

New Tools for Data Management: The Rise of Hadoop, NoSQL and NewSQL

Monday, 19 March 2012 15:00 (45 minutes)

For three decades, business users and scientists have chosen between two platforms for data storage and analysis: Simple file systems for complex or application-specific data, and relational databases for relatively simple, structured data that needed good query or transactional properties. Over the last five years, a variety of powerful but diverse new technologies has emerged: Columnar stores to speed up structured data access, so-called NoSQL or NewSQL engines for high performance storage and retrieval of single records, and Hadoop, for scale-out storage and massively parallel analysis of both structured and complex data.

In this talk, I will digest the technical, business and industry pressure that led to this sudden explosion in variability among platforms, and will explore some of the Darwinian pressures that apply to the market now. I'll make some predictions about the composition of the data center – including the infrastructure that will underlie big science – of the future, and will describe how these different kinds of products operate and interact.

Presenter: Mr OLSON, Mike (Cloudera)

Session Classification: Plenary