

May 7-11, 2006

Tampa Convention Center

Tampa, Florida, USA

IDUG* 2006

C02

CCSID 102 – What's a CCSID and why do I care?

North America

Christopher J. Crone
IBM/DB2 for z/OS Development

Monday, May 8th, 2006 • 1:00 p.m. – 2:10 p.m.

Platform: z/OS

GoFurther





Presentation Topics

- How does DB2 know the CCSID of My data
- When does DB2 convert
- Character Based Functions
- Queries with more than one Encoding
- Putting it all together – Examples
 - ▶ DDL
 - ▶ SQL
 - ▶ Utilities



Terminology

- For the purposes of this presentation
 - ▶ ASCII - all ASCII CCSIDs that DB2 currently supports
 - ▶ EBCDIC - all EBCDIC CCSIDs that DB2 currently supports
 - ▶ UNICODE - UTF-8 or UTF-16
 - ▶ Encoding Scheme
 - ASCII, EBCDIC or Unicode
 - ▶ CCSID – Coded Character Set Identifier (**used by DB2 to tag string data**)
 - Two-byte, unsigned binary integer identifying a specific set of encoding scheme and one or more pairs of character sets (CS) and code pages (CP)
 - ▶ CCSID set
 - The single byte CCSID value (SBCS), mixed CCSID value and double byte CCSID value (DBCS) associated with a particular encoding scheme
 - ▶ Multiple CCSID Sets
 - When two or more CCSID sets contain different CCSID values for one or more values in a set (SBCS, Mixed or DBCS).



IBM Software Group

How does DB2 know the CCSID of my data?

DB2 Information Management Software



@ business on demand software

Installation (continued)

- Information from DSNTIPF ends up in Job DSNTIJUZ

Mixed System

```

DSNHDECM  ASCCSID=1088,
           AMCCSID=949,
           AGCCSID=951,
           SCCSID=833,
           MCCSID=933,
           GCCSID=834,
           USCCSID=367,
           UMCCSID=1208,
           UGCCSID=1200,
           ENSCH=EBCDIC,
           APPENSCH=EBCDIC,
           MIXED=YES
END

```

Non-Mixed System

```

DSNHDECM  ASCCSID=819,
           AMCCSID=65534,
           AGCCSID=65534,
           SCCSID=37,
           MCCSID=65534,
           GCCSID=65534,
           USCCSID=367,
           UMCCSID=1208,
           UGCCSID=1200,
           ENSCH=EBCDIC,
           APPENSCH=EBCDIC,
           MIXED=NO
END

```

The information from panel DSNTIPF flows to Job DSNTIJUZ. In the case on the Left, we have a Mixed = Yes system that is set up to support Korea. The ASCII and EBCDIC system CCSIDs that actually would have been specified on panel DSNTIPF, to result in this specification, would have been 949 and 833.

For mixed systems, and for the Unicode CCSID, the Mixed CCSID is specified on install panel DSNTIPF and DB2 will pick the corresponding Single byte and Graphic (Double Byte) CCSIDs.

In the case on the Right, we have a Mixed = No system that is set up to support US English.

Note that the user specified 819 and 37 for the ASCII and EBCDIC Single byte CCSIDs and that DB2 used the value 65534 for the ASCII and EBCDIC Mixed and Graphic (Double byte) CCSIDs. 65534 is a reserved value that means no CCSID.

Also note that the Default Encoding and Default Application Encoding also flow to this job.

Note there is a bug in DSNTIJUZ and DSNHDECM - These ship with CCSID 500 as default.

Installation

Specification of CCSIDs is performed at installation via install Panel DSNTIPF

```

DSNTIPF INSTALL DB2 -APPLICATION PROGRAMMING DEFAULTS PANEL 1
====>_
Enter data below:
1 LANGUAGE DEFAULT      ==> IBMCOB  ASM,C,CPP,COBOL,COB2,IBMCOB,FORTRAN,PLI
2 DECIMAL POINT IS     ==> .        .or ,
3 MINIMUM DIVIDE SCALE ==> NO       NO or YES for a minimum of 3 digits
   to right of decimal after division
4 STRING DELIMITER     ==> DEFAULT  DEFAULT,"or '(COBOL or COB2 only)
5 SQL STRING DELIMITER ==> DEFAULT  DEFAULT,"or '
6 DIST SQL STR DELIMTR ==> '        'or "
7 MIXED DATA          ==> NO       NO or YES for mixed DBCS data
8 EBCDIC CCSID         ==> 0       CCSID of your SBCS or MIXED DATA
9 ASCII CCSID          ==> 0       CCSID of SBCS or mixed data.
10 Unicode CCSID       ==> 1208    CCSID of Unicode UTF-8 data.
11 DEF ENCODING SCHEME ==> EBCDIC  EBCDIC, ASCII, or UNICODE
12 LOCALE LC_CTYPE     ==>
13 APPLICATION ENCODING ==> EBCDIC  EBCDIC, ASCII, UNICODE ccsid (1-65533)
14 DECIMAL ARITHMETIC  ==> DEC15   DEC15,DEC31,15,31
15 USE FOR DYNAMICRULES ==> YES    YES or NO
16 DESCRIBE FOR STATIC ==> NO     Allow DESCRIBE for STATIC SQL.NO or YES.

```

Install panel DSNTIPF is used to specify CCSID information. Options 8, 9, and 10 are where the CCSIDs for the three encoding schemes are specified.

Notice that ASCII and EBCDIC CCSIDs are initialized to 0 and the Unicode CCSID is initialized to 1208.

The ASCII and EBCDIC CCSIDs are not pre-filled; these values need to be set by the customer.

The EBCDIC should be set to the CCSID that the customer's 3270 emulators, CICS, and IMS transactions use.

The ASCII value should be set to the CCSID that is most commonly used by workstations in the customer shop (1252 for example).

The Unicode value is pre-filled with 1208 and cannot be changed.

This value specifies the mixed CCSID for Unicode tables.

Other things to note on this page:

Option 11 - This specifies the default encoding scheme for Objects created in the DB2 subsystem.

Option 13 - This option specifies the default application encoding. Changing this value should be done with great care.

DB2 Metadata – Where DB2 looks for information

CCSIDs are stored in the following places

SYSIBM.SYSDATABASE (V5)
 SYSIBM.SYSCOLUMNS
 - FOREIGNKEY (V2.3) - subtype information
 - CCSID (V8)
 SYSIBM.SYSPACKAGE (V7) – Application Encoding
 SYSIBM.SYSPARMS (V6)
 SYSIBM.SYSPLAN (V7) – Application Encoding
 SYSIBM.SYSROUTINES (V8)
 SYSIBM.SYSTABLES (V8)
 SYSIBM.SYSTABLESPACE (V5)
 SYSIBM.SYSVTREE (V5)

Plans and Packages (SCT02 and SPT01) – No external

Directory (DSNDB01) (V5) – No external

DECP (V2.3)

BSDS (V8)

In ENCODING_SCHEME column of - Stored as 'A', 'E', 'U', or blank (default)

SYSIBM.SYSDATATYPES
 SYSIBM.SYSDATABASE
 SYSIBM.SYSPARMS
 SYSIBM.SYSTABLESPACE
 SYSIBM.SYSTABLES

Once we've specified CCSIDs for our system, what does DB2 do with them?

DB2 stores CCSIDs in the Catalog, the directory, in bound statements, the directory, and of course the DECP

The value stored in these areas depends on what release of DB2 was used to create the object, and the value in the DECP at the time the object was created.

If a value is 0, it is assumed that the object is EBCDIC.

In general, DB2 does not support changing of a CCSID once it is specified in a DECP

The exceptions are

Changing from 0 to a valid value

Changing from a CCSID that does not support the EURO symbol to a CCSID that supports the EURO symbol (37 -> 1140 for instance).

Note that this sort of change requires special, disruptive, changes and should be undertaken only after the documentation has been read and the process is thoroughly understood.

Encoding information is also stored in some catalog tables



IBM Software Group

When Does DB2 Convert?

DB2 Information Management Software



@ business on demand software

When does conversion occur?

Local

Generally, conversion does not occur for local applications

When dealing with ASCII/Unicode tables

When specified by application

- CCSID Override in SQLDA (V2.3)

- Declare Variable (V7)

- Application Encoding Bind Option (V7)

- Current Application Encoding Special Register (V7)

ODBC/JDBC/SQLJ

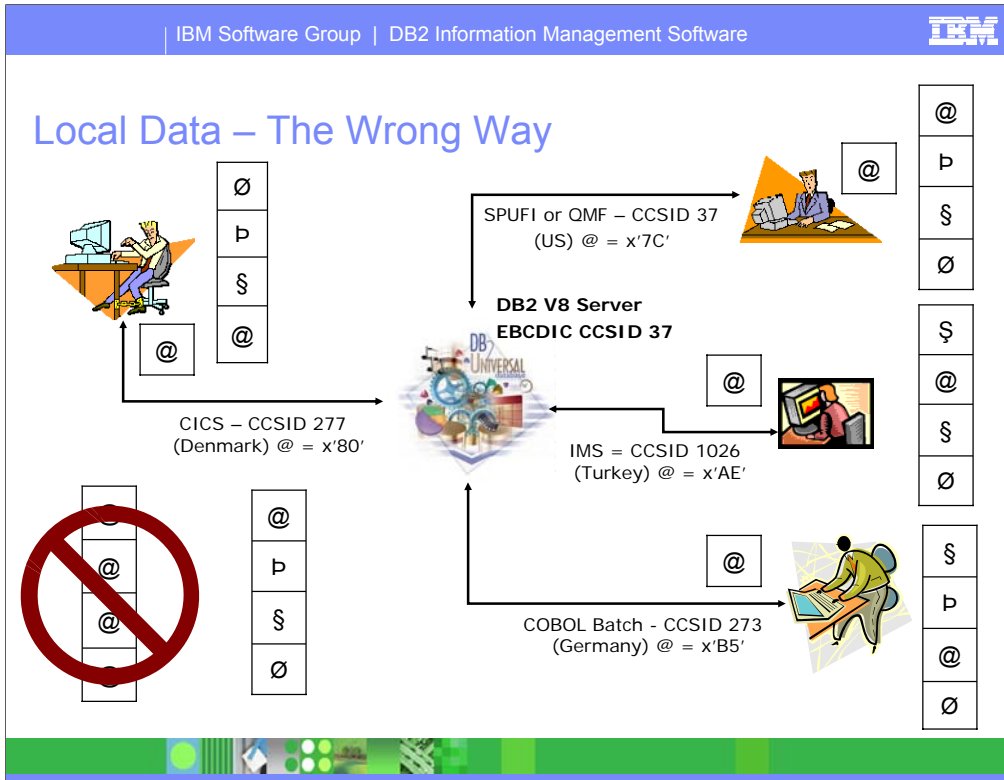
Remote

Automatically when needed

- DRDA Receiver Makes Right

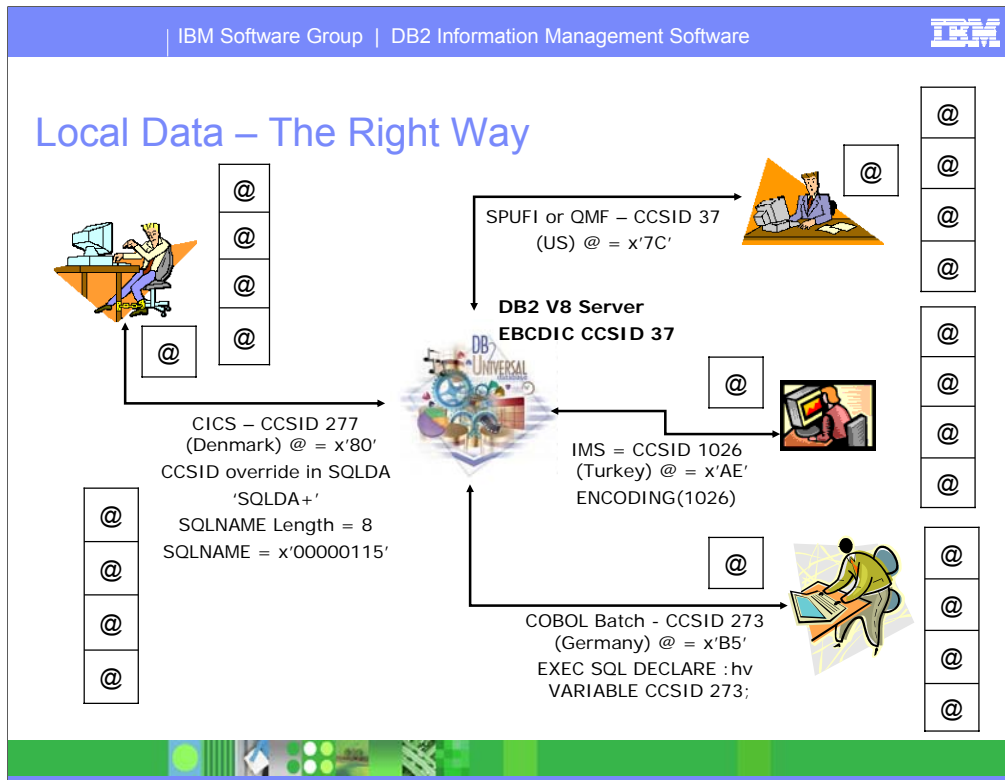
There are no hard and fast rules as to when a conversion occurs. The short answer is that conversion occurs when necessary.

Some of the cases when we do conversion are listed here.



In this example, assume that all the clients are referencing an EBCDIC table (CCSID 37 in this case). Each of the local applications is attempting to insert an “@” into a DB2 table. In general, the results will be incorrect unless DB2 “knows” that the data coming from the CICS, IMS, and COBOL applications is not CCSID 37. There are various ways of the application to tell DB2 the CCSID of the data that they are giving DB2 (in input host variables) or they want to receive from DB2 (in output host variables). These are:

- CCSID overrides in the SQLDA – when using a descriptor
- DECLARE VARIABLE – precompiler directive
- ENCODING bind option
- APPLICATION ENCODNG special register



In this example, assume that all the clients are referencing an EBCDIC table (CCSID 37 in this case). Each of the local applications is attempting to insert an “@” into a DB2 table. In general, the results will be incorrect unless DB2 “knows” that the data coming from the CICS, IMS, and COBOL applications is not CCSID 37.

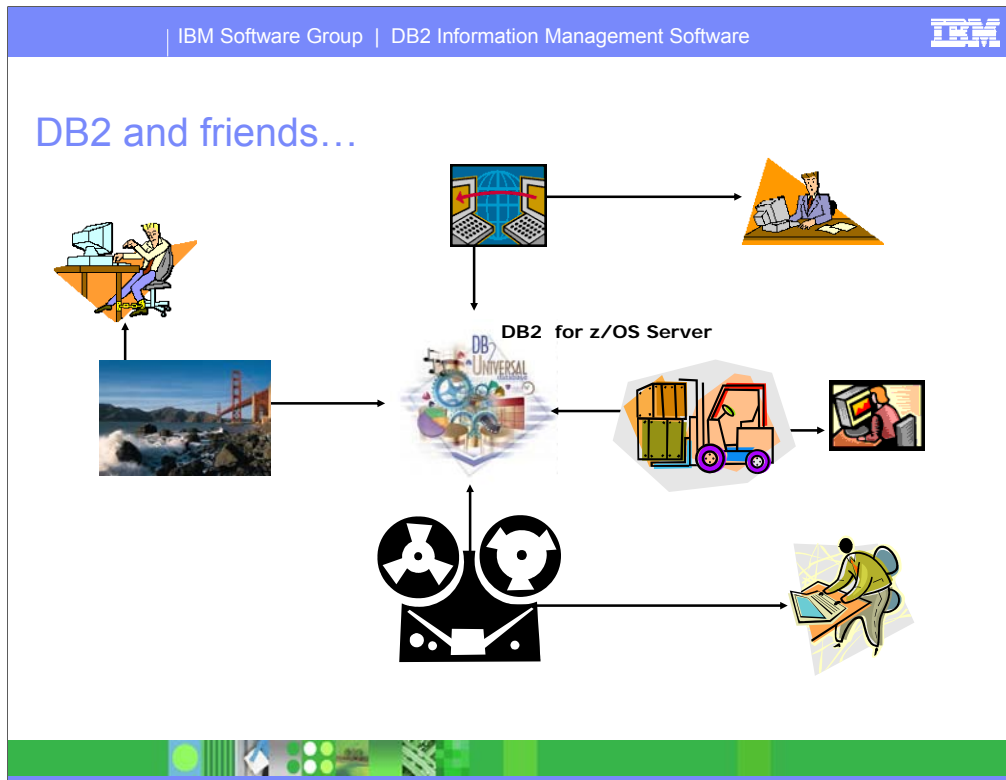
In this case:

The ENCODING bind option is used to let DB2 know the CCSID of the application inserting the data from Turkey.

DECLARE VARIABLE is used by the batch job inserting data from Germany

CCSID overrides are used in the CICS application inserting data from Denmark

The result is that the data is inserted correctly, and when each application fetches the data that was just inserted, it is represented correctly (each client sees the correct data).



This picture is showing that data does not have to come into DB2 through a terminal emulator to cause problems.

All character data has a CCSID. That CCSID may be implicit or it may be explicit, but there is a CCSID for the data.

In this case, you might have data coming through a (golden) gateway, a tape, electronic computer to computer transfer (such as FTP), or LOAD. In every case there is a CCSID associated with the data and if DB2 does not know what the CCSID is, it may treat the data incorrectly.



IBM Software Group

Character Based Functions

DB2 Information Management Software



@ business on demand software



Character Based Functions (PQ88784)

- New functions
 - ▶ CHARACTER_LENGTH
 - ▶ POSITION
 - ▶ SUBSTRING
- Updated Functions
 - ▶ CHAR
 - ▶ CLOB
 - ▶ DBCLOB
 - ▶ GRAPHIC
 - ▶ INSERT
 - ▶ LEFT
 - ▶ LOCATE
 - ▶ RIGHT
 - ▶ VARCHAR
 - ▶ VARGRAPHIC
- CAST Specification
 - ▶ Changes to enable specification of Code Units





Character Based Functions (sample syntax)

SUBSTRING

▶ SUBSTRING(*string-expression*, *start*, *length*, *CODEUNITS32* | *CODEUNITS16* | *OCTETS*)

CHAR

▶ CHAR(*string-expression*, *integer*, *CODEUNITS32* | *CODEUNITS16* | *OCTETS*)





Character Based Functions - Example

Assume that NAME is a VARCHAR(128) column, encoded in Unicode UTF-8, that contains 'Jürgen'. The following query:

```
SELECT CHARACTER_LENGTH(NAME, CODEUNITS32)
FROM T1 WHERE NAME = 'Jürgen';
```

or

```
SELECT CHARACTER_LENGTH(NAME, CODEUNITS16)
FROM T1 WHERE NAME = 'Jürgen';
```

returns the value 6. A similar query:

```
SELECT CHARACTER_LENGTH(NAME, OCTETS)
FROM T1 WHERE NAME = 'Jürgen';
```


or

```
SELECT LENGTH(NAME)
FROM T1 WHERE NAME = 'Jürgen';
```

returns the value 7.

Name	UTF-8 Representation	UTF-16 Representation	UTF-32 Representation
Jürgen	x'4AC3BC7267656E'	x'004A00FC007200670065006E'	x'0000004A000000FC0000007200000067000000650000006E'



IBM Software Group | DB2 Information Management Software 

LEFT and RIGHT Examples


LEFT

Statement	Result
LEFT('Jürgen',2,CODEUNITS32)	'Jü' -- x'4AC3BC'
LEFT('Jürgen',2,CODEUNITS16)	'Jü' -- x'4AC3BC'
LEFT('Jürgen',2,OCTETS)	'J ' -- x'4A20' a truncated string
LEFT('Jürgen',2)	'J?' -- x'4AC3' The letter 'J' and a Partial character*

RIGHT

Statement	Result
RIGHT('Jürgen',5,CODEUNITS32)	'ürgen' -- x'C3BC7267656E'
RIGHT('Jürgen',5,CODEUNITS16)	'ürgen' -- x'C3BC7267656E'
RIGHT('Jürgen',5,OCTETS)	' rgen' -- x'207267656E' a truncated string
RIGHT('Jürgen',5)	' ?rgen' -- x'BC7267656E' a partial character followed by 'rgen'*

*If conversion occurs with on a string with a partial character, SQLCODE -330 results.



LEFT and RIGHT operate in a similar manner.

In these examples, CODEUNITS32 and CODEUNITS16 return the same results.

When OCTETS is specified, an a partial character would result, the PAD character is used.

When no units are specified, the behavior is similar to that of OCTETS (byte based), but partial characters can result.

IBM Software Group | DB2 Information Management Software

SUBSTRING

Function ...	Returns ...
SUBSTRING('Jürgen',1,2,CODEUNITS32)	'Jü' -- x'4AC3BC'
SUBSTRING('Jürgen',1,2,CODEUNITS16)	'Jü' -- x'4AC3BC'
SUBSTRING('Jürgen',1,2,OCTETS)	'J' -- x'4A20' - a truncated string
SUBSTR ('Jürgen',1,2)	'J?' -- x'4AC3' -- a partial character
SUBSTRING('Jürgen',8,CODEUNITS16)	a zero-length string
SUBSTRING('Jürgen',8,4,OCTETS)	a zero-length string
SUBSTRING('ABCDEFG',-2,2,OCTETS)	a zero-length string
SUBSTRING('ABCDEFG',-2,4,OCTETS)	'A'
SUBSTRING('ABCDEFG',-2,5,OCTETS)	'AB'
SUBSTRING('ABCDEFG',-2,OCTETS)	'ABCDEFG'
SUBSTRING('ABCDEFG',0,1,OCTETS)	a zero-length string

Let C be the value of the *string-expression*, let LC be the length in characters of C, and let S be the value of the *start*. v If *length* is


The first and second examples return the same value because even though 'ü' is a multi-byte character in UTF-8, it is one CODEUNITS32 or CODEUNITS16 codepoint.

In the third example OCTETS is specified and since a partial character would have resulted, truncation occurs on the partial character and a pad character is returned in place of the partial character.

In the fourth example, we are using SUBSTR, not SUBSTRING. SUBSTR is byte based and will return partial characters. As mentioned previously if conversion occurs on a string with a partial character, -330 will result.

specified, let L be the value of <string length> and let E be S+L.

Otherwise, let E be the larger of LC + 1 and S. v If either C, S, or L is the null value, the result of the function is the null value. v If E is less than S, an exception condition is raised: data exception — substring error. v Otherwise: – If S is greater than LC or if E is less than 1 (one), the result of the function is a zero-length string. – Otherwise: - Let S1 be the larger of S and 1 (one). Let E1 be the smaller of E and LC+1. Let L1 be E1–S1. - The result of the function is a character string that contains the L1 characters of C starting at character number S1 in the same order that the characters appear in C.

IBM Software Group | DB2 Information Management Software 

Example where CODEUNITS16 and CODEUNITS32 differ

Unicode value \u1D400 - 'A' MATHEMATICAL BOLD CAPITAL A

UTF-8	UTF-16	UTF-32
X'F09D9080'	X'D835DC00'	X'0001D400'


Assume that C1 is a VARCHAR(10) column, encoded in Unicode UTF-8, and that table T1 contains one row with the value of the mathematical bold capital A (X'F09D9080').

The following similar queries return different answers:

```
SELECT CHARACTER_LENGTH(C1, CODEUNITS32) FROM T1; -- Returns 1
SELECT CHARACTER_LENGTH(C1, CODEUNITS16) FROM T1; -- Returns 2
SELECT CHARACTER_LENGTH(C1, OCTETS) FROM T1; -- Returns 4
```

The following similar queries return different answers:

```
SELECT HEX(SUBSTRING(C1, 1, 1, CODEUNITS32)) FROM T1; -- Returns X'F09D9080'
SELECT HEX(SUBSTRING(C1, 1, 1, CODEUNITS16)) FROM T1; -- Returns X'20'
SELECT HEX(SUBSTRING(C1, 1, 2, CODEUNITS16)) FROM T1; -- Returns X'F09D9080'
SELECT HEX(SUBSTRING(C1, 1, 1, OCTETS)) FROM T1; -- Returns X'20' (blank)
SELECT HEX(SUBSTR(C1, 1, 1)) FROM T1; -- Returns X'F0'
```



For many characters, specifying CODEUNITS16 and CODEUNITS32 on a BIF does not matter. However, when supplementary characters (sometimes also called surrogate characters) are used, the results can be different.

Multiple CCSID Sets per SQL Statement

While there are no syntax changes to allow multiple CCSID sets, the following SQL statements may be affected.

```
ALTER TABLE & ALTER TABLE ADD (materialized query table)
CREATE TABLE (materialized query table)
CREATE TABLE LIKE view-table
CREATE GLOBAL TEMPORARY TABLE LIKE view-table
CREATE VIEW
DECLARE GLOBAL TEMPORARY TABLE AS (fullselect) DEFINITION ONLY
DECLARE GLOBAL TEMPORARY TABLE LIKE view-table
DELETE
INSERT
SELECT
SELECT INTO
UPDATE
Scalar fullselect expression
```

The statements listed above may have their behavior affected when data from more than one encoding is processed as part of the execution of the statement

Multiple CCSID Sets - Example 1:

```

SELECT a.name, a.creator, b.charcol, 'ABC',
       :hvchar, X'C1C2C3'
FROM sysibm.systables a,
     ebcdictable b
WHERE a.name = b.name AND
      b.name > 'B' AND
      a.creator = 'SYSADM'
ORDER BY b.name;

```

In the above example, since both tables have the same system EBCDIC CCSID set, the comparisons are done in EBCDIC and the result data is EBCDIC.

To maintain compatibility with previous releases, statements that do not reference objects with more than one CCSID set, will continue to use the old rules.

- For compatibility with prior releases:
 - IF an SQL statement which references table objects with **only one CCSID set**
 - THEN **the results** will continue to have the same **result encoding scheme (CCSID set) as the table objects** and basic semantics.
- All string objects and special registers in the SQL statement are converted to this result encoding scheme (CCSID set) vs. using the application encoding scheme like multiple CCSID set SQL statements.

In this example, all the tables reference a single encoding scheme.

The results of this query will be EBCDIC.

Multiple CCSID Sets - Example 1 (continued)

```
SELECT a.name, a.creator, b.charcol, 'ABC',  
       :hvchar, X'C1C2C3'  
FROM sysibm.systables a,  
     ebcdictable b  
WHERE a.name = b.name AND  
      b.name > 'B' AND  
      a.creator = 'SYSADM'  
ORDER BY b.name;
```

Result or Evaluated:

EBCDIC

Unicode

Application Encoding Scheme

Assuming a Unicode catalog, the result will contain multiple CCSIDs and the comparisons and ordering will be dependent on the context.

In this example, where we have a Unicode catalog, some columns are EBCDIC and others are Unicode. Literal values are interpreted using the Application Encoding option in effect. Queries that cross encoding schemes are allowed once ENFM (Enabling New Function Mode) is entered. This is to ensure that applications that access the catalog continue to function.

SQL statements with multiple CCSID sets

Comparison and resulting data types for multiple CCSID sets...

If an expression or comparison involves two strings which contain columns with different CCSID sets,

Drive to Unicode if necessary

WHERE T1.C1 = T2.C1

If an expression or comparison involves two strings with different CCSID sets where only one of them contains a column,

Drive to the column's CCSID set

WHERE T1.C1 = X'C1C2'

If an expression or comparison involves two strings with different CCSID sets and neither contains a column,

Drive to Unicode

WHERE GX'42C142C2' = 'ABC' -- GX literal and 'ABC' are different CCSIDs

String constants and special registers in a context by themselves use the application encoding scheme

SELECT 'ABC' FROM T1 . . .

Here are some examples of how the new DB2 rules for multiple CCSID sets work

-In the first case we have two columns being compared – in this case, drive to Unicode if they are not already Unicode

-In the second case, we drive to the column CCSID set

-In the third case, neither side is a column, so drive to Unicode

-In the fourth case we use the APPLICATION ENCODING SCHEME – the thought here is that the application is providing the string, we should interpret the string in the same encoding as the application.

IBM Software Group | DB2 Information Management Software

Multiple CCSID Sets – A closer look

```

CREATE TABLE TCCSIDU (CU1 VARCHAR(12)) CCSID UNICODE;
CREATE TABLE TCCSIDE (CE1 VARCHAR(12)) CCSID EBCDIC;
INSERT INTO TCCSIDU VALUES ('Jürgen');
INSERT INTO TCCSIDE VALUES ('Jürgen');

SELECT LENGTH(A.CU1) AS L1, HEX(A.CU1) AS H1,
       LENGTH(B.CE1) AS L2, HEX(B.CE1) AS H2
FROM TCCSIDU A, TCCSIDE B WHERE A.CU1 = B.CE1;
    
```

Returns

	L1	H1	L2	H2
1_	7	4AC3BC7267656E	6	D1DC99878595

In this example, you can see that up until the data is returned to the application, the data from TCCSIDU is Unicode, and the data from TCCSIDE is EBCDIC. At the point the data is returned to the application (copied from DB2 to the applications host variables), conversion will be performed to the CCSID specified by the application (either implicitly or explicitly). In this case, since we are returning the HEX representation of the data, we really don't convert the data, we convert the hex representation of the data (which will always be composed of characters between 0-9 and A-F).

The result is the same for all Join methods. For Sort Merge Join (SMJ), we convert the TCCSIDE.CE1 to Unicode before sorting. In the sort workfile, we carry both the EBCDIC version of the string (to return to the application), and the Unicode version of the string (for comparison processing in the Join). As a result, the size of workfiles can increase when EBCDIC and Unicode tables are joined (as compared to joining two EBCDIC or two Unicode tables). Additionally, it is possible that SQLCODE -136 or -670 may be issued.

IBM Software Group | DB2 Information Management Software

Multiple CCSID – Performance Considerations

```

CREATE TABLE ET1(C1 VARCHAR(8) FOR SBCS DATA) CCSID EBCDIC;
CREATE TABLE ET2(C1 VARCHAR(8) FOR SBCS DATA) CCSID EBCDIC;
CREATE TABLE UT3(C1 VARCHAR(8) FOR SBCS DATA) CCSID UNICODE;

EXPLAIN ALL SET QUERYNO = 1 FOR SELECT ET1.C1 FROM ET1
  WHERE ET1.C1 IN (SELECT ET2.C1 FROM ET2 WHERE ET1.C1 = ET2.C1);

EXPLAIN ALL SET QUERYNO = 2 FOR SELECT ET1.C1 FROM ET1
  WHERE ET1.C1 IN (SELECT UT3.C1 FROM UT3 WHERE ET1.C1 = UT3.C1);
        
```

	QUERYNO	QBLOCKNO	METHOD	TNAME	TABNO	ACCESSTYPE
1	1	1	0	ET1	1	R
2	1	1	2	ET2	2	R
3	2	1	0	ET1	1	R
4	2	2	0	UT3	2	R

In this example, we have three tables ET1 and ET2 are EBCDIC tables, and UT3 is a Unicode table.

The two queries are identical, except the second query has UT3 where the first query has ET2.


Note that table UT3 isn't really even using Unicode data (it is defined as "FOR SBCS DATA" which means that the data will be 7-bit ASCII).

When two EBCDIC tables are involved, as in QUERYNO=1, we are able to transform the query into a join. In the second example, the transformation to a join did not take place, as a result the second query may perform slower than the first query.

This is a very simple example that demonstrates the ramifications to performance of multiple CCSID statements.


Column Name and Type	Description
TABLE_ENCODE CHAR(1)	Indicates the encoding scheme of the statement. If the statement represents a single CCSID set, then the column will contain 'E' for EBCDIC, 'A' for ASCII, or 'U' for Unicode. If the statement is a multiple CCSID set statement, then the column will be set to 'M' for multiple CCSID sets.
TABLE_SCCSID FIXED(16)	The SBCS CCSID value of the table or zero if the TABLE_ENCODE column is 'M'
TABLE_MCCSID FIXED(16)	The Mixed CCSID value of the table or zero if the TABLE_ENCODE column is 'M'
TABLE_DCCSID FIXED(16)	The DBCS CCSID value of the table or zero if the TABLE_ENCODE column is 'M'

The table above shows the changes to the PLAN_TABLE to support multiple CCSIDs in a statement


IBM Software Group | DB2 Information Management Software 

New EXPLAIN tables / columns -- DSN_STATEMNT_TABLE

Column Name and Type	Description
STMT_ENCODE CHAR(1)	Indicates the encoding scheme of the statement. If the statement represents a single CCSID set, then the column will contain 'E' for EBCDIC, 'A' for ASCII, or 'U' for Unicode. If the statement is a multiple CCSID set statement, then the column will be set to 'M' for multiple CCSID sets.




The table above shows the changes to the DSN_STATEMNT_TABLE to support multiple CCSIDs in a statement

IBM Software Group | DB2 Information Management Software 

Example of CAST to influence ordering

Given table names: TA, TB, T1, T2

<pre>SELECT NAME FROM SYSIBM.SYSTABLES WHERE NAME LIKE 'T%' ORDER BY NAME</pre> <p>In EBCDIC (V7), returns: TA, TB, T1, T2</p> <p>In Unicode (V8), returns: T1, T2, TA, TB</p>	<pre>SELECT CAST (NAME AS VARCHAR(128) CCSID EBCDIC) AS E_NAME FROM SYSIBM.SYSTABLES WHERE NAME LIKE 'T%' ORDER BY E_NAME</pre> <p>Returns: TA, TB, T1, T2</p>
--	--



The new CAST specification can be used to influence ordering.

The above example shows how SQL could be recoded to cause data to be returned in the same order as DB2 V7 would.

CAST – Changing a CCSID

- If both the *length* and CCSID clauses are specified, the data is first cast to the specified CCSID, and then the *length* is applied.

Example:

```
CAST ('Jürgen' as VARCHAR(6) CCSID UNICODE)
```

Returns 'Jürge'

- If either CODEUNITS32 or CODEUNITS16 is specified, the specification of *length* applies to the units specified.

Example:

```
CAST ('Jürgen' as VARCHAR(6 CODEUNITS16) CCSID UNICODE)
```

Returns 'Jürgen'

When a length is specified as part of a CAST that also specifies a CCSID, the length is applied after the data has been converted to the target CCSID, unless CODEUNITS16 or CODEUNITS32 is specified. If CODEUNITS16 or CODEUNITES32 is specified, then the length applies in the specified units, then the data is cast to the final CCSID.



Sample Schema - EBCDIC

```
CREATE TABLE DSN8XX0.EMP
(EMPNO      CHAR(6)                NOT NULL,
 FIRSTNME   VARCHAR(12)            NOT NULL,
 MIDINIT    CHAR(1)                NOT NULL,
 LASTNAME   VARCHAR(15)            NOT NULL,
 WORKDEPT   CHAR(3)                ,
 PHONENO    CHAR(4) CONSTRAINT NUMBER CHECK
 (PHONENO >= '0000' AND PHONENO <= '9999') ,
 HIREDATE   DATE                   ,
 JOB        CHAR(8)                 ,
 EDLEVEL    SMALLINT                ,
 SEX        CHAR(1)                 ,
 BIRTHDATE  DATE                   ,
 SALARY     DECIMAL(9, 2)           ,
 BONUS      DECIMAL(9, 2)           ,
 COMM       DECIMAL(9, 2)           ,
 PRIMARY KEY(EMPNO))
CCSID EBCDIC;
```




Sample Schema – Unicode UTF-16

```
CREATE TABLE DSN8ZZ0.EMP
  (EMPNO      GRAPHIC(6)                NOT NULL,
   FIRSTNME  VARGRAPHIC(12)            NOT NULL,
   MIDINIT   GRAPHIC(1)                 NOT NULL,
   LASTNAME  VARGRAPHIC(15)            NOT NULL,
   WORKDEPT  VARGRAPHIC(3)              ,
   PHONENO   VARGRAPHIC(4) CONSTRAINT NUMBER CHECK
     (PHONENO >= '0000' AND PHONENO <= '9999') ,
   HIREDATE  DATE                       ,
   JOB       VARGRAPHIC(8)              ,
   EDLEVEL   SMALLINT                   ,
   SEX       VARGRAPHIC(1)              ,
   BIRTHDATE DATE                       ,
   SALARY    DECIMAL(9, 2)              ,
   BONUS     DECIMAL(9, 2)              ,
   COMM      DECIMAL(9, 2)              ,
   PRIMARY KEY(EMPNO))
CCSID UNICODE;
```


COBOL

Enterprise COBOL for z/OS and OS/390 V3R1 Supports Unicode

NATIONAL is used to declare UTF-16 variables

```
MY-UNISTR pic N(10). -- declares a UTF-16 Variable
```

N and NX Literals

```
N'123'
```

```
NX'003100320033'
```

Conversions

NATIONAL-OF Converts to UTF-16

DISPLAY-OF Converts to specific CCSID

```
DECLARE Greek-EBCDIC pic X(10) value "ΕΣΦΛΘΖΔΓΩ".
```

```
UTF16STR pic N(10).
```

```
UTF8STR pic X(20).
```

```
Move Function National-of(Greek-EBCDIC, 00875) to UTF16STR.
```

```
Move Function Display-of(UTF16STR, 01208) to UTF8STR.
```

Cobol has recently added support for Unicode characters

Included in this support

- New NATIONAL data type
- N and NX literals
- Conversion operations

COBOL and DB2 (getting DB2 to convert for you)

```
EXEC SQL BEGIN DECLARE SECTION;
  01 HOST-VARS.
    05 GREEK-EBCDIC PIC X(10) VALUE "ΞΣΦΛΘΖΔΓΩ".
    05 UTF16STR PIC N(10).*
    05 UTF8STR PIC X(20).
  EXEC SQL DECLARE :UTF8STR VARIABLE CCSID 1208.
EXEC SQL END DECLARE SECTION;

INSERT INTO T1 (C1) VALUES(:GREEK-EBCDIC) END EXEC.**

EXEC SQL
  SELECT C1, C1 INTO :UTF16STR, :UTF8STR
END EXEC.
```

* COBOL will use an implicit DECLARE VARIABLE for PIC N data.
** This example assumes T1 is a table encoded in EBCDIC CCSID 875 and that the ENCODING bind option for this appl is also CCSID 875.

Using the previous example as a starting point, it is also possible to get DB2 to perform the needed conversions for you.

In this example, the EBCDIC 875 data is INSERTed into an EBCDIC CCSID 875 table. Then the COBOL application fetches the data into the UTF-16 and UTF-8 host variables.

Note that only one DECLARE VARIABLE statement is used needed because PIC N data is implicitly UTF-16, and the PIC X data is assumed to be in the CCSID specified by the ENCODING bind option.

IBM Software Group | DB2 Information Management Software

PL/I

```

%PROCESS CODEPAGE(277), WIDECHAR(BIGENDIAN);

DCL UTF16STR WIDECHAR(10) VARYING;
DCL uOneTwoThree WCHAR(3);
DCL eOneTwoThree CHAR(3);
uOneTwoThree = WX'003100320033';           /* UTF-16 '123' */
eOneTwoThree = '123';                       /* = x'F1F2F3' */

IF uOneTwoThree = eOneTwoThree THEN         /* FALSE */
...
IF uOneTwoThree = WIDECHAR(eOneTwoThree) THEN /* True */
...
UTF16STR = WIDECHAR('ABC@'); /* note '@' is assumed to be in CCSID
273 position (x'B5') because of the CODEPAGE(273) specification. UTF16STR
now = x'0041004200430040' */

```

PL/I has some support for Unicode UTF-16 data.

PL/I uses the WIDECHAR datatype to support UTF-16.

PL/I supports UTF-16 data as either BIG or LITTLE ENDIAN. For compatibility with DB2, BIGENDIAN data should be used.

PL/I has basic conversion support via the WIDECHAR function. This function does not use the z/OS conversion services at this time.

PL/I does not have any support for a UTF-16 literal, but does offer a WX literal which can be used to specify a UTF-16 codepoint (like the UX constant would be used in DB2).

PL/I Example with USING DESCRIPTOR

```

...
DCL STMT1 CHAR(100) VARYING INIT('INSERT INTO T1 VALUES (?,?) ');

DCL DA1 CHAR(16+(2*44)); /* ALLOCATE SPACE FOR 2 SQLDA ENTRIES */

EXEC SQL INCLUDE SQLDA;
SQLDAID = 'SQLDA+ '; /* Note the "+" */ /*
SQLN = 2; /* Allocated SQLVARS */
SQLD = 2; /* Used SQLVARS */
SQLVAR(1).SQLTYPE = 468; /* Graphic – not null */
SQLVAR(1).SQLLEN1 = 3; /* Length = 6 bytes */
SQLVAR(1).SQLDATA = ADDR(uOneTwoThree); /* Address of host var */
SQLVAR(1).SQLNAME = '0000048000000000'X; /* CCSID 1200 */
SQLVAR(2).SQLTYPE = 452; /* Character – not null */
SQLVAR(2).SQLLEN1 = 3; /* Length = 3 bytes */
SQLVAR(2).SQLDATA = ADDR(eOneTwoThree); /* Address of host var */
SQLVAR(2).SQLNAME = '0000011100000000'X; /* CCSID 273 */

EXEC SQL PREPARE S1 FROM :STMT1; /* Prepare Statement */

EXEC SQL EXECUTE S1 USING DESCRIPTOR :SQLDA; /* Execute Stmt */
...

```

This code snippet shows the code needed to use a descriptor (SQLDA) with a dynamic INSERT statement.

This is one example of how to perform this sort of programming and is intended to demonstrate the techniques associated with using a descriptor and setting CCSID values with a dynamic INSERT statement.



Conversion Considerations – Application –vs- DB2

- In many cases, conversion can be handled by the application or DB2 (as shown in previous slides).
- Conversion by DB2
 - ▶ Pro
 - Less application code
 - ▶ Con
 - Performance impact may be greater than just conversion cost
- Conversion by Application
 - ▶ Pro
 - May perform better in some cases
 - ▶ Con
 - Requires more application code



ODBC/SQLJ/JDBC

V7 ODBC Support

Support for Wide Character API's (UCS2/UTF-16)

See ODBC Guide and Reference (SC26-9941-01)

SQLRETURN	SQLRETURN
SQLPrepare (SQLPrepareW (
SQLHSTMT hstmt,	SQLHSTMT hstmt,
SQLCHAR *szSqlStr,	SQLWCHAR *szSqlStr,
SQLINTEGER cbSqlStr);	SQLINTEGER cbSqlStr);

V8 ODBC Support

CURRENTAPPENSCH ini file setting

SQLJ/JDBC Support



Remove current support for converting to EBCDIC before calling engine. Let DB2 engine determine where conversion is necessary

ODBC support for Unicode is included as part of the effort to support ODBC 3.0

In V7, ODBC added support for Wide Character APIs. These Wide Character APIs support Unicode input, but the data is converted to EBCDIC before being passed to DB2

In V8, ODBC added support for specification of processing using the CURRENTAPPENSCH ini file setting. Additionally, in V8, the ODBC Wide Character APIs pass data to DB2 in Unicode.

SQLJ and JDBC already support Unicode, but changes have been made to exploit Unicode support in DB2 V7 and DB2 V8.

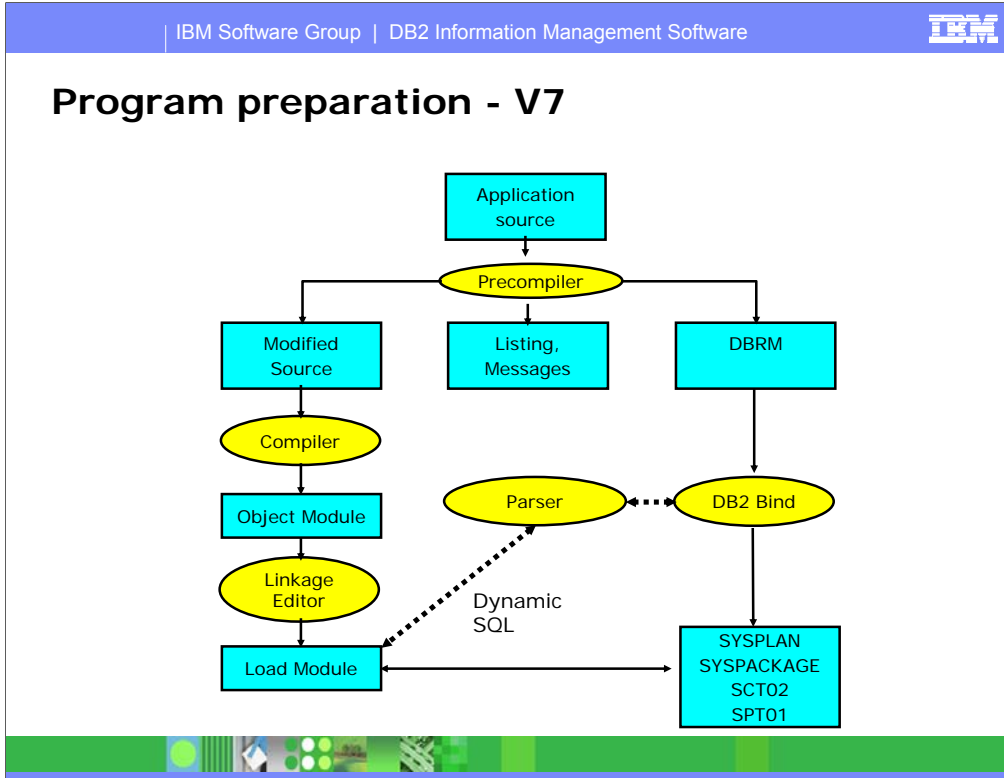
These changes are aimed at reducing the number of conversions that take place, and avoiding the potential for data loss.



SPUFI, DSNTEP2/4, DSNTIAUL (PQ98170 & PK06076)

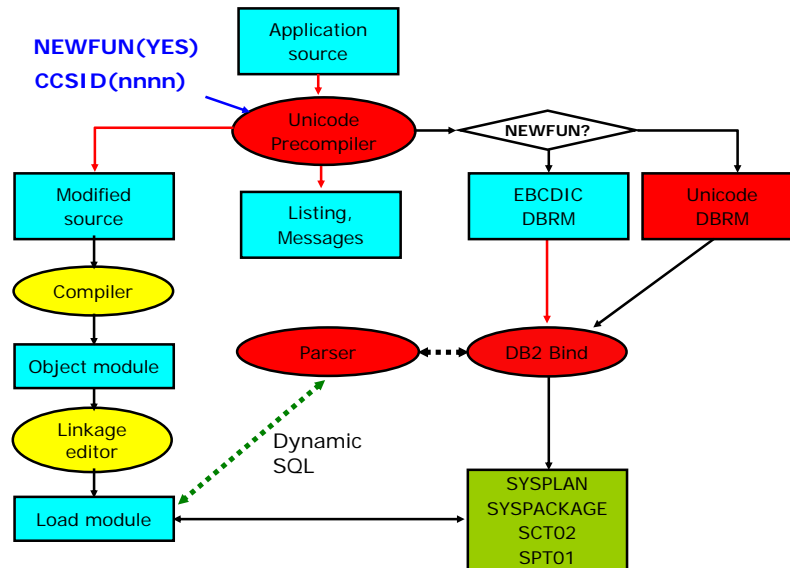
- Have been modified to support UTF-16 better. The SQLDA from DESCRIBE is changed to accommodate the data type of each graphic column having CCSID 1200 as follows:
 - ▶ GRAPHIC -> CHAR
 - ▶ VARGRAPHIC and LONG VARGRAPHIC -> VARCHAR - DBCLOB -> CLOB
- During FETCH, DB2 converts all UTF-16 characters that can be mapped to the EBCDIC SBCS CCSID, and otherwise sets SQLCODE 0 with SQLWARNING flag 8 (a character that could not be converted was replaced with a substitute character).





Program preparation in DB2 V7 and below assumes all data is in the SYSTEM EBCDIC CCSID

Program Preparation V8



CCSID(nnnn) input parameter to the Unicode Precompiler specifies the CCSID of the application source to ensure proper conversion to unicode for processing. The default value of the CCSID option is the EBCDIC system CCSID as specified on the panel DSNTIPF during installation.

The modified source program (an output of precompilation) remains in its original CCSID. If the DBRM is later bound to a server that does not support UTF-8, the SQL statements are then converted from CCSID 1208 (UTF-8) and sent in the EBCDIC system CCSID.

- NEWFUN(YES)
 - Accept V8 new syntax
 - Unicode DBRM
- NEWFUN(NO)
 - Reject V8 new syntax
 - EBCDIC DBRM
- CCSID(nnnn)
 - It is important to note that this CCSID applies to the source of the application program. It does not necessarily apply to host variables that are used by this application.



IBM Software Group

Utility Support

DB2 Information Management Software



@ business on demand software



Utility Unicode statements

Utility control statements may be specified in Unicode or EBCDIC

DB2 detects which encoding scheme is being used

Must be all UTF-8 or EBCDIC - no mixing!

Object names in messages will be in EBCDIC

New utility stored procedure interface DSNUTILU for Unicode

Identical to DSNUTILS except:

Inputs are in UNICODE

utility_name parameter dropped

Data set DYNALLOC keywords dropped

use TEMPLATE for all data sets

DB2 V8 adds support for Utility control statements in Unicode.

Additionally, a new Unicode Stored Procedure, DSNUTILU is added to compliment DSNUTILS

DSNUTILS is an updated version of DSNTILU that

- Has Unicode parameters
- Has fewer required parameters
- Allocates datasets based on TEMPLATES

Utility control statements

EBCDIC:

```
//SYSIN DD *  
  QUIESCE TABLESPACE A.B  
  COPY TABLESPACE A.B  
/*
```

UNICODE:

```
//SYSIN DD *  
"éíñáëää"è â<áë& äá" "â" "  
"ä|&ß"è â<áë& äá" "â" "  
/*
```

If you see what appears to be garbage where a utility statement should be, don't panic. It could be Unicode UTF-8. In V8, Utility control statements can now be in Unicode.

Control statements (in EBCDIC or Unicode) also allow long names (V8)



ISPF


- DISPLAY CCSID ccsid_number [LINE start_line end_line] [COLS start_col end_col]
CCSID
n
UTF8
UTF16
UTF32
UCS2
EBCDIC
UNICODE
- FIND C'xxxxxx' UTF8
- OA07685 (OA08496 fixes a problem with FIND)
- z/OS 1.4, 1.5, and 1.6



DSNUTILS –vs– DSNUTILU


<pre> CREATE PROCEDURE DSNUTILS (IN UTILITY_ID VARCHAR(16) , IN RESTART VARCHAR(8) , IN UTSTMT VARCHAR(32704) , OUT RETCODE INTEGER , IN UTILITY_NAME VARCHAR(20) , IN RECDSN VARCHAR(54) , IN RECDEVT CHAR(8) , IN RECSPACE SMALLINT , IN DISCDSN VARCHAR(54) , IN DISCDEVT CHAR(8) , IN DISCSpace SMALLINT , IN PNCHDSN VARCHAR(54) , IN PNCHDEVT CHAR(8) , IN PNCHSPACE SMALLINT , IN COPYDSN1 VARCHAR(54) , IN COPYDEVT1 CHAR(8) , IN COPYSPACE1 SMALLINT ... </pre>	<pre> CREATE PROCEDURE DSNUTILU (IN UTILITY_ID VARCHAR(16) CCSID UNICODE , IN RESTART VARCHAR(8) CCSID UNICODE , IN UTSTMT VARCHAR(32704) CCSID UNICODE , OUT RETCODE INTEGER) EXTERNAL NAME DSNUTILU LANGUAGE ASSEMBLE WLM ENVIRONMENT WLMENV1 NO COLLID RUN OPTIONS 'TRAP(OFF)' PROGRAM TYPE MAIN MODIFIES SQL DATA ASUTIME NO LIMIT STAY RESIDENT NO COMMIT ON RETURN NO PARAMETER STYLE GENERAL RESULT SETS 1 EXTERNAL SECURITY USER; </pre>
---	---

Items in green are parms which are not in DSNUTILU. Notice that DSNUTILU specifies CCSID UNICODE for it's character parameters.


IBM Software Group | DB2 Information Management Software 

Conversion support in Load and Unload

LOAD Utility
UTF-16 <-> UTF-8
SBCS/MIXED -> DBCS
DBCS -> SBCS/MIXED
ASCII/EBCDIC <->
UNICODE



UNLOAD Utility
ASCII/EBCDIC <->
UNICODE
No support for
SBCS/MIXED ->
DBCS
DBCS ->
SBCS/MIXED



The load utility has been extended to support conversion to and from Unicode.

Additionally, the load utility will support conversion between character and graphic as long as conversion exists.

Character in load dataset -> Graphic column

Graphic in load dataset -> character column

The unload utility, new for V7, supports conversion to/from Unicode, but does not support conversion between character and graphic



LOAD Example

```
CREATE TABLE TCCSIDU (CU1 VARCHAR(11)) CCSID UNICODE;
```

```
LOAD DATA INDDN SYSREC LOG YES CCSID(0037) INTO TABLE
  "USER1"."TCCSIDU" ( "CU1" POSITION( 00001:00011) CHAR(00011) )
```

Data is

```
Jürgen
Hegelstraße
```

```
DSNU000I      DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = LOAD65710
DSNU1044I     DSNUGTIS - PROCESSING SYSIN AS EBCDIC
DSNU050I      DSNUGUTC - LOAD DATA INDDN SYSREC LOG YES CCSID(37)
DSNU650I     ) DSNURWI - INTO TABLE "ADMF001"."TCCSIDU"
DSNU650I     ) DSNURWI - ("CU1" POSITION(1:11) CHAR(11))
DSNU334I     ) DSNURCON - INPUT FIELD 'CU1' INVALID FOR 'USER1.TCCSIDU',
  ERROR CODE '19'
DSNU302I     ) DSNURWBF - (RE)LOAD PHASE STATISTICS - NUMBER OF INPUT
  RECORDS PROCESSED=1
DSNU017I     DSNUGBAC - UTILITY DATA BASE SERVICES MEMORY EXECUTION
  ABENDED, REASON=X'00E40323'
```

To prevent this error table should be declared as follows:

```
CREATE TABLE TCCSIDU (CU1 VARCHAR(33)) CCSID UNICODE;*
```

*Recommendation based on Table 25 (*Result length of CCSID conversion*) in the SQL Ref

IBM Software Group | DB2 Information Management Software

UNLOAD Example

```

CREATE TABLE TCCSIDE (CE1 VARCHAR(11)) CCSID EBCDIC;
INSERT INTO TCCSIDE VALUES ('Jürgen');
INSERT INTO TCCSIDE VALUES ('Hegelstraße');

UNLOAD CCSID(1208) FROM TABLE USER1.TCCSIDE

DSNU000I      DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = AOAD65710
DSNU1044I      DSNUGTIS - PROCESSING SYSIN AS EBCDIC
DSNU050I      DSNUGUTC - UNLOAD CCSID(1208)
DSNU650I      DSNUGMS - FROM TABLE ADMF001.TCCSIDE
DSNU1233I      DSNULVA - DATA IS TOO LONG FOR FIELD CE1, TABLE ADMF001.TCCSIDE
DSNU1219I      DSNULVA - THE NUMBER OF RECORDS IN ERROR REACHED THE LIMIT 1
DSNU253I      DSNUNLD - UNLOAD PHASE STATISTICS - NUMBER OF RECORDS UNLOADED=1
                FOR TABLE USER1.TCCSIDE
DSNU252I      DSNUNLD - UNLOAD PHASE STATISTICS - NUMBER OF RECORDS UNLOADED=1
                FOR TABLESPACE DSNDB04.TCCSIDE
DSNU250I      DSNUNLD - UNLOAD PHASE COMPLETE, ELAPSED TIME=00:00:00
DSNU012I      DSNUGBAC - UTILITY EXECUTION TERMINATED, HIGHEST RETURN CODE=8

```

To prevent truncation, the UNLOAD statement should be modified as follows:

```

UNLOAD CCSID(1208) FROM TABLE USER.TCCSIDE (CE1 POSITION(*) VARCHAR(33))*

```

*Recommendation based on Table 25 (Result length of CCSID conversion) in the SQL Ref

When unloading data, if conversion is involved, Table 25 in the V8 SQL reference should be consulted and lengths adjusted accordingly to ensure that truncation problems do not occur.

In this case, 'Hegelstraße' expands from 11 bytes in EBCDIC to 12 bytes in Unicode UTF-8.

Also note that had 'Jürgen' been inserted as 'Jürgen ' (with 5 pad characters), that the unload of that row would also have failed. Unlike in SQL, where trailing blanks are insignificant, with the UNLOAD utility, trailing blanks are significant. The UNLOAD utility can remove trailing blanks if you use the STRIP clause.



UNLOAD Example

```
CREATE TABLE TCCSIDE (CU1 CHAR(1)) CCSID UNICODE;  
INSERT INTO TCCSIDE VALUES ('A');
```

```
UNLOAD EBCDIC FROM TABLE USER1.TCCSIDU
```

```
DSNU000I      DSNUGUTC - OUTPUT START FOR UTILITY, UTILID =  
              AOAD65710  
DSNU1044I      DSNUGTIS - PROCESSING SYSIN AS EBCDIC  
DSNU050I      DSNUGUTC - UNLOAD EBCDIC  
DSNU1250I )    DSNUULIA - ERROR IN CCSID TRANSLATION FOR  
              VERIFICATION, FROM CCSID 1208 TO 65534  
DSNU650I )    DSNUGMS - FROM TABLE ADMF001.TCCSIDU  
DSNU012I      DSNUGBAC - UTILITY EXECUTION TERMINATED, HIGHEST  
              RETURN CODE=8
```

To prevent truncation, the UNLOAD statement should be modified as follows:
UNLOAD CCSID(0037,0037) FROM TABLE USER1.TCCSIDU

In this example, the UNLOAD utility is trying to convert UTF-8 data to the MIXED EBCDIC CCSID. Since this system is a MIXED=NO system, there isn't a MIXED EBCDIC CCSID, so the reserved value of 65534 is used. To prevent this issue, explicit specification of CCSIDs should be used. If UTF-16 data is also involved, a similar issue could involve, and the UNLOAD statement should be changed to specify CCSID(0037,0037,0037).

REPAIR, DSN1COPY, and DSN1PRNT

- REPAIR

- ▶ LOCATE KEY, VERIFY DATA, and REPLACE DATA Keywords

- No conversions occur on character constants.
 - If table is ASCII or Unicode, hexadecimal constants must be used

Example on an ASCII or Unicode table use:

```
REPAIR OBJECT LOCATE TABLESPACE DSN8D81A.DSN8S81D
```

```
  PAGE X'02' VERIFY OFFSET 50 DATA X'313233' - Not '123'
```

```
  REPLACE OFFSET 50 DATA X'343536' - Not '456'
```

```
  DUMP OFFSET 50 LENGTH 4
```

- DSN1COPY and DSN1PRNT

- ▶ Option to PRINT data in ASCII, EBCDIC, and Unicode

- Explicit or Implicit specification
 - Default is EBCDIC if the first page of the input data set is not a header page

REPAIR can repair ASCII and Unicode data, as well as EBCDIC data, however, you must specify the data using hexadecimal constants, not literal constants.

IDUG® 2006 - North America

Session C02

CCSID 201 – What's a CCSID and why do I care?

Christopher J. Crone

IBM/DB2 for z/OS Development

cjc@us.ibm.com

GoFurther



References

DB2 UDB for z/OS Version 8:

Everything You Ever Wanted to Know , ... and More - SG24-6079

DB2 UDB for z/OS Internationalization Guide

<http://www.ibm.com/software/data/db2/zos/pdf/ccmstr.pdf>

DB2 Universal Database Administration Guide - SC09-2946

Appendix E - National Language Support

The Unicode Standard Version 4.0

The Unicode Consortium - Addison-Wesley - www.unicode.org

Character Data Representation Architecture: Reference & Registry


SC09-2190

National Language Design Guide Volume 2 - SE09-8002

eBusiness Globalization Solution Design Guide, Getting Started

SG24-6851-00

In appendix E of the DB2 UDB Admin Guide, there is a discussion of NLS issues and how to set/override the codepage at the client using the codepage keyword on NT, and LANG variable on AIX.

IBM Software Group | DB2 Information Management Software 


Appendix – z/OS Support for Unicode

z/OS support for Unicode (V7 & V8) - Conversion Services

Documentation :

Manual: *z/OS: Support for Unicode(TM): Using Conversion Services* (sc33-7050)
Additional configuration in information APARs I113048, I113049, and I113277
Requires OS/390 V2R8 and above + APAR OW44581
code and program directory
<http://www6.software.ibm.com/dl/os390/unicodespt-p>
documentation
<http://publibfp.boulder.ibm.com/pubs/pdfs/os390/cunpde00.pdf>
<http://publibfp.boulder.ibm.com/pubs/pdfs/os390/cunuge00.pdf>
Information APAR I113048 and I103049

z/OS Conversion Services (64 Bit enabled) (V8)
Requires z/OS V1R2 and above + OW56703 and OW56704
Documentation:
Pointers to new documentation contained in OW56703 and OW56704



The z/OS support for Unicode now offers

- Conversion
- Normalization
- Casing

These services are used by DB2 and are also available to Assembler and C/C++ applications.