



IBM Power Systems™



Agenda Key:
Session Number:

V6R1 DB2 for i5/OS Overview

Presented by Jarek Miszczyk
DB2 for i5/OS Technology Team
IBM Rochester, MN



© Copyright IBM Corporation, 2008. All Rights Reserved.
This publication may refer to products that are not currently
available in your country. IBM makes no commitment to make
available any products referred to herein.

IBM Power Systems



DB2 for i5/OS V6R1 Enhancements

Application Flexibility & Portability

- **SQL & RPG Integration**
- **Enhanced JDBC & .NET support**
- **Skip Locked Data**
- **Extended Indicator Variables**
- **VALUES on FROM**
- **Hidden Timestamp Columns**
- **Improved DB2 Family Compatibility**
 - OLAP Support – Cube & Rollup
 - INSERT on FROM
 - Unsupported Syntax Tolerance
 - AES Encryption

OnDemand & Availability

- **Enhanced, online Reorg**
- **Library-level Journaling**

Performance

- **SQL Query Engine enhancements**
 - Sort sequence support
 - Self-Learning Optimizer
 - EVI-Only Processing
- **Derived SQL Indexes**
- **Faster Full Opens**
- **Client Special Registers**

Usability

- **System i Navigator Enhancements**
 - Customizable Performance Analysis
 - Spreadsheet integration
 - Plan Cache Enhancements
 - Index Advisor Improvements
- **DB2 Web Query for System i**
- **OmniFind Text Search Server**

© 2008 IBM Corporation

Application Development Enhancements

Enhancements for i5/OS Application Development

- Improved RPG & SQL Integration
 - IFS source file support
 - Improved Variable Scoping
 - Enhanced LIKE supports variables based on SQLCA variables (eg, SQLSTATE)
 - Improved WDSC integration for SQL source code error resolution
 - SQL support in IBM Application Development tools
 - SQL syntax highlighting
 - SQL statement(s) templates
 - Formatting
- ILE COBOL SQL Precompiler support for Unicode data
- Enhanced SQL Scripting
 - RUNSQLSTM support for IFS files
 - Larger & “Wider” SQL scripts
 - Improved db2 qshell utility

IBM Power Systems

IBM

RPG SQL Precompiler Variable Scoping

```

PSubProc1      B              EXPORT
D              PI
D OutArray      ds              qualified dim(1000)
D customer      25A
D region        25A
/free
    exec sql DECLARE c1 CURSOR FOR SELECT customer,region FROM cust_dim;
    exec sql OPEN c1;
    exec sql FETCH NEXT FROM c1 FOR 100 ROWS INTO :OutArray;
    exec sql CLOSE c1;
return;
/end-free

P              E
PSubProc2      B              EXPORT
D              PI
D OutArray      ds              qualified dim(1000)
D part          55A
D mfgr          25A
D brand         10A
/free
    exec sql DECLARE c2 CURSOR FOR SELECT part,mfgr,brand FROM part_dim;
    exec sql OPEN c2;
    exec sql FETCH NEXT FROM c2 FOR 200 ROWS INTO :OutArray;
    exec sql CLOSE c2;
return;
/end-free

```

V5R4 PTFs planned!

© 2008 IBM Corporation

IBM Power Systems

IBM

Industry Standard Application Interface Improvements

- JDBC
 - JDBC 4.0
 - Alias support in Metadata APIs
 - Returning DEFAULT
 - Return update counts
- ADO.NET
 - Exploitation of ADO.NET 2.0
 - Visual Studio integration
 - Distributed transactions
 - Multi-row Insert
 - Enhanced data type support
- CLI
 - Wide API support for Unicode data
 - Alias support in Metadata APIs
 - Row-wise array INSERT
 - Complete ISO timestamp support

© 2008 IBM Corporation

Improved Scalability with Skip Locked Data

- By default, all DB2 requests wait for conflicting locks to be released
- SKIP LOCKED DATA clause can alter default behavior
 - Clause only honored with Cursor Stability(*CS) and Read Stability(*RS) levels
 - Clause can also be specified on Insert & Update statements

Executed at 11:30, not yet committed

```
UPDATE flights
SET departTime='05:25'
WHERE departTime = '04:30'
AND flightNum=331
AND destCity='HNL'
```

FLIGHTS

flightNum	destCity	departTime
...
4388	RST	08:23
331	HNL	05:25
3044	MSP	03:07
1025	SYD	02:45
389	HNL	06:10
...

Executed at 11:32

```
SELECT * FROM flights
WHERE
  departTime >= '05:15'
  AND destCity='HNL'
SKIP LOCKED DATA
```

Wait or Skip?

© 2008 IBM Corporation

INSERT on FROM (Select From Insert)

- Simplifies access to values generated by DB2 for identity columns
 - Prior support (IDENTITY_VAL_LOCAL function) required overhead of an extra SQL statement
 - Prior support had no solution for blocked Insert statements

Examples:

```
CREATE TABLE orders( order_id INTEGER AS IDENTITY,
                      order_date DATE,
                      order_qty INTEGER,
                      order_item CHAR(4));
```

```
SELECT order_id FROM FINAL TABLE (
  INSERT INTO orders VALUES(DEFAULT,'11/03/2007',50,'JM12'));
```

Result 1

```
SELECT order_id FROM FINAL TABLE (
  INSERT INTO orders VALUES(DEFAULT,'11/05/2007',12,'JM09'),
  (DEFAULT,'11/05/2007',1,'PC01'))
```

Result 2 3

```
ORDER BY INPUT SEQUENCE;
```

© 2008 IBM Corporation

VALUES on FROM

- Allows programmer to dynamically generate & populate temporary tables as part of the query definition
 - Enable SQL access of in-memory tables maintained by application
 - Supports table-less queries

Examples:

```
SELECT deptnum, deptname FROM org WHERE deptnum <20
UNION ALL
SELECT * FROM (VALUES(77,'New Department')) AS tmp(c1,c2)
```

```
WITH proposedRates(prType, prRate) AS ( VALUES(?,?),(?,?),(?,?),(?,?) )
SELECT rmttype, prRate, ((prRate - rmRate)/rmRate)*100 AS RateChgPercent
FROM rooms, proposedRates WHERE rmttype = prType
ORDER BY RateChgPercent DESC
```

SQL & DB2 Enhancements

Data Access Advancements

- SQL OLAP Extensions – Grouping Sets & Super Groups
- More “Flexible” FROM
 - VALUES on FROM
 - INSERT on FROM
 - FULL OUTER JOIN
- Expanded SQL Function Toolset
 - Data Encryption: ENCRYPT_AES
 - String Processing: ASCII & CHR
 - Date and Time Processing
 - TIMESTAMP_FORMAT & VARCHAR_FORMAT
 - MONTHS_BETWEEN
 - ROUND_TIMESTAMP & TRUNC_TIMESTAMP
- Limit Improvements
 - 120 columns on GROUP BY
 - 128-byte cursor and statement names
 - Partial support for 64K Result Set width (V5R4)

© 2008 IBM Corporation

Grouping Sets and Super Groups (ROLLUP & CUBE)

- Many BI applications and OLAP tools involve hierarchical, multi-dimensional aggregate views of transaction data
 - Users need to view results at multiple levels
 - Users need to view result data from different perspective
 - Current grouping support only allows aggregation data of along a SINGLE dimension

EXAMPLE: *SELECT country region, store, product, SUM(sales) FROM trans*
GROUP BY country region, store, product

 - Limitations result in extra coding for programmers
- V6R1 grouping and OLAP capabilities allow data to be grouped in multiple ways with a single SQL request
 - GROUPING SETS
 - ROLLUP
 - CUBE



Less Coding
for Developers!

© 2008 IBM Corporation

ROLLUP

- ROLLUP on GROUP BY clause results in DB2 returning aggregates for each level of the hierarchy implicitly represented in the grouping columns
 - ROLLUP(Country, Region) will result in the data being summarized at the following levels
 - (Country, Region)
 - (Country)
 - () << represents Grand Total

- **Example Query:**

```
SELECT country, region, SUM(sales)
FROM trans
GROUP BY ROLLUP (country, region)
```

© 2008 IBM Corporation

ROLLUP Output Example

```
SELECT country, region, SUM(sales) FROM trans
GROUP BY ROLLUP (country, region)
```

GROUP BY
country, NULL

GROUP BY
NULL, NULL

Country	Region	Sum(Sales)
Canada	-	100,000
Canada	NW	100,000
U.S.A.	-	3,250,000
U.S.A.	NE	450,000
U.S.A.	NW	940,000
U.S.A.	SE	550,000
U.S.A.	SW	1,310,000
-	-	3,350,000

© 2008 IBM Corporation

CUBE

- CUBE on GROUP BY clause results in DB2 returning aggregates for all possible distinct combinations represented by the grouping columns
 - CUBE(Country, Region) will result in the data being summarized at the following levels
 - (Country, Region)
 - (Country)
 - (Region)
 - () << represents Grand Total
 - Returns results at multiple intersection points
- **Example Query:**

```
SELECT country, region, SUM(sales)
FROM trans
GROUP BY CUBE(country, region)
```

© 2008 IBM Corporation

CUBE Output Example

```
SELECT country,region, SUM(sales) FROM trans
GROUP BY CUBE (country, region)
```

GROUP BY NULL, region →

GROUP BY NULL, NULL →

GROUP BY country, NULL →

Country	Region	Sum(Sales)
-	NE	450000
-	NW	1040000
-	SE	550000
-	SW	1310000
-	-	3350000
Canada	-	100000
U.S.A.	-	3250000
Canada	NW	100000
U.S.A.	NE	450000
U.S.A.	NW	940000
U.S.A.	SE	550000
U.S.A.	SW	1310000

© 2008 IBM Corporation

GROUPING SETS

- GROUPING SET on GROUP BY clause enables DB2 to return aggregates for multiple sets of grouping columns
 - GROUPING SETS((Country, Region), (Country, Store)) will result in the data being summarized at the following levels
 - (Country, Region)
 - (Country, Store)
 - CUBE and ROLLUP can be used in combination with Grouping Sets
 - **CAUTION:** These types of combinations can result in an exponential growth in the number of grouping sets returned by a query, combine carefully
- **Example Query:**

```
SELECT country, region, SUM(sales)
FROM trans
GROUP BY GROUPING SETS((country, region), (country, store))
```

© 2008 IBM Corporation

GROUPING SETS Output Example

```
SELECT country, region, store, SUM(sales) FROM trans
GROUP BY GROUPING SETS ((country, region), (country, store))
```

GROUP BY
COUNTRY, REGION

GROUP BY
COUNTRY, STORE

Country	Region	Store	Sum(Sales)
Canada	NW	-	100,000
U.S.A.	NE	-	450,000
U.S.A.	NW	-	940,000
U.S.A.	SE	-	550,000
U.S.A.	SW	-	1,310,000
Canada	-	Dougs	100,000
U.S.A.	-	Mariahs	350,000
U.S.A.	-	KMs	770,000
U.S.A.	-	Jennas	400,000
U.S.A.	-	Adrians	500,000
U.S.A.	-	Joshs	300,000
U.S.A.	-	TZs	200,000
U.S.A.	-	Maddies	210,000
U.S.A.	-	Dylans	520,000

© 2008 IBM Corporation

New Column Attributes – Hidden Timestamp

- Implicitly Hidden and Row Change Timestamp attributes enable DB2 to track row-level changes for you
 - Clauses can be used independently
 - Attributes frequently combined to create a Hidden Timestamp column
 - Hidden Timestamp Columns used to implement Optimistic Locking scheme
 - Read values in a row without an update lock
 - Perform calculations on fetched values
 - Update same row with new values, check timestamp column to see if row has changed

```
CREATE TABLE tickets(
  ticket_ord  INTEGER,
  ticket_qty  INTEGER,
  ticket_event VARCHAR(10),
  ticket_ts   TIMESTAMP NOT NULL
              IMPLICITLY HIDDEN
              FOR EACH ROW ON UPDATE AS ROW CHANGE TIMESTAMP);
INSERT INTO tickets VALUES(1,11,'mvGAME1'),
                          (2,8,'ihGAME4');
```

NOTE: Only 3 column values passed on INSERT

© 2008 IBM Corporation

New Column Attributes – Hidden Timestamp

Table contents after INSERT statement

SELECT * FROM tickets



TICKET_ORD	TICKET_QTY	TICKET_EVENT
1	11	mvGAME1
2	8	ihGAME4

Non-SQL interfaces
automatically
include any
hidden columns

Table contents after INSERT & UPDATE statements

UPDATE tickets SET ticket_qty = 6 WHERE ticket_ord = 2;

SELECT ticket_ord, ticket_qty, ticket_event, ticket_ts FROM tickets



TICKET_ORD	TICKET_QTY	TICKET_EVENT	TICKET_TS
1	11	mvGAME1	2007-11-26 16:19:08.516671
2	6	ihGAME4	2007-11-26 16:20:27.256864

© 2008 IBM Corporation

Unsupported Syntax Tolerance

- SQL parser enhanced to tolerate unsupported syntax that is not needed on DB2 for i5/OS
 - Unique architecture of i5/OS enables some SQL statements and clauses to be ignored
 - Speeds up porting process

Examples:

CREATE TABLESPACE TS1 MANAGED BY DATABASE USING (device '/dev/rcont \$N' 20000)

SQLSTATE: 01505 / SQLCODE: +143

Message: Statement CREATE TABLESPACE ignored.

CREATE TABLE newtab1 (c1 **INT**) **INDEX IN** ts1

SQLSTATE: 01680 / SQLCODE: +20367


Message: Clause INDEX IN ignored.

CREATE TABLE t1 (c1 **INT**) **IN** ts1

SQLSTATE: 42704 / SQLCODE: -204

Message: TS1 in L1 type *NODGRP not found.

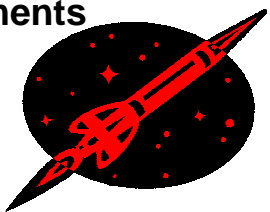
Performance Enhancements

IBM Power Systems



SQL Query Engine (SQE) Enhancements

- **Restriction Elimination**
 - National Language Sort Sequences
 - Translation support (UPPER, LOWER, etc)
 - User-Defined Table Functions
 - Remaining Restrictions
 - ICU 2.6.1 Sort Sequences
 - Non-SQL Interfaces (OPNQRYF, Query/400, QQQQRY api)
 - Logical File reference on FROM Clause
 - Select/Omit Logical Files defined on tables
- **Technological Advances**
 - Self-learning query optimization
 - Self-adjusting query execution
 - Faster optimization times
 - Smarter CASE processing

Default value for QAQQINI parameter:
IGNORE_DERIVED_INDEX
changed from *NO to *YES



© 2008 IBM Corporation


IBM Power Systems


SQL Derived Indexes

- SQL key definitions support expressions, functions and operators enabling more usage of indexes by query optimizer on complex queries
 - Fully supported by SQE optimizer (limited support by CQE)
 - EXAMPLES:


```
CREATE INDEX ix_TotalSalary ON employees (Sales + Bonus)
CREATE INDEX ix_FullName ON employees (CONCAT(CONCAT(FName, ' '), LName))
```
- Great for improving performance of case-insensitive searches


```
SELECT cust_id, cust_phone FROM customers
WHERE UPPER(company_name) = 'ACME'
```



```
CREATE INDEX ix_uCompName ON customers(UPPER(company_name))
```
- Create Index statement also supports the ability to create sparse indexes (ie, Select/Omit criteria) but there currently is no optimizer awareness
 - Short-term value: provide SQL replacement of DDS Select/Omit logicals
 - EXAMPLE: CREATE INDEX cust_ix1 ON customers(cust_id) WHERE activCust='Y'

© 2008 IBM Corporation



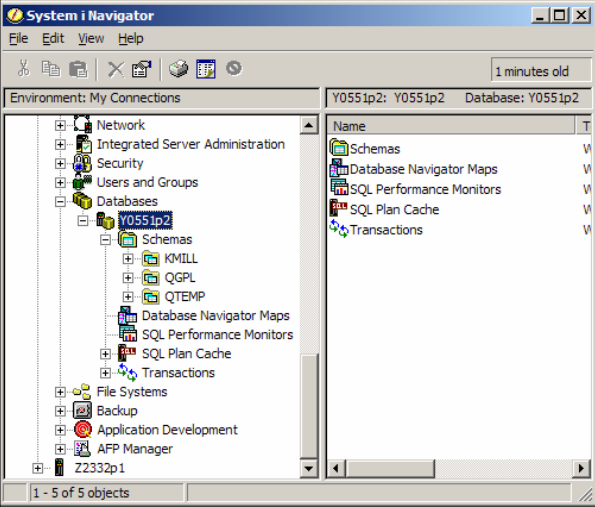
IBM Power Systems 

Ease of Use & Management Enhancements

© 2008 IBM Corporation

IBM Power Systems

System i Navigator Enhancements



OnDemand Performance Center

- Spreadsheet Integration
- Column Customization
- Index Advice Condenser (PTF)
- Fast monitor summary compare
- Plan Cache Resize & Event Monitor
- Enhanced Plan Cache Analysis
- Visual Explain While Running
- Special Client Register

Run SQL Scripts

- Improved Font Size
- SQL Syntax Flagger
- Source Member Integration
- UTF-8 support

Database Management

- Show object locks (WRKOBJLCK)
- Schema-level Index Evaluator
- Enhanced SQL Details for Job

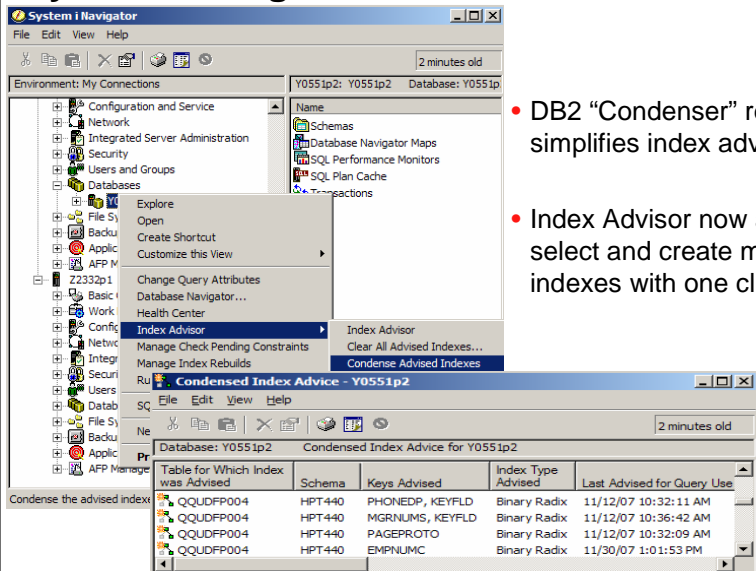
Health Center

- Environmental limits
- Journal and journal receivers
- Procedures and Packages
- Table activity

© 2008 IBM Corporation

IBM Power Systems

System i Navigator – Index Advisor



- DB2 “Condenser” reviews and simplifies index advice for you
- Index Advisor now allows you to select and create multiple indexes with one click

Table for Which Index was Advised	Schema	Keys Advised	Index Type Advised	Last Advised for Query Use
QQUDFP004	HPT440	PHONEDP, KEYFLD	Binary Radix	11/12/07 10:32:11 AM
QQUDFP004	HPT440	MGRNUMS, KEYFLD	Binary Radix	11/12/07 10:36:42 AM
QQUDFP004	HPT440	PAGEPROTO	Binary Radix	11/12/07 10:32:09 AM
QQUDFP004	HPT440	EMPNUMC	Binary Radix	11/30/07 1:01:53 PM

© 2008 IBM Corporation

IBM Power Systems

System i Navigator – Enhanced Index Evaluator

Environment: My Connections

Indexes for tables in QSYS2 - Y0551p2

SQL Name	Key Columns	Type	Last Query Use	Query Use Count
QSYS2.CHEMA...	Primary Key...	Primary Key...	12/21/07 10:46:55 AM	6488
QSYS2.SYSTEM...	Index	Index	12/21/07 1:01:51 PM	12587
QSYS2.CIFIC...	Index	Index	12/21/07 1:01:51 PM	182459

Index usage stats can be reset to initiate “new” performance analysis

© 2008 IBM Corporation

IBM Power Systems

System i Navigator – Enhanced Plan Cache Analysis

- Customized controls for systems pushing Plan Cache limits
 - Properties view allows you to change Cache size
 - Event Monitor can be started to “catch” statements removed from cache
 - Can manually delete or “pin” statements
- Additional data available when viewing Plan Cache contents
 - Average Processing Time
 - Plan Cache Score
 - Associated jobs

SQL Plan Cache Properties - Z2332p1(Z2332p1)

Description	Value
Time Of Summary	2007-08-13-20.15.02.83718
Active Query Summary	
Number of Currently Active Queries	21
Number of Queries Run Since Start	17322051
Number of Query Full Opens Since Start	3322485
Plan Usage Summary	
Current Number of Plans in Cache	425
Current Plan Cache Size	488 MB
Plan Cache Size Threshold	512 MB

SQL Plan Cache Event Monitor Wizard - Z2332p1(Z2332p1)

To limit the amount of data collected, specify which filters to use. When filters are provided, only statements that match the specified filter values will be captured.

☐ Minimum runtime for the longest execution of the statement:
[] Seconds

☐ Statements that ran on or after this date and time:
[] 2013/01/01 [] 11:52 PM

☐ Statements the following user has ever run:
[]

☐ Statements that are currently active

☐ Statements for which an index has been advised

☐ Statements for which statistics have been advised

☐ Include statements initiated by the operating system

☐ Statements that reference the following objects:
[]

☐ Statements that contain the following text:
[]

© 2008 IBM Corporation

Additional Information

- DB2 for i5/OS Websites
 - Home Page: ibm.com/systemi/db2
 - DeveloperWorks Zone: ibm.com/developerworks/db2/products/db2i5OS
 - Porting Zone: ibm.com/servers/enable/site/db2/porting.html
- Newsgroups
 - USENET: comp.sys.ibm.as400.misc, comp.databases.ibm-db2
 - System i Network DB2 Forum - <http://systeminetwork.com/isnetforums/forumdisplay.php>
- Education Resources - Classroom & Online
 - ibm.com/systemi/db2/gettingstarted.html
 - ibm.com/partnerworld/wps/training/i5OS/courses
- DB2 for i5/OS Publications
 - White Papers: ibm.com/partnerworld/wps/whitepaper/i5OS
 - Online Manuals: ibm.com/systemi/db2/books.html
 - DB2 for i5/OS Redbooks (<http://ibm.com/redbooks>)
 - [Getting Started with DB2 Web Query for System i](#) (SG24-7214)
 - [OnDemand SQL Performance Analysis ... in V5R4](#) (SG24-7326)
 - [Preparing for and Tuning the SQL Query Engine on DB2 for i5/OS](#) (SG24-6598)
 - [Modernizing iSeries Application Data Access](#) (SG24-6393)

Trademarks and Disclaimers

© IBM Corporation 1994-2007. All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Trademarks of International Business Machines Corporation in the United States, other countries, or both can be found on the World Wide Web at <http://www.ibm.com/legal/copytrade.shtml>.

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.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

The customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Prices are suggested U.S. list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Photographs shown may be engineering prototypes. Changes may be incorporated in production models.