

Contents

Part I An Inflection Point for Enterprise Applications

1	Desirability, Feasibility, Viability: The Impact of In-Memory	3
1.1	Information in Real Time: Anything, Anytime, Anywhere	3
1.1.1	Response Time at the Speed of Thought.	4
1.1.2	Real-Time Analytics and Computation on the Fly	6
1.2	The Impact of Recent Hardware Trends	6
1.2.1	Database Management Systems for Enterprise Applications	7
1.2.2	Main Memory is the New Disk	10
1.2.3	From Maximizing CPU Speed to Multi-Core Processors.	11
1.2.4	Increased Bandwidth Between CPU and Main Memory	13
1.3	Reducing Cost Through In-Memory Data Management	16
1.3.1	Total Cost of Ownership.	16
1.3.2	Cost Factors in Enterprise Systems.	17
1.3.3	In-Memory Performance Boosts Cost Reduction	18
1.4	Conclusion	19
2	Why Are Enterprise Applications So Diverse?	21
2.1	Current Enterprise Applications	22
2.2	Examples of Enterprise Applications	23
2.3	Enterprise Application Architecture	25
2.4	Data Processing in Enterprise Applications	26
2.5	Data Access Patterns in Enterprise Applications	27
2.6	Conclusion	27

3	SanssouciDB: Blueprint for an In-Memory Enterprise	
	Database System	29
3.1	Targeting Multi-Core and Main Memory	30
3.2	Designing an In-Memory Database System	31
3.3	Organizing and Accessing Data in SanssouciDB	33
3.4	Conclusion	35
Part II	SanssouciDB: A Single Source of Truth Through In-Memory	
4	The Technical Foundations of SanssouciDB	39
4.1	Understanding Memory Hierarchies	40
4.1.1	Introduction to Main Memory	40
4.1.2	Organization of the Memory Hierarchy	44
4.1.3	Trends in Memory Hierarchies	45
4.1.4	Memory from a Programmer's Point of View	46
4.2	Parallel Data Processing Using Multi-Core and Across Servers	53
4.2.1	Increasing Capacity by Adding Resources	54
4.2.2	Parallel System Architectures	56
4.2.3	Parallelization in Databases for Enterprise Applications	58
4.2.4	Parallel Data Processing in SanssouciDB	60
4.3	Compression for Speed and Memory Consumption	64
4.3.1	Light-Weight Compression	65
4.3.2	Heavy-Weight Compression	69
4.3.3	Data-Dependent Optimization	70
4.3.4	Compression-Aware Query Execution	70
4.3.5	Compression Analysis on Real Data	71
4.4	Column, Row, Hybrid: Optimizing the Data Layout	72
4.4.1	Vertical Partitioning	72
4.4.2	Finding the Best Layout	76
4.4.3	Challenges for Hybrid Databases	78
4.4.4	Application Scenarios	79
4.5	The Impact of Virtualization	79
4.5.1	Virtualizing Analytical Workloads	80
4.5.2	Data Model and Benchmarking Environment	80
4.5.3	Virtual Versus Native Execution	81
4.5.4	Response Time Degradation with Concurrent VMs	82
4.6	Summarizing the Technical Concepts	84
4.7	Conclusion	95

5	Organizing and Accessing Data in SanssouciDB	97
5.1	SQL for Accessing In-Memory Data	98
5.1.1	The Role of SQL	98
5.1.2	The Lifecycle of a Query	99
5.1.3	Stored Procedures	99
5.1.4	Data Organization and Indices	100
5.1.5	Any Attributes as Index	101
5.2	Increasing Performance with Data Aging	103
5.2.1	Active and Passive Data	104
5.2.2	Implementation Considerations for an Aging Process	105
5.2.3	The Use Case for Horizontal Partitioning of Leads	106
5.3	Efficient Retrieval of Business Objects	108
5.3.1	Retrieving Business Data from a Database	109
5.3.2	Object Data Guide	109
5.4	Efficient Execution of Business Functions	111
5.4.1	Separating Business Functions from Application Functions	111
5.4.2	Comparing Business Functions	112
5.5	Handling Data Changes in Read-Optimized Databases	114
5.5.1	The Impact on SanssouciDB	115
5.5.2	The Merge Process	116
5.5.3	Improving Performance with Single Column Merge	120
5.6	Append, Never Delete, to Keep the History Complete	123
5.6.1	Insert-Only Implementation Strategies	123
5.6.2	Minimizing Locking Through Insert-Only	125
5.6.3	The Impact on Enterprise Applications	128
5.6.4	Feasibility of the Insert-Only Approach	130
5.7	Enabling Analytics on Transactional Data	132
5.7.1	Aggregation on the Fly	133
5.7.2	Analytical Queries without a Star Schema	142
5.8	Extending Data Layout Without Downtime	149
5.8.1	Reorganization in a Row Store	149
5.8.2	On-The-Fly Addition in a Column Store	150
5.9	Business Resilience Through Advanced Logging Techniques	151
5.9.1	Recovery in Column Stores	152
5.9.2	Differential Logging for Row-Oriented Databases	154
5.9.3	Providing High Availability	155
5.10	The Importance of Optimal Scheduling for Mixed Workloads	156
5.10.1	Introduction to Scheduling	156
5.10.2	Characteristics of a Mixed Workload	159

5.10.3	Scheduling Short and Long Running Tasks	160
5.11	Conclusion	163

Part III How In-Memory Changes the Game

6	Application Development	167
6.1	Optimizing Application Development for SanssouciDB	167
6.1.1	An In-Memory Application Programming Model	168
6.1.2	Application Architecture	172
6.1.3	Moving Business Logic into the Database	173
6.1.4	Best Practices	175
6.1.5	Graphical Creation of Views	176
6.2	Innovative Enterprise Applications	178
6.2.1	New Analytical Applications	179
6.2.2	Operational Processing to Simplify Daily Business	182
6.2.3	Information at Your Fingertips with Innovative User-Interfaces	185
6.2.4	Combining Analytics and Textsearch	190
6.2.5	Basic Types of Search	191
6.2.6	Features for Enterprise Search	191
6.3	Conclusion	194
7	Finally, A Real Business Intelligence System is at Hand	195
7.1	Analytics on Operational Data	195
7.1.1	Yesterday's Business Intelligence	196
7.1.2	Today's Business Intelligence	199
7.1.3	Drawbacks of Separating Analytics from Daily Operations	201
7.1.4	Dedicated Database Designs for Analytical Systems	202
7.1.5	Analytics and Query Languages	205
7.1.6	Enablers for Changing Business Intelligence	206
7.1.7	Tomorrow's Business Intelligence	208
7.2	How to Evaluate Databases After the Game has Changed	209
7.2.1	Benchmarks in Enterprise Computing	210
7.2.2	Changed Benchmark Requirements for a Mixed Workload	211
7.2.3	A New Benchmark for Daily Operations and Analytics	213
7.3	Conclusion	216
8	Scaling SanssouciDB in the Cloud	219
8.1	What Is Cloud Computing?	220
8.2	Types of Cloud Applications	221

8.3	Cloud Computing from the Provider Perspective	223
8.3.1	Multi-Tenancy	223
8.3.2	Low-End Versus High-End Hardware.	228
8.3.3	Replication	229
8.3.4	Energy Efficiency by Employing In-Memory Technology	230
8.4	Conclusion	231
9	In-Memory Revolution has Begun	233
9.1	Risk-Free Transition to In-Memory Data Management	233
9.1.1	Operating In-Memory and Traditional Systems Side by Side	234
9.1.2	System Consolidation and Extensibility	235
9.2	Customer Proof Points	236
9.2.1	Charité—Universitätsmedizin Berlin.	237
9.2.2	Hilti	239
9.3	Conclusion	241
	About the Authors.	243
	Glossary	245
	References	255
	Index	263