

# BCV5™


## REDUCTION AND MASKING

The BCV5 Data Reduction and Masking Facility (BCV5 RM) is an add-on component to BCV5 that expands the BCV5 capabilities beyond making one-to-one clones of DB2 databases and tablespaces. BCV5 RM allows you to specify regular SQL queries for one or more tables that take part in a BCV5 copy process. Instead of copying the complete source data into the target environment, the target tables will be populated with the result sets from these queries.

When you want to provide production data for your test or development environment, it is not always necessary (or legal) to copy all of the contents of the tables involved. Take the case of where developers want to verify the correctness of a program change - they want to make a component test with only a subset of data. They don't want to test the fix with all the production data, which would significantly increase the runtime and project time of their automated unit-tests. In this case it is far more efficient to extract only a subset of the production data.

### Reduction

With RM you can specify expressions of arbitrary complexity to select your desired subset of data. Specify a regular SQL SELECT statement for those tables that you want to thin out. You have access to the complete feature set of DB2's SQL. It is possible to specify a simple WHERE clause, or to do a complex SUBSELECT that gets records from a related table to build the desired result set. The SELECT statements decide the data for the target table.



### Masking

Many shops fight with the problem that the data privacy laws do not allow them to just copy production data to a test or development environment. The data is too sensitive and it's highly illegal to share this data with unauthorized users. There are rules that define how the production data must be masked when used in another context outside of the production environment. For this purpose, the SELECT statements that you can specify in BCV5 RM also allow for masking of data. You can modify the contents of tables in a flexible and efficient way. Characters can be exchanged with a substitution character, contents can be replaced by random content, or a user defined exit can be called to modify the data using a more sophisticated algorithm.


### Ease of use

The BCV5 ISPF user interface guides you through the process of specifying your reduction and masking rules. When you specify the table name, BCV5 automatically obtains the structure of the table and generates a default SELECT statement. This statement can be modified by you in a way that provides the desired data. The validation check automatically checks if the SQL statement is a valid and well-formed statement.

Additionally BCV5 offers a powerful batch facility to specify masking and reduction rules for a set of tables. Consider a business case that requires the specification of a masking rule for all tables that contain the CLIENT column. In such a case it is not necessary to specify a separate rule for each matching table.



# BCV5™ – REDUCTION AND MASKING



## Benefits

- Compliance with data security regulations
- Systematic selection of test data to shorten runtime of component tests
- Saves CPU time with the fast integrated copy facility
- Less space usage of target tables

## Features

- Powerful and flexible methods to mask sensitive data
- Generic methods to specify rules for a set of tables
- Automatic inclusion of dependent tables by examining DB2 referential integrity
- Fully automated process that covers all required steps

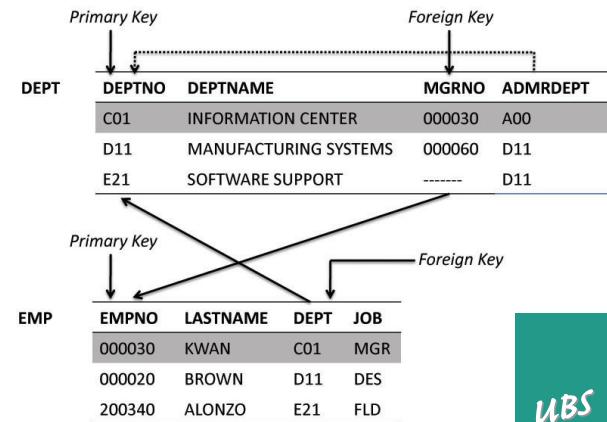
BCV5 handles this by specifying the selection criteria, which for example includes all tables that have a column called CLIENT. In addition you specify a rule that is executed for each matching table. This could be a rule that modifies and/or reduces the content of the column.

After specifying the statements, BCV5 generates the complete job chain to extract the data. This includes all steps that are necessary to transfer data from one DB2 subsystem to another. The complete process is automated - It generates the DDL for all objects involved and then executes it in the target environment. If the objects already exist in the target, BCV5 will perform a compatibility check. The process also includes the transfer of the data from the source to the target environment in the fastest possible way. Tables that are not processed by a masking or reduction rule are copied with a saving of 90% CPU and elapsed time – with masking and reduction other factors come into play.


## Handling Referential Integrity

The specification of a reduction or masking rule for one table may also affect other tables. In many data models tables are related to other tables using referential integrity constraints. Rows from one table must have a corresponding row in a parent table. If the parent table lacks rows that ought to exist, the referential integrity is violated.

In the following figure you see how rows of the DEPT and EMP table relate to each other. The primary and foreign keys of the table indicate these relationships. If you specify a reduction or masking rule for the DEPT table then you may also want to specify a corresponding rule for the EMP table too. In today's data models this may affect multiple tables that are part of this network of relationships.



## BCV5™ – REDUCTION AND MASKING



» *BCV5 RM integrates the provision of test-case data and the provision of preproduction-data-refresh in one product. And it saves significantly on those recurring costs.* «

BCV5 automatically handles this RI job for you. It examines the DB2 catalog for relationships to other tables that are related to the tables for which rules exist, and generates default rules for these tables as well. It generates a rule that extracts only the subset of the dependent child or parent table that does not violate the referential integrity constraints. For instance if you specify a reduction rule for the DEPT table that extracts only the records for department C01, BCV5 will automatically generate a rule that extracts all employees of the EMP table that are related to department C01. In our example this would be the employee KWAN.

This intelligent mechanism within BCV5 makes it very easy to process tables that are part of a referential integrity relationship.

**BCV5** stands for fast, efficient and convenient copying of DB2 data. Compared with UNLOAD/LOAD and most other methods, BCV5 works up to ten times faster. BCV5 requires only 10% of the resources (CPU time) required by conventional copying methods.



### The DB2 High-End Product Line:

#### BCV4™

Full DB2 Subsystem Clones in minutes versus days

#### BCV5™

Save 90% CPU & Run Time with each DB2 copy

#### BCV6™

Log enhanced copies for 24x7 DB2 refresh/migration

#### BPA4DB2™

Premier advisor for DB2 buffer pool optimization

#### ULT4DB2™

Easily identify & restore unwanted changes of DB2 data

#### TUC4DB2™

RTS & Policy driven automation of DB2 data maintenance

#### XM4DB2™

Pro-active surveillance for a greater DB2 availability

---

### Contact Us For More Information

We offer a free 30-day trial evaluation as well as private web demo. Learn more about BCV5 and our complete line of DB2 z/OS products at: [www.ubs-hainer.com](http://www.ubs-hainer.com)

For more information please send an email to:

[info@ESAIGroup.com](mailto:info@ESAIGroup.com)

PRESENTED BY

