

Testing of complex, large-scale distributed storage systems: a CERN disk storage case study

Thursday, July 12, 2018 11:15 AM (15 minutes)

Complex, large-scale distributed systems are more frequently used to solve extraordinary computing, storage and other problems. However, the development of these systems usually requires working with several software components, maintaining and improving large codebases, and also a relatively large number of developers working together. Therefore, it is inevitable to introduce faults to the system. On the other hand, these systems often perform important if not crucial tasks so critical bugs, performance-hindering algorithms are not acceptable to reach the production state of the software and the system. Also, the larger number of developers can work more liberated and productively when they receive constant feedback that their changes are still in harmony with the system requirements and other people's work which also greatly helps scaling out manpower, meaning that adding more developers to a project can actually result in more work done.

In this paper we will go through the case study of EOS, the CERN disk storage system and introduce the methods and possibilities of how to achieve all-automatic regression, performance, robustness testing and continuous integration for such a large-scale, complex and critical system using container-based environments. We will also pay special attention to the details and challenges of testing distributed storage and file systems.

Primary authors: MAKAI, Jozsef (CERN); PETERS, Andreas Joachim (CERN); BITZES, Georgios (CERN); SINDRILARU, Elvin Alin (CERN); SIMON, Michal Kamil (CERN); MANZI, Andrea (CERN)

Presenter: MANZI, Andrea (CERN)

Session Classification: T5 - Software development

Track Classification: Track 5 – Software development