DB2 for i: Autonomic Computing
Frequently Asked Questions

1. What is DB2 for i?

DB2 for i is the relational database management system (RDBMS) that is built into the Operating System (IBM i) running on Power Systems™ servers. It is an open, 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. See http:www.ibm.com/systems/i/db2.

2. What is Autonomic Computing

Autonomic Computing is a concept that applies to a self managing organism that has these characteristics:

- Self Configuring
- Self Healing
- Self Tuning
- Self Protecting

3. How does DB2 for i address Autonomic Computing?

DB2 for i is uniquely positioned to leverage the self managing concepts inherent in Autonomic Computing due to the fundamental architecture of IBM i. The following describes four key areas of autonomic computing:

- Self Configuring

  Unlike many Relational Database Management Systems, DB2 for i requires limited configuration work prior to use. Because it is pre-loaded with IBM i, and doesn't require any specific database installation procedures, the database itself is ready to go! Because of the unique Single Level Store concepts of the system, database management tasks such as creation and monitoring of table spaces or complex partitioning of data across disk subsystems are not required. DB2 for i provides automatic data spreading and automatic storage allocation.

- Self Healing

  DB2 for i contains a number of self-healing qualities. Systems Managed Access Path Protection (SMAPP) automatically journals (logs) high risk indexes (including Encoded Vector Indexes) to minimize the recovery time in the case of a system failure. Self-managed journal receivers allow the system to automatically monitor the capacity of a journal receiver (database transaction log), creating a new receiver and auto switching transaction logging to it should it become close to maximum capacity. Automatic journaling of Sequence Objects and auto restart of journal processing after a copy or restore further increase the autonomic capabilities of DB2!

  Index rebalancing, a common DBA task on other systems (because indexes can become unbalanced over time), is handled automatically by DB2 for i. DB2 will also automatically rebuild catalog views when necessary, eliminating the need for manual intervention.

- Self Tuning

  While IBM i comes with an "auto-tuner" that adjusts resources based on predictive analysis, DB2 for i also contains several self tuning attributes. One key component of DB2 for i is the cost-based optimizer. This component processes all SQL/Query database activity and automatically recognizes resource changes on the system and without user intervention rebuilds access plans (if desirable) to optimize the available resources.

Statistics are used by the optimizer to assist in determining the least cost method of accessing the data. Statistics are collected and maintained automatically by DB2 for i during normal database processing.

The DB2 On Demand Performance Center (included with IBM i) allows you to easily analyze and improve your indexing strategy by reviewing the indexes advised system-wide by the optimizer. This new support allows you to identify indexes that are not being used, as well as managing index creation. Index advice is provided for both traditional indexes (Keyed Logical Files or SQL Indexes), as well as Encoded Vector Indexes (EVIs).

The 'always-available' diagnostic information in the SQL plan cache allows you to analyze query performance and compare performance behavior from one point in time to another; all without the overhead of a Database Monitor. DB2 for i can also automate the creation of indexes when required, and keep them around for use by other queries or other jobs.

- Self Protecting

DB2 for i is very well positioned in this area because of the object based structure of IBM i, which DB2 leverages. In addition to the high degree of security that is made possible by the object based kernel, other self protecting attributes that are built into the operating system and available to use include:

- Digital signing of objects to help prevent unauthorized access
- 128 bit data encryption and SSL
- US Government C2 security compliance including object auditing

4. 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.

5. 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 ibm.com/developerworks/db2/library/techarticle/db2common/.

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.

6. Are Health Center and Configuration Advisor (DB2 V8.x components) available for DB2 for i?

No, these functions are either not required for IBM i or provided through other tools. This is a good example of administrative facility differences pointed out in question #5. These utilities are specific to DB2 V8.x. Health Center is unique to the Linux/Unix/Windows operating environment and provides functionality to address issues that generally do not exist with DB2 for i. Tools built into IBM i allow for collection of performance data similar to what is found in Configuration Advisor.

7. 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