

DataModel

Introduction

Cassandra is a partitioned row store, where rows are organized into tables with a required primary key.

The first component of a table's primary key is the partition key; within a partition, rows are clustered by the remaining columns of the PK. Other columns may be indexed independent of the PK.

This allows pervasive denormalization to "pre-build" resultsets at update time, rather than doing expensive joins across the cluster.

Patrick McFadin's data modeling series:

1. The Data Model is Dead; Long live the Data Model: [Video](#), [Slides](#)
2. Become a Super Modeler: [Video](#), [Slides](#)
3. The World's Next Top Data Model: [Video](#), [Slides](#)
4. Apache Cassandra 2.0: Data Model on Fire: [Video](#), [Slides](#)
5. Real Data Models of Silicon Valley: [Video](#), [Slides](#)
6. [The most important thing to know in Cassandra data modeling: The primary key](#). (Blog post)

<https://c.statcounter.com/9397521/0/fe557aad/1/> | stats