

Contents

Why You Should Read This Book by <i>Tom Ramey</i>	i
About the Authors	v
Introduction by <i>Surekha Parekh</i>	xv
DB2 12 for z/OS: Technical Overview and Highlights	
<i>by John Campbell and Gareth Jones</i>	1
Introduction	1
Why You Should Read This Paper.....	1
Highlights.....	1
Application Enablement	2
DBA Function	2
OLTP Performance.....	2
Query Performance.....	2
Quick Hits.....	3
Performance for Traditional Workloads.....	3
In-memory Computing	3
INSERT Performance.....	6
Access Path/Access Plan Stability.....	10
More Granular Global Commit LSN and Read LSN.....	15
Avoid Scheduling Unnecessary Prefetch.....	15
Miscellaneous Performance Enhancements for Traditional Workloads.....	15
System Scaling Enhancements	16
DB2/DASD Synergy	17
High-level Performance Expectations	17
Instrumentation Enhancements.....	18
Performance Enablers of Modern Applications	19
Performance Focus: Enabling New Applications.....	19
High-level Performance Focus	20
Query Performance Focus	21
Application Enablement.....	25
DB2 for z/OS—Adapting to Modern Application Development Paradigms	25
RAS—Reliability, Availability, Scalability	33
Lifting Partition Limits.....	33

Online Schema.....	35
Security and General Enhancements	37
Utilities	38
Data Sharing Improvements	42
Migration and Prerequisites	44
Migration Prerequisites: Hardware and Operating System.....	44
Migration and the Catalog	44
Online Migration Improvements	45
Questions & Answers.....	45

DB2 12 for z/OS: What’s the Latest from the Optimizer for Improved Query Performance?

<i>by Terry Purcell</i>	53
Introduction.....	53
Why Read Further?	54
What Is the Performance Focus in DB2 12 for z/OS?.....	54
Performance Expectations	55
UNION ALL and Outer Join Performance	59
Reorder Outer Join Tables to Avoid Materializations	60
Push Predicates Inside UNION ALL Legs or Outer Join Query Blocks	60
Trim Unnecessary Columns and Tables	64
Push Down of ORDER BY and FETCH FIRST	67
Summary of Targeted UNION ALL and Outer Join Performance Enhancements.....	70
Runtime Optimizations	71
Runtime Adaptive Index.....	71
Internal Runtime and Complex Expression Optimization	74
Sort, Workfile, and Sparse Index Improvements.....	75
Partial Sort-Avoidance for FETCH FIRST	75
Extensions to Sort Avoidance for OLAP Functions	76
Sort Space Reductions and In-memory Exploitation.....	77
Sparse Index Improvements	78
Predicate Optimizations	79
Sort for Stage 2 Join Expressions	79
User-Defined Table Function Predicate Optimizations.....	80
VARBINARY Indexability Improvements	80
Row Permission Indexability for Correlated Subqueries.....	81

Optimizer Cost Model Enhancements.....	82
Extending NPGTHRSH to Default Statistics	82
List Prefetch and Hybrid Join Cost Improvements.....	83
Improved Filter Factor for CURRENT DATE/TIMESTAMP Predicates.....	84
Improved Resolution of Filter Factors at Bind/Prepare Utilizing Index Probing	84
Improved Bind/Prepare Performance with a Large Number of Indexes.....	85
RUNSTATS and Optimizer-generated Statistics Profiles.....	86
Static and Dynamic Plan Stability.....	88
Static Plan Stability Usability Improvements.....	89
Dynamic Plan Stability.....	89
Summary	90

Build a DB2 12 for z/OS Mobile Application Using IBM MobileFirst

<i>by Jane Man and Clement Leung</i>	91
Abstract	91
Introduction.....	91
Installation and Configuration.....	92
Installation	92
Configure JNDI for DB2 for z/OS.....	94
Introducing Scenario 1: Retrieve a string from a string ID from the SYSIBM.SYSXMLSTRINGS table using regular SQL statements	95
Client-Side Development	95
Server-Side Development	97
Create an SQL Adapter.....	98
Deploy and Test an Adapter	100
Use the Result from the Server Side	102
Create New Environments	102
Build and Deploy	102
Test Using MobileFirst Console	103
Run Android Emulator (Optional)	106
Introducing Scenario 2: Execute RUNSTATS for a particular database and table space by calling a stored procedure.....	106
Summary	107
Acknowledgments.....	108
Resources	108

IBM DB2 Analytics Accelerator: A Revolution in Performance

by Namik Hrle, Ruiping Li, and Wolfgang Hengstler..... 109

Executive Overview 109

Introduction..... 110

The Changing Role of Analytics 110

Challenges with the Traditional Approach to Analytics 111

Business Usage Scenarios for Improved Insight..... 111

 Rapid Acceleration of Business-Critical Queries 113

 Derive Business Insight from z/OS Transaction Systems 114

 Reduce IT Sprawl for Analytics Initiatives 115

DB2 Analytics Accelerator Key Design and Operational Features..... 116

DB2 Analytics Accelerator Query Process and Features 118

 DB2 Analytics Accelerator Content Maintenance..... 119

 DB2 Analytics Accelerator Table Definition and Deployment 119

 High Performance Storage Saver..... 120

 Connectivity Options and Workload Balance..... 120

 Disaster Recovery..... 121

DB2 Integration..... 122

 Query Acceleration Criteria..... 123

 Query Acceleration Control Knobs 123

 EXPLAIN Function 123

 Accelerator Modeling..... 124

 Accelerator-related Instrumentation 125

 Accelerator-only Tables 126

 In-database Transformation in the Accelerator..... 127

 In-database Analytics in the Accelerator 129

 DB2 Analytics Accelerator Administrative Stored Procedures 130

 DB2 Analytics Accelerator DB2 Command..... 131

 Performance Considerations..... 132

PureData System for Analytics, powered by Netezza 132

 The PureData System for Analytics S-Blade, powered by IBM Netezza
 technology 134

 Applying Data Stream Processing to DB2 Queries 134

 Improved Disk Failover..... 135

 Drive Redistribution for S-Blade Failures 135

Comparison of PureData System for Analytics Appliance Models 136

Summary	136
For More Information	137
Acknowledgments.....	137
Making Data Simple and Accessible: The Role of Technology in Delivering Analytic Results	139
Overview	139
A Technology Revolution on z Systems	139
Understanding User Roles and Data	140
Use Cases	141
Avoid Data Movement	141
Integrate Multiple Data Types	142
Run Analytics on Operational Data	143
Summary	143
For More Information	144
Notes	144