

Challenges with Continuous Deployment of NoSQL-backed Database Applications

Meike Klettke¹, Stefanie Scherzinger², Uta Störl³,
Stephanie Sombach², and Katharina Wiech²

¹ University of Rostock, Germany

² OTH Regensburg, Germany

³ Darmstadt University of Applied Sciences, Germany

We address a practical challenge with the continuous deployment of database applications, which actually constitutes a data integration problem: Upon a new deployment of the application code, entities already persisted in the production database no longer match what the application code expects. Apart from migrating all legacy entities *eagerly* at the time of the release, *lazy* migration is an alternative popular with NoSQL data stores: A schema-flexible database stores entities with legacy structure, as well as up-to-date entities. When a legacy entity is loaded into the application, all pending structural changes are applied. Thus, from the viewpoint of the application, entities are always up-to-date.

Yet lazily migrating legacy data from several releases back, involving more than one entity at-a-time, is not a trivial task. At LWA 2015 [3], we presented our vision of a schema management unit for NoSQL data stores that carries out schema evolution lazily: This involves an internal, Datalog-based model for reading, writing, and migrating data [2]. However, we use Datalog not only to specify the semantics of schema evolution operations, but Datalog is our actual vehicle for carrying out data migrations: In this overview talk, we introduce *Datalution* [1], a tool that alternatively evaluates our Datalog rules bottom-up (for eager data migration) or top-down (for lazy data migration). In particular, our tool allows for an easy comparison of both approaches in terms of the number of physical writes to the data store.

We demonstrate *Datalution*, provide insight into its mechanics, and outline our next steps in integrating the *Datalution* engine with an industrial-strength NoSQL data store.

Keywords: Schema evolution, NoSQL data stores, Datalog

References

1. Scherzinger, S., Sombach, S., Wiech, K., Klettke, M., Störl, U.: Datalution: A Tool for Continuous Schema Evolution in NoSQL-backed Web Applications. In: Proceedings QUDOS'16 (2016), Tool Demo.
2. Scherzinger, S., Störl, U., Klettke, M.: A Datalog-based Protocol for Lazy Data Migration in Agile NoSQL Application Development. In: Proceedings DBPL'15 (2015)
3. Störl, U., Klettke, M., Scherzinger, S.: Kontrolliertes Schema-Evolutionsmanagement für NoSQL-Datenbanksysteme. In: Proceedings LWA 2015 Workshops: KDML, FGWM, IR, and FGDB. (2015)