Multi-Stakeholder Policy Enforcement for Distributed Systems

Robert Walther
robert.walther@tu-dresden.de
TU Dresden
Germany

Carsten Weinhold
carsten.weinhold@barkhauseninstitut.org
Barkhausen Institut
Germany

Peter Amthor
peter.amthor@tu-ilmenau.de
TU Ilmenau
Germany

Michael Roitzsch
michael.roitzsch@barkhauseninstitut.org
Barkhausen Institut
Germany

Abstract

Cloud environments, comprising both virtual and physical servers, are complex distributed systems that require clear and expressive configuration descriptions. While human-readable formats like Kubernetes YAML are common, they often lack the granularity needed for fine-grained control and advanced policy enforcement. To address these limitations, we propose an abstract system description approach that incorporates additional application properties, enabling more sophisticated policy decision-making beyond basic resource and port restrictions.

Our approach implements an end-to-end policy enforcement mechanism that verifies high-level system descriptions against security policies and automatically translates them into concrete configurations. This mechanism not only ensures compliant deployment, but also enables communication partners to assess trustworthiness with greater granularity before interacting.

1 Problem Statement

In cloud environments, applications rely on the operating system for resources like networking and storage while interacting with programs owned by diverse, partially untrusted, stakeholders. Managing these interactions requires evolving deployment and security policies to meet regulatory and operational demands. For instance, policies may prevent applications from sharing a processor core or demand that a database is configured to encrypt all stored data. Manually configuring systems to enforce such multi-layered and often dynamic rules is not only cumbersome but prone to human error.

The Need for Automation While existing cloud solutions provide run-time access control, they lack automated policy enforcement during system configuration and deployment. This gap increases the risk of insecure or otherwise non-compliant deployments. Moreover, trust assessments often rely on certificate checks or remote attestation, which commonly reduces available configuration details to an opaque hash sum, thereby preventing nuanced policy evaluation. An automated solution is needed to ensure policy compliance from deployment through operation, thereby minimizing misconfigurations and strengthening overall security.

End-to-End Enforcement We propose an approach that bridges the gap between the intention of stakeholders and the actual implementation by automatically propagating high-level security policies from abstract configuration to deployment and then to run-time verification. First, a policy enforcement engine assesses a platform-independent system description — annotated with metadata about applications, services, and communication channels — against predefined security requirements. If the system's description complies with all policy constraints, the engine automatically generates a platform-specific configuration. To guarantee that subsequent verification accurately reflects the system's untampered state, configurations must be bound to the platform using, for example, trusted execution environments (to be done in future work).

Abstract Policy Specifications By abstracting away low-level configuration details, our method enables nuanced, context-aware policy decisions that are automatically translated into concrete, system-specific settings. This separation not only simplifies policy specification but also empowers communication partners such as users or other services that interact with the system. They can independently verify the system's trustworthiness against their own security policies without having to concern themselves with platform details. In this way, our solution extends robust trust guarantees from the moment of deployment to the operational phase.

2 Contribution

With our work, which is described in more detail in a paper [1], we propose a comprehensive, end-to-end policy enforcement mechanism for distributed systems. Our solution introduces a straightforward scenario and policy description language that captures the communication requirements and constraints of multi-program deployments. The policy enforcement engine automatically translates high-level, metadata-rich descriptions into policy-conformant, platform-specific configurations and supports independent verification of the scenarios against security policies prior to communication. We designed and implemented the translation step in an easily extendable way. We demonstrate the suitability of the approach for three different platforms, ranging from an embedded system to state-of-the-art container runtimes, namely Kubernetes and Docker Compose. Our approach to describing, enforcing, and verifying application deployments empowers all stakeholders — including regulatory bodies and infrastructure providers — to express fine-grained policies for distributed systems without needing to manage the complexities of the underlying cloud infrastructure.

References

[1] Robert Walther, Carsten Weinhold, Peter Amthor, and Michael Roitzsch. Multi-stakeholder policy enforcement for distributed systems. In *Proceedings of the 10th International Workshop on Container Technologies and Container Clouds*, WoC '24, page 7–12, New York, NY, USA, 2024. Association for Computing Machinery.