

# Preface

---

**I**n the world of information technology today, it is more and more difficult to keep up with the skills required to be successful on the job. This book was developed to minimize the time, money, and effort required to learn DB2 for Linux, UNIX, and Windows. The book visually introduces and discusses the latest version of DB2, DB2 9.5. The goal with the development of DB2 was to make it work the same regardless of the operating system on which you choose to run it. The few differences in the implementation of DB2 on these platforms are explained in this book.

## WHO SHOULD READ THIS BOOK?

This book is intended for anyone who works with databases, such as database administrators (DBAs), application developers, system administrators, and consultants. This book is a great introduction to DB2, whether you have used DB2 before or you are new to DB2. It is also a good study guide for anyone preparing for the IBM® DB2 9 Certification exams 730 (DB2 9 Fundamentals), 731 (DB2 9 DBA for Linux, UNIX and Windows), or 736 (DB2 9 Database Administrator for Linux, UNIX, and Windows Upgrade exam).

This book will save you time and effort because the topics are presented in a clear and concise manner, and we use figures, examples, case studies, and review questions to reinforce the material as it is presented. The book is different from many others on the subject because of the following.

- 1. Visual learning:** The book relies on visual learning as its base. Each chapter starts with a “big picture” to introduce the topics to be discussed in that chapter. Numerous graphics are used throughout the chapters to explain concepts in detail. We feel that figures allow for fast, easy learning and longer retention of the material. If you forget some of the concepts discussed in the book or just need a quick refresher, you will not need to read the entire chapter again. You can simply look at the figures quickly to refresh your memory. For your convenience, some of the most important figures are provided in color on the book’s Web site ([www.ibmpressbooks.com/title/0131580183](http://www.ibmpressbooks.com/title/0131580183)). These figures in color can further improve your learning experience.
- 2. Clear explanations:** We have encountered many situations when reading other books where paragraphs need to be read two, three, or even more times to grasp what they are describing. In this book we have made every effort possible to provide clear explanations so that you can understand the information quickly and easily.
- 3. Examples, examples, examples:** The book provides many examples and case studies that reinforce the topics discussed in each chapter. Some of the examples have been

taken from real life experiences that the authors have had while working with DB2 customers.

- 4. Sample exam questions:** All chapters end with review questions that are similar to the questions on the DB2 Certification exams. These questions are intended to ensure that you understand the concepts discussed in each chapter before proceeding, and as a study guide for the IBM Certification exams. Appendix A contains the answers with explanations.

## GETTING STARTED

If you are new to DB2 and would like to get the most out of this book, we suggest you start reading from the beginning and continue with the chapters in order. If you are new to DB2 but are in a hurry to get a quick understanding of DB2, you can jump to Chapter 2, DB2 at a Glance: The Big Picture. Reading this chapter will introduce you to the main concepts of DB2. You can then go to other chapters to read for further details.

If you would like to follow the examples provided with the book, you need to install DB2. Chapter 3, Installing DB2, gives you the details to handle this task.

## A Word of Advice

In this book we use figures extensively to introduce and examine DB2 concepts. While some of the figures may look complex, don't be overwhelmed by first impressions! The text that accompanies them explains the concepts in detail. If you look back at the figure after reading the description, you will be surprised by how much clearer it is.

This book only discusses DB2 for Linux, UNIX, and Windows, so when we use the term DB2, we are referring to DB2 on those platforms. DB2 for i5/OS<sup>®</sup>, DB2 for z/OS<sup>®</sup>, and DB2 for VM and VSE are mentioned only when presenting methods that you can use to access these databases from an application written on Linux, UNIX, or Windows. When DB2 for i5/OS, DB2 for z/OS, and DB2 for VM and VSE are discussed, we refer to them explicitly.

This book was written prior to the official release of DB2 9.5. The authors used a beta copy of the product to obtain screen shots, and to perform their tests. It is possible that by the time this book is published and the product is officially released, some features and screenshots may have changed slightly.

## CONVENTIONS

Many examples of SQL statements, XPath/XQuery statements, DB2 commands, and operating system commands are included throughout the book. SQL keywords are written in uppercase bold. For example: Use the **SELECT** statement to retrieve data from a DB2 database.

XPath and XQuery statements are case sensitive and will be written in bold. For example:  
**/employee/DEPT/Id**

DB2 commands are shown in lowercase bold. For example: The **list applications** command lists the applications connected to your databases.

You can issue many DB2 commands from the Command Line Processor (CLP) utility, which accepts the commands in both uppercase and lowercase. In the UNIX operating systems, program names are case-sensitive, so be careful to enter the program name using the proper case. For example, on UNIX, **db2** must be entered in lowercase. (See Appendix B, Use of Uppercase versus Lowercase in DB2, for a detailed discussion of this.)

Database object names used in our examples are shown in italics. For example: The *country* table has five columns.

Italics are also used for variable names in the syntax of a command or statement. If the variable name has more than one word, it is joined with an underscore. For example: **CREATE SEQUENCE *sequence\_name***

Where a concept of a function is new in DB2 9, we signify it with the following icon:



When a concept of a function is new in DB2 9.5 we signify it with the following icon:



Note that the DB2 certification exams only include material of DB2 version 9, not version 9.5

## CONTACTING THE AUTHORS

We are interested in any feedback that you have about this book. Please contact us with your opinions and inquiries at [udb2book@ca.ibm.com](mailto:udb2book@ca.ibm.com).

Depending on the volume of inquires, we may be unable to respond to every technical question but we'll do our best. The DB2 forum at [www-128.ibm.com/developerworks/forums/dw\\_forum.jsp?forum=291&cat=81](http://www-128.ibm.com/developerworks/forums/dw_forum.jsp?forum=291&cat=81) is another great way to get assistance from IBM employees and the DB2 user community.

## WHAT'S NEW

It has been a while since we've updated this book, and there have been two versions of DB2 for Linux, UNIX, and Windows in that time. This section summarizes the changes made to the book from its first edition, and highlights what's new with DB2 9 and DB2 9.5.

The core of DB2 and its functionality remains mostly the same as in previous versions; therefore, some chapters required minimal updates. On the other hand, some other chapters such as Chapter 15, The DB2 Process Model, and Chapter 16, The DB2 Memory Model, required substantial changes. Chapter 10, Mastering the DB2 pureXML Support, is a new chapter in this book and describes DB2's pureXML technology, one of the main features introduced in DB2 9.

We have done our best not to increase the size of the book, however this has been challenging as we have added more chapters and materials to cover new DB2 features. In some cases we have moved some material to our Web site ([www.ibmpressbooks.com/title/0131580183](http://www.ibmpressbooks.com/title/0131580183)).

There have been a lot of changes and additions to DB2 and to this edition of the book. To indicate where something has changed from the previous version of the book, or was added in DB2 9 and/or DB2 9.5, we have used the icons shown below.

**V9**

**V9.5**

The following sections will briefly introduce the changes or additions in each of these new versions of DB2.

## **DB2 9**

DB2 9 introduced a few very important features for both database administrators as well as application developers. Developers have been most interested in pureXML, DB2's introduction of a native data type and support for XML documents. This allows you to store, index, and retrieve XML documents without the need to shred them, or store them in a CLOB and do unnatural things to access specific elements in the document. In addition, you can use SQL or XQuery to access the relational and XML data in your tables. The new developer workbench is a development environment that is integrated into the eclipse framework. It helps you to build, debug, and test your applications and objects such as stored procedures and UDFs.

From the DBA's point of view, there are a number of enhancements that came in DB2 9 that made life easier. Automatic storage allows you to create paths that your databases will use, and then create table spaces in the database that will automatically grow as data is added in these table spaces. Automatic self-tuning memory provides you with a single memory-tuning knob. You can simply tell DB2 to "manage yourself" and DB2 will adjust the memory allocations based on the system configuration and the incoming workload, and it will adapt to changes in the workload automatically.

Table or range partitioning allows you to logically break up your tables for easier maintenance and faster roll in and roll out of data. This, combined with MDC and DB2's database partitioning, provides the ultimate in I/O reduction for large, data-intensive queries. Data compression automatically replaces repeated strings, which can even span multiple columns with a very small symbol, resulting in even more I/O reduction. When you add data compression on top of the data partitioning, you get much smaller databases, with much faster performance than ever before possible. In DB2 9, the compression dictionary, where the symbols and the values they replace are stored, must be built offline using the REORG tool. This dictionary is then used for subsequent inserts and updates, but it is static, meaning new repeated values are not added automatically.

Another important feature of DB2 9 that we discuss in detail in this book is Label Based Access Control, or LBAC. LBAC allows you to control access to the rows and columns of data in your tables so that you can comply with both business and regulatory requirements.

## DB2 9.5

If you were excited by DB2 9, you are going to be even more excited about DB2 9.5. One thing those of you running on Linux or UNIX will notice immediately is that DB2 9.5 uses a threaded engine, not a process-based engine as in previous versions. This will help to increase performance and make administration even easier. Because there is only one main engine process now, the memory configuration task has been simplified. The default **AUTOMATIC** setting on most memory-related configuration parameters makes memory management much easier, and much more efficient. You will not have to worry about those pesky db2 agent processes any more. DB2 9.5 is also introducing a new locking mechanism so that readers do not block writers.

While DB2 9 introduced automated statistics gathering, DB2 9.5 takes that a major step further. If there has been a significant amount of changes to the table since the last **RUNSTATS** was run, DB2 will automatically run a **RUNSTATS** before it optimizes a new access plan that hits that table. Along the same vein, if you have a table which has been enabled for compression, DB2 will now create a dictionary, or update the existing dictionary as data is added to, or updated in the table, and if you are **LOADING** data into the table. You no longer have to run a **REORG** or **INSPECT** to get a dictionary, and it is no longer a static dictionary.

If you tried HADR with DB2 V8.2, or DB2 9, it is now even easier, and that is saying something. You no longer have to worry about setting up the scripts and the clustering software to automatically detect server failures and manage server failovers. This setup is now done for you automatically when you set up HADR.

Roles are another important feature added in DB2 9.5. Roles allow you to mimic the job functions within your organization, and grant privileges or rights to roles instead of to specific users or groups. This makes the management of who can see what data much easier and more intuitive since it provides a single point of security and privilege management.

Both of these new versions of DB2 include features and functions that will make your life easier. We hope that we have described them in a manner that is easy to understand, and that shows you how and when you might want to take advantage of the new features of the different versions of DB2.