The Database Kairometer

Chronological Years

Event Years

Events
Stonebraker
Data management research
Data management industry

Events	Start	End
Preface	2018	
Introduction	1971	2018
I. 2014 ACM A. M. Turing Award Paper and Lecture		
The Land Sharks Are on the Squawk Box	1984	2016
II. Mike Stonebraker's Career		
1. Make It Happen: The Life of Michael Stonebraker	1943	2018
The Career of Mike Stonebraker: The Chart: A. Pavlo	1970	2015
Mike Stonebraker's Student Genealogy	1973	2018
III. Mike Stonebraker Speaks Out: An Interview by Marianne Winslett		
2. Mike Stonebraker Speaks Out: An Interview	1965	2017
Interview Video	1965	2017
IV. The Big Picture		
3. Leadership and Advocacy	1973	2018
4. Perspectives: The 2014 ACM Turing Award		2015
5. Birth of an Industry; Path to the Turing Award	1973	2018
6. A Perspective of Mike from a 50-Year Vantage Point	1970	2018
V. Startups		
7. How to Start a Company in F ye (Not so Easy) Steps	2013	2018
8. How to Create and Run a Stonebraker Startup: The Real Story	2004	2018
9. Getting Grownups in the Room: A VC Perspective	2002	2018
VI. Database Systems Research		
10. Where Good Ideas Come From and How to Exploit Them	1971	2017
11. Where We Have Failed	1971	2017
12. Stonebraker and Open Source	1979	1993
13. The Relational Database Management Systems Genealogy	1970	2018
VII. Contributions By System		
14. Research Contributions of Michael Stonebraker: An Overview	1971	2018
VI.A Research Contributions by System		
15. The Later Ingres Years	1971	1984
16. Looking Back at Postgres	1983	1993
17. Databases Meet the Stream Processing Era	2000	2017
18. C-Store: Through the Eyes of a Ph.D. Student	2002	2011
19. In-Memory, Horizontal, and Transactional: The H-Store OLTP DBMS Project	2007	2016
20. Scaling Mountains: SciDB and Scientific Data Management	1992	2017
21. Data Unification at Scale: Data Tamer	2012	2018
22. BigDAWG Polystore System	2012	2017
23. Data Civilizer: End-to-End Support for Data Discovery, Integration, and Cleaning	2014	2018
VII.B Contributions From Building Systems VII.B		
24. The Commercial Ingres Codeline	1980	1990
25. The Postgres and Illustra Codelines	1989	1997
26. The Aurora/Borealis/StreamBase Codelines: A Tale of Three Systems	2000	2018
27. The Vertica Codeline	2005	2013
28. The VoltDB Codeline	2008	2014
29. The SciDB Codeline: Crossing the Chasm	2013	2016
30. The Tamr Codeline	2013	2018
31. The BigDAWG Codeline	2014	2018
VIII. Perspectives		
32. IBM Relational Database Code Bases	1991	1997
33. Aurum: A Story About Research Taste	2016	2018
34. Nice: Or What It Was Like to Be Mike's Student	1983	1994
35. Michael Stonebraker: Competitor, Collaborator, Friend	1970	2015
36. The Changing of the Database Guard	1974	1986
IX. Seminal Works of Michael Stonebraker and his Collaborators		
OLTP Through the Looking Glass, and What We Found There	1980	2008
"One Size Fits All": An Idea Whose Time Has Come and Gone	1970	2005
The End of an Architectural Era (It's Time for a Complete Rewrite)	1970	2007
C-store: A Column-Oriented DBMS	1990	2005
The Implementation of Postgres	1984	1989
The Design and Implementation of Ingres	1971	1976
Bibliographies		
The Collected Works of Michael Stonebraker	1972	2018
References		

Very Large Data Bases

ACM SIGMOD

Berkeley Database Research Group

INGRES — SQL — Microsoft SQL Server — ORACLE — IBM IMS — freeBSD — ILLUSTRA — PostgreSQL — ORACLE TimesTen — goby — NoSQL Not Only SQL — hadoop — Apache HBASE — HIVE — Apache Spark — streambase — VERTICA — VOLTDB — paradigm4 — SciDB — tamr — IBM DB2 — NewSQL — VOLTDB — Clustrix — xeround — nuoDB — memsql

The Entity-Relationship model