

Infrastructure-as-Code: How to make the Textual Representation of Scientific Infrastructures FAIR

Tuesday, 2 April 2019 14:00 (30 minutes)

Background:

The approach of infrastructure-as-code allows to efficiently manage large infrastructures, for instance to support FAIR data management. A canonical and machine-actionable description of these infrastructures can itself be an item of research and an essential component in handling reproducibility challenges for the results achieved on the infrastructures. Such a description would include all necessary steps and dependencies in order to provision the infrastructure described in a machine-actionable and human-readable way.

Objective:

The textual representations of the infrastructures need both to be handled in a FAIR spirit and to be compliant to common DevOps quality management standards: On the one hand, they should be Findable, Accessible, Interoperable and Reusable. This can be achieved by annotating the infrastructure-as-code with rich metadata, including a permanent identifier that can be used to retrieve the textual representation. On the other hand, they should be integrated into a workflow of continuous integration and deployment and managed by a version control system.

Method:

Guidelines and best practices to meet these requirements are presented, derived from experience with several research infrastructures hosted at the Leibniz Supercomputing Centre. The infrastructures considered span from small, but specialized systems to large-scale HPC clusters. Where beneficial, interviews with administrators were conducted to enrich the guidelines with their experiences.

Results:

The major recommendation is to stick to those standards and tools that are already existent and relevant for the continuous and scalable management of infrastructures in an industrial context. Ansible is an example for a well-established configuration and provisioning tool that finds application in both scientific and industrial service proliferation. DataCite as a generic metadata scheme is well-equipped to meet the domain-specific requirements with regard to infrastructure description. Beyond that the crucial part is to integrate both domains with qualified references and show their potential with a proof-of-concept.

Evaluation:

The guidelines are exemplified by a proof-of-concept to provision a generic research data infrastructure which runs on a kubernetes cluster. This infrastructure includes services to

harvest data providers, a search engine and several backend services to facilitate common workflows such as bookmarking of search results, data staging and platforms to process and analyze the data. The infrastructure as well as the code deployed on it are open source and can be used to reproduce the presented findings. Conclusion: The adoption of the recommendations presented cannot only make infrastructure setups citeable, but might boost best practices and facilitate the federation of distributed services in the context of science.

Summary

The approach of infrastructure-as-code allows to efficiently manage large infrastructures, for instance to support FAIR data management. The canonical description of these infrastructures can itself be an item of research and an essential component in handling reproducibility challenges with regard to the results achieved on the infrastructures. Such a canonical description typically includes a machine-actionable set of instructions to setup machines, services and seed data. The textual representations of the infrastructures need both to be handled in a FAIR spirit and to be compliant to common DevOps quality management standards. Guidelines and best practices to meet these requirements are presented, derived from experience with large-scale research infrastructures hosted at the Leibniz Supercomputing Centre. The major recommendation is to stick to those standards and tools that are already existent and relevant for the continuous and scalable management of infrastructures in an industrial context. The only modifications necessary of these can be found where scientific context adds requirements that typically cannot be found in a business context. The adoption of these recommendations cannot only make infrastructure setups citeable, but might boost best practices and facilitate the federation of distributed services in the context of science.

Primary authors: Dr KRANZLMULLER, Dieter (LMU Munich); Mr WEBER, Tobias (Leibniz Supercomputing Centre)

Presenter: Dr KRANZLMULLER, Dieter (LMU Munich)

Session Classification: Networking, Security, Infrastructure & Operations

Track Classification: Network, Security, Infrastructure & Operations