

David Dykstal

Architect, Rational Developer for Power
david_dykstal@us.ibm.com

Session: 480208
Agenda Key: 14FB

What's New with RPG & COBOL


XML, Extended I/O and More


ibm.com/power




- A bit of history
- Modern RPG
- Modern COBOL
- Tools for IBM i
- Summary









RPG IV

Performance

Productivity

Open and Modern

Mixed Environment Support

Maintainability

V5R1 - Free Form C Specs

- Free Form code allows you to indent the code to show the control structure
- Makes control form errors much more apparent
- Makes finding missing code easier

```
C      READ      orderFile
C      DOW       NOT %eof
C CUSTID CHAIN    custFile
C      IF        %found
C      CALLP     doOrder()
C      ENDIF
C      READ      orderFile
C      ENDDO
```

```
READ orderFile;
DOW NOT %eof;
  CHAIN CUSTID custFile;
  IF %found;
    doOrder();
  ENDIF;
  READ orderFile;
ENDDO;
```

Missing code for the
"record not found" case

V5R1 - Qualified Names

- These make the code clearer by explicitly showing which data structures are affected
 - Does %EOF(orderFile) check the end-of-file status for the READ to orderRec?
 - Is the DISCOUNT related to the order, or to the customer?
 - Were all the parameters in doOrder set by the READ or CHAIN operations?

```
READ orderRec;  
DOW NOT %eof(orderFile);  
  CHAIN custId custRec;  
  doOrder(custId: ordId:  
          desc: qty:  
          price: discount:  
          ordDate);  
  READ orderRec;  
ENDDO;
```

```
READ orderFile.orderRec order;  
DOW NOT %eof(orderFile);  
  CHAIN order.custId custFile cust;  
  doOrder(cust.id: order.id:  
          order.desc: order.qty:  
          order.price: cust.discount:  
          ordDate);  
  READ orderFile.orderRec order;  
ENDDO;
```

I ain't gonna do all that extra stinkin' typing!

You don't have to if you use Content Assist in RD Power.

V5R2 - Result Data Structures for I/O Operations

- More control over which records are used in these operations
 - Program has several versions of the Order record: old, updated, most-recent
 - Left side uses I and O specs.
 - The orderDs data structure has the fields from the Order record which are also in the screen record.
 - All the I/O uses the orderDs subfields.
 - Right side specifies the data structures directly on the I/O.
 - EVAL-CORR is used to transfer data from the order data structure to the screen data structure.
 - The second program may be easier to maintain because it is more clear how the data is flowing during the I/O operations.

```
CHAIN(N) key orderRec;  
oldDs = orderDs;  
EXFMT screenRec;  
updateDs = orderDs;  
CHAIN key orderRec;  
if oldDs <> orderDs;  
    ... record was changed during EXFMT  
else;  
    orderDs = updateDs;  
    update orderRec;  
endif;
```

```
CHAIN(N) key orderRec oldDs;  
EVAL-CORR screenDs = oldDs;  
EXFMT screenRec screenDs;  
EVAL-CORR updateDs = screenDs;  
CHAIN key orderRec newDs;  
if oldDs <> newDs;  
    ... record was changed during EXFMT  
else;  
    update updateDs;  
endif;
```

6.1 - Local Files in Subprocedures

- Limit where a file can be used. Prevent unexpected access.
 - The QPRINT printer file is only intended to be used within the print procedure. The original programmer wants to ensure that each line begins with a timestamp.
 - Another programmer has used the QPRINT file in the main procedure, against the rules stated in the comment for the file.
 - Prior to 6.1, the file had to be defined as a global file, so the compiler would allow the “illegal” red line using the file in the main procedure
 - In 6.1, the compiler will not allow the illegal line since the file is invisible to the main procedure.

```
* Note! Only use this file in
*           the print procedure
Fqprint o 132 printer

    line = 'Name is ' + name;
    write qprint line;

P print      b

msgDs = %timestamp + ':' + msg;
write qprint msgDs;
```

```
    line = 'Name is ' + name;
    write qprint line;

P print      b
Fqprint o 132 printer
msgDs = %timestamp + ':' + msg;
write qprint msgDs;
```



7.1 - Support for ALIAS names

- Background
 - Fields in externally described files can have a standard name up to 10 characters and an alternate (ALIAS) name up to 128 characters.
 - RPG III only allowed 6 characters, so many customers have files with cryptic names like CUSNAM, CUSADR. The files often have alternate names such as CUSTOMER_NAME and CUSTOMER_ADDRESS, that can be used in SQL queries.
- New ALIAS keyword for RPG
 - When ALIAS is specified, RPG will use the alternate name instead of the 10-character standard name.
 - Supported on F specs for any file that will not have Input or Output specs generated. Used for LIKERECD data structures.
 - Supported on D specs for any externally-described data structure.

```
A          R CUSTREC
A          CUSTNM      25A      ALIAS (CUSTOMER_NAME)
A          CUSTAD      25A      ALIAS (CUSTOMER_ADDRESS)
A          ID          10P 0

-----
D custDs          e ds          ALIAS
D                                     QUALIFIED EXTNAME(custFile)
/free
  custDs.customer_name = 'John Smith';
  custDs.customer_address = '123 Mockingbird Lane';
  custDs.id = 12345;
```

7.1 - %LEN(varying: *MAX)

- Problem solved by %LEN(*MAX):
 - Starting in 6.1 calculations involving the size of the varying-length prefix became error-prone because the size might be either 2 or 4 bytes in length.
 - In 6.1, %ADDR was enhanced with %ADDR(*DATA) to simplify getting the address of the data part.
 - Prior to 6.1 programmers added 2 to %ADDR.
 - %LEN(*MAX) simplifies getting the maximum length, in characters, of the data part.
 - Prior to 7.1 programmers subtracted 2 or 4 from %SIZE, and then divided by 2 for UCS-2 and DBCS.

```
D char_varying      S          5000A  VARYING
  /free
    // Pass address and length of the data
    conv (%addr(char_varying : *data)
          : %len(char_varying : *max));
```

7.1 - %SCANRPL

- Gets rid of error prone code that everyone has to write at least once.
 - Example: Replace all occurrences of '%NAME%' with the value of custName

```
msg = msgTemplate;  
p = %scan('%NAME%' : msg);  
dow p > 0;  
  msg = %replace(custName :  
                msg :  
                p :  
                6);  
  s = p + %len(custName);  
  if s + 5 > %len(msg);  
    leave;  
  endif;  
  p = %scan('%NAME%' : msg : s);  
enddo;
```

```
msg = %scanrpl ('%NAME%' :  
               custName :  
               msgTemplate);
```

◆ 7.1 PTF & 6.1 PTF - Exceptions on CCSID conversions

❖ Example

- The Thai word for "house" is "บ้าน".
- The Thai word for "house" is "■■■".

❖ RPG status code 50 is set when the conversion has to use substitution characters, but

- it's awkward to check for status code 50 after every statement with a CCSID conversion
- it's not always easy to tell which statements have CCSID conversions

❖ New H spec keyword CCSIDCVT(*EXCP) causes exception to be thrown instead

- New status code 00452

❖ New H spec keyword CCSIDCVT(*LIST) lists all CCSID conversions in the module

- shows source statements in the conversion
- shows if the conversion might result in a substitution

❖ Can code both as CCSIDCVT(*EXCP:*LIST)



- RNF7357 Conversion from UCS-2 to Alpha might not convert all data.
- RNF7358 Conversion from UCS-2 to DBCS might not convert all data.
- RNF7359 Conversion from DBCS to DBCS might not convert all data.
- RNF7360 Conversion from Alpha to DBCS might not convert all data.
- RNF7361 Conversion from DBCS to Alpha might not convert all data.

C C S I D C o n v e r s i o n s						
	From CCSID	To CCSID	References			
RNF7361	834	*JOB RUN	15	25		
RNF7357	1200	*JOB RUN	27	921	1073	
	*JOB RUN	1200	28	12	321	426
			552	631		
RNF7359	835	834	41	302	302	
RNF7360	*JOB RUN	834	242	304	305	
* * * *	E N D	O F	C C S I D	C O N V E R S I O N S	* * * *	

Performance

Productivity

Open and Modern

Mixed Environment Support

Maintainability

7.1 - Sort and Search Data Structure Arrays

- Sort and Search data structure arrays

- To sort or search a data structure array, one subfield is identified as the "key" for the array, using the syntax *ARRAY-NAME(*).KEY-SUBFIELD-NAME* in the SORTA or %LOOKUP operations.

Sort a data structure array using one subfield as a key

```
// sort by name  
SORTA info(*).name;
```

```
// sort by due date  
SORTA info(*).dueDate;
```

Search a data structure array using one subfield as a key

```
// search for a name  
pos = %LOOKUP('Jack' : info(*).name);
```

```
// search for today's date  
pos = %LOOKUP(%date() : info(*).dueDate);
```

7.1 - Sort Ascending and Descending

- Non-sequenced arrays can be sorted either ascending or descending
- (D) Extender indicates a descending sort
- (A) Extender indicates ascending
- No extender indicates ascending as in previous releases

```
D meetings          S          D  DIM(100)

/free
// sort descending, with the most recent date first
sorta(d) meetings;
...
```

Performance

Productivity

Open and Modern

Mixed Environment Support

Maintainability

7.1 - Performance when returning large values

- RTNPARM keyword greatly improves performance when a procedure returns a large value
- The speed of using a parameter with the convenience of using a return value
- Especially noticeable when the prototyped return value is a large varying length value

```
D center          pr          100000a  varying
D                                     rtnparm
D   text          50000a  const varying
D   len           10i  0  value
D title          s          100a  varying
/free
  title = center ('Chapter 1' : 60);
  ...
```

◆ 7.1 & 6.1 PTF - New H spec VALIDATE(*NODATETIME)

- ❖ Date, time, timestamp operations are costly since DTZ fields are validated on each use
- ❖ VALIDATE(*NODATETIME) tells RPG to treat DTZ as alphanumeric when possible
 - incorrect data will not always be detected and could be propagated
- ❖ Only use this in modules where you **NEVER** have date, time, or timestamp errors.
- ❖ Use the TEST operation to check a field, TEST is independent of VALIDATE keyword
- ❖ Affects moves where fields have the same format and separator
- ❖ Affects comparisons as above and
 - where dates are formatted as yyyy mm dd
 - where times are formatted as hh mm ss

Performance

Productivity

Open and Modern

Mixed Environment Support

Maintainability

7.1 - Implicit CCSID Conversion

- Implicit CCSID conversion is now supported for prototyped parameters passed by VALUE and by read-only reference (CONST)
- Reduces the number of code changes that have to be made when a database field is changed from alphanumeric or DBCS to Unicode (UCS-2 or UTF-16).
- Allows you to write a single procedure to handle string operations for alphanumeric, UCS-2 and DBCS strings

In the example below, there is only one “makeTitle” procedure with a UCS-2 parameter and return value. If the passed parameter is alpha or DBCS, it will be converted to UCS-2 on the call. The procedure will work with the UCS-2 parameter and return a UCS-2 value. This returned value can then be converted on assignment to alpha or DBCS, if necessary.

```
// makeTitle() upper-cases the parameter
// and centers it within the provided length

alphaTitle = makeTitle(alphaValue : 50);
ucs2Title = makeTitle(ucs2Value : 50);
dbcsTitle = makeTitle(dbcsValue : 50);
```

7.1 (6.1 PTF)- datasubf option of XML-INTO

- XML-INTO matches child elements and attributes to subfields.
- Specify datasubf to retrieve both text and attributes in the same operation
- Without datasubf no where to put text

```
<emp id="13573" type="regular">John Smith</emp>
```

```
D emp  
D id          7p 0  
D type       10a  
D value      100a
```

```
/free
```

```
XML-INTO emp %xml('emp.xml' : 'datasubf=value doc=file');
```

```
// emp.id = 13573
```

```
// emp.type = 'regular'
```

```
// emp.value = 'John Smith'
```

7.1 & 6.1 PTF - countprefix option of XML-INTO

- XML-INTO places repeating contained elements into an array subfield of the receiving data structure
 - The programmer declares the maximum number of elements, but the XML document may have fewer
 - Prior to 7.1, the allowmissing=yes option would be required and the RPG programmer would have to figure out how many XML elements there were by looking for some default value or by writing an XML handler.
 - the countprefix option specifies a subfield in which to place the count of elements actually read

```
<team>
  <manager>Jack Spratt</manager>
  <emp>John Smith</emp>
  <emp>Mary Jones</emp>
</team>
```

```
D  team          DS
D  manager       25A
D  emp           25A   dim(100)
D  numemp        10i  0
```

numexp = countprefix & contained element id
numemp not in the XML document.
allowmissing=yes not required.
has a value of 2 after the operation.

```
/free
  xml-into team %xml('teams.xml' : 'doc=file countprefix=num');
```

7.1 PTF & 6.1 PTF - ns option of XML-INTO

- brings XML namespace support to XML-INTO
- **ns=remove**: remove the namespace
- **ns=merge**: merge the namespace and the name using underscore.

```
<emp employee:type="regular" employee:id="13573">  
  <standard:name>John Smith</standard:name>  
</emp>
```

This attribute can't match an RPG subfield.

with ns=remove, matches "type"

with ns=merge, matches "employee_type"

7.1 PTF & 6.1 PTF - nsprefix option of XML-INTO

- What if the namespace is different between docs?
- **ns=remove**: remove the namespace
- **nsprefix=<subfield-prefix>**: specify the prefix of the field to receive the namespace

```
<emp>  
  <standard:type>manager</standard:type>  
</emp>
```

```
-----  
D emp                DS          qualified  
D  type              25A  
D  ns_type           25A
```

```
xml-into emp %xml('emp.xml' : 'ns=remove nsprefix=ns_');  
// emp.type = 'manager'  
// emp.ns_type = 'standard'
```

remove the namespace
standard:type => type

save the namespace name in
ns_<subfield>

can now be used to select processing based on
namespace

7.1 PTF & 6.1 PTF - case=convert option of XML-INTO

- Accented characters in XML can't be used in RPG variables
- **case=convert** matches accented characters according to *LANGIDSHR
 - other characters mapped to _

Étudiant is not a valid RPG variable name

```
<Étudiant Pre-nom="Élise" Âge="12">  
  <École>Collège Saint-Merri</École>  
</Étudiant>
```

```
-----  
D etudiant      ds              qualified  
D  age          3p 0  
D  pre_nom      25a  varying  
D  ecole        50a  varying
```

```
xml-into etudiant %xml('info.xml'  
                      : 'case=convert') ;  
// etudiant.age = 12  
// etudiant.pre_nom = 'Élise'  
// etudiant.ecole = 'Collège Saint-Merri'
```

case=convert matches "etudiant" and "ecole"

Performance

Productivity


Open and Modern

Mixed Environment Support

Maintainability


7.1 - Rational Open Access: RPG Edition

- Open Access opens the RPG I/O interface
- Allows programmers to reach out and touch resources beyond what the native interfaces would ordinarily allow.
- Handlers direct RPG IV I/O operations to devices other than workstations, printers, and disk.
 - ▶ Web browsers
 - ▶ mobile devices
 - ▶ Web services
 - ▶ XML
 - ▶ Spreadsheets
- Possibilities are limitless once programmer imagination and modern coding techniques are applied.




When an RPG program performs an I/O operation for a “normal” file, a system data management function is called to handle the operation

Traditional I/O




Open Access I/O

When the RPG program performs an I/O operation for an Open Access file, the Open Access handler is called.



The handler receives a data structure parameter with subfields that enable the handler to perform the correct I/O operation, and provide information back to RPG.



- ◆ Write a general handler some time after the application and retrofit
- ◆ Built for more than one application
- ◆ Built around device characteristics
- ◆ Example: a 5250 handler
 - ❖ An existing application that uses 5250 display files modified to use Open Access for the WORKSTN files.
 - ❖ The RPG program is modified by adding the HANDLER keyword to the WORKSTN files
 - ❖ The handler must handle all the operations and requirements of the existing RPG program.
 - ❖ This type of handler will normally be provided by an outside expert such as a software tool vendor or business partner.
- ◆ This can be REALLY hard
 - ❖ Depends on application usage, conventions used in applications
 - ❖ Mixed applications - pieces not under your control?

- ◆ Write a handler with or before the application
- ◆ Usually specific to an application or service
- ◆ Built within the context of an application's design
- ◆ Application has knowledge of handler and can exploit
- ◆ Example: exploiting a web service that returns the current weather for a particular city
 - ◆ Can be conceptually like a keyed database file, with the web service request = key
 - ◆ Web service data = record matching the key
 - ◆ The handler provider creates a keyed database file matching the web service.
 - ◆ The handler provider can tell the RPG programmer what I/O operations that the handler will support. For example, only OPEN, CHAIN, CLOSE.
 - ◆ The RPG programmer codes that file as an externally described keyed DISK file.
- ◆ This type of handler may be written by the same RPG programmer who uses the Open Access file or it may be provided by an outside expert.

- ◆ Open Access does not provide handlers
- ◆ Who provides?
 - ❖ Software tool vendors
 - ❖ Business partners
 - ❖ Services organizations
 - ❖ Programmers
- ◆ Anyone can write the handlers that extend RPG IV's I/O capabilities to new resources and devices

- Any RPG device type can be defined as an Open Access file: DISK, PRINTER, or WORKSTN
- The provider of the handling procedure can choose the RPG device-type that best fits the function that the handler provides
 - ▶ Based on operations
 - ▶ Based on data
 - ▶ Based on interaction
- For example
 - ▶ User interface device: WORKSTN file
 - ▶ Creating an spreadsheet: PRINTER file
 - ▶ Accessing a Web service: Keyed DISK file

Then

- 5733-OAR debuted in April 2010
- Trial licenses made available in July 2011
- 5733-OAR is no longer available

Now

- Open Access is integrated into
 - ▶ 5761-WDS for 6.1
 - ▶ 5770-WDS for 7.1
 - ▶ and into the base operating systems for each of those releases.

Removes price and configuration barriers to adoption

COBOL 7.1

- ◆ COMP-4 binary type does not support the full range of the binary values
 - ❖ For example PIC S9(4) COMP-4 is 2 bytes and supports values between -9999 and 9999
- ◆ Can change this with NOSTDTRUNC PROCESS option
 - ❖ forces all COMP-4 values to be treated as full binary
 - ❖ A 2 byte PIC S9(4) COMP-4 supports values between -32767 and 32767
 - ❖ A 2 byte PIC 9(4) COMP-4 with NOSTDTRUNC supports 0 to 32767
 - ❖ NOSTDTRUNC is “cracking an egg with a sledge hammer”
- ◆ COMP-5 uses all the bits
 - ❖ A 2 byte PIC S9(4) COMP-5 supports values between -32768 and 32767
 - ❖ A 2 byte PIC 9(4) COMP-5 with supports values 0 to 65535
- ◆ The COMP-5 is true binary integer but with decimal positions
- ◆ COMP-5 is also supported by System z COBOL
 - ❖ eases porting applications that use COMP-5 between System z and IBM i

◆ National data items introduced in V5R1

- ◆ Supports UCS-2 data
- ◆ PIC N(nn) and USAGE NATIONAL
- ◆ Only national hexadecimal literals (e.g. NX"0020") were allowed on the VALUE clause

◆ 7.1 allows non-hex (i.e. "normal") literals on VALUE clause

- ◆ 99 CompanyName PIC N(30) USAGE NATIONAL VALUE "ACME Sprockets Inc."

XML GENERATE is faster when

- XML is repeatedly generated for the same data structure or record format
- the data for the XML never contains special characters
 - specify process option XMLGEN(ASSUMEVALIDCHARS) to bypass checking XML data for (<,>,&,'", and characters outside the normal alpha-numeric range)
- The XML is generated to an IFS file
 - specify process option XMLGEN(KEEPFILEOPEN) to keep the stream file open after an APPEND
 - code XML GENERATE FILE-STREAM without the OVERWRITE or APPEND option to close the file.

```
PROCESS XMLGEN(KEEPFILEOPEN) .  
...  
  * write out XML records to the stream file  
  perform until DONE = 1  
    read MYFILE next record into CUST-INFO  
  ...  
    xml generate file-stream 'cust.xml' append from cust-info  
  end-perform.  
  * close the stream file  
  xml generate file-stream 'cust.xml' from cust-info.
```

- ◆ Inhibit optimization with
 - ❖ OPTIMIZE(*NEVER) compiler option
 - ❖ NEVEROPTIMIZE PROCESS option
- ◆ OPTIMIZE(*NEVER) and NEVEROPTIMIZE
 - ❖ Prevent the COBOL compiler from generating the information used to optimize the program.
 - ❖ Reduces the size of the generated code
 - ❖ Allows larger programs to be generated
- ◆ As the option implies programs created with this option cannot be optimized later. They must be recompiled without these options. This may not be possible for very large programs.

Both RPG and COBOL

◆ The problem:

- ❖ You want to ship a debuggable version of your application to your customers, but you don't want them to be able to read your source code through the debug view

◆ The solution:

- ❖ Encrypt the debug view so that the debug view is only visible if the person knows the encryption key. For example:

```
CRTBND CBL MYPGM DBGENCKEY ( 'my secret code' )
```

- ❖ Then either

```
STRDBG MYPGM DBGENCKEY ( 'my secret code' )
```

or

```
STRDBG MYPGM
```

and wait to be prompted for the encryption key

◆ The problem:

- ❖ 16MB automatic storage limits with the single-level storage model, for a single procedure, and for all the procedures on the call stack

◆ The solution: use the teraspace storage model

- ❖ New STGMDL compiler option supports values *SNGLVL, *TERASPACE, *INHERIT
- ❖ Ridiculously higher limits for automatic storage.
- ❖ STGMDL(*TERASPACE) defaults to activation group QILETS (rather than QILE)
- ❖ Can compile *CALLER programs with STGMDL(*INHERIT) so they can be called from either single-level or teraspace programs
 - More flexible
 - Caveat emptor - don't make assumptions about pointer sharing
- ❖ The ACTGRP (activation group) compiler option now defaults to ACTGRP(*STGMDL) on the CRTBNDCBL command
- ❖ All programs in an activation group must have compatible storage models
- ❖ See “ILE Concepts” for more info on teraspace

Rational Developer for Power

◆ Modern development tools for RPG, COBOL, CL, DDS, and SQL

- ◆ SEU and PDM still supported but no longer being enhanced
- ◆ WDSC went out of support in April 2011
- ◆ Replaces Rational Developer for i (RDi)

◆ Features


- ◆ Lightweight, graphical, development tools for RPG, COBOL, CL, DDS, and SQL
- ◆ Integrated object access & filtering, search, edit, compile and debug
- ◆ Rich editing features
 - outline view
 - content assist
 - formatting
 - color tokenizing
 - searching
 - compare
- ◆ Visualize program structure: Application Diagram
- ◆ Debug batch, interactive, and web applications with a common visual debugger
- ◆ Visual DDS design tools: Screen and Report Designer
- ◆ Integration with Rational Team Concert

- ◆ Improved productivity from
 - ❖ Integrated search / edit / compile / debug
 - ❖ Rich source editing features
 - ❖ Visual analysis tools
 - ❖ Visual DDS designers
- ◆ Improved developer skills
 - ❖ Lower learning curve for other technologies
 - ❖ Learn Rational Developer for Power (Eclipse) for RPG or COBOL development
 - ❖ Use similar Eclipse-based tools for Java, Web, EGL, XML, or PHP development
- ◆ Common development tools across multiple platforms and languages
 - ❖ Rational Developer for Power
 - ❖ Rational Application Developer for J2EE
 - ❖ Rational Business Developer for EGL
 - ❖ Rational Software Architect for UML modeling
- ◆ Integration with Rational Team Concert for Power Systems Software

And In Conclusion...

- ◆ IBM i applications represent a large investment and asset to businesses
- ◆ IBM offers a complete development solution for IBM i development to help leverage and maintain these assets
- ◆ There are other modern and productive development tools from IBM Rational including application lifecycle management
- ◆ RPG and COBOL language enhancements are keeping the languages current
- ◆ Rational Open Access: RPG Edition leverages new devices and resources (AND IT IS FREE!)

Check out the documents section of the RPG and COBOL Cafés to keep up with language news.



The screenshot shows the IBM Rational Cafés website. The main header includes the IBM logo, language selection (English), and a sign-in/register link. Below this is a navigation bar with links to developerWorks, Technical topics, Evaluation software, Community, and Events, along with a search bar. The main content area is titled "IBM Rational Cafés Connecting Communities". A sub-navigation bar shows links to various cafés: Café home, C/C++, COBOL, EGL, and RPG. The RPG Café page is displayed, featuring a welcome message, a list of community links (Discussions, Documents, Blog posts), and sections for discussing RPG and reading bloggers' content. The "Documents" link in the RPG Community section is circled in red.

<https://www.ibm.com/developerworks/rational/community/cafe/rpg.html>

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquiries, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Revised September 26, 2006

IBM, the IBM logo, ibm.com AIX, AIX (logo), AIX 5L, AIX 6 (logo), AS/400, BladeCenter, Blue Gene, ClusterProven, DB2, ESCON, i5/OS, i5/OS (logo), IBM Business Partner (logo), IntelliStation, LoadLeveler, Lotus, Lotus Notes, Notes, Operating System/400, OS/400, PartnerLink, PartnerWorld, PowerPC, pSeries, Rational, RISC System/6000, RS/6000, THINK, Tivoli, Tivoli (logo), Tivoli Management Environment, WebSphere, xSeries, z/OS, zSeries, Active Memory, Balanced Warehouse, CacheFlow, Cool Blue, IBM Systems Director VMControl, pureScale, TurboCore, Chiphopper, Cloudscape, DB2 Universal Database, DS4000, DS6000, DS8000, EnergyScale, Enterprise Workload Manager, General Parallel File System, , GPFS, HACMP, HACMP/6000, HASM, IBM Systems Director Active Energy Manager, iSeries, Micro-Partitioning, POWER, PowerExecutive, PowerVM, PowerVM (logo), PowerHA, Power Architecture, Power Everywhere, Power Family, POWER Hypervisor, Power Systems, Power Systems (logo), Power Systems Software, Power Systems Software (logo), POWER2, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, POWER6+, POWER7, System i, System p, System p5, System Storage, System z, TME 10, Workload Partitions Manager and X-Architecture are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries.

A full list of U.S. trademarks owned by IBM may be found at: <http://www.ibm.com/legal/copytrade.shtml>.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

AltiVec is a trademark of Freescale Semiconductor, Inc.

AMD Opteron is a trademark of Advanced Micro Devices, Inc.

InfiniBand, InfiniBand Trade Association and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.

Microsoft, Windows and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries or both.

NetBench is a registered trademark of Ziff Davis Media in the United States, other countries or both.

SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECapc, SPECchpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).

The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.

TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).

UNIX is a registered trademark of The Open Group in the United States, other countries or both.

Other company, product and service names may be trademarks or service marks of others.

Revised December 2, 2010

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM Power Systems Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html.

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, the latest versions of AIX were used. All other systems used previous versions of AIX. The SPEC CPU2006, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C for AIX v11.1, XL C/C++ for AIX v11.1, XL FORTRAN for AIX v13.1, XL C/C++ for Linux v11.1, and XL FORTRAN for Linux v13.1.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

TPC	http://www.tpc.org
SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
VolanoMark	http://www.volano.com
STREAM	http://www.cs.virginia.edu/stream/
SAP	http://www.sap.com/benchmark/
Oracle, Siebel, PeopleSoft	http://www.oracle.com/apps_benchmark/
Baan	http://www.ssaglobal.com
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
Ideas International	http://www.ideasinternational.com/benchmark/bench.html
Storage Performance Council	http://www.storageperformance.org/results

Revised December 2, 2010

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM Power Systems Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html.

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, the latest versions of AIX were used. All other systems used previous versions of AIX. The SPEC CPU2006, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C for AIX v11.1, XL C/C++ for AIX v11.1, XL FORTRAN for AIX v13.1, XL C/C++ for Linux v11.1, and XL FORTRAN for Linux v13.1. Linpack HPC (Highly Parallel Computing) used the current versions of the IBM Engineering and Scientific Subroutine Library (ESSL). For Power7 systems, IBM Engineering and Scientific Subroutine Library (ESSL) for AIX Version 5.1 and IBM Engineering and Scientific Subroutine Library (ESSL) for Linux Version 5.1 were used.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
STREAM	http://www.cs.virginia.edu/stream/
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
AMBER	http://amber.scripps.edu/
FLUENT	http://www.fluent.com/software/fluent/fl5bench/index.htm
GAMESS	http://www.msg.chem.iastate.edu/gamess
GAUSSIAN	http://www.gaussian.com
ANSYS	http://www.ansys.com/services/hardware-support-db.htm
ABAQUS	http://www.simulia.com/support/v68/v68_performance.php
ECLIPSE	http://www.sis.slb.com/content/software/simulation/index.asp?seg=geoquest&
MM5	http://www.mmm.ucar.edu/mm5/
MSC.NASTRAN	http://www.mscsoftware.com/support/prod%5Fsupport/nastran/performance/v04_sngl.cfm
STAR-CD	www.cd-adapco.com/products/STAR-CD/performance/320/index/html
NAMD	http://www.ks.uiuc.edu/Research/namd
HMMER	http://hmmer.janelia.org/ http://powerdev.osuosl.org/project/hmmerAltivecGen2mod

Click on the "Benchmarks" icon on the left hand side frame to expand. Click on "Benchmark Results in a Table" icon for benchmark results.

Revised December 2, 2010

Notes on performance estimates

rPerf for AIX

rPerf (Relative Performance) is an estimate of commercial processing performance relative to other IBM UNIX systems. It is derived from an IBM analytical model which uses characteristics from IBM internal workloads, TPC and SPEC benchmarks. The rPerf model is not intended to represent any specific public benchmark results and should not be reasonably used in that way. The model simulates some of the system operations such as CPU, cache and memory. However, the model does not simulate disk or network I/O operations.

- rPerf estimates are calculated based on systems with the latest levels of AIX and other pertinent software at the time of system announcement. Actual performance will vary based on application and configuration specifics. The IBM eServer pSeries 640 is the baseline reference system and has a value of 1.0. Although rPerf may be used to approximate relative IBM UNIX commercial processing performance, actual system performance may vary and is dependent upon many factors including system hardware configuration and software design and configuration. Note that the rPerf methodology used for the POWER6 systems is identical to that used for the POWER5 systems. Variations in incremental system performance may be observed in commercial workloads due to changes in the underlying system architecture.

All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, and application sizing guides to evaluate the performance of a system they are considering buying. For additional information about rPerf, contact your local IBM office or IBM authorized reseller.

=====

CPW for IBM i

Commercial Processing Workload (CPW) is a relative measure of performance of processors running the IBM i operating system. Performance in customer environments may vary. The value is based on maximum configurations. More performance information is available in the Performance Capabilities Reference at: www.ibm.com/systems/i/solutions/perfmgmt/resource.html

Revised April 2, 2007