERFILE and
DOCKER images

DOCKER [WHAT]

DOCKER images and
DOCKER COMPOSE

DOCKER [WHERE]

KUBERNETES and
DOCKER SWARM

- lighter than virtual machines (if sharing the same OS kernel)
- the DOCKERFILE still has to be written at some point
- DOCKER COMPOSE to write an ordered list of DOCKER images to deploy locally
- KUBERNETES and DOCKER SWARM to manage set of containers, their placement, their replicas

Example of DOCKER COMPOSE LAMP deployment on GITHUB.