

New PL/I and C/C++ releases

Spring 2015

Peter Elderon

elderon@us.ibm.com

Insert
Custom
Session
QR if
Desired.



SHARE is an independent volunteer-run information technology association that provides **education, professional networking and industry influence.**



Enterprise PL/I 4.5



#SHAREorg



SHARE is an independent volunteer-run information technology association that provides **education, professional networking and industry influence.**

Insert
Custom
Session
QR if
Desired.



Enterprise PL/I 4.5 highlights

- Improved performance
- Enhanced middleware support
- Increased string length limit
- Introduced support for JSON
- Increased functionality
- Added features to enforce code quality
- Satisfied 28 RFE's

Improved performance



#SHAREorg



SHARE is an independent volunteer-run information technology association that provides **education, professional networking and industry influence.**

Insert
Custom
Session
QR if
Desired.



ARCH(11)

- Under ARCH(11), the compiler will exploit
 - ▶ The new Load Halfword Immediate on Condition instruction
 - ▶ The new vector hardware instructions (and registers)
- The latter significantly improves the performance of the code generated for (MEM)SEARCH and (MEM)VERIFY of CHAR and WIDECHAR strings
- The vector facility consists of:
 - ▶ Support instructions
 - ▶ Integer instructions
 - ▶ String instructions
 - ▶ Float instructions

ARCH(11)

- PL/I currently exploits only the support and string instructions
- In particular, to implement SEARCH and VERIFY, these the vector instructions are very useful
 - ▶ find_any_element_equal
 - ▶ find_any_element_not_equal
 - ▶ string_range_compare

ARCH(11)

- For example, this simple code that tests if a UTF-16 string is numeric

```
wnumb: proc( s );
```

```
    decl s    wchar(*) var;  
    decl n    wchar value( '0123456789' );  
    decl sx   fixed bin(31);
```

```
    sx = verify( s, n );  
    if sx > 0 then ...
```

- Is done with an expensive library call with ARCH ≤ 10

ARCH(11)

- With ARCH(11), the vector instruction facility is used to inline it as

E700	E000	0006		VL	v0,+CONSTANT_AREA(,r14,0)
E740	E010	0006		VL	v4,+CONSTANT_AREA(,r14,16)
			@1L2	DS	0H
A74E	0010			CHI	r4,H'16'
4150	0010			LA	r5,16
B9F2	4054			LOCRL	r5,r4
B9FA	F0E2			ALRK	r14,r2,r15
E725	E000	0037		VLL	v2,r5,_shadow1(r14,0)
E722	0180	408A		VSTRC	v2,v2,v0,v4,b'0001',b'1000'
E7E2	0001	2021		VLGV	r14,v2,1,2
EC5E	000D	2076		CRJH	r5,r14,@1L3
A74A	FFF0			AHI	r4,H'-16'
A7FA	0010			AHI	r15,H'16'
EC4C	000E	007E		CIJNH	r4,H'0',@1L4
A7F4	FFE5			J	@1L2
0700				NOPR	0
			@1L3	DS	0H

ARCH(11)

- You can learn a lot more about the vector facility at this talk

Session 16897: IBM z Systems z13 Vector Extension Facility (SIMD)

Tuesday, March 3: 4:30 PM - 5:30 PM

Other performance improvements

- Much faster code is now generated for MOD and REM of large FIXED DEC
- Previously, calls to a library routine were used for this
- Now inline code using DFP makes the calculation much faster
- ARCH(11) is not required for this

Other performance improvements

- ARCH(11) not required is also not required for improvements to
- INLIST of CHAR(1)
- BETWEEN for CHAR(1) and WCHAR(1)
- SEARCH and VERIFY of WCHAR(1)

Other performance improvements

- A SELECT statement of the form

```
select( x );  
  when( '..' ) ..  
  when( '..' ) ..  
  ...
```

- Was turned into a (fast) branch table if x was CHAR(1) and into a (slow) series of string compares if x had length > 1
- Now it will also be turned into a branch table if x is CHAR(2) or CHAR(4)

Other performance improvements

- One user has some code with this SELECT statement

```
SELECT(PLAUS.PLZ);  
  WHEN('0000') ...  
  WHEN('9999') ...  
  WHEN('1000') ...  
  WHEN('1004') ...
```

- With more than 2500 WHEN clauses
- With 4.4, it takes 90 seconds to compile under OPT(2)
- With 4.5, it takes 5 seconds, and the generated code is much better, too!

EXEC CICS statements

- The code generated for every EXEC CICS statement consists of a call to a CICS entry point with a first parameter that is an unprintable character string often longer than 100 bytes
- This string encodes for CICS what the statement is requesting
- Since the call appears to be an ordinary call to the compiler, it allocates a temporary on the stack and generates a MVC instruction to copy the 100+ bytes from the constant area to that temporary

EXEC CICS statements

- The 4.5 compiler now marks the first parameter of the CICS entry point with the INONLY attribute
- This eliminates the allocation of the temporary on the stack and the (slow) MVC instruction to copy to it
- This means the code will run faster, and your DSA will be smaller too

Enhanced middleware support



#SHAREorg



SHARE is an independent volunteer-run information technology association that provides **education, professional networking and industry influence.**

Insert
Custom
Session
QR if
Desired.



Structures as indicator variables

- Previously, if you wanted to use indicator variables with a structure in an EXEC SQL statement, you had to name each element of the structure and an associated indicator variable. For example, given

```
dc1 1 h3, 2 a fixed dec, 2 b fixed dec, ..., 2 z fixed dec;
```

- You had to code a matching structure and then name everything

```
dc1 1 in3, 2 a fixed bin(15), 2 b fixed bin(15), ..., 2 z fixed dec;
```

```
exec sql insert into mytable
```

```
values( :h3.a:in3.a, :h3.b:in3.b, ..., :h3.z:in3.z );
```

Structures as indicator variables

- Not fun and not easy to maintain or enhance. But now you can use a structure as indicator variable. So, given

```
dc1 1 h3, 2 a fixed dec, 2 b fixed dec, ..., 2 z fixed dec;
```

- you can use the matching indicator structure and name only the structures

```
dc1 1 in3, 2 a fixed bin(15), 2 b fixed bin(15), ..., 2 z fixed dec;
```

```
exec sql insert into mytable
```

```
values( :h3:in3 );
```

Structures as indicator variables

- But it gets better: with the new INDFOR attribute (it's like LIKE except the copied names all get the FIXED BIN(15) attribute), the matching indicator structure is easy to declare. So, given

```
dc1 1 h3, 2 a fixed dec, 2 b fixed dec, ..., 2 z fixed dec;
```

- you can use INDFOR and code simply

```
dc1 1 in3 indfor h3;  
exec sql insert into mytable  
values( :h3:in3 );
```

Structures as indicator variables

- And multi-row fetch (and dimacross) work here, too! So, given

```
dc1 1 h3(3) dimacross, 2 a fixed dec, 2 b fixed dec, ..., 2 z fixed dec;
```

- you can use INDFOR and code the very simple

```
dc1 1 in3(3) dimacross indfor h3;
```

```
exec sql insert into mytable  
values( :h3:in3 );
```

Named constants as host variables

- You can now use named constants as SQL host variables if
 - ▶ DB2 allows a simple, unnamed constant at that place in the EXEC SQL statement
- And
 - ▶ the named constant has either the attribute
 - CHARACTER, in which case its VALUE attribute must specify a character string
 - ▶ or
 - FIXED, in which case its VALUE attribute must specify a decimal number or an expression that can be reduced to an integer constant

Statement validation

- Previously, when the SQL preprocessor scanned an EXEC SQL statement, it would report only the first error in the statement
- Now, it will report all the errors in every EXEC SQL statement

SQL CODEPAGE option

- Using the new SQL preprocessor option (NO)CODEPAGE, you can control the preprocessor's use of the compiler's CODEPAGE option when it sets the CCSID of a host character variable
- When CODEPAGE is in effect, the compiler's CODEPAGE option is always used as the CCSID for SQL host variables of character type
- When NOCODEPAGE is in effect, the compiler's CODEPAGE option is used as the CCSID for SQL host variables of character type only if the SQL preprocessor option NOCCSID0 is also in effect
- NOCODEPAGE is the default for compatibility with previous releases

SQL WARNDECP option

- Using the new SQL preprocessor option (NO)WARNDECP, you can control whether the SQL preprocessor issue a warning message if it uses the DB2-provided DSNHDECP module
- `PP(SQL('WARNDECP'))` matches what the previous releases did
- But NOWARNDECP is now the default (since this message is almost meaningless to most users and hence is just distracting noise)

Less noise from the SQL preprocessor

- Message DSNH4760I is now suppressed – this shows up as

**IBM3024I I DSNH4760I DSNHPSRV The DB2 SQL Coprocessor is
using the level 2 interface under DB2 V9**

- This is almost always uninteresting
- This has been removed for 4.3 as well as 4.4

Increased string length limit



#SHAREorg



SHARE is an independent volunteer-run information technology association that provides **education, professional networking and industry influence.**

Insert
Custom
Session
QR if
Desired.



From 32K to 128M

- Previously no string could be 32K or longer, and VARYING strings had only a 2-byte length prefix
- The new VARYING4 attribute will let you declare strings as having a 4-byte length prefix
- The new STRING suboption of the LIMITS option will let you set the threshold for string lengths to be 32K, 512K, 8M, or 128M

From 32K to 128M

- The maximum length is one less than the threshold
- The default is 32K
- The threshold values may be specified using K or M suffices or as decimal numbers, i.e. as 32K or 32768 or as 128M or 134217727
- But ...

From 32K to 128M

- A threshold bigger than 32K may be specified only if the CMPAT(V3) option is also in effect – because an expanded string descriptor is needed
- And mixtures of code compiled with CMPAT(V2) and CMPAT(V3) face the same restrictions as mixtures of code compiled with CMPAT(V1) and CMPAT(V2) faced 30-years ago
- Fwiw, the new VARYING4 attribute may be used with CMPAT(V2) - but then the length must still be less than 32K

Long string considerations

- Specifying `LIMITS(STRING(n))` where $n > 32K$ may also greatly increase the amount of stack storage used by your code
- For example, in `PUT LIST(A || B)` where A and B are `CHAR(*) VARYING4`, the compiler will allocate a temporary on the stack equal to the maximum string size – so under `LIMITS(STRING(128M))` this would be 128M off the stack!!
- The `MAXTEMP` compiler option will alert you to such statements

Long string considerations

- The STRING limit applies to all kinds of strings: BIT, CHAR, GRAPHIC and WIDECHAR
- The VARYING4 attribute is also supported for all kinds of strings
- ALIGNED VARYING strings are halfword-aligned – ALIGNED VARYING4 strings will be fullword-aligned

CMPAT(V3) considerations

- The descriptors generated under CMPAT(V3) are different than those generated under CMPAT(V2) – and not just for strings
- The offsets in structure descriptors and the bounds etc in array descriptors are all 8-byte integers

Support for JSON



#SHAREorg



SHARE is an independent volunteer-run information technology association that provides **education, professional networking and industry influence.**

Insert
Custom
Session
QR if
Desired.



JSON overview

- A series of built-in functions provide the ability to
 - ▶ Generate JSON text
 - ▶ Parse JSON text
 - ▶ Validate JSON text

JSON generation

- All JSON written out will be in UTF-8 with the compiler and library handling any necessary conversions from EBCDIC
- a series of "put" functions are provided and all have a buffer address and buffer length as their first 2 arguments, and all return the number of bytes written
- attempts to write variables containing data types incompatible with JSON will be flagged at compile time
- escaped characters will be created as needed

JSON parsing

- All JSON to be parsed must be in UTF-8 with the compiler and library handling any necessary conversions to EBCDIC
- a series of “get” functions are provided and all have a buffer address and buffer length as their first 2 arguments, and all return the number of bytes read
- attempts to read variables containing data types incompatible with JSON will be flagged at compile time
- whitespace characters will be skipped over when found

JSON

- For example, suppose we have this sample JSON text

```
{ "passes" : 3,  
  "data" :  
  [  
    { "name" : "Mather", "elevation" : 12100 }  
    , { "name" : "Pinchot", "elevation" : 12130 }  
    , { "name" : "Glenn", "elevation" : 11940 }  
  ]  
}
```

- And that it is in a buffer at address p and of length n

JSON

- If we had a corresponding PL/I structure

dc1

```
1 info
2 passes fixed bin(31),
2 data(3),
3 name char(20) varying,
3 elevation fixed bin(31);
```

- Then *jsonGetValue(p, n, info)* will by itself fill in the whole structure
- This one simple function call would do all the work for you

JSON

- But if we did not know how many data instances we would get, our PL/I structure might instead look like

dc1

```
1 info based(q)
2 count fixed bin(31),
2 data( passes refer(count) ),
3 name char(20) varying,
3 elevation fixed bin(31);
```

- And it would have to be dynamically allocated – but this is still easy:

JSON

- Four built-in references would suffice:

```
rd = jsonGetObjectStart(p,n);          /* read over {          */
rd = jsonGetMember(p+rd,n-rd,passes);  /* read "passes":3 and assign it */
allocate info;
rd = jsonGetComma(p+rd,n-rd);          /* read over ,          */
rd = jsonGetValue(p+rd,n-rd,info.data); /* read "data" ... and assign it */
```

- And this works, of course, no matter how much whitespace is present

Increased functionality



#SHAREorg



SHARE is an independent volunteer-run information technology association that provides **education, professional networking and industry influence.**

Insert
Custom
Session
QR if
Desired.



INLIST

- This built-in function is useful in determining if a value belongs to a set of values and allows you to put a SELECT in the middle of an IF
- It requires a minimum of 3 arguments and accepts a maximum of 64
- `INLIST(x, a, b, c, ...)` is equivalent to `(x = a) | (x = b) | (x = c) ...`
- All the arguments must have computational type
- The compiler will optimize this when possible

INLIST

- The arguments can be numbers, strings, or arbitrary expressions
 - ▶ if inlist(stadt, 'Berlin', 'Bern', 'Rom', 'Wien') then
 - ▶ if inlist(uppercase(stadt), 'BERLIN', 'BERN', 'ROM', 'WIEN') then
 - ▶ if inlist(x, 2, y, 7, z) then
 - ▶ if inlist(b, b1 | b2, '1100'b, b3 & b4) then
- These would all be converted to a short-circuited series of compares

INLIST

- But if the first argument is “nice” and the rest are all similar values, then the compiler will turn the inlist reference into a branch table. For example,
 - ▶ `inlist(x, 2, 3, 5, 7, 11, 13, 17, 19)`
- would become a branch table if `x` is `FIXED BIN(p,0)` with `p <= 31` or if `X` is `FIXED DEC(p,0)` with `p <= 9`
- The values 2, 3, 5, etc don't have to be literals – they can be named constants (VALUEs) or restricted expressions
- And if all are `CHAR(1)`, a simple table look-up is generated

INLIST

- Branch tables would also be built if
 - ▶ X is BIT(n) with $1 \leq n \leq 16$ and the other arguments are bit constants
 - ▶ X is CHAR(1) and the other arguments are CHAR(1) constants
- Again the second and subsequent arguments don't have to be literals – they can be named constants (VALUES) or restricted expressions

BETWEEN

- This built-in function is useful in determining if a value is in an interval
- It requires exactly 3 arguments
- `BETWEEN(x, a, b)` is equivalent to `(x >= a) & (x <= b)`
- All the arguments must be ordinals or have real numeric type
- The compiler will optimize this when possible
 - ▶ For example, if `x`, `a`, and `b` are all `FIXED BIN(p,0)` with `p <= 31`, then the compiler will turn `BETWEEN(x, a, b)` into one comparison (not two!)
 - ▶ `OORDINAL`, `CHAR(1)`, and `WCHAR(1)` are optimized in the same way

NULLENTRY

- This built-in function allows you to initialize an entry variable with a null value – including static variables
- `Dcl function_pointer limited entry static init(nullentry());`
- You can also use it to test an entry value to see if it is null

PLISTCK, PLISTCKE, and PLISTCKF

- These built-in subroutines will generate the corresponding instructions
- PLISTCK(x) sets an UNSIGNED FIXED BIN(64)
- PLISTCKE(x) sets a CHAR(16) NONVARYING
- PLISTCKF(x) sets an UNSIGNED FIXED BIN(64)

- Each returns a FIXED BIN(31) that is the condition code set by the corresponding hardware instruction

SMFTOJULIAN and JULIANTOSMF

- These built-in functions convert between the Julian and SMF date formats
- SMFTOJULIAN(d) converts a CHAR(4) SMF date to a CHAR(7) YYYYDDDD
- JULIANTOSMF(d) converts a CHAR(7) YYYYDDDD to a CHAR(4) SMF
- No error checking is done at run-time for these functions – the conversions are done in-line, and the input data must be valid

PLISAX

- These built-in functions now allow the event function pointers to be null
- When null, the event will not be called
- This allows you to write smaller, faster XML parsing code
- It requires a library PTF (but not the 4.5 compiler)

REINIT

- This statement allows you to reset a variable with its INITIAL values

- The variable must be level-one, unsubscripted with storage class
 - ▶ AUTOMATIC
 - ▶ BASED
 - ▶ CONTROLLED
 - ▶ STATIC

SYSDIMSIZE

- This preprocessor built-in function returns the number of bytes needed to hold the largest array bound
- Under CMPAT(V1), SYSDIMSIZE returns 2
- Under CMPAT(V2), SYSDIMSIZE returns 4
- Under CMPAT(V3), SYSDIMSIZE returns 8

SYSPOINTERSIZE and SYSOFFSETSIZE

- These preprocessor built-in functions return the size (in bytes) of a **POINTER** and an **OFFSET**
- As of now, they always return a 4
- But when 64-bit code is supported they could return an 8, and they will be useful in writing code that will compile and run correctly both in 32-bit and in 64-bit mode

SYSPONTERSIZE and SYSOFFSETSIZE

- For example, you could use syspointersize to declare C's malloc

```
%if syspointersize = 8 %then %do;  
    define alias size_t    fixed bin(63);  
%end; %else %do;  
    define alias size_t    fixed bin(31);  
%end;
```

```
dc1 malloc    ext('malloc')  
              entry( type size_t byvalue )  
              returns( byvalue pointer )  
              options( linkage(optlink) nodescriptor );
```

- And this would be correct for 32- and 64-bit

XML compiler option

- This option now supports an XMLATTR suboption with suboptions of APOSTROPHE or QUOTE
- It determines whether XML attributes are enclosed as '...' or “....”
- It requires a library PTF (as well as the 4.5 compiler)

Quotes in pictures

- The apostrophe symbol is now accepted in PICTURE specifications
- With the same usage as the comma or period symbol
- This requires only a library PTF

Added features to enforce code quality



#SHAREorg



SHARE is an independent volunteer-run information technology association that provides **education, professional networking and industry influence.**

Insert
Custom
Session
QR if
Desired.



RULES(NOLAXRETURN)

- Previously RULES(NOLAXRETURN) caused the compiler to generate code to raise the ERROR condition if
 - ▶ RETURN; was hit in a PROC with the RETURNS attribute
 - ▶ RETURN(...); was hit in a PROC without the RETURNS attribute
- Now the options has been enhanced so that the compiler will raise the ERROR condition if code falls through to the END statement in a PROC with the RETURNS attribute

RULES(NOLAXRETURN)

- This makes it much easier to detect the lack of the desired statement after the else in this segment of real customer code

```
func: proc( . . . ) returns( fixed bin(31) );
```

```
. . .
```

```
if flags = 'b' then;
```

```
else
```

```
    return(0);
```

```
end;
```

- This problem would probably now be resolved within minutes rather than days it took to resolve it previously via the PMR process

MAXBRANCH option

- The new MAXBRANCH compiler option lets you find blocks (PROCEDURES and BEGIN-blocks) that are perhaps too complex. More precisely, it flags any block that has too many conditional branches
- A statement of the form "if a then ...; else ..." adds 1 to the total number of branches in its containing block, and a statement of the form "if a = 0 | b = 0 then ..." adds 2.
- SELECT statements and conditional DO loops also to the total
- The default is MAXBRANCH(2000)

NONASSIGNABLE option

- The new NONASSIGNABLE suboption of the DEFAULT compiler option now supports INONLY and STATIC as suboptions
- NONASSIGNABLE(INONLY) specifies that parameters declared with the INONLY attribute are given the NONASSIGNABLE attribute.
- NONASSIGNABLE(STATIC) specifies that STATIC variables are given the NONASSIGNABLE attribute.
- These suboptions have no effect on variables that are explicitly given the ASSIGNABLE or NONASSIGNABLE attribute or on structure members that have inherited the ASSIGNABLE or NONASSIGNABLE attribute from a parent where it was explicitly specified.

NONASSIGNABLE option

- BYVALUE parameters are given the INONLY attribute after the resolution of the (NON)ASSIGNABLE attribute, and hence the NONASSIGNABLE(INONLY) suboption has no effect on BYVALUE parameters (unless, of course, they are explicitly given the INONLY attribute).
- To specify that both STATIC and INONLY variables are to be given the NONASSIGNABLE attribute, then you must specify the suboption NONASSIGNABLE(STATIC INONLY).
- The NONASSIGNABLE attribute may be specified without any suboptions, in which case it has the same meaning as in previous releases, namely NONASSIGNABLE(STATIC).

RULES(NOLAXQUAL)

- The NOLAXQUAL suboption of RULES now accepts another choice of suboptions:
- Under the default NOLAXQUAL(ALL), the compiler will flag all violations of the qualifications rules (either STRICT or LOOSE)
- But under NOLAXQUAL(FORCE), the compiler will flag only those violations when the element belongs to a structure with the new FORCE(NOLAXQUAL) attribute
- This gives you the ability to enforce mandatory qualification on a structure-by-structure basis

RULES(NOLAXNESTED and NOPADDING)

- The new ALL | SOURCE suboptions to the RULES(NOLAXNESTED) and RULES(NOPADDING) compiler options provide finer control over when the compiler flags questionable coding
- Under ALL, all violations are flagged
- Under SOURCE, only those violations that occur in the primary source file are flagged
- This matches what is already supported for RULES(NOUNREF)
- For each of these options, ALL is the default

DEPRECATE(NEXT) in the MACRO preprocessor

- DEPRECATE and DEPRECATENEXT are now accepted as MACRO preprocessor options.
- They currently have only one suboption: ENTRY
- Any use of macros with the names specified in DEPRECATE(ENTRY(...)) will be flagged with an E-level message (or W-level for DEPRECATENEXT)

FILEREF option

- The new FILEREF compiler option lets you get an even smaller listing:
- Under NOFILEREF, if the compilation produces no messages, then the file reference table is not generated into the listing
- For compatibility, the default is NOFILEREF

SUBSCRIPTRANGE checking

- Previously in an array assignment, the array bounds were not checked (except at compile-time if they were constants) even if SUBRG was enabled. So, given

```
dc1 a(1:x)    fixed bin ct1;  
dc1 b(1:y)    fixed bin ct1;
```

- SUBRG would have been raised in the loop, but not in the first assignment

```
x = 2; y = 4;  
alloc a,b;  
a = b;  
do jx = 1 to hbound(a);  
    a(jx) = b(jx);  
end;
```

SUBSCRIPTRANGE checking

- Now when the array bounds in an array assignment are not constant, if SUBRG is enabled, then they will be checked once before the array assignments occur. So, given

```
dc1 a(1:x,1:z)    fixed bin ct1;  
dc1 b(1:y,1:z)    fixed bin ct1;
```

- SUBRG would now be raised by either of the array assignments below

```
x = 2; y =4; z = 6;  
alloc a,b;  
a = b;  
a(*,1) = b(*,1);
```

ATTRIBUTES listing

- If a FIXED or FLOAT variable has the VALUE attribute, then its value will now be included in the ATTRIBUTES listing. So, given

```
dc1 a  fixed dec(3) value(17);  
dc1 b  fixed dec(3) value(47);  
dc1 c  fixed dec(9,7) value( b/a );  
dc1 d  float bin(21) value( acos(-1) );
```

- The ATTRIBUTES listing now shows

```
A          VALUE ( 17 ) FIXED DEC(3,0)  
B          VALUE ( 47 ) FIXED DEC(3,0)  
C          VALUE ( 2.7647058 ) FIXED DEC(9,7)  
D          VALUE ( 3.141593E+00 ) FLOAT BIN(21)
```

AGGREGATE listing

- Arrays of VARYING (or VARYING4) strings may contain padding if they are ALIGNED since each string must halfword (or fullword) aligned. So, given

```
dc1 a(10) char(31) varying aligned;
dc1 b(10) char(30) varying4 aligned;
```

- The AGGREGATE listing now shows

Aggregate Length Table

Line	File	Dims	Offset	Total Size	Base Size	Identifier
3.0		1	0	339	33	A /* padding */
4.0		1	0	358	34	B /* padding */

Compatibility



#SHAREorg



SHARE is an independent volunteer-run information technology association that provides **education, professional networking and industry influence.**

Insert
Custom
Session
QR if
Desired.



LIMITS(FIXEDBIN)

- The default has changed from
 - ▶ LIMITS(FIXEDBIN(31))
- to
 - ▶ LIMITS(FIXEDBIN(31,63))
- The new default was already required to use the SQL preprocessor and to use various other functions
- However, if you have not previously made this change, it can cause some differences in how some (esoteric) code runs

CMPAT(V3)

- All code using CMPAT(V3) must be recompiled
- Fortunately, there was no reason to use this option before this release

XL C/C++ V2R1M1



#SHAREorg



SHARE is an independent volunteer-run information technology association that provides **education, professional networking and industry influence.**

Insert
Custom
Session
QR if
Desired.



XL C/C++ V2R1M1 highlights

- Support for the new hardware
- Asm and hardware model support
- Two new high performance libraries
- Improved functionality

Support for the new hardware



#SHAREorg



SHARE is an independent volunteer-run information technology association that provides **education, professional networking and industry influence.**

Insert
Custom
Session
QR if
Desired.



ARCH(11)

- To enable the compiler with the z13 feature we are introducing new sub options for ARCHITECTURE and TUNE options
- ARCH(11) will prompt the compiler to utilize the list of hardware instructions new in z13
- TUNE(11) will prompt the compiler to improve run time performance according to the micro-architecture characteristics of z13

Support for Load Halfword Immediate on Condition

- The new instructions LOCHI and LOCGHI are like LHI and LGHI except that the load depends on the conditional code and the mask in the instruction
- The immediate value, as indicated in the name of the instruction, is halfword in size, so this applies only to values that fit in 16 bits
- The compiler will exploit this under OPTimize and ARCH(11)

Support for the Vector Facility

- Vector programming support provides programmers direct access to SIMD instructions from the z13 Vector Facility for z/Architecture
- Language extensions based on the AltiVec Programming Interface specification with suitable changes and extensions

Support for the Vector Facility

- A new VECTOR option
- macro: `__VEC__`
- Vector data types:
 - {vector, `__vector`} {bool, signed, unsigned} {char, short, int, long long}
 - {vector, `__vector`} double
- Various language extensions on the vector types for natural usage just like for the native types
 - Ex.
 - Assignment operator (=)
 - Address operator (&)
 - Pointer arithmetic
 - Unary operators (++ , -- , + , - , ~)
 - Binary operators (+ , - , * , / , % , & , | , ^ , << , >>)
 - Relational operators (== , != , < , > , <= , >=)

Support for the Vector Facility

- Comprehensive set of vector built-in functions for access and manipulation of individual vector elements
 - In the following high-level categories:
 - Arithmetic
 - Compare
 - Compare ranges
 - Find any element
 - Gather and scatter
 - Generate Mask
 - Isolate Zero
 - Load and Store
 - Logical
 - Merge
 - Pack and unpack
 - Replicate
 - Rotate and shift
 - Rounding and conversion
 - Test
 - All Predicates
 - Any Predicates

Support for the Vector Facility

- See z/OS XL C/C++ V2R1M1 programming guide, Ch. 35 Using vector programming support for more details
 - <http://publibz.boulder.ibm.com/epubs/pdf/cbc1p201.pdf>

Support for the Vector Facility

- Example:

```
#include <builtins.h>
#include <stdio.h>
int main() {
    vector signed int a = {-1, 2, -3, 4}; // declare and initialize a vector with 4 signed integer elements
    vector signed int b = {-5, 6, -7, 8};
    vector signed int c, d;           // declare vectors with 4 signed integer elements

    c = a + b;           // Generates VAF
    d = vec_abs(c);     // Generates VLPF

    printf("d[0] = %d\n",d[0]); // prints 6 -- d[0] extract the 1st element from the vector
    printf("d[1] = %d\n",d[1]); // prints 8
    printf("d[2] = %d\n",d[2]); // prints 10
    printf("d[3] = %d\n",d[3]); // prints 12

    return 0;
}
```

One more reminder:

- You can learn a lot more about the vector facility at this talk

Session 16897: IBM z Systems z13 Vector Extension Facility (SIMD)

Tuesday, March 3: 4:30 PM - 5:30 PM

Asm and hardware model support



#SHAREorg



SHARE is an independent volunteer-run information technology association that provides **education, professional networking and industry influence.**

Insert
Custom
Session
QR if
Desired.



Ability to insert asm statements

- Clients have asked us to support HLASM statements for all C or C++ programs and not just with Metal C
- With V2R1M1 release we are rolling in support for HLASM statements in LE enabled C or C++ programs as well
- Now you can inline hardware instructions in your C or C++ program using the `__asm` keyword

Ability to insert asm statements

- New options and sub-options
 - **ASM**: causes `__asm`, `__asm__` to be statements
 - **Keyword(ASM)**: gives `asm` the same semantic as `__asm`
 - **ASMLIB**: specifies the macro libraries to be used when assembling the inline assembler source code
- New messages will be reported by the compiler under message CCN1148

Mixing ARCH levels in one compile

- Before, you had to put code using different ARCH levels into separate functions – and you had to pay for the expense of calling the appropriate function
- You can now have code using different ARCH levels within one compilation unit

Mixing ARCH levels in one compile

- Now this code can be inlined with the new
 - **#pragma arch_section(<architecture>)**
- The pragma indicates the start of a section of the source intended for the machine indicated by <architecture>
- The compiler switches to architecture specified, and at the end the section switches back to the previous architecture

Identifying the hardware model and features

- Programs may need to check the machine model before doing some processing
- Three new builtins help with this:
 - `__builtin_cpu_init(void)`
 - Runs the CPU detection code, and saves the CPU information in a compiler defined/managed buffer
 - Must be called at least once before either of the following is invoked

Identifying the hardware model and features

- `__builtin_cpu_is(const char* cpumodel)`
 - Returns 1 if the CPU is of type *cpumodel* (“5”, “6”, ..., “11”)

- `__builtin_cpu_supports(const char* feature)`
 - Returns 1 if the CPU supports one of the features indicated

 - *feature* values are
 - `"longdisplacement", "etf2", "etf3", "dfp", "prefetch", "storeclockfast", "loadstoreoncond", "popcount", "interlocked", "tx", "dfpzoned", "vector128", "5", ..., "11"`

New high performance libraries



#SHAREorg



SHARE is an independent volunteer-run information technology association that provides **education, professional networking and industry influence.**

Insert
Custom
Session
QR if
Desired.



New high performance math library

- MASS (Mathematical Acceleration Sub-System) library
- A comprehensive set of elementary/special mathematical functions (e.g. exp, log, sin, etc.) tuned for high performance on zEC12 and z13
- 3 kinds of libraries:
 - 59 scalar functions
 - Easiest to use in existing code since names match existing runtime library functions
 - 77 vector functions
 - Generally provides the highest performance, provided `vector_length` is sufficient (approximately >2 to >10 depending on the function)
 - 8 SIMD (z13 only) functions
 - Convenient for code written to use vector datatypes and built-in functions
- Single- and Double-precision FP, in IEEE

MASS performance

- z/OS MASS (Mathematical Acceleration Subsystem) vector functions on z13 demonstrate **up to 2.9x higher throughput** than the corresponding z/OS V2.1 XL C/C++ runtime library math functions on zEC12.
- A key subset of MASS vector functions are heavily used in Analytics and are accelerated using SIMD instructions. These functions demonstrate **up to 6.8x higher throughput** than the corresponding z/OS V2.1 XL C/C++ runtime library math functions on zEC12.

Disclaimer

This claim is based on results from internal lab measurements. A subset of the MASS vector functions is accelerated using SIMD instructions on z13. The SIMD benefit is demonstrated using this subset. The performance improvements achieved will vary depending on the workload and other factors.

New high performance linear algebra library

- ATLAS (Automatically Tuned Linear Algebra Software) library
- A high-performance versions of all the BLAS (basic linear algebra subprograms) routines, and a subset of the LAPACK (linear algebra package) routines
- Tuned for high performance on zEC12/zBC12 and z13
- Both single and multi-threaded versions available, with IEEE floating point
- Supplied libraries
 - ATLAS main libraries
 - ATLAS specific variants of the BLAS, CBLAS, and LAPACK routines
 - CBLAS libraries
 - C interface versions of the BLAS routines
 - LAPACK libraries
 - C interface versions of the LAPACK routines
 - Fortran BLAS libraries
 - Fortran 77 interface versions of the BLAS routines
 - Supports 31-bit C linkage, 31-bit XPLINK, and 64-bit XPLINK

ATLAS performance

- A key subset of z/OS ATLAS 3.10.0 (Automatically Tuned Linear Algebra Software) double precision functions on z13 demonstrate **up to 44% higher throughput** than the corresponding functions on zEC12.
- Selected key z/OS ATLAS 3.10.0 (Automatically Tuned Linear Algebra Software) functions are accelerated using SIMD instructions. These functions demonstrate **up to 80% higher throughput** on z13 than the corresponding functions on zEC12.

Disclaimer:

This claim is based on results from internal lab measurements. The double precision function improvement is derived from comparisons of a select set of commonly used z/OS ATLAS 3.10.0 functions executing on z13 to the equivalent functions executing on zEC12. A subset of these functions is accelerated using SIMD instructions on z13. The SIMD benefit is demonstrated using this subset. The performance improvements achieved will vary depending on the workload and other factors.

Increased functionality



#SHAREorg



SHARE is an independent volunteer-run information technology association that provides **education, professional networking and industry influence.**

Insert
Custom
Session
QR if
Desired.



Improved make

- 'make' dependency file generation through -M did not include non-existent headers or allow named targets
- New make options address these and similar concerns:
 - -MT allows setting the dependent target name
 - -MQ is -MT but also escapes 'make' special characters for easier dependency file usage
 - -MG allows missing header files to be included in the dependency list
 - The q-makedep=pponly suboption will run the include preprocessing only
 - A dependency file will be generated
 - But no object code will be generated

Redefining macro values (C++)

- `LANGVL(REDEFMAC)` allows redefining a macro to a different value
- Without the option, an error is generated when a macro value is redefined
- With the option, a warning is emitted, and the macro value is redefined

Expanded user-reserved space on the Metal C stack

- This is an enhancement to an existing feature
- The option DSAUSER, available only with MetalC, allows the user to reserve a pointer on the stack
- However, the user may want to leave more space on the stack
- DSAUSER has been enhanced to allow you to specify how much space to reserve on the stack

–xlc -qDSAUSER=12

- This will reserve a space equal to the size of 12 words on the stack

Eliminating a null pointer check

- The C++11 ANSI Standard requires
 - checking the pointer returned from the placement new operator
 - and performing the initialization only when the pointer is not null
- The check for null pointer returned from other operators new and new[] is not required but performed
- New sub-option `LANGlvl (NOCHECKPLACEMENTNEW)` will remove the check of the pointer returned by the placement new operator
 - And can thus speed up the program by eliminating unnecessary null pointer checks

Questions?

- Connect with us
 - Email me at elder@us.ibm.com
 - Rational Café - the compilers user community & forum
 - C/C++: <http://ibm.com/rational/community/cpp>
 - PL/I: <http://ibm.com/rational/community/pli>
 - RFE community – for feature requests
 - C/C++:
http://www.ibm.com/developerworks/rfe/?PROD_ID=700
 - PL/I: http://www.ibm.com/developerworks/rfe/?PROD_ID=699
 - Product Information
 - C/C++: <http://www-03.ibm.com/software/products/us/en/czos>
 - PL/I: <http://www-03.ibm.com/software/products/en/plicompfam>

Thank You!



Learn more at:

- [IBM Rational software](#)
- [IBM Rational Software Delivery Platform](#)
- [Process and portfolio management](#)
- [Change and release management](#)
- [Quality management](#)
- [Architecture management](#)
- [Rational trial downloads](#)
- [developerWorks Rational](#)
- [IBM Rational TV](#)
- [IBM Rational Business Partners](#)

© Copyright IBM Corporation 2008. All rights reserved. The information contained in these materials is provided for informational purposes only, and is provided AS IS without warranty of any kind, express or implied. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, these materials. Nothing contained in these materials is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software. References in these materials to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in these materials may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. IBM, the IBM logo, the on-demand business logo, Rational, the Rational logo, and other IBM products and services are trademarks of the International Business Machines Corporation, in the United States, other countries, or both. Other company, product, or service names may be trademarks or service marks of others.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval