

Dependable Cloud Computing with OpenStack

OpenStack – The “Linux of the Cloud”

OpenStack is steadily gaining popularity as a rising standard for cloud management stacks. As a young and dynamic open source project with short feedback loops and release cycles, OpenStack is representative of recent trends in cloud computing, and therefore interesting to study.

Major companies such as HP and Huawei are looking into new application areas for OpenStack, such as telecommunications and the thorough virtualization of network functions. While such applications could profit enormously from a flexible and scalable cloud solution, these domains also strictly demand *dependability* on all levels of the software stack.

Dependability Concerns in Cloud Management Stacks

Experience from previous projects has taught us that the deployment of OpenStack is a complex and error-prone task. Nevertheless, failures in the cloud infrastructure induce high financial companies and certain security and availability violations may be unacceptable. Anecdotes from practice suggest that cloud failures arise from complex error propagation chains. Since they are frequently owed to design and software flaws, such failures cannot be avoided by simply replicating everything.

Project Goal: Studying the Dependability of OpenStack

In this master project, we want to tackle the question how a desired level of dependability can be achieved within OpenStack. We will approach the question from two sides.

First, we want to create a fault tolerance test bed installation of OpenStack, which allows for fault injection experiments and the testing of different fault tolerance strategies. This test bed may be virtualized, so that repeated deployment, fault injection and reproducible experiments can be scripted and simplified. There have been pioneering approaches to fault injection in cloud environments such as ChaosMonkey¹, which we will study and apply.

Second, we will analytically study the dependability of an OpenStack architecture. We will develop a specification of what the expected behaviour of OpenStack under faultload is. Further, we will discuss an OpenStack-specific fault model, so that quantitative dependability modelling techniques can be applied to evaluate and compare different OpenStack deployments. Such a fault model needs to take into account failures on all levels of a complex and intertwined cloud software stack. A vast body of anecdotal evidence² from modern cloud companies can be consulted to develop such a fault model.

Contact

Prof. Dr. Andreas Polze andreas.polze@hpi.de

Christian Neuhaus christian.neuhaus@hpi.de

Lena Herscheid lena.herscheid@hpi.de

Operating System and Middleware Group

¹ <https://github.com/Netflix/SimianArmy>, accessed 03/02/2015

² http://failure.wiki/Category:Post_Mortem, accessed 03/02/2015