

GoogleSQL: A SQL Language as a Component

(Invited Talk)

David Wilhite

Senior Staff Software Engineer & Manager

Google

wilhite@google.com

Abstract

Google has built many systems using SQL: public cloud products like BigQuery and Spanner, internal systems like F1 Query and Procella, open source ZetaSQL, and other systems that use SQL as a declarative API for data processing and management outside traditional SQL databases and query engines. Building GoogleSQL as a component has enabled easy reuse across many systems, with full language compatibility.

Sharing a language definition and parser alone is not sufficient to achieve real consistency or easy reuse. The GoogleSQL libraries include shared language analysis, shared implementations of common functions and operators, a compliance testing framework for validating consistency across engines, and other shared subcomponents.

The consistency that results from sharing the language and implementation across tools has enabled the growth of a large ecosystem of interoperable tools that can easily leverage the full power of SQL, while still allowing domain-specific extensions and customizability.

Biography

David Wilhite has worked as a Software Engineer at Google for 10 years, where he currently leads a team building componentized infrastructure for large-scale analytics use cases, which includes the GoogleSQL component. Prior to joining Google, David worked across the DBMS stack (compiler, optimizer, execution) at various DBMS companies including Red Brick Systems and ParAccel.