

DB2 for i Frequently Asked Questions

1. What is DB2 for i ?

DB2 for i is the relational database management system (RDBMS) that is built into the IBM i Operating System that runs on IBM's Power Systems Servers. It is an extensible, high performance, scalable database that adheres to many industry standards while leveraging the IBM i architecture to maintain its value proposition of lower total costs to own.

2. Do I need to buy DB2 for i?

DB2 for i is built into IBM i at no additional charge. This includes the core database management system, the SQL/Query engine, storage management, security, and more. There are several utilities that you can add onto the core database functionality, including utilities for data replication, parallel processing, query tools and more. A list of DB2 for i utilities can be found at http://www.ibm.com/servers/eserver/series/db2/prod_admin.html.

3. What level of IBM i supports DB2 for i?

DB2 for i is not relevant to a specific OS (Operating System) level, but rather, built into the OS at all levels. There is only one RDBMS within IBM i. The current level of DB2 for i is Version 6 Release 1 (6.1). The functionality of DB2 for i will differ with each version/release of IBM i, as new enhancements are introduced with each new version.

4. How does DB2 for i compare to DB2 for Linux/Unix/Windows (LUW)?

All the DB2 family members (DB2 runs on over twenty platforms) share common attributes (like SQL Standard syntax support), but also differ in their optimization for the specific platform and operating environments, and differ in packaging and in some administrative areas. At any point in time, feature/function differences may exist as well, depending on delivery schedules and market requirements.

SQL Syntax is very common across the DB2 products. To see more details about the common SQL supported by DB2, refer to <http://www.ibm.com/servers/enable/site/db2/db2common.html>.

Version/Release naming will differ because, as mentioned above, DB2 for i follows the IBM i version/release numbering scheme, and not the DB2 LUW version/release scheme (i.e., you cannot assume a comparison of version/release is a reflection on functional equivalency). For instance, in the above document detailing common SQL Syntax, V5R1 of DB2 for i is compared favorably to DB2 LUW V7, and V5R3 is mapped to DB2 LUW V8.

Because DB2 is included as part of IBM i, packaging is obviously different, whereas DB2 for Linux/Unix/Windows is a common code base that can be installed on many different operating systems and platforms.

Graphical tools to administer the database are provided as part of IBM i Navigator component of IBM i, whereas DB2's Control Center would be used to administer DB2 UDB for LUW, however, other "centers", such as Replication Center, Development Center, and Warehouse Center can be used across all DB2 products.

5. What is "Viper"?

Viper is a code name for Version 9 of DB2 for Linux/Unix/Windows (LUW). DB2 V9 has a planned availability date of July 28, 2006. DB2 9 can run on Power System servers in any of these environments: Linux or AIX in a logical partition, or the Integrated x-Series adapter/server or via the Blade Center support. However, administration and management of DB2 9 running in Linux, Unix, or Windows operating systems would be different since DB2 9 is a different code base than DB2 for i .

Version 9 of DB2 for LUW includes specific enhancements in the area of XML data storage and access, improved development environments, data compression, security, and enhanced installation processes.

6. How does DB2 for i equate to DB2 LUW, specific to the DB2 9 enhancements?

DB2 9 offers a capability to store XML documents in DB2, and access them via XML-like interfaces (an emerging standard called xQuery) that are similar to SQL. DB2 for i is taking a phased approach to XML integration. DB2 XML Extender, part of the DB2 Extenders product for IBM i, has been available for many releases and provides interfaces to store, parse, shred and generated XML documents to/from data stored in DB2. Recent performance enhancements to the DB2 XML Extenders support have been delivered in V5R3 and V5R4 and make the XML Extenders a viable solution for many customers and ISVs. In addition, XML integration into existing RPG and COBOL applications has been simplified with new built in support for XML into these languages (RPG in V5R4, COBOL in V5R3).

Native XML storage and xQuery syntax support are future directions being considered for DB2 for i. For more information on XML integration with DB2 for i, refer to the following redbook:
<http://www.redbooks.ibm.com/abstracts/sg247258.html?Open>

Data Compression was available in DB2 for i for several years but had little uptake in the marketplace and is not planned for DB2 for i in the future at this point. The reason it had little uptake was that data compression trades off I/O utilization for CPU utilization. For example, data needs to be uncompressed on reads, and compressed on writes. This is great for I/O bottlenecked systems with spare CPU, because you are streamlining I/O but adding work to the CPUs. With IBM i's single level storage and parallel I/O processor architecture, I/O bottlenecks are not that prevalent with a balanced configuration.

7. Is DB2 for i “optimized” for SAP?

Yes. For many years, there has been a joint development team between DB2 for i and SAP. In addition to this joint team, the SAP and DB2 development teams work very closely together through a single IBM focal point, ensuring requirements from SAP are translated into compatible features across our DB2 portfolio. This ensures that requirements that are implemented in DB2 for Linux/Unix/Windows will be implemented in a consistent manner in DB2 for i and vice versa.

Some of the specific features added to DB2 for i to optimize for SAP include:

- Enhanced SQL Performance Analysis
- SAP specific SQL optimization techniques
- Increased Table Limits
- High Availability enhancements to improve SAP upgrade processes
- Improved Memory Management
- Improved Data Warehousing features (allowing DB2 for i to maintain leadership in the SAP BW benchmark during 2005 and 2006)

More DB2 features that are designed to meet SAP future needs are already being developed for the next release and still more items for are planned for the following release.

8. Do I need DB2 Connect to connect my Linux, Unix, or Windows clients to DB2 for i?

The answer is that DB2 Connect can be used, but it is not necessarily a requirement. DB2 Connect can be used to connect applications from those environments to any DB2 server via ODBC, .NET, JDBC, or using CLI/DRDA interfaces. So if the only connectivity required is from Windows or Linux, to DB2 for i (no other DB2 family member), you might consider the IBM i Access for Windows middleware (ODBC, OLE DB, .NET) or Java Toolbox JDBC drivers available at no charge with the IBM i operating system. However, DB2 Connect also includes additional programmer productivity features through .NET (Visual Studio) or PHP (Zend Studio) plug-ins for DB2. DB2 Connect can also be used for DB2 data “federation”,

such as the ability to join data from two DB2 (or Informix) databases with a single SQL statement. DB2 Connect's mobility on demand feature includes a synchronization server to replicate (and synchronize) data from DB2 Servers with DB2 products running in a mobile device.

9. Can you run DB2 InfoSphere Warehouse on an Power System Server in an IBM i partition?

Infosphere Warehouse is a packaging of DB2 for Linux/Unix/Windows (LUW) with some additional data warehouse oriented tools. The elements of Infosphere Warehouse do not support IBM i, but can certainly run in a Linux or AIX partition on a Power System server.

10. Other database systems tout the ability to do tasks in parallel or to partition data for better performance. Can DB2 for i do that?

Yes. DB2 for i introduced parallel database functionality, also known as SMP (Symmetric Multi-Processing) capability in V3R1 (1995), and has enhanced those functions significantly since then. The DB2 for i SMP feature is an Operating System add-on that can significantly improve performance of database tasks by breaking the processing across multiple CPUs. By doing this, SMP optimizes the available resources to complete the database task in a cost-effective manner.

In addition, DB2 for i, with the DB2 Multisystem feature, can also partition data across multiple servers/partitions, or partition data locally (within the same server/partition). One of the key benefits of this is to expand the size of a single database table/file, without application changes. However, it should be noted that other databases use the notion of "data partitioning" to address performance issues. Because these databases do not employ **single-level store** (single-level store is a fundamental design point of IBM i), they are restricted on how much space they can address. Furthermore, they cannot effectively use all the processors on a single box for a single unit of work. Database partitions enable them to overcome both problems.

Because of Single-Level Store, and the ability for near-linear scalability in an SMP environment, "data partitioning" as it is called in other database management systems is not required, and this also reduces the amount of technical resource required to manage the partitioned environment. For IBM i customers interested in data partitioning, a technical review is highly recommended.

11. DB2 for i DBAs are not readily available in the inventory of available skills in the marketplace. I can find lots of DBAs for other database products. Wouldn't an investment in one of these other databases be a better solution?

Availability of DBA skills on the marketplace is a subject that requires some further investigation. For instance, keep in mind that (a) many typical DBA tasks required in other database products are not required on DB2 for i, (b) DBAs have many skills that are transferable to DB2 for i, and (c) the DBA "community" has evolved because of the need for this specific skill on other database management systems.

As noted, many of the skills of a DBA are common across all relational databases. This includes knowledge of SQL, relational data modeling, indexing, and query optimization. With the exception of query optimization, learning these skills on DB2 for i is a matter of learning new user interfaces. Online tutorials exist for this via the DB2 for i home page: ibm.com/systems/i/db2.

Query optimization on DB2 for i will require some education for the DBA. IBM has developed a four day class called "DB2 for i SQL Performance Tuning and Monitoring Workshop." For details, refer to: ibm.com/series/service/igs/db2performance.html.

Lastly, it is important to note that many of the tasks to manage other databases are not required with DB2 for i. These tasks may include creation and monitoring of tablespaces, logs, and cache area disk usage, re-balancing of indexes, running utilities to gather statistics, and partitioning of data. The IBM i architecture and sophisticated cost-based optimizer of DB2 for i automates these processes.

12. Does DB2 for i require a DBA?

Maybe! Again, a “traditional” DBA will spend more of their time with DB2 for i doing data modeling, SQL coding, developing indexing strategies, and SQL optimization. If the application is SQL based, this is a critical element to ensuring appropriate design and SQL optimization. A person with an understanding of SQL optimization is highly recommended to fully leverage the benefits of DB2 for i.

But as traditionally defined, typical tasks required by other database systems will NOT be required due to the advanced architecture of IBM i. Many of the typical DBA tasks are handled automatically by the OS (data striping, collection of statistics for use by the SQL Optimizer, managing log or buffer space).

For more details, refer to: <ftp://ftp.software.ibm.com/systems/i/db2/pdf/acfaq.pdf>

13. What are DDS and RPG? Am I really using DB2 or a flat file system?

DB2 for i has many programming and operational interfaces available for application developers, DBAs, and operations/support personnel. However, the database engine is still the same regardless of the interface being used. With this architecture, you actually have MORE flexibility in how you use the database and what skills are available in the marketplace to help develop or support.

DDS (Data Description Specifications) provide one option to define database tables and indexes. RPG is a common programming language in the IBM i world. These interfaces, commonly referred to as “native” interfaces, were in place before SQL became an industry standard and continue to be used heavily in an IBM i environment.

DB2 for i also supports SQL interfaces, and you can create tables and indexes through SQL, generally considered the “open” standard for database development. DB2 for i also supports many different high level languages (e.g., Cobol, JAVA), 4GL and CASE tools, and web or client/server based interfaces (JDBC, ODBC, Microsoft’s .NET framework) to allow a tremendous amount of flexibility while still leveraging the underlying advantages of the single RDBMS!

14. Where can I find more education on DB2 for i?

There is a wealth of information available via the web, at the DB2 for i home page:
<http://www.ibm.com/systems/i/db2>.

In addition, educational roadmaps and course descriptions can be found at:
http://www.ibm.com/systems/i/db2/db2educ_m.html.

For information on porting applications from other databases to DB2 for i , check out:
<http://www.ibm.com/servers/enable/site/db2/porting.html>