Sodalite

An Approach to Support Automated Deployment of Applications on Heterogeneous Cloud-HPC Infrastructures

Elisabetta Di Nitto, Jesus Gorronogoitia,
Indika Kumara, Georgios Meditskos,
Dragan Radolovic, Karthee Sivalingamk,
and Roman Sosa Gonzlez

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 825480.

An example of complex application

Sodalite

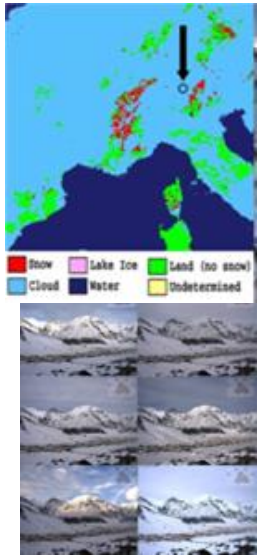

The goal is to exploit publicly available images to obtain a water availability indicator.

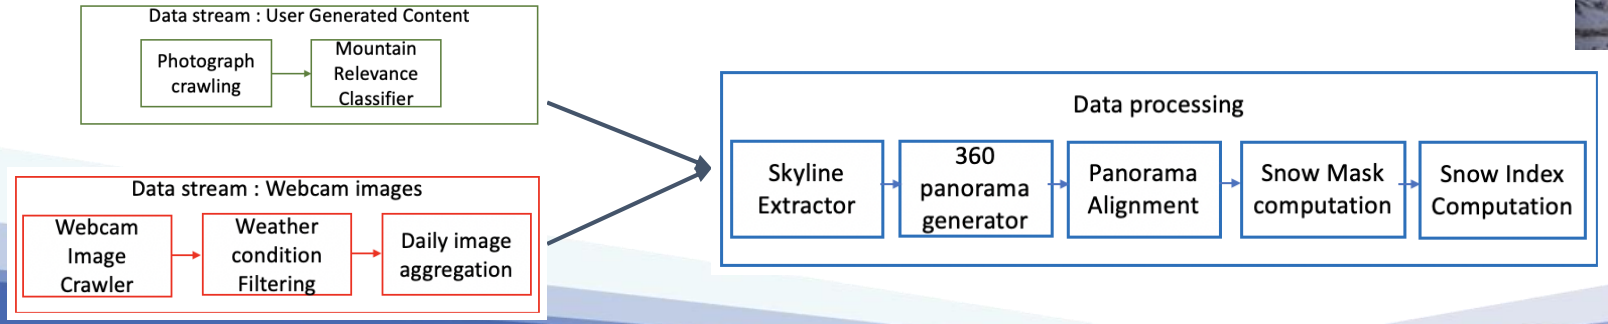

A closer look - multiple configurations

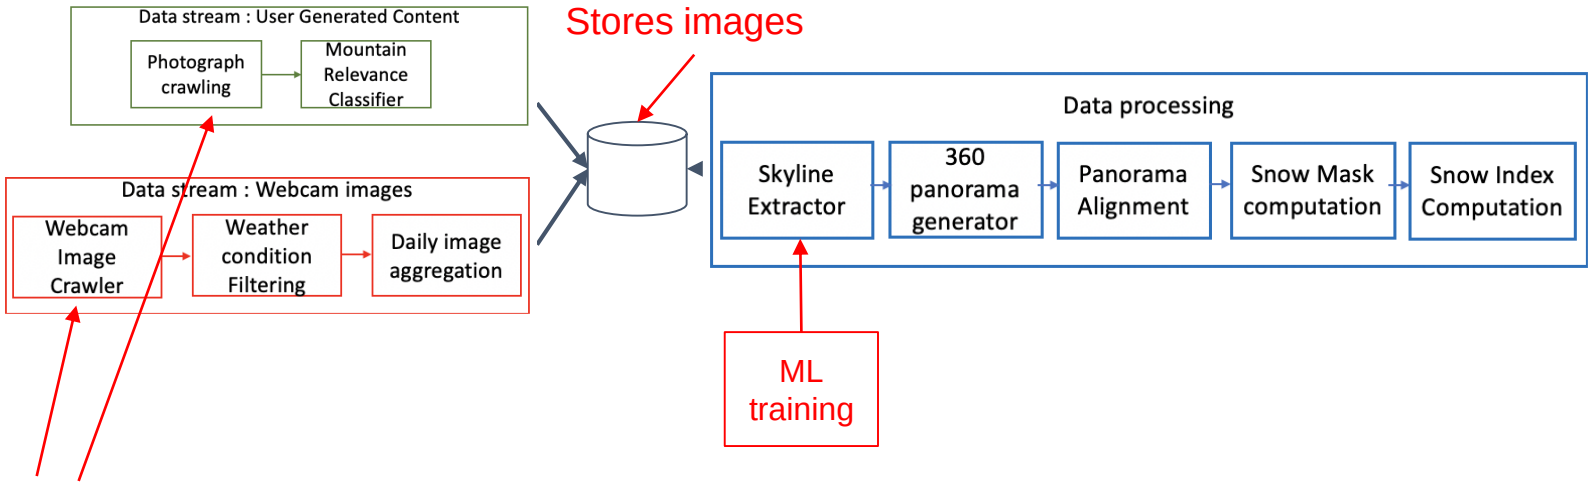

Chron jobs calling the other components through a pipe & filter approach

A closer look - multiple configurations

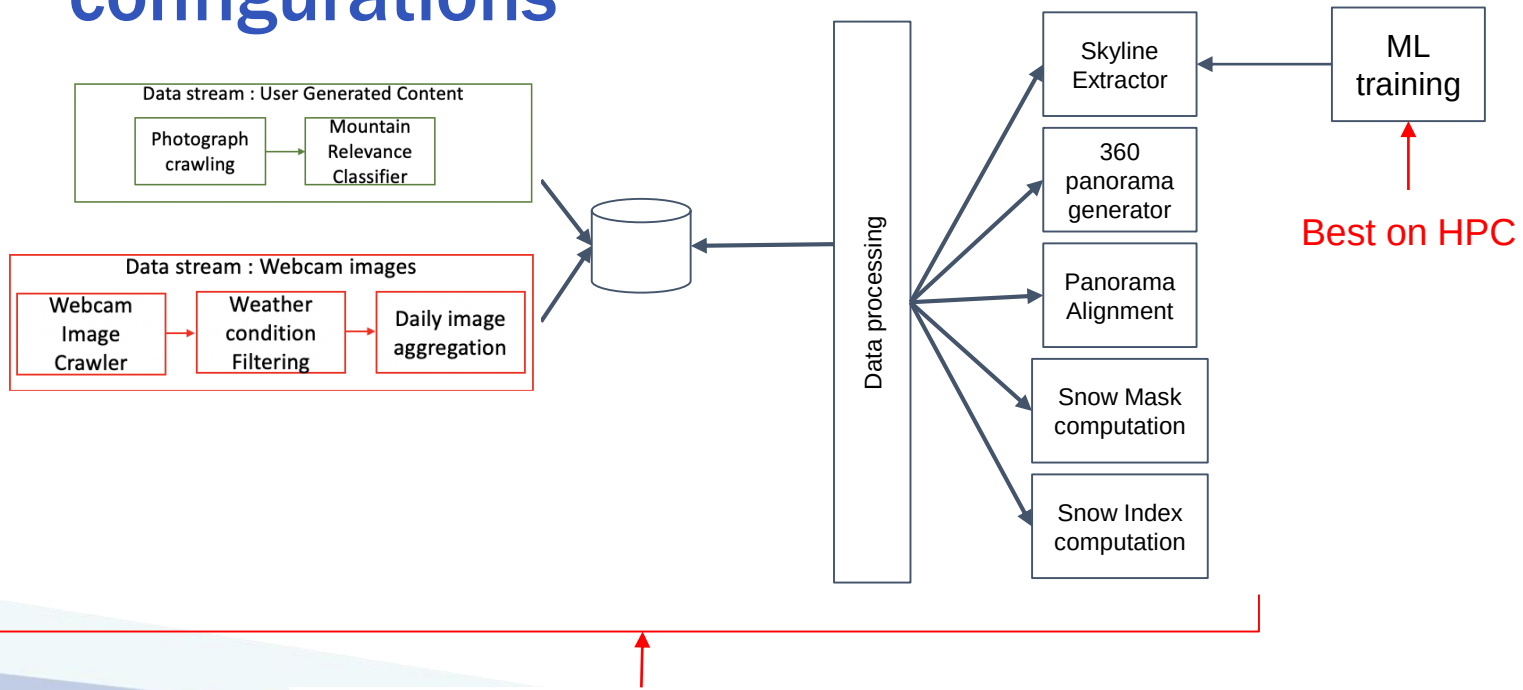

The problem

How expensive is for a non-IT intensive company

- to handle the deployment of such application and
- to make the process replicable?

The realm of DevOps tools

Sodalite

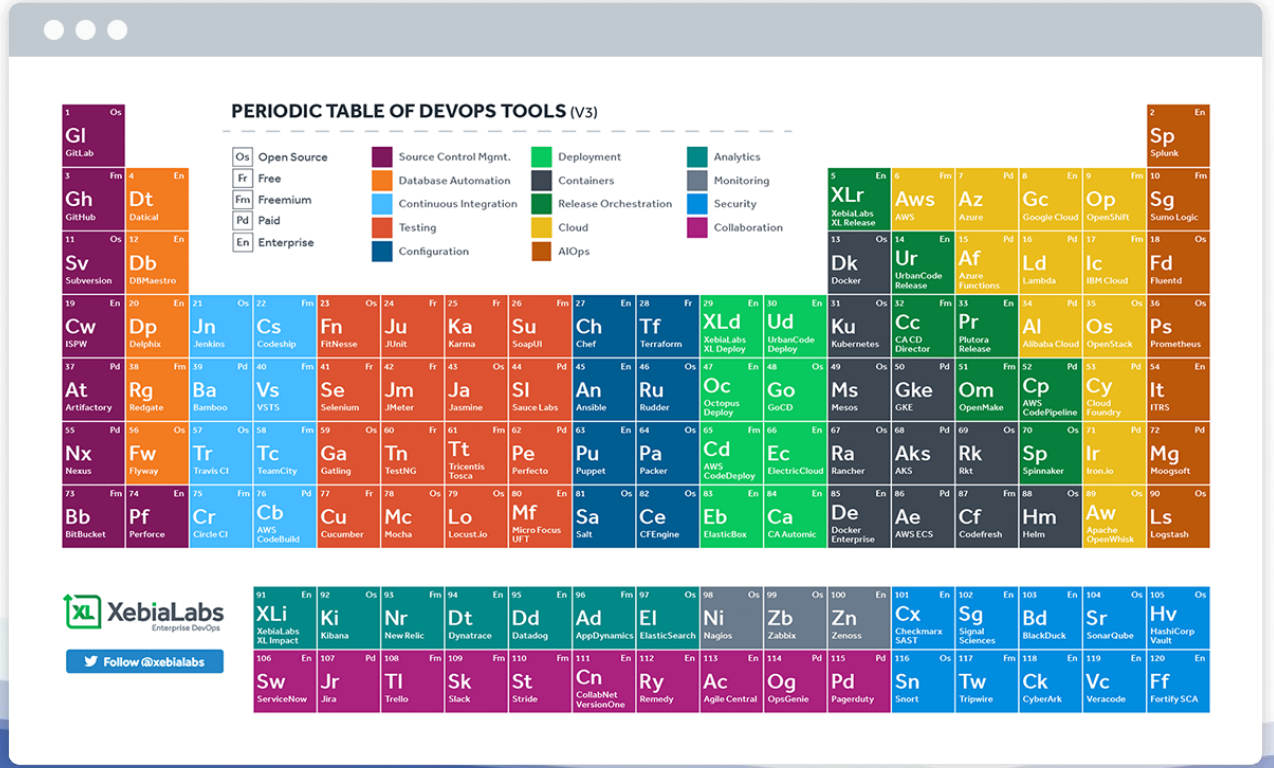

Towards standard Infrastructure-as-Code

Sodalite

➔ **An Application Deployment Topology**, i.e., “a graph of physical artefacts that need support for several lifecycle phases (e.g., procurement, installation, configuration, deployment, undeployment, teardown, etc.)”

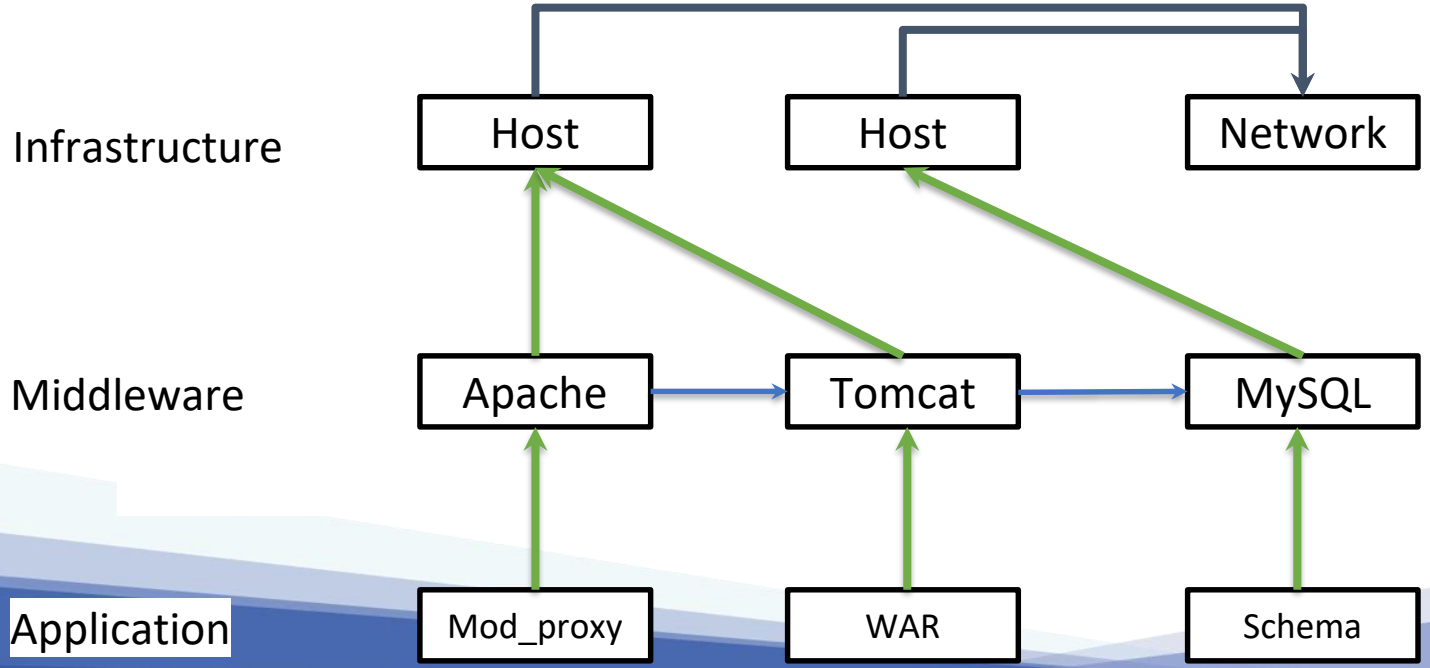

Towards standard Infrastructure-as-Code

Sodalite

➔ **Infrastructure-as-code**, i.e., “a blueprint detailing physical artefacts, all scripts for all lifecycle phases and all artefacts needed for deployment”

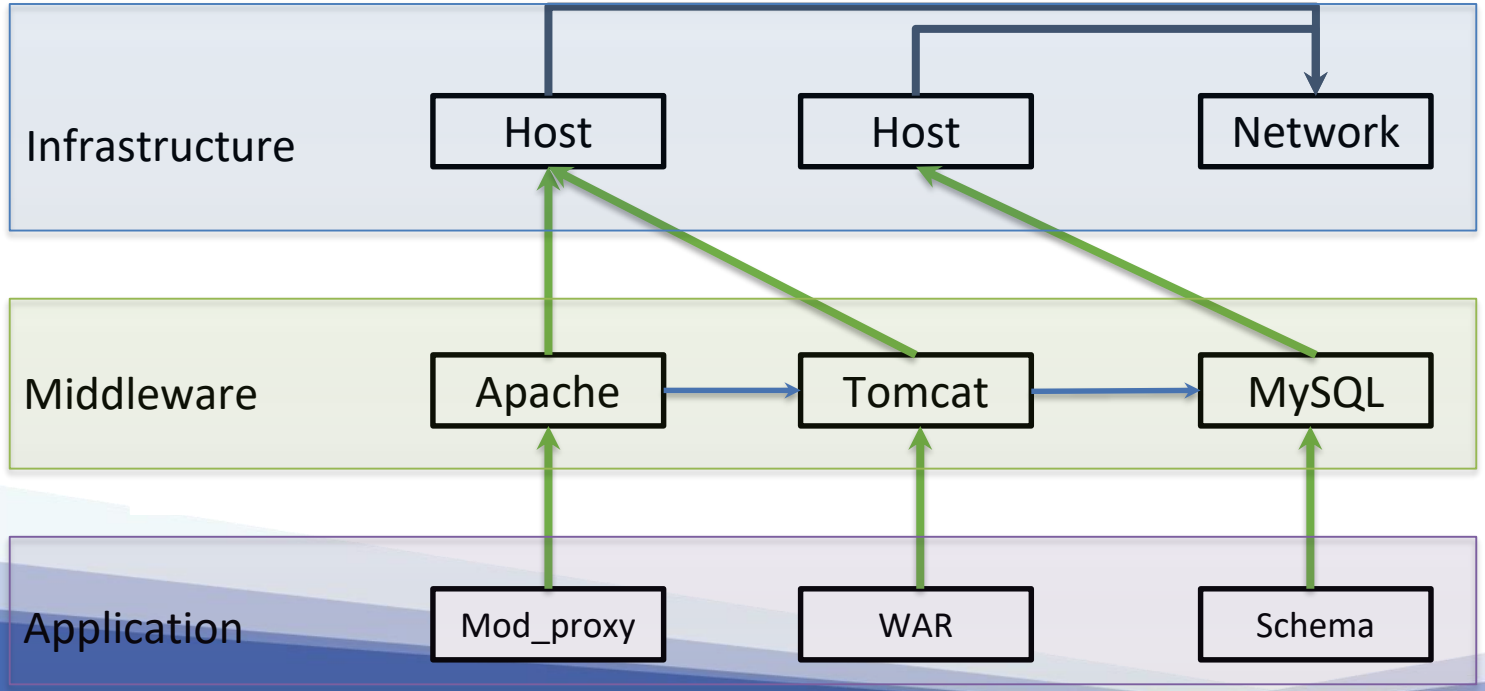

IaC Blueprint

Referenced in

IaC Middleware Scripts (e.g., Chef, Puppet, etc.)

Included in

Deploy artefacts (e.g., JARs, etc.)

Where Does TOSCA fit into?

Sodalite

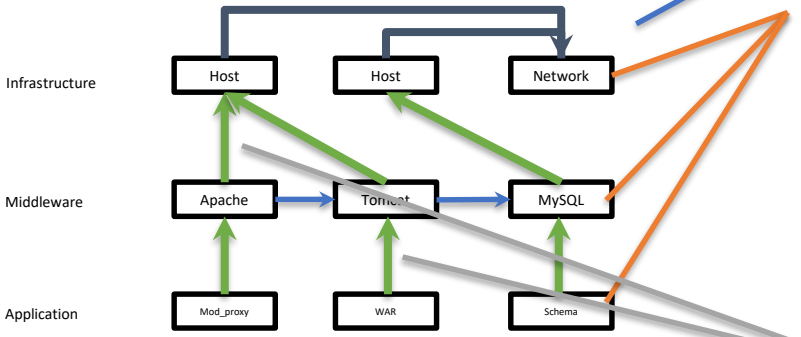

- An application topology
- 3 layers
  - Infrastructure (Cloud or DC objects)
  - Platform or Middleware (App containers)
  - Application modules, schemas and configurations
- Relationships between components:
  - What's hosted on what or installed on what
  - What's connected to what

Topology and Orchestration Specification for Cloud Applications (TOSCA) - an OASIS standard

Issues

Sodalite

- Complexity of the specification
- If special-purpose resources are used, they need to be specified as well (this may be difficult for application experts)
- Different types of resources often offer different APIs and access control mechanisms
  - Sometimes even a different programming style

What SODALITE offers

Sodalite

- Smart creation of deployment models through a textual and graphical DSL
- Editing is supported by an ontology-based reasoning mechanism that
  - Checks the semantic validity of a model
    - E.g., it signals a problem if a requirement of a source node is not satisfied by a capability of the target node
  - Enables the development of decision making tools, e.g.:
    - context-aware assistance of user at design-time
    - model enrichment taking into account domain knowledge

What SODALITE offers

Sodalite

- Automatic discovery and modeling of new infrastructural resources

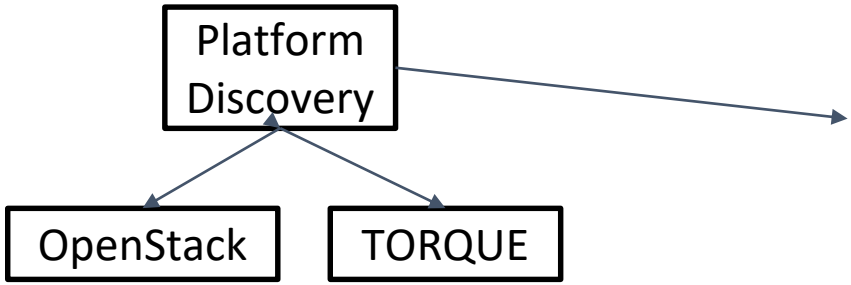


```
sodalite.nodes.OpenStack.VM:
  derived_from: tosca.nodes.Compute
  properties:
    name:
    …
  attributes:
    id:
      type: string
      description: OpenStack id of the VM
    private_address:
      type: string
      description: Private ipv4
      ….
```

What SODALITE offers

Sodalite

Support to design time application optimization for HPC

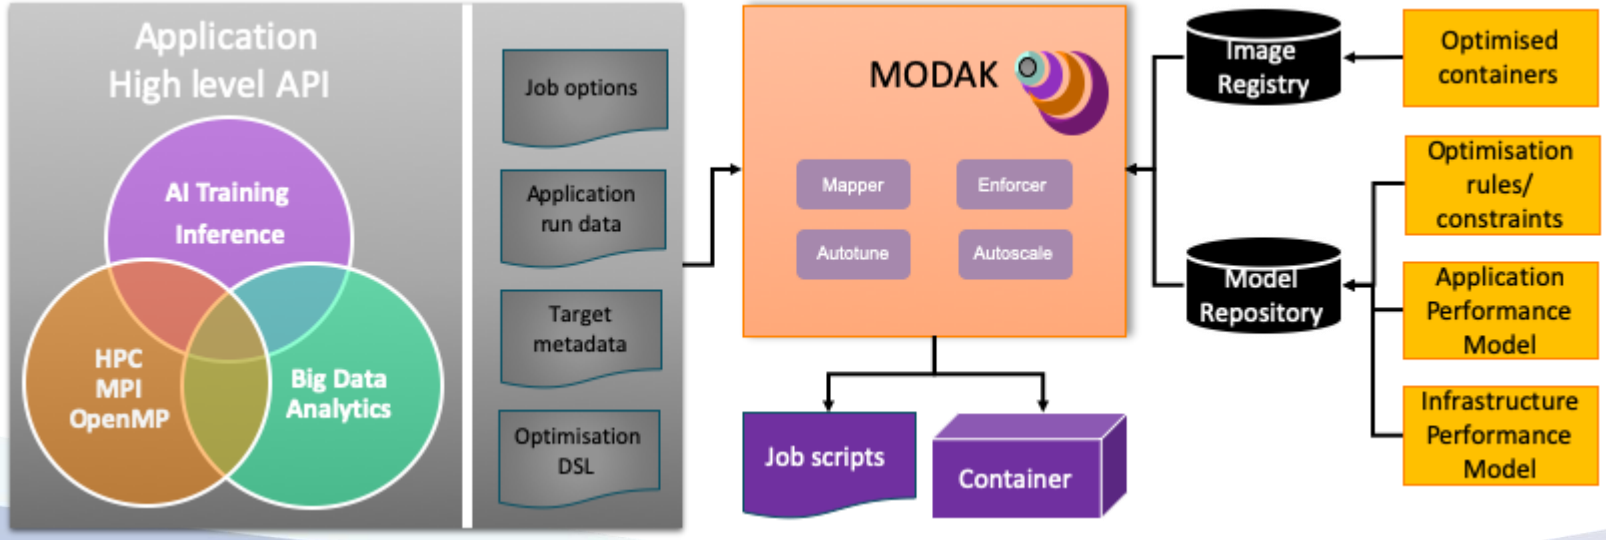

What SODALITE offers

- Supports the identification of bug smells in deployment models
- Identify potentially advantageous reconfigurations in running applications

The SODALITE Deployment

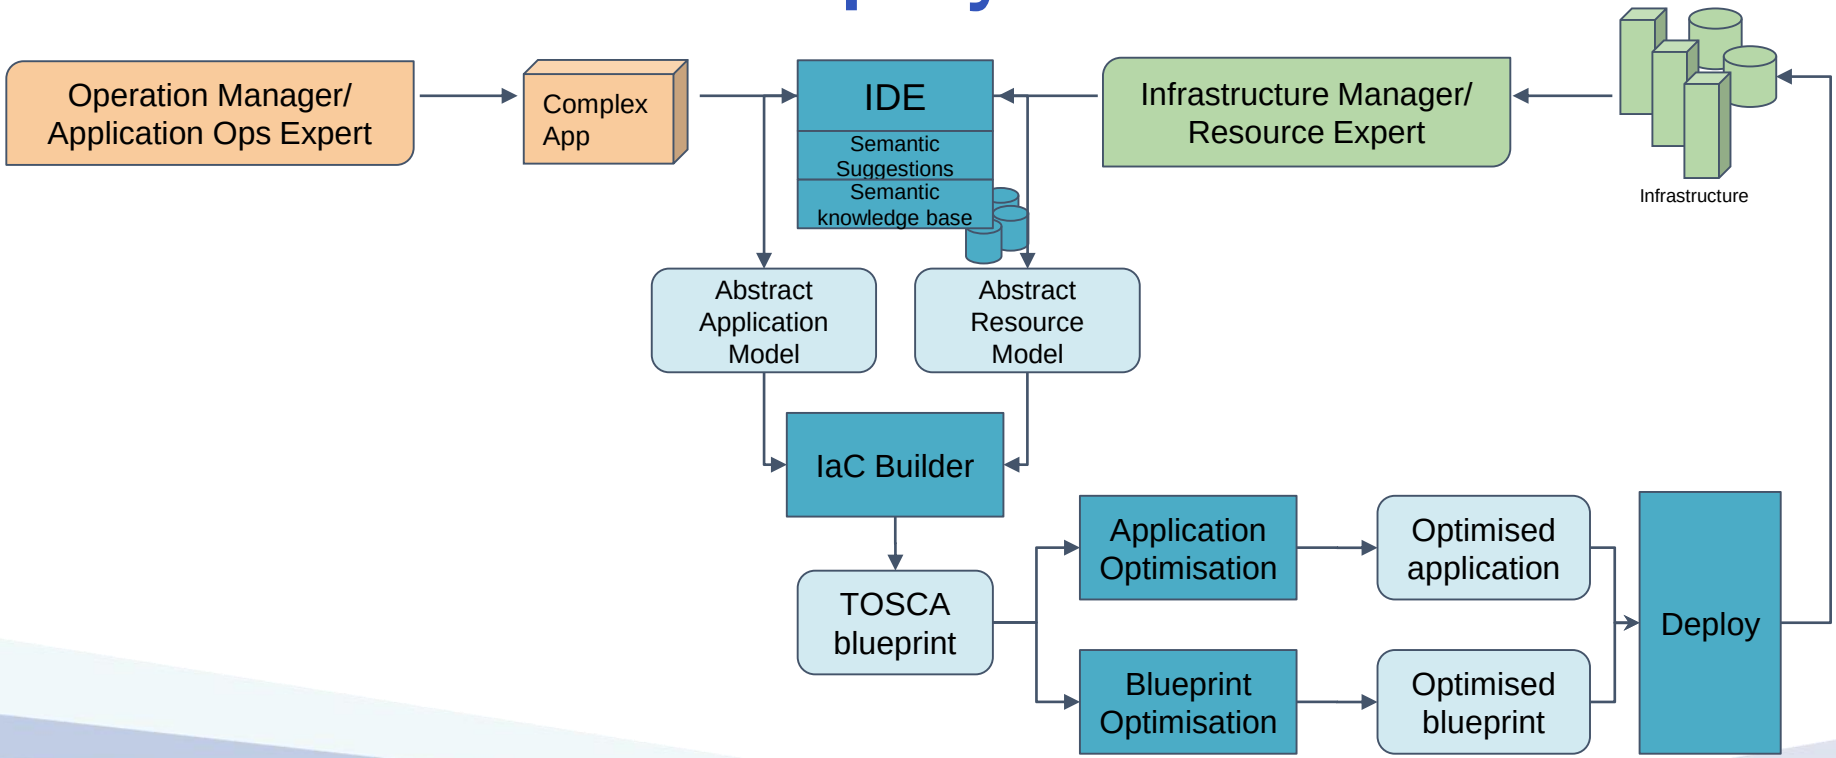

Evaluation

Sodalite

- Three case studies from three different domains:
  - A computationally-intensive scientific workflow aiming at simulating the behavior of a screw-rod fixation bone implant system on virtual patients – **Clinical Trial Simulation**
  - An adaptive system for acquiring and elaborating information from moving vehicles, able to reconfigure based on privacy preferences and rules holding in specific countries – **Vehicle IoT**
  - A system able to collect images concerning mountains and to compute the quantity of snow and, therefore, of water available in the area **– Snow**

Conclusion

Sodalite

- SODALITE provides tools to enable simpler and faster development, deployment and execution of heterogeneous apps in HPC, Cloud, Edge, & SW defined computing environments.
- **Future works:**
  - Incorporation of edge resources
  - Dynamic self-adaptation of application deployments
  - Enhancement of the reasoning capabilities to reduce even further the modeling effort by end users

Questions?

Sodalite

Thank You