

---

# DB2 and Chargeback

---

*TOPICS DISCUSSED:*

- What is DB2?
- Advantages of DB2
- Overview of DB2 Subsystem
- How SMF Data Interacts with DB2
- DB2 Chargeback Problems
- Accounting for DB2 Overhead
- DB2 User Identification
- Conclusion

## What is DB2?

DB2 is IBM's relational database. A relational database is often viewed as a lot of files that are tied together and can talk to each other. This is true to a degree. It is a database perceived by the users as a collection of tables and nothing but tables. All of the data is viewed as rows and columns, otherwise referred to as "records" and "fields". In DB2 terminology they are called "tuples" and "attributes". A DB2 relation or table may consist of rows of data showing company name, state, and status (i.e., active, inactive). In each column or field there is only one piece of information, i.e., state. You do not have more than one state listed for a company in the same row. This table of information is called a "relation" which is where the term relational database comes from.

## Advantages of DB2

DB2 is becoming more popular because it is a sound theoretical base. In 1970 Dr. E. F. Codd wrote the paper "*A Relational Model Of Data For Large, Shared Databases*", which gave us the first set of rules and regulations for databases. With the exception of one set of minor updates, this set of rules has held up since 1970. These rules make the system predictable. Data will always react the same way. Because of this, user interfaces are easy to document, learn, teach, use and remember. Data is represented in only one way—column values within rows of tables.

DB2 has a dual-mode principle, which means SQL (structure query language) is used for both programming and interactive access to the database. Instead of the database administrator, programmer, and end-user using different languages, everyone is speaking the same language (SQL). Being a relational database, DB2 has considerable flexibility in coding. For example, in a customer database containing customer first name, last name and telephone number, one is able to add fields for street address, state and zip code while the database is running. It is not necessary to bring down the database in order to make these additions.

DB2 has set-level operators which provide automatic navigation. This means the end-user merely has to specify what is wanted and not how to get there. The database does all the navigating. This makes DB2 a lot more user friendly. It is easy to use, however, there is a cost. DB2 incurs high resource utilization. There is a lot of overhead related to all the features provided by DB2. DB2 is a CPU and memory "hog".

## Overview of DB2 Subsystem

An address space is a "neighborhood", for example, where my batch job, started task, or TSO session lives. DB2 is a little different in that it uses three address spaces:

- System services
- IMS resource lock manager
- Database services

System services control the connections to other MVS subsystems. DB2 is always accessed from a fourth address space, i.e., IMS, CICS, or a batch job. DB2 uses an attachment facility. For example, it determines that TSO wants to talk to DB2. It

establishes a path of communication, otherwise known as a “thread”. When that thread disconnects and that path disappears, a DB2 transaction record is cut for logging purposes. System services manage the system log. Everything going on within DB2 is recorded in the activity logs and is archived. System services control system startup, shutdown, and operator communications. It allows the operator to talk to the system and to shutdown all three address spaces. It also gathers all system-wide statistics for SMF (system management facility) or GTF (general trace facility) records in order to capture the resources consumed within a transaction.

The IMS resource lock manager controls concurrent access to data. For example, if I am updating a record, someone else cannot come in and update the same record.

Database services implement the functions of SQL. It has the following components:

	<b>Function</b>
Precompiler	Preprocessor for the host programming language (i.e., COBOL, PL/1) and constructs a Database Request Module (DBRM).
Bind	Produces an application plan (directions how to best access the data) for the DBRMs and acts as a compiler for the SQL statements.
Runtime supervisor	Interfaces to the application program and uses information from the application plan to make requests of the data manager.
Data manager	A sophisticated access method that manages the physical databases on disk or in memory.
Buffer manager	Performs the actual I/O against the data and reads the required “page” into memory from external storage if the data is not in memory.

It takes .030 seconds to access data from external storage and only .000075 seconds to access data from memory, so for DB2 to provide very efficient processing, it tries to move as much data as possible into memory. At this point, if the system is not efficiently tuned, DB2 could turn into a “resource hog”.

Within the DB2 subsystem, all three address spaces run as started tasks. DB2 uses cross-memory service facility (XMS) to allow these address spaces to interact.

## How SMF Data Interacts with DB2

SMF data consists of records created by the operating system that provides an audit trail of events and/or utilization. Within SMF there are 256 possible record types (0-255). Record types 0-127 are reserved by IBM to create a record of anything that happens within the operating system. Record types 128-255 are available to the user to use user exits to write their own SMF records. We are going to discuss the following record types:

- Type 30 Measures batch jobs, TSO sessions, started tasks
- Type 110 Measures CICS resources at the transaction level
- Type 101 Measures DB2 resources at the transaction level

A CICS region may involve a variety of different transactions. For chargeback purposes, one can record the total resource utilization of the entire CICS region using type 30 records and allocate the resource cost in proportion to the number of transaction. Implicit

in this approach is that all transactions use the same amount of resources and therefore, cost the same as all other transactions. The more equitable approach would be to track each individual transaction with the 110 records in order to record the actual resource utilization and cost of each type of transaction.

A CPU hour for type 30 records is not the same as a CPU hour for type 110 records. The problem is that the total CPU time for all of the 110 records within the CICS region is less than the CPU time reported by the type 30 records. There are certain system overheads that are not carried in the 110 records, which can account for 20% or more of the region's total utilization. CPU cycle time for type 30 records contain TCB (task control block) time, SRB (system request block) time, initiator TCB time, and initiator SRB time. If you have ESA, it also contains hiperspace CPU time, I/O interrupt time and region control task time. There can be seven fields of CPU time and various overheads for type 30 records. The CPU time for 110 records contains only two fields of CPU time—TCB time and SRB time. For chargeback purposes, a different rate (unit cost) can be developed for the two types of records.

Type 101 records measure DB2 resources at the transaction level. There are six types of SMF 101 records: 100 Statistical, 101 Accounting, 102 Audit, 102 Performance, 102 Monitor, and 102 Global. For chargeback purposes, you want to use the 101 Accounting records. Within 101 Accounting records, there are 32 classes of records. The fewer classes of data captured the lower the overhead cost of capturing the data. Class 1 records are standard accounting data and class 2 provides elapsed and CPU time in DB2. These are among the recommended classes to be used for chargeback. Class 1 and class 2 account for a 5% CPU overhead according to the IBM DB2 Performance Monitor Usage Guide. The QWACAJST field in DB2 measures the TCB time spent in DB2. This measure of CPU time is repeatable and should be used for chargeback. The QWACASRB field measures SRB time spent in DB2 and is unrepeatable, so it is not recommended for chargeback purposes.

## **DB2 Chargeback Problems**

DB2 does have a problem with double counting CPU time if one is not careful. For example, suppose you have a TSO session or batch job which incur DB2 transactions. Type 30 records are captured for the TSO session or batch job. Type 101 records are captured for the DB2 transactions. The problem is that almost all (92-96%) of the DB2 CPU time found in the type 101 records are also found in the type 30 records. If you use both types of records, you will double charge DB2 CPU time. The same double counting problems exists for IMS if you are using type 7 IMS records.

On the other hand, if you have a CICS transaction that submits DB2 transactions, type 110 records are captured for the CICS transactions and type 101 records are captured for the DB2 transactions. In this case, there is no duplication of CPU time within the 101 and 110 records. Capturing both types of records will not result in double charging for DB2 CPU time. If most of your DB2 transactions are coming from CICS, and only type 110 records are captured, you are not currently charging anything for DB2. You also need to capture type 101 records for DB2 transactions. Another solution to avoid the double counting is to capture only type 30 records for CICS, which contain both CICS and DB2 utilization data. A third alternative to avoid the double counting is to drop the data in the connection ID (QWHCCN) field in DB2 records.

## Accounting for DB2 Overhead

How about recovering the cost of DB2 overhead that is not directly attributed to a specific DB2 user, for example, memory, DASD space, and staff? The solution depends upon the degree of unbundling you choose to use for chargeback purposes. The cost of memory could be bundled into the CPU rate. To chargeback separately for memory, the number of pages captured in the QBACGET field in DB2 records could be used. DASD used for DB2 can be charged separately based on the actual usage of individual users if you are using a DASD accounting package. DB2 staff-related expenses, i.e., system support, can be charged separately if you are using a manpower accounting system or the cost can be bundled in the CPU rate.

## DB2 User Identification

User account numbers may be derived from such fields as Authorization ID, Connection ID, Correlation ID, Plan Name, and Operator ID. The connection ID identifies where the transaction came from. Based on the connection ID, there are an authorization ID and correlation ID which provide further clarification of the specific user. For example, if the connection ID identified CICS as the origin of the transaction, the authorization ID might provide the user ID, terminal ID, or transaction ID from CICS.

## Conclusion

Chargeback for DB2 is no different than any other subsystem. Yes, there are CPU differences between DB2 and MVS, but there are differences in CPU time recorded for CICS, IMS, and MVS. There are significant overhead costs associated with DB2 that must be tracked for equitable enterprise wide accounting.