WHITE PAPER #3

SODALITE

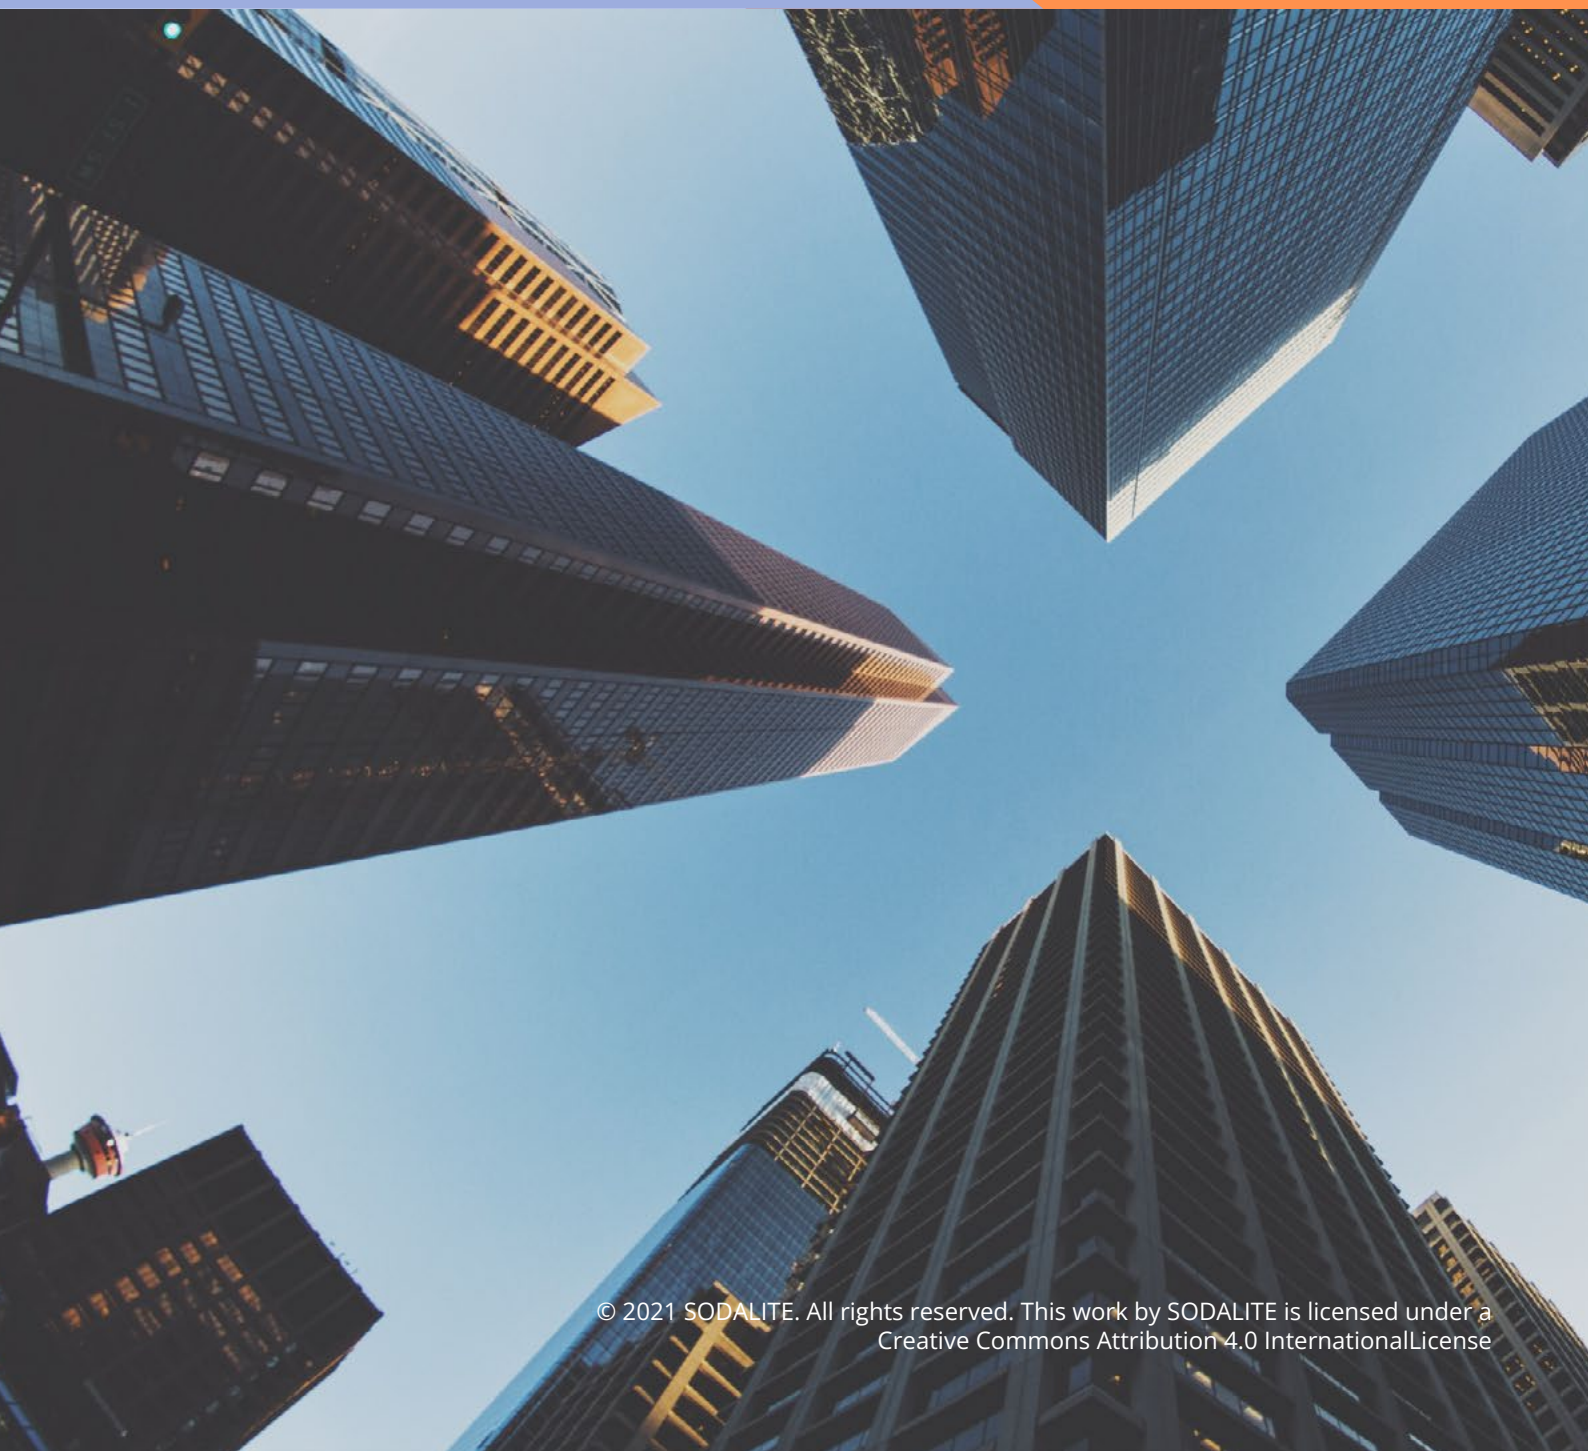

Sodalite

SOftware Defined AppLication Infrastructures managementT and Engineering

Abstract

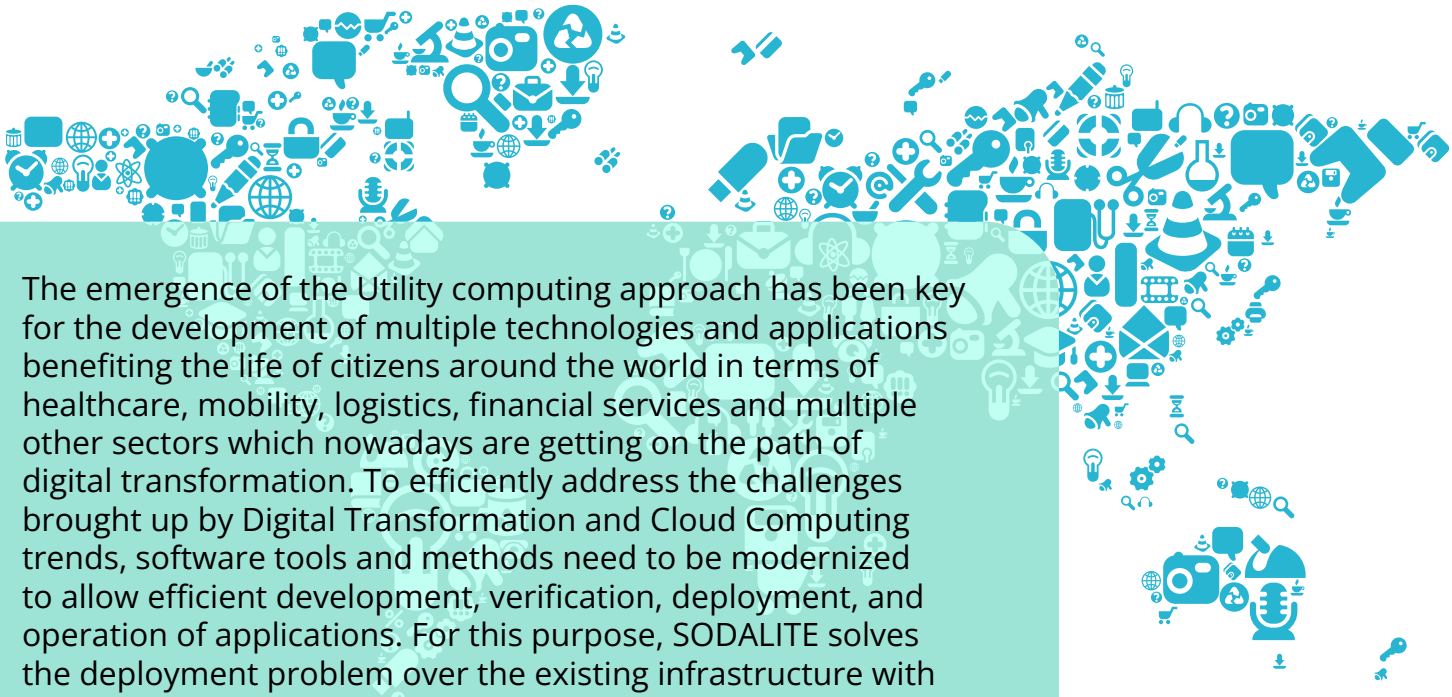

The emergence of the Utility computing approach has been key for the development of multiple technologies and applications benefiting the life of citizens around the world in terms of healthcare, mobility, logistics, financial services and multiple other sectors which nowadays are getting on the path of digital transformation. To efficiently address the challenges brought up by Digital Transformation and Cloud Computing trends, software tools and methods need to be modernized to allow efficient development, verification, deployment, and operation of applications. For this purpose, SODALITE solves the deployment problem over the existing infrastructure with focus on application management and performance. SODALITE provides a powerful and robust end-to-end toolkit supporting Digital Transformation of European Industry by decreasing design and runtime effectiveness of software-defined infrastructures, ensuring high-performance execution over dynamic heterogeneous execution environments, increasing simplicity of modelling applications and infrastructures, and improving manageability, collaboration, and time-to-market.

Keywords

HPC, IaC, Runtime, DevOps, Heterogeneous applications, Cloud Computing, software-defined, innovation

Motivation

In recent years, the global ICT market has seen a tremendous rise in utility computing, which serves as the backend for practically any new technology, methodology or advancement from healthcare to aerospace. We are entering a new era of heterogeneous, software-defined, high performance computing environments. In this context, SODALITE creates order into chaos, as the way for all these technologies to be harmonized in a single ecosystem extending their focus to other types of hardware and architectures.

SODALITE provides application developers and infrastructure operators with tools that a) abstract their application and infrastructure requirements to b) enable simpler and faster development, deployment, operation, and execution of heterogeneous applications reflecting diverse circumstances over c) heterogeneous, software-defined, high-performance, cloud infrastructures, with a particular focus on performance, quality, manageability, and reliability.

In doing so, it produces several tangible results:

1) A pattern-based abstraction library, including application, infrastructure and an absolute novum, performance abstractions.

2) A design and programming model for full-stack application and infrastructure descriptions, using the abstraction library.

3) A deployment framework, enabling static optimization of the so-abstracted applications onto specific infrastructures.

4) An automated runtime optimization and management of so-deployed applications.

This toolset will directly support Digital Transformation of European Industry through i) increasing design and runtime effectiveness of software-defined infrastructures, to ensure high-performance execution over dynamic heterogeneous execution environments; and ii) increasing simplicity of modelling applications and infrastructures, to improve manageability, collaboration, and time-to-market.

SODALITE Stack

SODALITE solution is based on four pillars:

1) **Semantic Intelligence** for describing, deploying and optimizing applications: Provides meaningful support to developers/operators and optimize their applications. By meaningful support, we denote not just support on how to construct the TOSCA (through a simplified TOSCAesque DSL), but also on the deployment of the actual application on the targeted infrastructure and networking of the application. The user gets feedback on what can be improved. The support is also provided with regards to the SLAs under which the application should run (these are the so-called static performance guarantees).

2) **Quality of service** based on guaranteed SLAs: Runtime performance optimization and adherence to the design time SLAs, guaranteeing running within the SLA envelope, provided during design time and offer performance improvements dynamic, based on the system state.

3) **Heterogeneity** support of underlying infrastructures: In order to deploy and optimize an application practically on any underlying system.

4) **Security/Privacy & Policy:** Considered as a cross-cutting concern in all layers.

These pillars support the SODALITE layers, which indeed represents the final offering to end users.

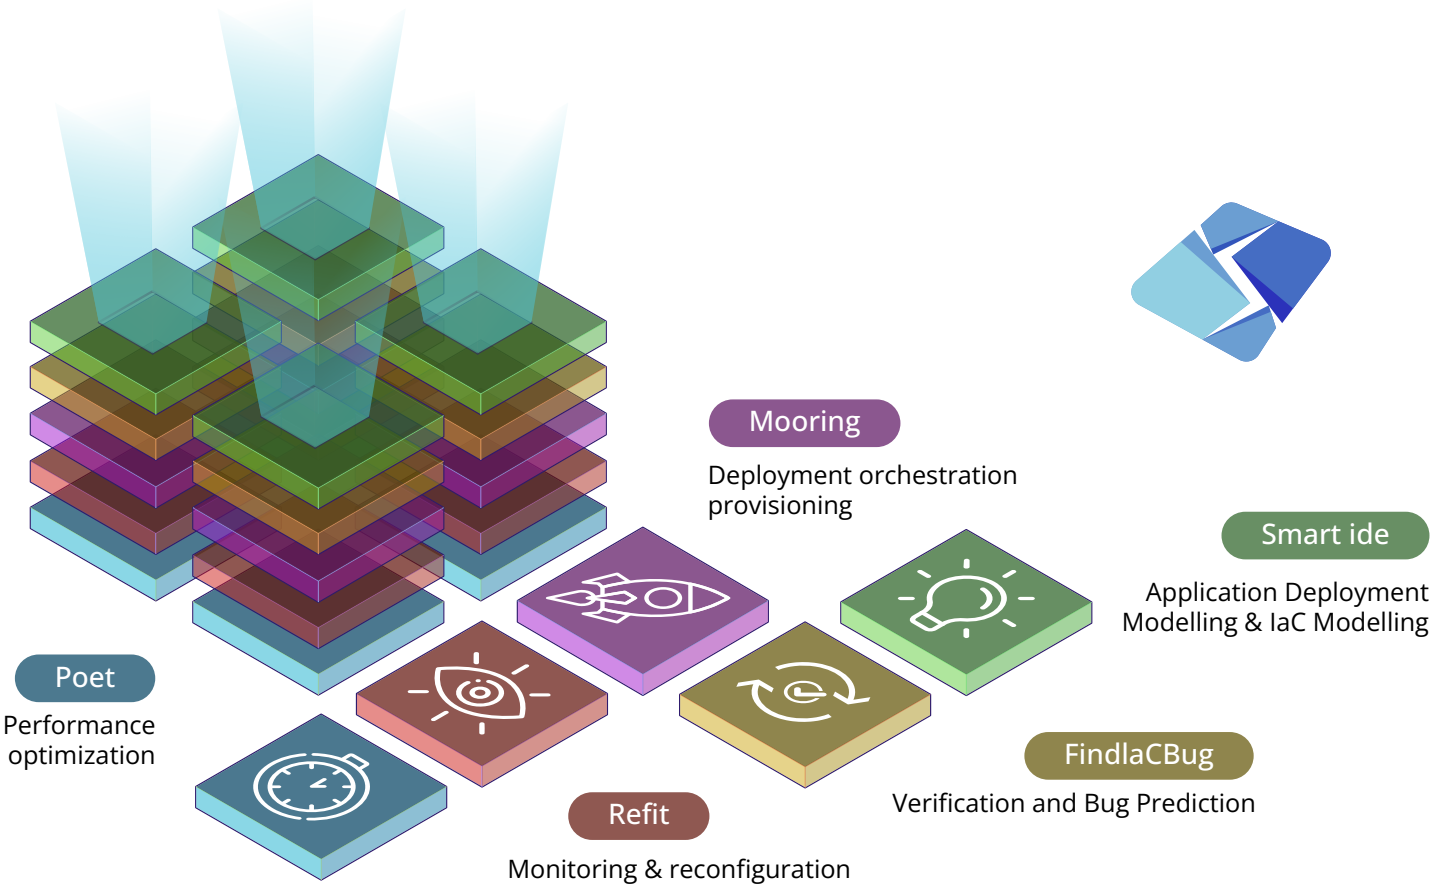

Smart IDE

Nowadays, the rise of new technologies and the digitalization of firm-related processes are highlighting the need of managing applications over heterogeneous infrastructures. SODALITE addresses this challenge providing an integrated environment for developers to model complex application deployment schemes.

Smart IDE provides an integrated programming environment for developing applications to be deployed over heterogeneous infrastructures. Through the DSL editor, complex application deployment models can be developed, assisted by the Reasoner which provides recommendations and validations, taking into account additional considerations such as performance and security. In this way, modelers, even less experienced ones, are able to develop applications to be developed using more complex deployment schemes without increasing the learning curve.

Provides the user with an IDE and a semantic reasoner API based on sophisticated knowledge graphs and semantic reasoning.

An IDE that requires the least information from the experts by getting suggestions and validation from the Knowledge Base

- **Problem**: Complex deployment model on heterogeneous infrastructures (too many considerations to be taken into account) without further assistance based on previous experiences
- **Solution**: DSL editor to develop a complex application deployment model, Reasoner to provide recommendations/validations, TOSCA application deployment model.
- **Value**: Intellisense integrated process to develop complex deployment schemas taking into account different considerations such as performance and security

Benefits

- Smart: It requires the least information from the experts by getting suggestions and validation from the Knowledge Base
- Comprehensive: Based on a wide range of knowledge-base semantic reasoning-driven results
- Configurable: The corresponding Experts can author AADMs, RMs, by getting validation and intelligent recommendations during design time
- Meaningful: Provides inferred knowledge for application modelling possibilities, performance models, security policy, deployment policies, monitoring and refactoring

FindIaCBug

The increased usage of DevOps makes software development more agile and with a higher quality. One of the main practices is Infrastructure as Code (IaC), where infrastructure is treated in the same way as the application code does. However, this is not an easy job as misconfiguration is one of the major causes of cloud incidents on enterprises due to the high amount of insecure IaC templates used.

It is responsible for the verification of the IaC code and the detection of various quality issues, such as errors, smells, antipatterns and bugs in IaC artifacts.

FindIaCBug allows users to develop high-quality detect-free error-free IaC codes. This is done by the provided tool, able to detect main quality issues using data-driven techniques for detecting antipatterns and misconfigurations, as well as semantic reasoning and rule-based models for detecting smells. It also provides recommendation fixes for detected ones thanks to a unified catalog of IaC best/bad practices, smells and bugs. These allow users to secure even more IaC templates avoiding cloud incidents and their associated costs.

DevOps experts can find and fix the syntactical/structural and semantic errors in IaC artifacts prior to deploying and executing them

- **Problem**: To build high-quality IaC artifacts, the users need to follow the recommended best practices of developing IaC scripts, and avoid applying the bad practices. Still, they can inadvertently introduce errors, smells and bugs to the IaC code.
- **Solution**: The users need a tool that can help them to easily and interactively check the quality of the IaC code they develop, and to get recommendations on how to fix any detected quality issue
- **Value**: The ability to develop high-quality defect-free error-free IaC code, based on the most advanced machine learning methods, semantic reasoning and rule-based models

Benefits:

- Progressive: Detection of linguistic anti-patterns and misconfigurations in IaC using data-driven techniques such as machine learning, deep learning, and natural language processing
- Complete: It has the sufficiently complete taxonomies of IaC best/bad practices, smells, and bugs
- Capable: can verify IaC codes for some errors, can find some smells, and can find linguistic anti-patterns
- Scalable: in time it will support more verification cases, detect more smells and misconfiguration errors, and recommend fixes for some detected smells

MOORING

Deploying applications along the continuum means dealing with a set of heterogeneous infrastructures from the edge up to the cloud, including HPC facilities. The decision-making process is not so simple, identifying when and where applications must be deployed in a simple manner without experts intervening manually.

MOORING provides the means to automate the deployment of different applications over heterogeneous infrastructures, from HPC to Kubernetes clusters, including support for OpenFaaS functions. Through an orchestration engine, infrastructure provisioning and configuration as well as application deployment, redeployment, configuration and refactoring are possible.

Consists of a framework for the usage of IaC and concepts of modelling application deployment, ensuring a smooth resource operation.

Offering through IaC the choice between various HPC, Cloud & Edge environments and scalability of applications

- **Problem**: Difficulty of manual deployments, lack of TOSCA compliance for deployment of complex application on heterogeneous infrastructure including HPC, cloud, edge
- **Solution**: Providing TOSCA-compliant means to the automation of the provisioning, deployment and orchestration of applications on target infrastructure
- **Value**: Agnostic to the choice between various HPC, private/public Cloud and Edge environments and the appropriate scalability of applications

Benefits:

- Compliant: Builds runtime images supporting full or partial redeployment and refactoring of the TOSCA blueprints
- Inclusive: Supports the heterogeneity of the building process using IaC and TOSCA with the orchestrator for building heterogeneous runtime images
- Equipped: Enables the truly concurrent provisioning and configuration of infrastructure and deployment of applications, with lightweight orchestration in HPC
- Robust: Ensure deployment, orchestration and provisioning, relying on mature technologies including the xOpera orchestrator and the application lifecycle management ALDE

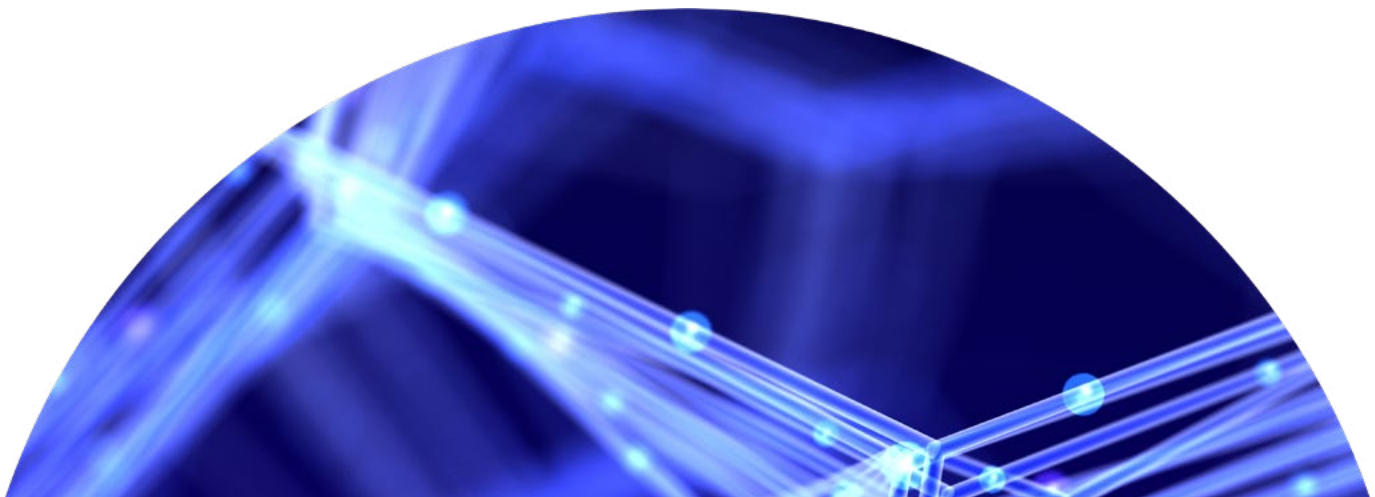

REFIT

Working with heterogeneous infrastructures involves monitoring data coming from different sources to select the most appropriate deployment options and to detect any malfunctioning of deployed applications. Current tools do not provide a single unified view to analyse traces of functional and non-functional requirements. SODALITE provides a set of tools to simplify decision-making processes based on monitored data.

REFIT provides a single entry-point for dynamic monitoring of runtime data in multi-cloud and HPC clusters with an alerting mechanism, that can be used to trigger the refactoring of deployments on several infrastructures for ensuring the highest QoS level. User will also benefit of the deployment configuration selection methodology, based on benchmarking, runtime monitoring, machine learning and software product line techniques to ensure proper deployment and/or re-deployment of their applications.

Ensures the proper system and network monitoring, collecting application-level and infrastructure-level metrics and events, further used for reconfiguration and

Dynamic monitoring, runtime resource discovery and autonomous refactoring of application deployments will change the way we manage heterogeneous (multi-cloud and HPC) environments

- **Problem**: It is difficult to monitor data management coming from different sources, pinpoint the application that is causing the problem and provide an accurate selection of appropriate deployment options for a given context
- **Solution**: Appropriately collect, store and aggregate data to simplify access and use per-connection network metrics instead of just per network interface
- **Value**: Unified monitoring data collection and presentation to simplify the monitoring of heterogeneous infrastructures and can ease decision-making processes

Benefits:

- Vigilant: Can detect the violations in the application performance SLAs based on runtime monitoring data, and refactor the current deployment of an application to maintain or improve their performance
- Flexible: Balancing different non-functional requirements such as performance, cost, thermal energy, security risks, and privacy risks
- Adjustable: Dynamic discovery and use of deployment options as new options become available/ unavailable/ changed
- Comprehensive: Allows for a coordinated management of CPU and GPU resources for multiple concurrent and containerized applications

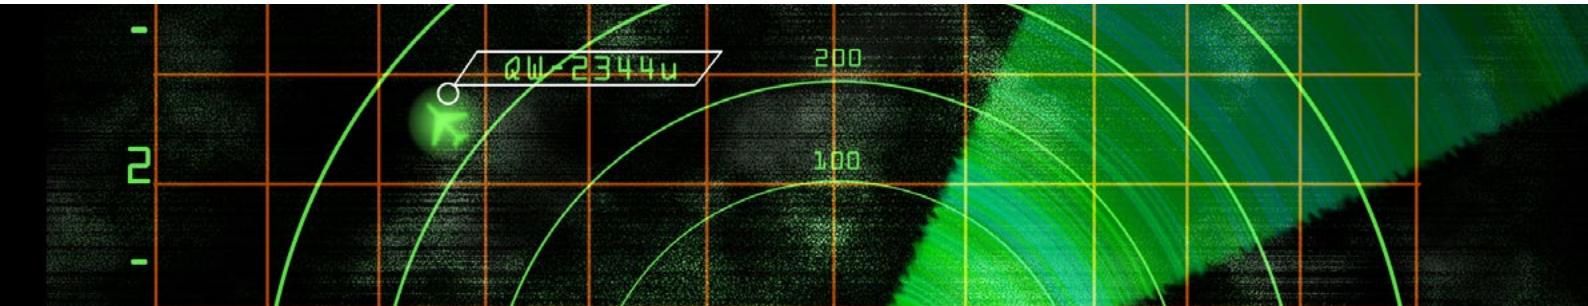

POET

Migrating applications over heterogeneous infrastructures requires re-starting the optimization processes again, as they are not portable, and involves a huge investment in system configuration. SODALITE addresses this challenge providing an abstraction layer for application experts with limited hardware or optimization knowledge to be able to use diverse targets in an optimal way.

POET provides a set of tools for scaling applications across infrastructures in the most optimal way as possible, without the need of undertaking an exhaustive study of both application and underlying infrastructure. At the same time, it can cater for additional knowledge, profiling data or autotuning of application parameters for AI and Big Data deployments.

Static and dynamic application optimization, during design and runtime.

Enabling application experts with limited hardware or optimisation knowledge to use diverse targets in an optimal way

- **Problem**: Application experts with limited hardware or optimisation knowledge using diverse targets in non optimal ways and in the absence of automation
- **Solution**: Autotune and improve application runtime parameters for better application performance, using the pre-built, optimised containers, modified to build an optimised container for the application deployment
- **Value**: Automating experience of manual optimisation for an IaC environment, porting application optimisation to cloud and edge deployments based on HPC

Benefits:

- Proactive: Can make performance decisions based on the available target's application inputs and configuration, and on the performance modelling
- Resourceful: Provides the user with tools to optimise an application deployment during design (status) and runtime (runtime)
- Prepared: Uses the pre-built, optimised containers from the Image Registry and modifies them to build an optimised container for the application deployment
- Adjustable: The optimisation process is adjusted to cater for the wide diversity of usage patterns and requirements in HPC, Cloud and Edge

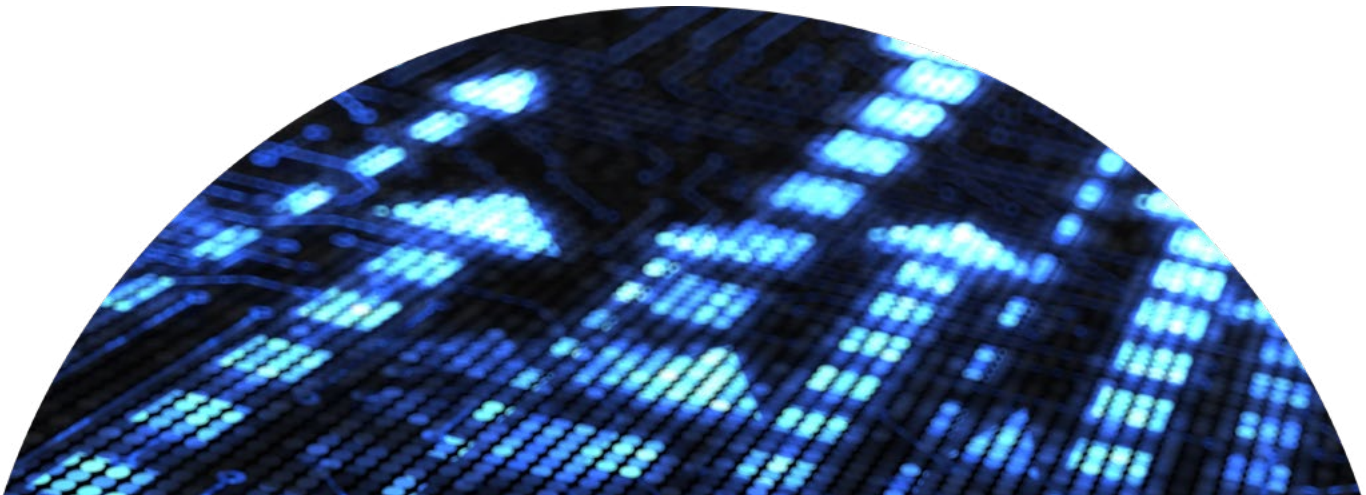

Conclusions & Recommendation

SODALITE brings the vast knowledge of performance optimization accrued by the HPC industry into the cloud computing arena. So, it exploits a model-driven approach, empowered by ontological reasoning, to support experts operating applications or resources over heterogeneous infrastructures.

SODALITE offers its layered solution based on four pillars and five stacks. Each of the SODALITE stacks represents a product itself but can be combined for an enhanced experience.

SODALITE tools are available for download in https://github.com/SODALITE-EU. This site contains detailed information about all tools, including source code access and installation and use instructions.

Acknowledgements

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

Authors

- Lara López (ATOS)
Senior Business Consultant

- Joao Pita Costa (XLAB)
Exploitation Leader

Annex 1: Glossary

Acronym	Description
Devs	Developments
DevOps	Development and Operations
DSL	Domain Specific Language
HPC	High Performance Computing
IaC	Infrastructure as Code
Ops	Operations
SLA	Service Level Agreement

About SODALITE

Computer application developers and infrastructure operators are always searching for new tools to improve their performance. The EU-funded SODALITE project supplies tools for developers and infrastructure operators that abstract their applications and infrastructure requirements to facilitate simpler and faster development and operation and execution of diverse applications over diverse software-defined, high-performance cloud infrastructures. It generates a model-based abstraction library and develops a design and programming model for full-stack application and infrastructure descriptions. It also establishes a deployment framework, enabling static optimization of the so-abstracted applications onto specific infrastructures, and automated runtime optimization and management of applications. SODALITE supports the digital transformation of the European industry by enhancing the design and effectiveness of software-defined infrastructures and simplicity of modelling applications.

Find more about us


sodalite.eu


sodalite-eu


sodalitesw


projectinfo@sodalite.eu


github.com/SODALITE-EU

Let's start a discussion together