SGD Migration to AWS Architecture

Ajay Shrivatsav Vichanthangal Prathivadhib, Travis K. Sheppard, Gail A. Binkley, Felix Gondwe, Kalpana Karra, Shuai Weng, J. Michael Cherry and the SGD Project

Stanford University, School of Medicine, Department of Genetics, Stanford, CA

Over its twenty-five years of existence, the *Saccharomyces* Genome Database (SGD; https://www.yeastgenome.org) has become an indispensable service for the yeast community. During that time, SGD has also experienced a proliferation of services, data types, and web pages which has made it more expensive to maintain. In order to make a more sustainable service for its users, SGD was recently re-architected for Amazon Web Services (AWS) cloud infrastructure to be consistent with industry best practices. Migrating SGD to AWS had many benefits including flexibility to provision and decommission servers as needed, while maintaining high reliability and robustness in service. The architecture is multi-layered. User requests are handled by the load balancer that reroutes them to one of the productions web servers, which in turn query the database slave servers for data replicated from the master database. These architecture changes have left SGD a leaner, more sustainable service consistent with industry best practices. SGD is funded by the NIH-NHGRI [U41HG001315].

Previous Architecture

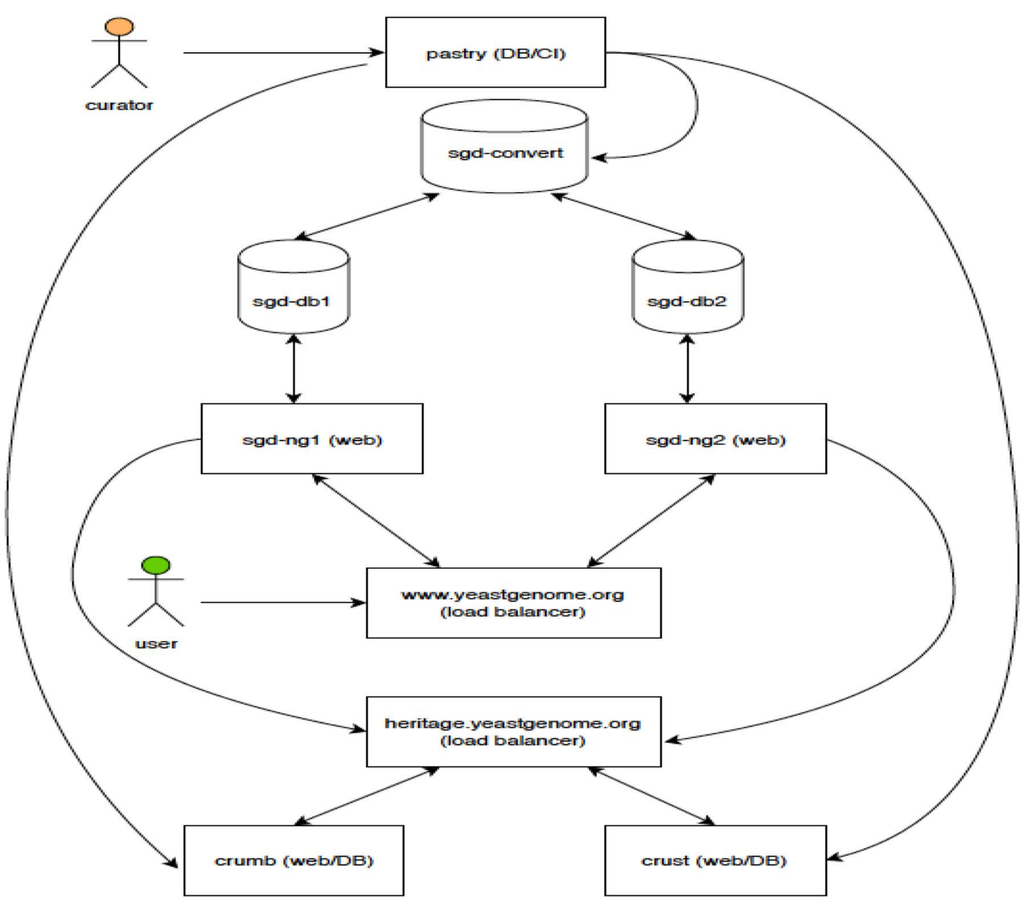

SGD Architecture at AWS

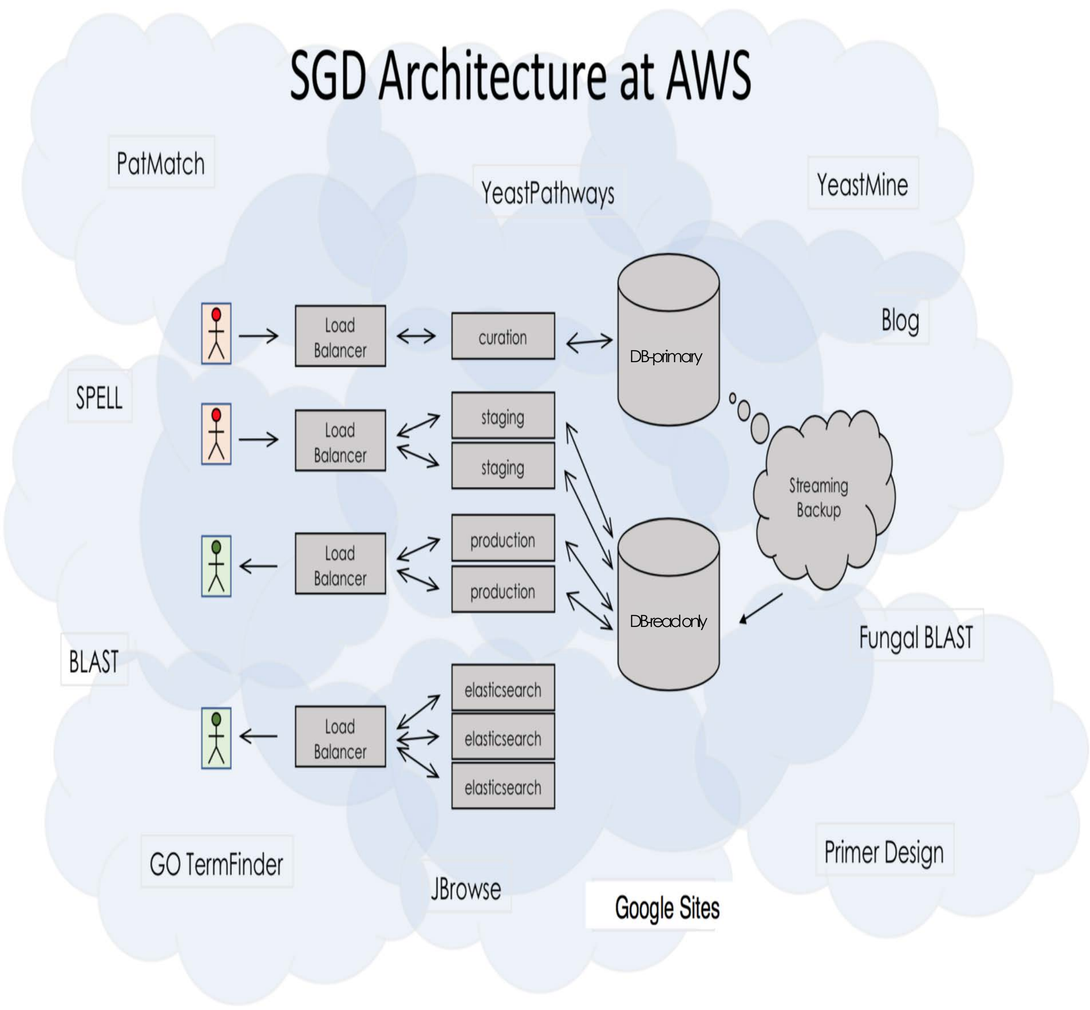

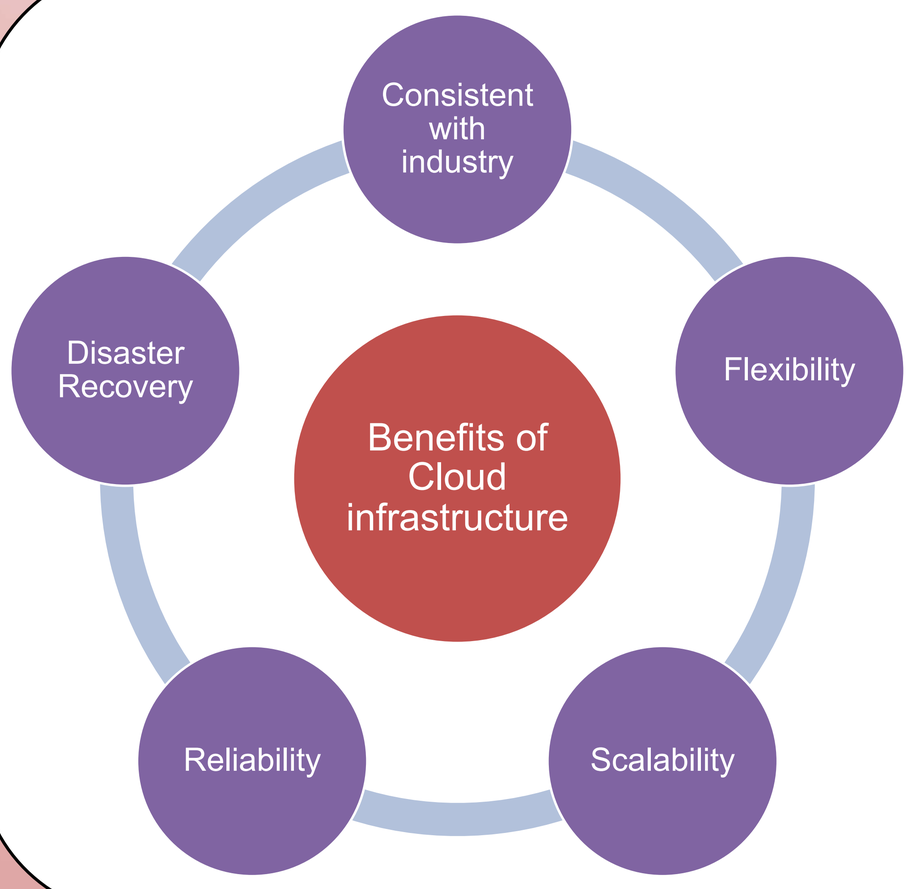

Industry best practice recommends having average response time < 200 ms

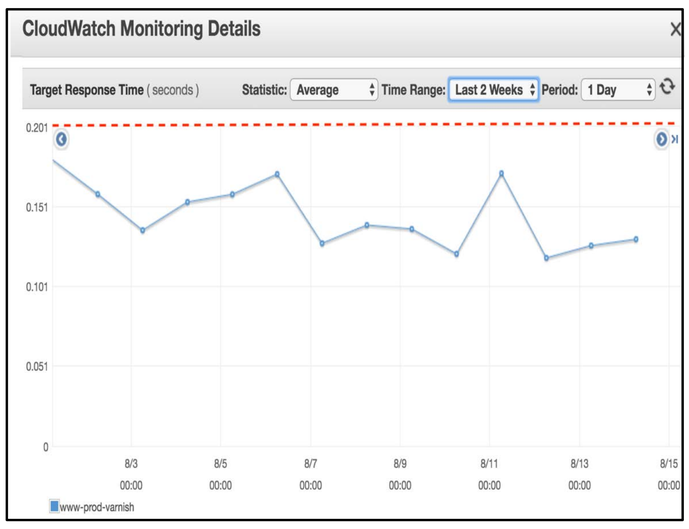

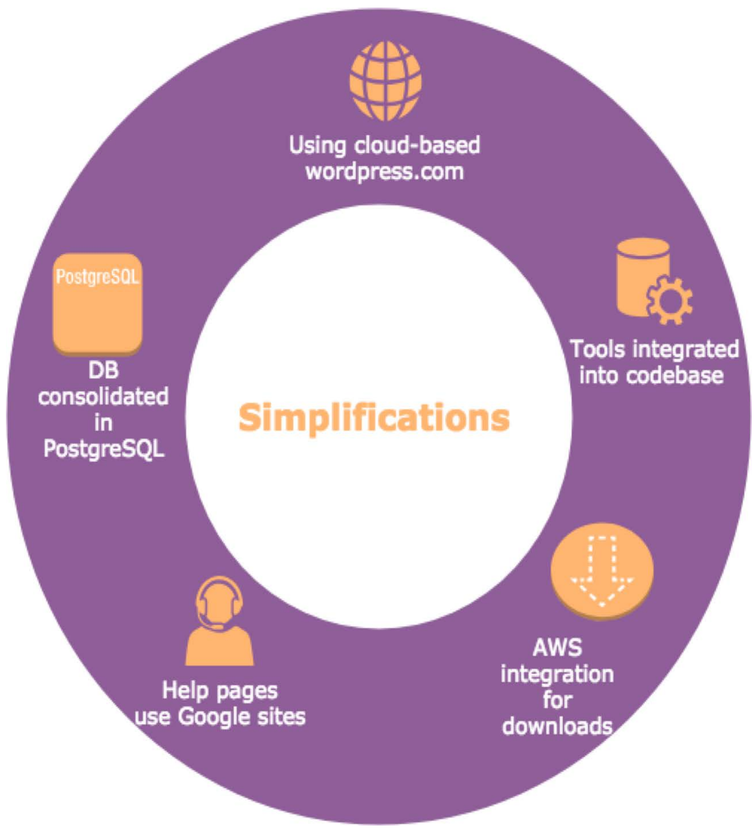

@yeastgenome

sgd-helpdesk@lists.stanford.edu

https://www.facebook.com/yeastgenome/

https://www.youtube.com/SaccharomycesGenomeDatabase

ALLIANCE of GENOME RESOURCES FOUNDING MEMBER

department of GENETICS

Stanford University