Thesis for The Degree of Doctor of Philosophy

Pre-deployment Description Logic-based Reasoning for Cloud Infrastructure Security

Claudia Cauli

Presentation:
June 17th, 2022, 09:30
Online & Room EC, 4th Floor, EDIT Building
Hörsalsvägen 11
University of Gothenburg, Campus Johanneberg

Faculty Opponent:
Prof. Piero Bonatti
Dipartimento di Ing. Elettrica e Tecnologie dell'Informazione
Università degli Studi di Napoli Federico II

Division of Computing Science
Department of Computer Science & Engineering
University of Gothenburg
Gothenburg, Sweden, 2022

Abstract

Ensuring the security of a cloud application is exceptionally challenging. Not only is cloud infrastructure inherently complex, but also a precise definition of *what is secure* is hard to give. Business context, regulatory compliance, use cases, intent, and human interpretation influence this definition, and what is considered secure in one setting may not be in another. This thesis aims to improve the extent to which automated techniques support manual security reviews and, by doing so, to aid users of all levels in designing infrastructure compliant with their security standards. To achieve this objective, we investigate the application of provable techniques to security analyses amenable to early design phases. In particular, we study description logic-based semantic reasoning for the pre-deployment modeling and verification of cloud infrastructure.

The body of this thesis is based on three published papers. In the first paper, we encode AWS CloudFormation deployment language into the expressive description logic $\mathcal{ALCOIQ}$. We verify configuration checks with ad-hoc reasoners and sketch an axiomatization of security knowledge to reason about system-level properties. We find that expressive logics can simulate partial closed-world reasoning, vulnerabilities, and mitigations to threats but trigger high complexity of the reasoning tasks and require cumbersome modeling. To overcome these, in the second paper, we define a novel lightweight logic and a query language for security threats. The logic mixes open- and closed-world assumptions to succinctly encode complete and incomplete knowledge. The query language embeds optimistic and pessimistic reasoning to express *vulnerabilities that may be present* versus *mitigations that must be in place*. Lightweight logics enable tractability: knowledge base satisfiability and query answering become decidable in AC^0 and LogSpace data complexity, respectively. Lastly, in the third paper, we build on this new formalism by introducing a language to encode *mutating actions* (that create, delete, or modify cloud resources) and defining the transition system generated from an initial configuration when all possible actions are applied. In the transition system, states represent alternative configurations, and transitions represent changes induced by the actions. By focusing on the planning problem, we search for sequences of actions that mitigate the potential vulnerabilities of the initial configuration. Due to the practical decision procedures of the underlying formalism, we do so in PTime data complexity.

Keywords: Description Logic, Security, Cloud, Automated Reasoning

ISBN 978-91-8009-839-7 (PRINT)
ISBN 978-91-8009-840-3 (PDF)