



iSeries

Memorandum to users

Version 5 Release 3 (May 2008 update)





iSeries

Memorandum to users

Version 5 Release 3 (May 2008 update)

Note

Before using this information and the products it supports, be sure to read the information in "Notices" on page 65.

Fifth Edition (May 2008)

This edition applies to version 5, release 3, modification 0 of the IBM Operating System/400® (Program 5722-SS1) and iSeries licensed programs. This version does not run on all reduced instruction set computer (RISC) models nor does it run on CISC models.

© Copyright International Business Machines Corporation 2004, 2008. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About iSeries Memorandum to Users **vii**

Who should read this memorandum	vii
Additional incompatibility information	vii
What's new	vii
Installing V5R3 over V5R1	ix
Discontinued support for certain software and hardware	ix
Installation considerations.	ix
PTF numbers in this memorandum	ix
Memorandums for previous releases	x
Prerequisite and related information	x
iSeries Navigator	x
How to send your comments	x

Chapter 1. Read this first **1**

Current customers only — before you install V5R3	1
V5R3 installation requirements	2
Plan for your Operations Console installation or upgrade	3
Upgrading to V5R3 using image catalog support	5
I/O software prerequisites.	5
Support for I/O expansion units.	6
Enterprise Storage Server (ESS) SAN attachment prerequisites	6

Chapter 2. Operating system. **7**

Programming considerations	7
Output file (OUTFILE) changes	7
Security audit record changes.	7
Programs that use customized versions of IBM-supplied commands	7
Clustering does not support switchable independent ASP (IASP) with IOP-less disk units	7
Modification levels for licensed internal code and operating system can be different	7
Prepare for Install required during installation process	8
Server IPL might be required	8
Collection Services and performance database file changes	8
XML procedural parser service programs removed	9
Performance Explorer statistics change.	9
Default values for pool tuning options changed on systems with 100% interactive capacity	10
Active user profile now stored with each message	10
User profiles to which you are not authorized are not shown for DSPLIB of QSYS library	10
Tape management system changes.	10
SNTP client changes the system clock instead of the software clock	10
Integrated file system changes	11
*TYPE2 directory conversion.	11
Network Server Description Considerations	11
Reclaim Storage (RCLSTG) restriction.	11

Convert Directory (QP0FCVT2) API and Convert Directory (CVTDIR) command	12
Management Central changes	12
QYPSSRV job no longer exists	12
Default DNS lookup frequency changed.	12
Hardware and software inventory no longer published to LDAP database	12
Management Central user data moved from QUSRSYS library	12
Inventory information from prior releases not migrated	12
Graph history changes.	12
Security update to the /tmp, /var, /QOpenSys/var, and /QOpenSys/var/preserve directories	13
Limit of one Java virtual machine (JVM) per process	13
Meaning of the console type on the Select Console	14
Type DST screen changed	14
Program QSPGETF to be removed in future release	14
Coded character set identifiers (CCSIDs) with bidirectional encoding schemes.	15
Virtual IP Address automatic agent selection behavior	15
TCP/IP server values *DCE and *NSMI removed.	15
Network Quality of Service (QoS) policy	15
If you configured a Network Quality of Service (QoS) policy in V5R1 or V5R2	15
Running multiple iSeries servers using QoS	16
Enable ECN support added	16
Required authorities changed for security tools CL commands.	16
Cryptographic product and DCM needed for IBM Universal Connection	16
Directory Server (LDAP) changes	16
Configuration changes.	16
LDAP Kerberos service name changed	17
Replication stores passwords in a validation list	18
System SSL default cipher specification list changed	18
Using Service Tools Device IDs in System Service Tools (SST)	19
Obsolete command menus removed	19
Applications using the iSeries SSL_ C APIs for secure sockets	20
Changes in C language runtime	20
Saving and restoring journal receivers to a previous release	23
Mark counters changed from 32-bit integers to 64-bit integers.	23
Universal Connection replaces IGN on the Work with PM eServer iSeries customization display	24
New security auditing for special authority violations	24
Restricting access to audit settings.	24
Java Database Connectivity (JDBC) coded character set identifier (CCSID) requirement.	24
Remote command and distributed program call server	25

Saving UIM panels to a previous release.	25	New DB2 tables and views added to the QSYS2 library	39
Reported utilization for shared partitions	25	DYNAMIC keyword handled as if SENSITIVE	
High availability clustering considerations	25	DYNAMIC has been specified	41
Command and API changes	25	New EXTRACT keyword.	41
Authority change for CHGDDMTCPA command	25	BINARY data type has a release-to-release	
Authority change to remove *LOCAL RDB		compatibility issue	41
directory entry	25	Stricter enforcement of ROWID compatibility	
CPYFRMIMPF and CPYTOIMPF command		rules.	41
changes.	26	SQL DELETE and journal entries	42
Create Validation List (CRTVLDL) command		Reduction of scale no longer performed on SQL	
changes.	26	query comparisons involving numeric literals	42
Digital certificate management APIs require DCM		Enable concurrent write (ECW) default behavior	
option	26	on record inserts.	42
Journaling command and API changes	27	New CACHE_RESULTS Query QAQQINI option	43
adjtime(), gettimeofday(), and settimeofday()		Rules for specifying SQLSTATE values enforced	
APIs updated to access the system clock.	29	in SQL routines	43
Open List of ASPs (QYASPOL) API changes	29	Authority to create a view	43
Changes to profile handle and token APIs might		Error message change when comparing numeric	
require source code changes	30	and string fields.	43
Reply handling exit point changes.	30	ILE RPG SQL precompiler enforces unique	
Activation group default changing on Create		variable names	44
Program (CRTPGM) command	31	Updating rows with NULL value fails with	
Integrated file system byte stream file APIs now		SQL0407 error when NULL values are not	
use the 1200 CCSID for storing path names.	31	allowed.	44
Add Mounted File System (MOUNT, ADDMFS)		Format level identifiers changed for record	
command changes	31	formats of new SQL tables that have one or more	
Change to path name format for Qlg		TIMESTAMP columns.	44
(NLS-enabled) APIs.	31	Machine Interface (MI) instruction changes	45
Display DLO Auditing Level (DSPDLOAUD)		Materialize Machine Attributes (MATMATR) MI	
command changes	31	instruction.	45
Work with Object Links (WRKLNK) and Display		Materialize Resource Management Data	
Object Links (DSPLNK) command changes.	31	(MATRMD) MI instruction	45
Get Attributes (Qp0lGetAttr) API changes	31	Materialize Journal Port Attributes (MATJPAT)	
FRCACCPH(*YES) parameter on the CHGLF		MI instruction no longer reports a local journal	
command	32	as inactive.	45
Removal of support for ATM and ISDN	32	Materialize System Object (MATSOBJ) MI	
Default value on AUT parameter changed for		instruction changed.	45
some commands.	35	Process access group to be eliminated	45
Change TCP/IP Attributes (CHGTCPA)		Optical support changes	45
command changes	36	Optical device support changes.	45
Perform Hardware Configuration Operation		Optical command changes	46
(QYHCHCOP) API restriction	36	System value changes	46
WAITFILE parameter on CRTPF and CRTLF		Retrieve System Values (RTVSYSVAL) against	
commands changed	36	QAUDLVL system value might not return the	
Backup and recovery changes	37	complete list of audit values.	46
Save command and API changes	37	New system value (QTHDRSCAFN) and a new	
SAVLIB LIB(*NONSYS), SAVLIB LIB(*ALLUSR),		job attribute for changing and displaying the	
and SAVCHGOBJ LIB(*ALLUSR) commands save		main storage affinity	47
libraries in a different order	37	New system value QENDJOBMT.	47
Output file will contain additional records when		Change System Value (CHGSYSVAL) command	
INFTYPE(*MBR) is specified.	38	requires *ALLOBJ authority to change the date	
Cannot save a file restored with partial		and time system values	47
transactions to an earlier release	38	Changes to system QIGC value.	47
Save-while-active processing changes.	38	QPRCMLTTSK system value changed	47
PTF adds additional error checking to the restore		Default for QPWRDWNLMT system value	
function	38	changed	47
DB2 UDB for iSeries changes	38	QUTCFFSET and QTIMZON system value	
Software prerequisite for DB2 UDB for iSeries		changes.	48
inter-operation with DB2 Connect and DB2 UDB		Work management changes	48
for Linux/UNIX/Windows Version 8.	38	New default configuration	48

Fewer "Job started" and "Job ended" messages sent to the history log	49		
Test request key no longer recognized by operating system	49		
Run priority changed for new batch immediate jobs	49		
Jobs that provide system services for independent ASPs are now system jobs	49		
Work management enforcement of object authorities	49		
End Job (ENDJOB) command change	50		
l Host print transform changes	50		
Chapter 3. Options	51		
S/36 and S/38 Migration (Option 4) and S/36 Migration Assistant (Option 11)	51		
Host Servers (Option 12)	52		
Open List APIs moved to library QSYS	52		
Ultimedia System Facilities (Option 16)	52		
OS/400 Portable Application Solutions Environment (Option 33)	54		
Qp2RunPase API change	54		
OS/400 PASE program name on job displays	54		
International Components for Unicode (Option 39)	54		
Chapter 4. Licensed programs	55		
Backup, Recovery and Media Services (5722-BR1)	55		
BRMS product initialization required	55		
Special values for ASPDEV parameter changed	55		
DB2 UDB XML Extender (5722-DE1, Option 2)	56		
IBM HTTP Server (5722-DG1)	56		
Java CGI no longer supported	56		
Configuration update for Tomcat 3.2.4	56		
Error log file expanded	56		
HTTP Server (original) withdrawn	56		
Developer Kit for Java (5722-JV1)	57		
Support for option 3 (JDK 1.2), option 4 (JDK 1.1.8), and JPDC removed	57		
Support for Remote AWT withdrawn	57		
Performance Tools (5722-PT1)	57		
Changes to interactive and batch workload reporting by Performance Tools in V5R3	57		
		l	Performance report percentage values 58
			IBM Content Manager OnDemand for iSeries (5722-RD1) 58
			STRASMOND command changed 58
			DB2 Query Manager and SQL Development Kit for iSeries (5722-ST1) 58
			CRTSQLC command removed 58
			Precompiler commands changed 58
			WebSphere Development ToolSet (5722-WDS) 58
			Application Development Manager (Option 22) and Application Dictionary Services (Option 23) licensed programs withdrawn 58
			New reserved words added to ILE COBOL 60
			IBM Open Class - source and samples (Option 55) withdrawn 60
			ILE C and ILE C++ *PRV Compilers not separate options 60
			Independent ASP support enhancements 60
			iSeries Integration for Windows Server (5722-WSV) 60
			iSeries Access for Windows (5722-XE1) 61
			Support for Application Development folder and Interprocess Communication function in iSeries Navigator 61
			Migration of PC5250 workstation profiles 61
			iSeries Navigator changes 61
			iSeries Access POP server options adjusted 62
			Windows Operating Systems support changes 62
			AFP and SCS Printer Driver considerations 62
			Upgrading from a previous product or release 63
			Distributing iSeries Navigator plug-ins and add-ins 63
			Data Transfer considerations 63
			Ultimedia System Facilities 63
			iSeries Access for Web (5722-XH2) 63
			Notices 65
			Trademarks 66
			Terms and conditions for downloading and printing publications 67

About iSeries Memorandum to Users

This information describes the changes in Version 5 Release 3 Modification 0 (V5R3M0 or V5R3) that could affect your programs or system operations. Use the information in this memorandum to prepare for changes on your current release and to use the new release.

Who should read this memorandum

The Memorandum to Users contains information that is critical for several audiences. This memorandum has four chapters:

- **Read this first** provides information that should be considered before you install V5R3. This chapter is intended for system and application programmers and for the person responsible for system management.
- **Operating system** contains new release changes to basic operating system functions. This chapter includes changes to systems management functions, such as configuration and tailoring the system, as well as changes that could affect the way things operate or appear in the new release. This chapter is intended for all users of the iSeries® servers.
- **Options** provides information about new release changes that affect specific program options of the operating system. This chapter is intended for all users of the iSeries servers.
- **Licensed programs** contains new release changes that might affect existing applications. These changes might also affect applications saved on a V5R3 server to be restored on a previous release server. This chapter is intended for application programmers and system programmers who use the iSeries server and its licensed programs, as well as for businesses with complex networks or application development businesses that have systems at different releases.

Additional incompatibility information

After the publication of the Memorandum to Users, updates to this document will be available in the English Internet version of the iSeries Information Center at this Web site:

<http://www.ibm.com/eserver/iseries/infocenter>

Technical changes to the text are indicated by a vertical line to the left of the change.

- | For other incompatibility information that was not available when this memorandum was published, see
- | PTF cover letters, preventive service planning (PSP), and authorized problem analysis reports (APAR)
- | information at the following Web site:

- | <http://www.systems/support/i> 

- | Click **Technical Databases**.

What's new

The following revisions or additions have been made to the Memorandum to Users since the first V5R3 publication:

- **May 2008 update:**
 - “Host print transform changes” on page 50 was added.
- **December 2007 update:**
 - The link to the Support for System i™ Web site was updated throughout this document.

- Instructions have been provided for locating the readme.txt file for the V5R3M2 version of the iSeries Information Center CD, SK3T-4091.
- **June 2006 update:**
 - “PTF adds additional error checking to the restore function” on page 38 was added.
 - “Application Development Manager (Option 22) and Application Dictionary Services (Option 23) licensed programs withdrawn” on page 58 was modified.
- **April 2006 update:**
 - “Format level identifiers changed for record formats of new SQL tables that have one or more TIMESTAMP columns” on page 44 was added.
- **February 2006 update:**
 - “V5R3 installation requirements” on page 2 was modified to include information about server firmware requirements.
 - “Clustering does not support switchable independent ASP (IASP) with IOP-less disk units” on page 7 was added.
 - “High availability clustering considerations” on page 25 was added.
 - “WAITFILE parameter on CRTPF and CRTLF commands changed” on page 36 was added.
 - “Software prerequisite for DB2 UDB for iSeries inter-operation with DB2 Connect and DB2 UDB for Linux/UNIX/Windows Version 8” on page 38 was modified.
- **October 2005 update:**
 - “V5R3 installation requirements” on page 2 was modified.
 - “Plan for your Operations Console installation or upgrade” on page 3 was modified.
 - “Upgrading to V5R3 using image catalog support” on page 5 was modified.
 - “Modification levels for licensed internal code and operating system can be different” on page 7 was added.
 - “Security update to the /tmp, /var, /QOpenSys/var, and /QOpenSys/var/preserve directories” on page 13 was modified.
 - “Meaning of the console type on the Select Console Type DST screen changed” on page 14 was modified.
 - “Using Service Tools Device IDs in System Service Tools (SST)” on page 19 was modified.
 - “Saving UIM panels to a previous release” on page 25 was added.
 - “Reported utilization for shared partitions” on page 25 was added.
 - “adjtime(), gettimeofday(), and settimeofday() APIs updated to access the system clock” on page 29 was modified.
 - “ILE RPG SQL precompiler enforces unique variable names” on page 44 was added.
 - “Updating rows with NULL value fails with SQL0407 error when NULL values are not allowed” on page 44 was added.
 - “DB2 UDB XML Extender (5722-DE1, Option 2)” on page 56 was modified.
 - “Performance report percentage values” on page 58 was modified.
- **June 2004 update:**
 - IBM® i5/OS® is the next generation of OS/400®. The information in this document may refer to i5/OS as OS/400.
 - “Current customers only — before you install V5R3” on page 1 was modified.
 - “Plan for your Operations Console installation or upgrade” on page 3 was modified.
 - “Server IPL might be required” on page 8 was added.
 - “Virtual IP Address automatic agent selection behavior” on page 15 was modified.
 - “Java Database Connectivity (JDBC) coded character set identifier (CCSID) requirement” on page 24 was added.


- “Remote command and distributed program call server” on page 25 was added.
- “Perform Hardware Configuration Operation (QYHCHCOP) API restriction” on page 36 was added.
- “Error message change when comparing numeric and string fields” on page 43 was added.
- “QUTCOFFSET and QTIMZON system value changes” on page 48 was modified.
- “Independent ASP support enhancements” on page 60 was added.
- “iSeries Access for Web (5722-XH2)” on page 63 was added.

Installing V5R3 over V5R1

Be sure to read the *iSeries Read This First and Memorandum to Users* for V5R2. This document contains incompatibility-related information about the new functions and enhancements incorporated into V5R2. You can order this document by typing the following command:

```
SNDPTFORD SF98076
```

It is also available in the PSP information at this Web site:

| [http://www/systems/support/i](http://www.systems/support/i) 

| Click **Technical Databases > Preventive Service Planning (PSP) > All Preventive Service Planning Documents by Release.**

Discontinued support for certain software and hardware

| It is important that, as a customer, you review and understand all new software release considerations.
| This is especially true of discontinued support for selected software and hardware products or features.
| This information is contained within the iSeries announcement material. To get the most current
| information on discontinued products or features along with suggested replacements, go to the Support
| for System i Planning Web site:

| [http://www/systems/support/i/planning/](http://www.systems/support/i/planning/)

| Select the **Migration & upgrades** link for upgrade planning information.

Installation considerations

Information about the installation and installation-related tasks is included in the book *Install, upgrade, or delete OS/400 and related software*. Information regarding planning and preparing for software installation along with licensed program conceptual and reference information can be found in the iSeries Information Center at:

| <http://www.ibm.com/eserver/iseries/infocenter> 


Click **OS/400 and related software > Install, upgrade, or delete OS/400 and related software.**

PTF numbers in this memorandum

Program temporary fix (PTF) numbers in this memorandum might have been superseded.

Memorandums for previous releases

In addition to ordering previous releases of the *Read this First* and the *Memorandum to Users* using the SNDPTFORD command, you can also view these documents at the following Web site:

<http://www.systems/support/i> 

Click **Technical Databases > Preventive Service Planning (PSP) > All Preventive Service Planning Documents by Release.**

Prerequisite and related information

Use the iSeries Information Center as your starting point for looking up iSeries technical information.

You can access the Information Center two ways:

- From the following Web site:
<http://www.ibm.com/eserver/series/infocenter>
- From CD-ROMs that ship with your operating system order:
iSeries Information Center, SK3T-4091-04.

The iSeries Information Center contains advisors and important topics such as Java™, TCP/IP, Web serving, secured networks, logical partitions, clustering, CL commands, and system application programming interfaces (APIs). It also includes links to related IBM Redbooks® and Internet links to other IBM Web sites such as the IBM home page.

With every new hardware order, you receive the *iSeries Setup and Operations CD-ROM*, SK3T-4098-02. This CD-ROM contains IBM @server iSeries Access for Windows and the EZ-Setup wizard. iSeries Access Family offers a powerful set of client and server capabilities for connecting personal computers to iSeries servers. The EZ-Setup wizard automates many of the iSeries setup tasks.

iSeries Navigator

IBM iSeries Navigator is a powerful graphical interface for managing your iSeries servers. iSeries Navigator functionality includes system navigation, configuration, planning capabilities, and online help to guide you through your tasks. iSeries Navigator makes operation and administration of the server easier and more productive. It also includes Management Central for managing multiple servers from a central system.

You can find more information on iSeries Navigator in the iSeries Information Center and at the following Web site:

<http://www.ibm.com/eserver/series/navigator/>

How to send your comments

Your feedback is important in helping to provide the most accurate and high-quality information. If you have any comments about this memorandum or any other iSeries documentation, fill out the readers' comment form at the back of this memorandum.

- If you prefer to send comments by mail, use the readers' comment form with the address that is printed on the back. If you are mailing a readers' comment form from a country other than the United States, you can give the form to the local IBM branch office or IBM representative for postage-paid mailing.
- If you prefer to send comments by FAX, use either of the following numbers:
 - United States, Canada, and Puerto Rico: 1-800-937-3430
 - Other countries: 1-507-253-5192

- If you prefer to send comments electronically, use one of these e-mail addresses:
 - Comments on books:
RCHCLERK@us.ibm.com
 - Comments on the iSeries Information Center:
RCHINFOC@us.ibm.com

Be sure to include the following:

- The name of the book or iSeries Information Center topic.
- The publication number of a book.
- The page number or topic of a book to which your comment applies.

Chapter 1. Read this first

Current customers only — before you install V5R3

Each of the following publications contains additional information you should read and understand before you install this release. All of the sources referenced in this list are located either on the *iSeries Information Center*, SK3T-4091-04, CD-ROM or on the Internet at the following Web site:

<http://www.ibm.com/eserver/series/infocenter>

Note: After the release of the Information Center, updates will be available in the English Internet version. To review these updates, click **Updates since V5R3 announcement** on the Information Center home page.

| The readme.txt file is missing from the V5R3M2 version of the iSeries Information Center CD,
| SK3T-4091. To access the readme.txt file, run the **install_win.exe** script under the root of the CD.

Access the Internet version or place the CD in the CD-ROM drive of your PC to view its contents. The instructions refer you to several sources:

- | • The iSeries Information Center Install, upgrade, or delete OS/400 and related software topic contains software preinstallation information and information on installing or upgrading the operating system release, part of the release, or related licensed programs. You can also order a printed version of this PDF (SC41-5120; feature code 8004) with your software upgrade orders or new hardware orders.
- | • The Preventative Service Planning (PSP) Information provides information about software problems you may encounter as you install the new release. You can download PSP by using electronic customer support from the Support for System i Web site:

| <http://www/systems/support/i>

Alternatively, you can obtain PSP from your software service provider. The following items are included in PSP:

- The PSP identifier for information that pertains to installing V5R3 is SF98080. Information within the PSP is grouped by product area. To receive this PSP information by using electronic customer support, type the following command on an iSeries command line:

```
SNDPTFORD SF98080
```

- The PSP identifier for information that pertains to problems discovered since the availability of the current cumulative PTF package is SF98530. Information within this PSP describes all PTFs released since the current cumulative PTF package began shipping. It also contains information on all known high-impact and pervasive problems that are not included in the latest cumulative PTF package. To receive this PSP information by using electronic customer support, type this command on an iSeries command line:

```
SNDPTFORD SF98530
```

- The PSP identifier for information that pertains to installing V5R3 hardware is MF98530. Review this PSP information before you install either new iSeries servers or hardware devices. To receive this PSP information by using electronic customer support, use this command:

```
SNDPTFORD MF98530
```

- The PSP identifier for information that pertains to server upgrades and data migrations is SF98167. Information within this PSP describes upgrade and migration corrections. Review this PSP information before you upgrade your server model or migrate your data between servers. To receive this PSP information by using electronic customer support, type the following command on an iSeries command line:

- The iSeries PTF maintenance strategy. A PTF maintenance strategy is recommended for all iSeries customers. This may reduce the impact to iSeries system operations that result from unplanned outages or program failures. For more information on the iSeries maintenance strategy, do the following:
 1. Go to: <http://www/systems/support/i>.
 2. Click **Fixes**, and then click **Fix Central**.
- The iSeries Operations Console overview (**Connecting to iSeries > Operations Console**).

Notice

Operations Console support is available on V5R1 and later releases of the operating system. In V5R3, the only type of PC console that iSeries Models 270, 800, 810, 820, 825, 830, 840, 870, and 890 support is Operations Console.

V5R3 installation requirements

Successful installation of V5R3 requires the following for each server or logical partition:

- The V5R3 Licensed Internal Code requires more storage space on the load source than in previous releases. To upgrade to the V5R3M0 Licensed Internal Code, you must have a load-source disk unit of 4 GB or larger on each server or logical partition. For optimum performance, at least 8 GB is recommended. To upgrade to the V5R3M5 LIC (licensed internal code), you must have a load-source disk unit of 17 GB or larger. Before you upgrade, ensure that the server meets disk storage requirements for upgrades; in the Information Center, click **Install, upgrade, or delete i5/OS and related software > Upgrading or replacing i5/OS and related software > Preparing to upgrade or replace software > Performing initial upgrade or replacement tasks > Ensuring the server meets disk storage requirements for upgrades**.
- In addition, all server models with V5R2 or earlier installed require additional reserved storage before V5R3 can be installed. All server models with V5R3M0 Licensed Internal Code installed require additional reserved storage before the V5R3M5 LIC (licensed internal code) can be installed. Your upgrade will stop during the installation if you do not allocate the additional space. In the Information Center, click **Install, upgrade, or delete i5/OS and related software > Upgrading or replacing i5/OS and related software > Preparing to upgrade or replace software > Performing initial upgrade or replacement tasks > Required: Allocating additional space for Licensed Internal Code** for instructions.
- A minimum memory size of 128 megabytes. Systems with less than this minimum memory fail the installation process for Licensed Internal Code. Memory requirements for partitions are dependent on partition configuration, I/O resources assigned, and applications used. The primary partition needs a minimum of 256 megabytes of memory. Depending on configuration values used, a primary partition may require more than 256 megabytes. Secondary partitions running V5R1 and V5R2 require a minimum of 128 megabytes of memory. Also depending on configuration values used, a secondary partition may require more than 128 megabytes. For more information on memory requirements for logical partitions, see the Logical partition concept: memory topic in the iSeries Information Center.
- Server firmware requirements before you install V5R3M5 Licensed Internal Code: To determine the minimum firmware requirements for your system, you need to understand both the hardware and software requirements. The IBM eServer™ Prerequisite Web page provides the prerequisite information for features that you currently have or plan to add to your system. If you have an IBM System i5™, System p5™, or IBM eServer i5 or p5 system, the server firmware must be at a minimum level of SF235_160 for support of V5R3M5 Licensed Internal Code. The minimum level of firmware is included in the V5R3M5 Licensed Internal Code.
- Some products have specific requirements of their own. Review all product requirement information that you receive with your order.

| To determine the appropriate V5R3 level of the operating system for your server and partitions, see the
| **Installation & setup** or **Migration & upgrades** topics at the following Web site:

| <http://www.systems/support/i/planning/migrationupgrade.html> 

| Also, click **Technical Databases > Authorized Program Analysis Reports APARS** at the following Web
| site:

| <http://www.systems/support/i> 

| For upgrade planning information, go to the following Web site:

| <http://www.systems/support/i/planning/upgrade> 

| More information on partition release requirements for all iSeries models is in the Logical partition
| release support topic in the iSeries Information Center.

Process for ordering cumulative PTF packages

You will not receive a cumulative PTF package with your software order. To ensure consistent cumulative PTF packaging and to be able to get the latest PTFs, IBM has established a single delivery source. This single source is Corrective Service.

When you order the Cumulative PTF Package for V5R3 (SF99530) from Corrective Service, you will receive the latest cumulative PTF plus DB (Database) Group PTFs and HIPER (High Impact Pervasive) Group PTFs to install with your software upgrade. The cumulative PTF package that in the past may have been delivered with your software order did not include the HIPER or Database Group PTFs.

| If an extended time has elapsed since you received your order, a later cumulative PTF package might be
| available. You can order cumulative PTF packages either by using the SNDPTFORD command, by using
| Fix Central, or by contacting your software service provider. You are advised to check to see if you have
| the latest cumulative PTF package. To view the latest available cumulative PTF package identifier for
| your software release, go to the Support for System i Web site:

| <http://www.systems/support/i> 

| Click **Technical Databases > Preventative Service Planning - PSP**.

Refer to the software installation preventative service planning (PSP) document, SF98080, for the release of the operating system software that you are installing.

| If you do not have the latest cumulative PTF package for your operating system software installation,
| place an order for the package 7 to 10 working days before you plan to install the software to allow
| sufficient delivery time. You can view PTF ordering options through the Internet at the Support for
| System i Web site. Click **Technical Databases > PTF Ordering**. Also, you can download the cumulative
| PTF package using FTP by specifying that option in Fix Central; this option provides faster turnaround
| than the 7-10 days required for the physical package to arrive.

Plan for your Operations Console installation or upgrade

| If you are upgrading to V5R3 and you want to replace an existing console with Operations Console,
| upgrade the system before migrating the console. This will prevent any conflict between the existing
| console and the Operations Console. For instructions on upgrading your operating system, click **i5/OS
| and related software > Install, upgrade, or delete i5/OS and related software > Upgrade or replace
| i5/OS and related software**.

| **Prerequisite information for Operations Console users upgrading to, or installing, V5R3:**

| You must comply with the following before upgrading or installing your software (operating system, Licensed Internal Code) to V5R3:

- | 1. If you use the Operations Console configuration, Local console directly attached to the server with a server that has no logical partitions, and the adapter used for this connection is a 2771, then use the following table to verify the correct location of this adapter. This does not pertain to any POWER5™ systems or servers.

| *Table 1. Card location*

iSeries model	Operations Console asynchronous card location for the cable
170 or 250	C08
270	C07
720	C09
730 or 740	See cabling information.
800 or 810	C07
820	C06
825	C06
830 or SB2	C02
840 or SB3	C02
870 or 890	C02

- | 2. For all upgrades and installations, you need to establish a connection between the server and Operations Console PC using the 11111111 (eight 1's) service tools user ID. The default password for this user ID is 11111111; however, this password might have been changed following a previous installation. This default user ID ensures a successful re-authentication of the client connection to the server. When you receive the operating system release upgrade, the shipped service tools user IDs (except 11111111) are expired. To re-authenticate the client connection to the server, enter the service tools user ID of 11111111 (eight 1's) and either the default password of eight 1's or the password you might have previously created for this user ID. This is especially important for automatic installations.
- | 3. It is recommended that you update iSeries Access for Windows® to V5R3 before you upgrade the operating system. For more information, see the Install iSeries Access for Windows topic in the iSeries Information Center.

| **Note:** Failure to comply with the above actions may prevent the console from working correctly during the upgrade or installation.

| **Important:** During a manual IPL of the server, and if no console has been specified before, you will receive two additional screens to confirm the setting of the console mode. The first will require a F10 to accept your current console type and the second will show that a value did not previously exist (a zero will be present for the old value) and the new value will be shown. Pressing Enter will exit and set the console mode automatically. The IPL will then continue to the IPL or Install the System screen. This condition is most likely to occur during the installation of a new partition but could happen on your first manual IPL of V5R3; for example, the A-mode IPL following the restore of Licensed Internal Code during the upgrade or install when a console value of zero is found.

| **Migrating to Operations Console before your server model upgrade**

| If you are going to be utilizing Operations Console on your new iSeries server (migrating from a different console type), it is important that you configure the new Operations Console PC before the beginning of the server model upgrade. At the point in the upgrade instructions where console functions are required

| on the new iSeries server, you will be able to perform any required functions without the need for your current console device. The Operations Console features matching the connectivity you plan to use should be specified as part of the order for your new iSeries server.

| **New models**

| If you are using a model 5xx that uses a Hardware Management Console (HMC), you can change from HMC to Operations Console or from Operations Console to HMC. For details on how to do this, see the Changing the console that manages i5/OS topic in the @server Hardware Information Center (see **Managing consoles, interfaces, and terminals > Changing consoles, interfaces, and terminals**).

| Embedded ports on models 5xx cannot be used for Operations Console. Currently, the model 825 is the only model using an embedded port for Operations Console.

| **Switching console control to another user**

| If you use Operations Console and run a 5250 emulation console, you can use a new option on the Select Console window, called **Allow console to be taken over by another console**. This option allows users of local consoles on a network to take control of the console from another user, if needed.

| **Recovery**

| The **Allow console to be taken over by another console** option also controls a new recovery function that allows the recovery of the console without the loss of data or the job. For details of what this option does, see the Operations Console topic in the iSeries Information Center.

| **Upgrading to V5R3 using image catalog support**

| When you use image catalog support to upgrade to V5R3 from V5R1 or V5R2, you must apply the following PTFs before you prepare your image catalog for the upgrade:

- | • For V5R1, use SI10553
- | • For V5R2, use SI10563

| The Virtual media PTFs required for the corresponding releases are listed below:

- | • V5R1M0: SI17565
- | • V5R2M0: SI17564
- | • V5R3M0: SI17566

| You need to apply the V5R3M0 PTF if you are installing a later version of V5R3M0.

| **I/O software prerequisites**

| For the latest information on software prerequisites for new I/O features and new I/O expansion units that are available, see Information APAR II13440 at the Support for System i Web site:

| <http://www/systems/support/i> 

| Click **SEARCH Technical Databases > Authorized Program Analysis Reports APARS**.

Support for I/O expansion units

iSeries models announced in January 2003 (800, 810, 825, 870, and 890 models) support PCI-attached I/O expansion units rather than SPD-attached features. More information on current support for I/O features is available at the following Web address:

| <http://www/systems/support/i/planning/upgrade/future.html> 

Enterprise Storage Server (ESS) SAN attachment prerequisites

The IBM TotalStorage® Enterprise Storage Server® (ESS) must be upgraded to a minimum level of 2.2.0 before you install or upgrade to V5R3. For additional information related to iSeries attachment to ESS, see the *ESS interoperability matrix* and the *ESS host systems attachment guide* at the following Web site:

| <http://www/systems/storage>

Chapter 2. Operating system

This chapter describes changes to the operating system and its functions. Changes to systems management functions, such as configuring and tailoring the system, are also included.

Programming considerations

Output file (OUTFILE) changes

Applications using LVLCHK(*YES) may be affected by changes made to IBM-shipped system output files in this release. IBM commands and APIs that generate database output files added new fields to the end of record formats for additional information returned this release. Adding these new fields to the record format changed the level check value for the file. This might cause an application with LVLCHK(*YES) to fail with a level check error. If a level check error does occur, please review the application to determine which system file it uses. New fields have been added to IBM-shipped database files in each release of AS/400® and iSeries.

Security audit record changes

Changes made to security auditing for this release may affect applications that read those audit records. Actions that were not audited in previous releases may now be audited. Existing audit records may have been changed by the addition of new fields in a reserved area of the audit record or at the end of the audit record. Existing fields may contain new values. Applications that read the audit records should be changed to tolerate these types of changes.

Programs that use customized versions of IBM-supplied commands

Some operating system functions that use IBM-supplied CL commands that are not library-qualified in this release may be changed in a future release to specify a specific library of *NLVLIBL or *SYSTEM for the library qualifier. Applications that depend on using their own version of commands instead of the IBM-supplied commands might not work as they had on earlier releases. These applications should be changed to use the retrieve command exit point (QIBM_QCA_RTV_COMMAND) or the change command exit point (QIBM_QCA_CHG_COMMAND) that allows your exit program to get control and possibly change the command that is used.

Clustering does not support switchable independent ASP (IASP) with IOP-less disk units

Switchable IASPs with IOP-less DASD require V5R4. The model 520 with advanced POWER5 processors running with V5R3M0 of i5/OS and V5R3M5 Licensed Internal Code and IOP-less DASD will not support switchable IASPs. IASPs would need to have IOPs in the IASP configuration, or the system would have to be upgraded to V5R4M0 i5/OS and V5R4M0 Licensed Internal Code.

Modification levels for licensed internal code and operating system can be different

Prior to V5R3M0, the licensed internal code (product ID 5722-999) and the *BASE option of the operating system (product ID 5722-SS1) have been required to be the same version, release, and modification (VRM) level.

With the V5R3M5 LIC (licensed internal code), the version and release level numbers must be the same, but the modification level numbers can be different. This could affect you if, for example, you retrieve the VRM of the operating system and assume the VRM of the licensed internal code is exactly the same.

| There is not a separate cumulative PTF package for the V5R3M5 LIC. The V5R3M0 cumulative PTF package contains PTFs for both V5R3M0 and V5R3M5. You should order the cumulative PTF package using the VRM of the operating system, SF99530.

| The Save System (SAVSYS) command saves to optical media three additional files that distinguish the level of the licensed internal code from the level of the operating system. To use this function, you must install PTF SI17733. The following new files indicate the licensed internal code information:

- | • /QLANGID/SVvRrMm
- | • /QLANGID/S_L29nn
- | • /QLANGID/S_LVLnn

| The following files continue to save the operating system information:

- | • /QLANGID/QVvRrMm
- | • /QLANGID/Q_L29nn
- | • /QLANGID/Q_LVLnn

Prepare for Install required during installation process

Prior to V5R3, the Prepare for Install function (Option 5 from the GO LICPGM menu) was an optional step. In V5R3, the online software agreement acceptance step makes the Prepare for Install step mandatory. You must accept software agreements before you upgrade to V5R3, because the upgrade overlays the previous release, making it very difficult to return back to a previous release if you do not want to accept the terms of the agreement. The Prepare for Install function allows you to review and accept these agreements before installing V5R3. If you do not complete this step, the upgrade process to V5R3 will be stopped by the Licensed Internal Code install.

Server IPL might be required

If you have an @server i5, you might be required to perform a server IPL after you apply or remove a PTF on the service partition, if that PTF affects the server firmware portion of the Licensed Internal Code. You also may be required to perform a server IPL after you upgrade or install a new level of the Licensed Internal Code on the service partition.

A server IPL is an IPL in which all logical partitions on the server are shut down at the same time; this allows, for example, a new level of the server firmware to be activated on the server.

Collection Services and performance database file changes

Two major changes have occurred within collection services to better support dynamic and shared processor partition environments. If you use the performance database files or the performance collector API QPMLPFRD, and work with shared processor partitions or partitions whose CPU resources are changed without an IPL, you need to be aware of the following:

- Collection services will no longer cycle its collection when a partition's configuration is changed. The values reported for virtual processors and processor units (found in the QAPMCONF and QAPMSYSCPU files) represent the current value at the time sampled. These values, if changed, will not be representative of conditions throughout an interval or collection. Consequently, calculations based on those values (for example, CPU percent) will be incorrect for periods involving change.
- In a shared processor partition, collection services no longer scales partial processor CPU data to whole virtual processors. In addition, the HVLPTASK entry will no longer appear in the QAPMJOBMI/QAPMJOB files.

Prior to V5R3, the CPU reported for virtual processors (QAPMSYSCPU/QAPMSYSL) and the CPU reported for secondary workload and database (QAPMSYSTEM/QAPMSYSL) was scaled (increased) to what the utilization would be if a whole virtual processor was allocated, rather than a partial processor. Scaling was accounted for in the list of jobs by the reporting of an additional task called

HVLPTASK. This allowed system utilization to be calculated in the same manner as previously done (that is, time consumed divided by elapsed interval time).

In V5R3, all CPU reported will be actual time consumed. (Note that interactive CPU and individual job CPU have always been reported as actual CPU time consumed.)

In V5R3, a new method for calculating CPU utilization eliminates the problems associated with dynamic configuration changes and shared (partial) processor partitions. Although this methodology can be used in all environments, no changes are required if you do not use shared processor partitions or if you do not change CPU resources.

All system processor utilization (time consumed) is now being reported in QAPMSYSTEM/QAPMSYSL. This includes total system CPU and interactive CPU. You no longer need to sum individual processor utilization to obtain system CPU.

The time available for resources like system CPU and interactive CPU is now reported in QAPMSYSTEM/QAPMSYSL. By using these new available capacity metrics, you no longer need to be concerned with configuration issues, such as the number of processors, processor units, and interactive thresholds, to properly calculate percent utilized. These metrics already account for changes in configuration over time.

Existing (unchanged) tools that calculate CPU utilization will not show the correct results for shared processor partitions or partitions that have had configuration changes during data collection. This includes those that use the performance database as well as those that use the QPMLPFRD API.

When performance data from a prior release is moved to V5R3 (a management collection object *MGTCOL or V5R2 database files) and that data is converted or generated into V5R3 database files, it will be modified to be consistent with the V5R3 reporting. You might see differences between the V5R3 view of the data and the prior release view of the data, especially when significant scaling (HVLPTASK activity) occurred in the prior release. These differences will also affect the percentage values reported in performance reports generated by the Performance Tools product (5722-PT1). See "Performance report percentage values" on page 58 for more information.

A V5R3 management collection object can still be copied to a prior release, and the database files generated. However, if that object is from a shared processor partition, the reported QAPMSYSCPU data will remain unscaled. This means that total system CPU reporting by tools using virtual processors will be incorrect. Also, if that *MGTCOL contains data that spans partition configuration changes, CPU percent calculations might be incorrect during intervals when a configuration change occurred, or they might be incorrect for all intervals where the configuration is not the same as it was at the beginning of the collection.

XML procedural parser service programs removed

The XML for C++ and procedural parsers, previously shipped in service programs QXML4C310 and QXML4PR310 in library QSYS, have been removed in V5R3. XML parser technology is evolving at a rapid pace with new release levels being delivered more frequently than operating system releases. Also, new releases are not necessarily compatible with prior XML parser releases. XML parser technology is available through the XML Toolkit for iSeries product (5733-XT1).

Performance Explorer statistics change

In V5R3, the QAYPESTATS file contains one record for each program, module, or procedure for each job on the system, regardless of the value specified for the MERGE parameter on the Add Performance Explorer Definition (ADDPEXDFN) command.

In previous releases, if MERGE(*YES) was specified, information for the same program, module, or procedure in multiple jobs was merged into a single record.

Default values for pool tuning options changed on systems with 100% interactive capacity

On the Display tuning data screen (F11) of the Work with Shared Pools (WRKSHRPOOL) command, the default values for *Priority* and *Faults/Second* changed on systems with 100% interactive capacity on V5R3. The *BASE pool and the *INTERACT pools both have a priority of '1' and the faults per second values have also changed. The tuning options on this display affect how the QPFRADJ performance adjuster makes adjustments to the pools.

Active user profile now stored with each message

In V5R3, the active user profile is stored with each message. The *active user profile* is the user profile that the job is doing work for. This can be different than the user name from the qualified job name if the job had swapped to run under another user profile. Another name for active user is *current user*.

For job logs, this can mean that the job log will wrap or fill up with fewer messages. For standard message queues, this means that a message queue might hold fewer messages before it is full.

Many messages will not increase in size. This is because, for standard message queues, the active user name is saved only if it is different than the name in the qualified job name.

The *From* user shown on the Display message (DSPMSG) command, the Work with message (WRKMSG) command, and additional message information screens is now the actual sending user, rather than the user name from the qualified job name of the sending job.

User profiles to which you are not authorized are not shown for DSPLIB of QSYS library

In V5R3, when you display the QSYS library with the Display Library (DSPLIB) command, user profile (*USRPRF) objects to which you are not authorized are not returned in the list of objects.

Tape management system changes

Tape cartridges are no longer ejected when a cartridge is mounted in a device and in the *INSERT category.

The non-printable characters in the message replacement text that is passed to the Tape Management exit program are no longer blanked out, so the exit program processes the actual binary data in the replacement text.

SNTP client changes the system clock instead of the software clock

In V5R3, the SNTP client changes the system clock instead of the software clock. If you want to have the software clock changed, or if you install V5R3 on a server that does not support system clock adjustment (iSeries models 7xx, 170, and 250), you must create a data area (*DTAARA) object named QTOTSNTP. Use the following CL command to create this data area object:

```
CRTDTAARA DTAARA(QUSRSYS/QTOTSNTP) TYPE(*CHAR) LEN(9) VALUE(*SOFTWARE) AUT(*USE)
```

To switch back to having the SNTP client change the system clock, run the following CL command to delete the data area object:

```
DLTDTAARA DTAARA(QUSRSYS/QTOTSNTP)
```

Integrated file system changes

*TYPE2 directory conversion

The "root" (/), QOpenSys, and user-defined file systems (UDFS) in the integrated file system support the *TYPE2 directory format as of OS/400 V5R1. The *TYPE2 directory format is an enhancement of the original *TYPE1 directory format. *TYPE2 directories have a different internal structure from *TYPE1 directories and provide improved performance and reliability.

Shortly after V5R3 is installed, the system will automatically begin converting any file systems that currently support *TYPE1 directories to *TYPE2 directories. This conversion should not significantly impact your system activity, as it will run in a low priority background job. The conversion will run until all directories have been converted to the *TYPE2 format. This means it will be restarted after every IPL until all eligible file systems have been converted. The order of conversion will be the "root" (/) file system, QOpenSys file system, and user-defined file systems in auxiliary storage pools (ASPs) 1 through 32, in that order. This might cause additional CPU usage and an increased number of input and output operations on the iSeries server. This is normal. The CPU utilization and the I/O rate will return to normal after the directory conversion function has completed. To check the status of the directory conversion function, use the Convert Directory (CVTDIR) command with OPTION(*CHECK). The priority of the directory conversion function can be controlled using the CVTDIR command with OPTION(*CHGPTY). For more information, see Convert directories from *TYPE1 to *TYPE2 in the Integrated file system topic in the iSeries Information Center.

Network Server Description Considerations

If you have a varied-on network server description (*NWSD) that was created on a release prior to V5R1M0, the automatic *TYPE2 directory conversion may wait until the *NWSD is varied off. This is normally not a concern, although the directory conversion will not finish as long as the *NWSD is varied-on. If you have such a *NWSD, and the Convert Directory (CVTDIR) command with OPTION(*CHECK) specified shows that the directory conversion is not progressing, consider running the Vary Configuration (VRYCFG) command to vary off the *NWSD, which will allow directory conversion to resume. Then, after waiting for a short time, vary it back on.

You can use the Display Object Links (DSPLNK) command to check if your *NWSD is one that may affect directory conversion as described above.

1. Run the DSPLNK command with DSPOPT(*ALL) specified on the /QFPNWSSTG/xxxxx directory, where xxxxx is the name of your *NWSD. If no directory is found, then your *NWSD will not affect directory conversion and there is nothing more to do. If the directory is found, proceed with the next step.
2. Select option 8 (Display attributes) from the Display Object Links main panel next to the directory name. Page down on the Display Attributes panel to see the Directory format field. If it has a value of *TYPE1, then directory conversion is running. Proceed with the next step. If it has a value of *TYPE2, the conversion of your *NWSD directory is complete and there is nothing more to do.
3. Return to the Display Object Links main panel and select option 5 (Display) next to the directory name. Select option 8 (Display attributes) next to the QFPCLTSTG1 file name. Page down on the Display Attributes panel to see the 'File format' field. If it has a value of *TYPE1, then this *NWSD may affect directory conversion as described above. If it has a value of *TYPE2, this *NWSD will not affect directory conversion and there is nothing more to do.

Reclaim Storage (RCLSTG) restriction

In V5R3, while the *TYPE2 directory conversion is running, the Reclaim Storage (RCLSTG SELECT(*ALL)) command requires that directory processing be omitted by specifying the OMIT(*DIR) parameter. If you want to perform a reclaim of directories, use the CVTDIR command with the OPTION(*CHECK) parameter to display the current status of the directory conversion. When the conversion has completed for all file systems, you can run a RCLSTG that includes directories.

Convert Directory (QP0FCVT2) API and Convert Directory (CVTDIR) command

The Convert Directory API (QP0FCVT2) that was used in previous releases to convert directories is no longer supported. When called, it will now return escape message CPF9890 (Function not supported). The Convert Directory (CVTDIR) command is still supported, but it can no longer request a directory conversion. Instead, it provides status information about the directory conversion that is taking place and allows you to change the priority of that directory conversion function.

Management Central changes

QYPSSRV job no longer exists

Starting in V5R3, the Management Central server job QYPSSRV no longer exists.

Default DNS lookup frequency changed

Starting in V5R3, the default domain name server (DNS) lookup frequency value for Management Central is *always* rather than *never*, if the value was never explicitly set.

Hardware and software inventory no longer published to LDAP database

From V4R4 through V5R2, Management Central published hardware and software inventory information to the LDAP (Lightweight Directory Access Protocol) database. Starting in V5R3, Management Central no longer publishes this information to LDAP. This inventory information can be exported to a file on the PC of a Management Central user by using the Inventory function in Management Central. This export support has been available since V4R4 of Management Central.

Management Central user data moved from QUSRSYS library

In prior releases, Management Central user data was stored in the QUSRSYS library. Beginning with V5R3, this data will be stored in two new libraries, QMGTC and QMGTC2. Any existing Management Central user data, except for inventory data in the QUSRSYS library, will be automatically moved to the QMGTC and QMGTC2 libraries. If you have backup and recovery programs that specify individual IBM user libraries to be saved or restored, you should add QMGTC and QMGTC2 to the list of libraries. The QMGTC and QMGTC2 libraries are automatically included if LIB(*ALLUSR) is specified for any command, such as Save Library (SAVLIB) or Save Changed Objects (SAVCHGOBJ).

Inventory information from prior releases not migrated

In V5R3, the record formats for database files where Management Central inventory information is stored for certain Management Central functions (Hardware, Software, System Values, Users and Groups, and PTFs) changed. Inventory information from a prior release will not be migrated. If needed, this same information can be refreshed the first time you run the Inventory Collection function in Management Central.

Graph history changes

Graph history no longer dynamically creates detailed graph data at the time of the query. For a V5R3 endpoint system, you now must explicitly select the option to create the graph data. For V5R3 iSeries Navigator clients, this option is available on a collection object. For pre-V5R3 iSeries Navigator clients, select the option to create the graph data when the collector cycles.

Security update to the /tmp, /var, /QOpenSys/var, and /QOpenSys/var/preserve directories

Starting in V5R3, the /tmp, /var, /QOpenSys/var, and /QOpenSys/var/preserve directories might have the Restricted rename and unlink attribute set to Yes. (Note: this attribute is equivalent to the S_ISVTX mode bit for a directory.) This is being done for operating system commonality and security purposes. If this attribute value is Yes, you cannot rename or unlink objects within these directories unless one of the following is true:

- You are the owner of the object.
- You are the owner of the directory.
- You have all object (*ALLOBJ) special authority.

If the attribute value is Yes, and you do not have the appropriate authorities, you will see error number 3027 (EPERM) or message CPFA0B1 (Requested operation not allowed. Access problem.) for rename or unlink failures. The following commands and APIs are affected:

- Remove Link (RMVLNK, DEL, and ERASE) command
- Remove Directory (RMVDIR, RD, and RMDIR) command
- Rename Object (RNM and REN) command
- Move Object (MOV and MOVE) command
- Rename File or Directory (rename) API
- Remove Directory (rmdir) API
- Remove Link to File (unlink) API

You can modify this attribute (or mode bit) using the Change Attribute (CHGATR) command or the Qp0lSetAttr or chmod APIs, provided you are the owner of the object, or you have all object (*ALLOBJ) special authority. But, if the attribute is changed to No, then you will lose the operating system commonality and security benefits the Yes setting provides.

Whenever the /tmp directory must be created during an IPL, the attribute will be set to Yes. If the /tmp directory already exists at IPL time, the attribute will not be changed. The remaining directories will have the attribute changed to Yes during an i5/OS Option 33 (i5/OS PASE) installation, regardless of whether or not they have to be created during the installation.

Limit of one Java virtual machine (JVM) per process

In V5R2, a special Java property was required in order to create multiple Java virtual machines for a single operating system process. In V5R3, this property is ignored and it is not possible to create multiple Java virtual machines for a single process.

If JNI_CreateJavaVM is called more than once in a single process, an error code (JNI_EEXIST -5 VM Already Created) is returned. This change was made to the iSeries JVM in order to make it consistent with established JVM implementations for other operating systems.

Prior to V5R3, the JNI_GetCreatedJavaVMs function could return more than one Java virtual machine in its list of Java virtual machines. In V5R3, JNI_GetCreatedJavaVMs returns at most one Java virtual machine.

Meaning of the console type on the Select Console Type DST screen changed

The meaning of the console type on the Select Console Type DST screen has changed. In previous releases, the value represented the maximum allowed type. For example, a value of 3 meant LAN console was allowed, Direct console was allowed, and twinaxial console was allowed.

In V5R3, the meaning of console type was changed to indicate that only the selected type of console is allowed. When option 3 (LAN) is selected, only LAN console is allowed.

In V5R3, a new function shows data for a console type other than the specified console, which resolves the display of blank screens that had been a problem in earlier releases. This function works only on 5250-based consoles (Operations Console and HMC); twinax consoles are not supported.

When a 5250-based device connects, it is either going to be the console or eligible to become the console; however, this does not imply that an eligible console can take over the console session without change. For example, if a LAN console is defined and two devices connect, the first device becomes the console and displays what the console would normally present for the state of the system. The second device displays either a special DST signon screen (if you have enabled the **Allow console to be taken over by another console** option), or the Console Information Status screen (if the **Allow console to be taken over by another console** option is not enabled).

In another example, if a console mode is anything other than HMC, an HMC displays a message during an attempt to connect to a logical partition that the partition is not configured for HMC console.

No console-capable device should have a blank screen on initial connection.

For more details regarding the behavior of connecting console devices see the Take over or recover an Operations Console connection topic in the Operations Console collection in the iSeries Information Center (click **Connecting to iSeries > Operations Console > Manage Operations Console > Manage your multiple consoles > Take over or recover an Operations Console connection**).

Program QSPGETF to be removed in future release

The QSPGETF program has never been documented or supported as an operating system API. In the release following V5R3, this program will be removed.

In preparation for this, you should remove any references to this program from your application programs. You should replace calls to QSPGETF with calls to documented and supported APIs for working with spooled files:

- QUSRSPLA
- QSPCRTSP
- QSPOPNSP
- QSPGETSP
- QSPPUTSP
- QSPCLOSP
- QUSLSPL

Appendix A of the IBM Redbook *IBM AS/400 Printing III* (GG24-4028-00), demonstrates how you can use the supported APIs to perform the same function as the QSPGETF program.

IBM Backup Recovery and Media Services for iSeries and IBM Content Manager OnDemand for iSeries are examples of products that provide support for saving and restoring spooled files.

Coded character set identifiers (CCSIDs) with bidirectional encoding schemes

If you use coded character set identifiers (CCSIDs) with bidirectional encoding schemes, you may be affected by the change described here.

Prior to V5R3, integrated file system commands would incorrectly process path names when they were supplied in certain CCSIDs, including Hebrew EBCDIC (424), and Arabic EBCDIC (420, 12708). This condition may be difficult to detect if you only use CL commands to access objects in the integrated file system because all of the commands use the same incorrect path name processing. The problem becomes apparent when using commands to access objects created through APIs, or vice-versa. The incorrect path name processing is being corrected in V5R3M0.

This change may cause operations to fail that previously succeeded. The most common errors would involve indications that an object could not be found. This is because the order of characters in affected path names may appear to be the reverse of the intended order. Affected objects must be renamed to correct these problems.

The following is an example on how to rename an affected object:

The affected object is named FILE1. The path to this object is /directory1. Run the following CL command:

```
WRKLNK OBJ('/directory1/*')
```

The affected object will appear as 1ELIF. On the option field next to the affected object, enter option 7 (Rename). Enter the correct object name and press the F5 (Refresh) key to verify the object has been renamed.

Virtual IP Address automatic agent selection behavior

In previous releases, when a Virtual IP Address (VIPA) interface that was configured to use proxy address resolution protocol (ARP) was started, the server always selected the first local area network (LAN) adapter that had an active IP address configured whose network address range included the VIPA. If more than one VIPA was configured for Proxy ARP on the same subnet, they would use the same LAN adapter.

Starting with V5R3, the server will perform a round-robin search of all appropriate LAN adapters that have an active IP address configured whose network address range includes the VIPA. If more than one VIPA is configured, they will be balanced across the adapters.

TCP/IP server values *DCE and *NSMI removed

For V5R3, values *DCE and *NSMI are not valid for the SERVER parameter on the Start TCP/IP Server (STRTCPSVR) and End TCP/IP Server (ENDTCPSVR) commands.

Network Quality of Service (QoS) policy

If you configured a Network Quality of Service (QoS) policy in V5R1 or V5R2

After upgrading to V5R3, you must run the V5R3 iSeries Navigator QoS Configuration Wizard to convert the existing policy to V5R3 format so that the V5R3 QoS server can load the rule set into the QoS Manager. This means that the prior QoS policy is not active until you have run the Configuration Wizard. If you start the QoS server before you run the Configuration Wizard, the rules are not loaded and a

message is sent to the QSYSOPR message queue indicating that you must run the wizard. You can then run the wizard to either restart the QoS server or select the server update function from iSeries Navigator.

You can store the QoS configuration data in either the LDAP directory or a file on the server.

Running multiple iSeries servers using QoS

If you run multiple iSeries servers using QoS, see the QoS topic in the iSeries Information Center for information about migration options for multiple systems.

Enable ECN support added

In V5R3, a record in the TCP/IP attributes file (QUSRSYS/QATOCTCPIP) was added to support explicit congestion notification (ECN). The record will be added to the TCP/IP attributes file the first time the file is accessed. After conversion takes place, the V5R3M0 version record in the file shows that the file has been updated with the new record.

Required authorities changed for security tools CL commands

Prior to V5R3, the following commands required all object (*ALLOBJ) special authority and were shipped with no public authority (*EXCLUDE).

Starting in V5R3, the commands are shipped with public authority of *USE, and only a user who has *AUDIT special authority (as well as users who have *ALLOBJ and any other required special authorities) can run the commands.

- DSPSECAUD (Display Security Auditing)
- PRTADPOBJ (Print Adopting Objects)
- DSPAUDJRNE (Display Audit Journal Entries)
- PRTPVTAUT (Print Private Authorities)
- PRTPUBAUT (Print Publicly Authorized Objects)
- PRTCMNSEC (Print Communications Security)
- PRTJOBDAUT (Print Job Description Authority)
- PRTQAUT (Print Queue Authority)
- PRTSBSDAUT (Print Subsystem Description)
- PRTSYSSECA (Print System Security Attributes)
- PRTRGPGM (Print Trigger Programs)
- PRTUSROBJ (Print User Objects)
- PRTUSRPRF (Print User Profile)

Cryptographic product and DCM needed for IBM Universal Connection

If you had previously configured operating system service functions to use the AT&T or Multihop connectivity options of the IBM Universal Connection, in V5R3 you will need operating system option 34 (Digital Certificate Manager) and one of the Cryptographic Access Provider products (5722-AC1, 5722-AC2, or 5722-AC3) installed to continue to use those service functions. In prior releases, these were only required when the IBM Universal Connection configuration options established a direct Virtual Private Network (VPN) through a customer-provided Internet Service Provider (ISP).

Directory Server (LDAP) changes

Configuration changes

In V5R3, several changes related to the IBM Directory Server on iSeries require your attention:

- A Web Administration tool is provided for configuring the IBM Directory Server (IDS) on iSeries. You must use this tool to configure the replication, authority, and ACL group functions. In prior releases, you could configure these functions using iSeries Navigator. In V5R3, the Directory Management Tool (DMT) is no longer provided.
- If you are using replication today, you must perform additional post-installation configuration steps before replication will work in V5R3. The iSeries Information Center Directory Server (LDAP) topic, Migration considerations, describes the steps you must take. In addition to these steps, you must also set the QRETSVRSEC system value to 1.
- You must use the WebSphere® Application Server - Express product with the Web Administration tool. If your iSeries server is too small to run this server with V5R3, you can use an alternate version of the Web Administration tool. This alternate version, called Web Administration 5.1, is available as a download on the iSeries LDAP Web site at the following URL:
<http://www.ibm.com/eserver/iseries/ldap>

LDAP Kerberos service name changed

In V5R3, Directory Server changed the service name used by the directory server and client APIs for GSSAPI authentication (Kerberos). This change (also provided for V5R2 servers with PTF SI08487) makes the service name compliant with the standards that define GSSAPI authentication, which state that the principal name should start with lowercase ldap.

In previous releases, the directory server and client APIs used the following service name when the GSSAPI mechanism (Kerberos) is used for authentication:

```
LDAP/dns-host-name@Kerberos-realm
```

The uppercase LDAP created the potential for the directory server and client APIs to not interoperate with other vendor's products, particularly if the Kerberos key distribution center (KDC) had case sensitive principal names. The LDAP service provider for JNDI, a commonly used Java LDAP client API, is an example of a client included with the operating system that uses the correct service name.

In V5R3, the service name complies with the standards. You might encounter the following compatibility problems:

- A directory server configured to use GSSAPI authentication will not start after installing this release. This is because the key tab file used by the server has credentials using the old service name (for example, LDAP/mysys.ibm.com@IBM.COM), and the V5R3 server is looking for credentials using the new service name (for example, ldap/mysys.ibm.com@IBM.COM).
- A directory server or LDAP application using the LDAP APIs at V5R3 might not be able to authenticate with operating system servers or clients running on a prior release of OS/400.

To correct this, take the following actions:

1. If the KDC uses case sensitive principal names, create an account using the correct service name (ldap/mysys...).
2. Update the key tab file used by the directory server to contain credentials for the new service name. You might also want to delete the old credentials. You can use the Qshell keytab utility to update the key tab file. By default, the directory server uses the /QIBM/UserData/OS400/NetworkAuthentication/keytab/krb5.keytab file. The V5R3 Network Authentication Service (Kerberos) wizard in iSeries Navigator also creates key tab entries using the new service name.
3. Apply PTF SI08487 to update OS/400 systems running on V5R2 that use GSSAPI.

Alternately, you can choose to have the directory server and client APIs continue to use the old service name. You might want to do this when you are using Kerberos authentication in a mixed network of systems running with and without the PTFs. To do this, set the LDAP_KRB_SERVICE_NAME environment variable. You can set this for the entire system (required to set service name for the server) using the following command:

```
ADDENVVAR ENVVAR(LDAP_KRB_SERVICE_NAME)
```

In QSH (to affect LDAP utilities run from this QSH session), run the following command:

```
export LDAP_KRB_SERVICE_NAME=1
```

Replication stores passwords in a validation list

In V5R3, Directory Server changed to store replication passwords in a validation list when simple authentication is used. In order for these passwords to be stored and retrieved correctly, the QRETSVRSEC system value must be set to 1 before the password is stored. An error message will be sent to the Directory Server's job log if the QRETSVRSEC system value is not 1, and replication will fail. If this happens, set the QRETSVRSEC system value to 1 and set the password in the replication credentials again.

System SSL default cipher specification list changed

The System SSL default cipher specification list in V5R3 now includes two Transport Layer Security (TLS) Version 1 Advanced Encryption Standard (AES) ciphers:

- TLS_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_AES_256_CBC_SHA

The TLS_RSA_WITH_AES_128_CBC_SHA cipher has been supported by system SSL on the iSeries since the V5R1M0 release.

In prior releases, in order for this supported cipher to be used, an application was required to alter code to specifically allow the use of TLS_RSA_WITH_AES_128_CBC_SHA. Applications designed and coded to use the system SSL default cipher specification list rather than their own configured list would not use cipher TLS_RSA_WITH_AES_128_CBC_SHA in these previous releases.

In V5R3, those same applications using the system SSL default cipher specification list will now support both cipher TLS_RSA_WITH_AES_128_CBC_SHA and cipher TLS_RSA_WITH_AES_256_CBC_SHA with no code change. V5R3 applications that do not use the system SSL default cipher specification list, but instead configure the list, will be required to make a code or configuration change to allow the application to support either TLS_RSA_WITH_AES_128_CBC_SHA or TLS_RSA_WITH_AES_256_CBC_SHA.

The AES ciphers are valid only with the TLS Version 1 protocol and when the Cryptographic Access Provider (5722-AC3) product is installed.

The new system SSL default cipher specification list for this combination is:

- TLS_RSA_WITH_RC4_128_MD5
- TLS_RSA_WITH_RC4_128_SHA
- TLS_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_AES_256_CBC_SHA
- TLS_RSA_WITH_3DES_EDE_CBC_SHA
- TLS_RSA_WITH_DES_CBC_SHA
- TLS_RSA_EXPORT_WITH_RC4_40_MD5
- TLS_RSA_EXPORT_WITH_RC2_CBC_40_MD5

The previous system SSL default cipher specification list for this combination was:

- TLS_RSA_WITH_RC4_128_MD5
- TLS_RSA_WITH_RC4_128_SHA
- TLS_RSA_WITH_3DES_EDE_CBC_SHA
- TLS_RSA_WITH_DES_CBC_SHA

- TLS_RSA_EXPORT_WITH_RC4_40_MD5
- TLS_RSA_EXPORT_WITH_RC2_CBC_40_MD5

The order of the ciphers in the list has meaning when the iSeries acts as the server side of the SSL handshake. The server will prefer to negotiate the highest listed cipher in the list that is supported in the peer's client hello. When the iSeries is the client side for the SSL handshake, this default list will be sent in the client hello and the peer server will determine the cipher used for the negotiation based on the cipher specification list it is using.

The AES ciphers are not positioned first in the system SSL default cipher specification list. In order to guarantee that only an AES cipher is negotiated, you should code or configure a cipher specification list in an application that contains only the AES ciphers. In order to increase the likelihood that an AES cipher is negotiated for a server side application, without removing support for the other ciphers, configure a cipher specification list that positions the AES ciphers at the top of the list.

Using Service Tools Device IDs in System Service Tools (SST)

Starting with V5R3, you can change Operations Console configuration data from System Service Tools (SST) using the Work with service tools user IDs and Devices option. From the Work with Service Tools User IDs And Devices display, select the Service tools device IDs option . By default, this option is locked to prevent unauthorized changes to existing service tools device IDs, the creation of new IDs, or the deletion of IDs. To unlock this SST option, you must use a macro in Dedicated Service Tools (DST). Follow the steps below to run this macro.

Important: By default, the Service tools device IDs option is locked. If you receive the message **The user cannot perform the option selected**, it indicates that the option has not been unlocked.

Note: This function requires PTF MF32320 for V5R3M0; if you are running V5R3M5 LIC (Licensed Internal Code), this function is integrated into the LIC.

1. Access Dedicated Service Tools (DST).
2. Select Start a service tool.
3. Select Display/Alter/Dump.
4. Select Display/Alter storage.
5. Select Licensed Internal Code (LIC) data.
6. Select Advanced analysis.
7. Page down until you find the FLIGHTLOG option. Place a 1 next to the option and press Enter. You should be on the Specify Advanced Analysis Options display. The command should show as FLIGHTLOG.
8. Enter the option SEC UNLOCKDEVID in the Options field.

Note: If you want to lock this option to prevent later use, enter the option SEC LOCKDEVID.

Obsolete command menus removed

In V5R3, the following menu objects are no longer shipped as part of the operating system:

- CMDACN
- CMDADM
- CMDADSM
- CMDATM
- CMDCNNLE
- CMDDKT
- CMDDKTF

- CMDDOM
- CMDIDLC
- CMDIDX
- CMDIDXE
- CMDISDN
- CMDITM
- CMDLANG
- CMDOCL
- CMDPART
- CMDPRJ
- CMDRTL
- CMDRTLFL
- CMDS34
- CMDS38
- CMDTPL
- CMDTXT
- CMDTYPE
- CMDUSF

Application programs that attempt to use the Go to Menu (GO) command for any of these menus will fail with message CPF6AC7.

Applications using the iSeries SSL_ C APIs for secure sockets

If you maintain an application written to the SSL_ C API, you should update your source code to include `qsossl.h` rather than `ssl.h`. The new include `qsossl.h` contains the same prototypes and other information previously contained in `ssl.h`.

In V5R3, `ssl.h` remains available on the server, but will contain only a pointer to `qsossl.h`. In a future release, include `ssl.h` might be deleted.

Changes in C language runtime

Several changes in the C language runtime have been made to better conform with ANSI and POSIX specifications:

- **wcsrtombs**. If NULL is passed for the address of the state variable, a global state variable is used. In prior releases, the global state variable was set to the initial state for every call to this routine. In V5R3, to match the ANSI and POSIX specifications, the global state variable is set to the initial state only once, at program startup.
- **mbsrtowcs**. If NULL is passed for the address of the state variable, a global state variable is used. In prior releases, the global state variable was set to the initial state for every call to this routine. In V5R3, to match the ANSI and POSIX specifications, the global state variable is set to the initial state only once, at program startup. Also, in prior releases, passing a value of 0 for the output buffer length gave a return value of -1. In V5R3, a return value of 0 is given.
- **wcstombs**. In prior releases, the *LOCALEUCS2 version of this routine could generate a multibyte string that did not end in initial state if the given buffer was smaller than was required to hold the entire multibyte string. In V5R3, the *LOCALEUCS2 version of this routine was changed to be consistent with the *CLD and *LOCALE versions of this routine, and will always generate a multibyte string that ends in initial state.
- **wcrtomb**. In prior releases, passing NULL as the first parameter to this routine gave a return value of 0. In V5R3, a return value of 1 is given. Also, in prior releases, if NULL was passed for the address of

the state variable, the `*LOCALE` version of this routine would generate a trailing SI for a double-byte character. In V5R3, the `*LOCALE` version of this routine was changed to be consistent with the other routines and will not generate a trailing SI for a double-byte character. If a trailing SI is desired, `wcrtomb` can be called a second time, and given a wide character code of 0. An alternative is to use the `wctomb` routine, which will generate a trailing SI.

- **mbrlen/mbrtowc.** In prior releases, these routines did not save enough state in the `mbstate_t` parameter when an incomplete multibyte character was seen (that is, a return code of -2 was given). In V5R3, these routines now save additional information to work as intended by the ANSI and POSIX specifications.
- **mblen.** In prior releases, the `*CLD` version of this routine allowed out-of-order SI characters within the stream (that is, SI characters without a previous SO character). In V5R3, this routine will catch this error condition; each SI character must be matched with an SO character.
- **btowc.** In prior releases, this routine allowed values above 255 to be passed in, and would return the value as a wide character. In V5R3, this routine was changed to match the ANSI and POSIX specifications, and will truncate the given value to the range of 0 through 255.
- **strtol/strtoll/strtoul/strtoull/wcstol/wcstoll/wcstoul/wcstoull.** In prior releases, some versions of these routines would incorrectly update the `endptr` parameter if nothing but white space was seen. In V5R3, these routines have been changed to update the `endptr` parameter only if a valid sequence is seen, as indicated by the ANSI and POSIX specifications. In addition, these routines now set `errno` to `EINVAL` in this case. In prior releases, when an invalid base was given, `errno` was set to `EDOM`. In V5R3, to match the ANSI and POSIX specifications, `errno` is now set to `EINVAL`. In prior releases, the unsigned versions of these routines were returning an incorrect value when passed a string starting with the minus ('-') character. In prior releases, the minus character was being ignored. In V5R3, the behavior was changed to match the ANSI and POSIX specifications.
- **atof.** In prior releases, out of range values caused an exception to occur. In V5R3, this routine now handles out of range conditions in the same manner as `strtod`. Rather than generating an exception on overflow, it returns `HUGE_VAL`.
- **ungetc/ungetwc.** In prior releases, the `ungetc` buffer was 4 bytes in size. In V5R3, the `ungetc` buffer is 8 bytes in size. As a result, these routines can now `ungetc` more bytes before failing. ANSI and POSIX specifications guarantee at least one character of `ungetc`.
- **scanf/fscanf/sscanf/wscanf/fwscanf/swscanf.** In prior releases, under certain error situations, these routines were returning a 0 where they should have been returning -1. In V5R3, this has been fixed.
- **fileno.** In prior releases, this routine always returned -1 for the standard files (`stdin`, `stdout`, `stderr`), even if they were overridden to integrated file system files. In V5R3, this routine returns the underlying file number if the file is an overridden integrated file system file.

Integrated file system input and output routines

In prior releases, the integrated file system routines sometimes used `*CLD` locale items and sometimes used `*LOCALE` locale items. In V5R3, the integrated file system routines now use strictly `*LOCALE` locale items.

Also, in prior releases, the integrated file system routines sometimes performed shift-out and shift-in processing even for single byte locales such as the default "C" locale. In V5R3, the integrated file system routines now only do shift-out and shift-in processing for mixed byte locales. Some examples of routines affected by this behavior change are `scanf`, `printf`, and `getwc`.

Wide character input and output functions

In prior releases, the wide character input and output routines sometimes did not check validity of the given wide characters. In V5R3, more consistent checking is done. One example of this is with the `printf("%1s")` function. In V5R3, if an invalid wide character is given to an input and output routine, the routine will fail and set `errno` to `EILSEQ`.

Calendar time functions

In prior releases, the time routines filled in the `tm_isdst` field with 0 (DST_OFF) rather than a negative value (DST_UNKNOWN) when the time zone information did not exist. In V5R3, this has been corrected.

localtime/localtime_r/gmtime/gmtime_r. In prior releases, the *CLD versions of these routines did not allow a time of -1 to be passed in. In V5R3, these routines were changed to be consistent with the *LOCALE version so that a time of -1 is valid.

_gcvt. In prior releases, this routine sometimes returned "0." as a return string. In V5R3, it does not return the trailing decimal point and will return "0" in those cases.

nl_langinfo. In prior releases, the CRNCYSTR string was not returned correctly. In V5R3, a preceding '-' or '+' character is returned as indicated by the POSIX specification.

strptime/strptime/strfmon. In prior releases, these functions were not compliant with the POSIX specification in several areas. In V5R3, they are now compliant with the POSIX specification.

strftime. In prior releases, this routine would always generate an integrated file system newline character (X'25') for "%n". In V5R3, this routine will generate the correct newline character (either integrated file system X'25' or database X'15') based on the compile option. In prior releases, the "%V" and "%W" formats incorrectly computed the week number. In V5R3, this has been corrected.

localeconv. In prior releases, this routine returned the information given within the locale, even if it was not valid. In V5R3, all information given within the structure will be valid and will be what the ILE C runtime is using. Some fields have a range of valid values; anything not in that range will be interpreted as 'Not Available'.

The following `lconv` fields are affected by this change:

- `decimal_point` is now truncated to 1 byte.
- `p_cs_precedes` and `n_cs_precedes` are now binary values; anything other than 0 or 1 will be set to CHAR_MAX (Not Available).
- `p_sep_by_space` and `n_sep_by_space` can be 0-2; anything else is set to CHAR_MAX (Not Available).
- `p_sign_posn` and `n_sign_posn` can be 0-4; anything else is set to CHAR_MAX (Not Available).

Changes for programs compiled with SYSIFCOPT(*IFSIO) and SYSIFCOPT(*IFS64IO)

In previous releases, when `fopen()` created a file with mode "w", the `o_inherit` mode was used to create the file. When the file being created was a member in a database file, this meant that the member inherited its authorities from the database file. When the file being created is an integrated file system file, this meant that the file inherited its authorities from the parent directory. This is the desired behavior for the database member, but not for the integrated file system file. This will be changed in V5R3 to `S_IRWXO`, `S_IRWXU`, or `S_IRWXG` for integrated file system files. The database members will not change. The above bit mask gives the creator, his group and public read, write, and execute. So in V5R3, integrated file system files created with "w" mode will no longer inherit authorities from the parent directory.

Before V5R3, when `fopen()` created a file with mode "w+" or "a", the file was created with the following authorities:

- `S_IRWXO`
- `S_IRWXU`
- `S_IRWXG`

This gives the creator, the creator's group, and the public read, write, and execute (*RWX) authorities. This is desired for integrated file system files but not when creating members in a database file. If you attempt to create a member in a database file that you have authority to, but public does not, then the fopen would fail. Starting in V5R3, the database member will be created while inheriting all authorities from the database file. The authority behavior with integrated file system files will not be changed. So in V5R3, modes "w", "w+", and "a" all behave the same.

Integrated file system interfaces that create a physical file member, such as the creat() and open() APIs, will no longer require the input mode to exactly match the mode of the physical file. Additionally, these operations will no longer require the user and group running the operation to exactly match the user and group profiles currently associated with the physical file.

Saving and restoring journal receivers to a previous release

For V5R3, to exchange journal receivers in a network between a V5R3 server and systems that are running earlier releases, you must apply the following PTFs to the earlier release systems:

- SI10197 for V5R1M0
- SI10111 for V5R2M0

You must apply these PTFs if you intend to employ remote journal support between a V5R3 server and a system that is running an earlier release.

Mark counters changed from 32-bit integers to 64-bit integers

A *mark* is an integer value that uniquely identifies and determines relativity (like a time stamp) for program activations, invocations, and activation groups. In V3R6, the Licensed Internal Code changed to support these marks as 64-bit integers instead of 32-bit integers. However, at the MI (Machine Interface) level, marks were still returned as 32-bit integers by returning the low 32 bits of the mark value. For very long-running jobs (the well-known reliability of OS/400 has allowed jobs to run for weeks or even months), mark values could exceed the maximum 32-bit integer value. At the MI level, these 32-bit mark values could wrap. If above-MI code tried to use a wrapped mark value in an MI instruction, unexpected results could occur, including having the job terminate abnormally.

In V5R3, MI instructions and callable APIs have been added or changed to return the full 64-bit mark values. The operating system was changed to use the longer activation and invocation marks. Any applications that use 32-bit activation and invocation marks should be changed to use the longer 64-bit mark values to avoid unexpected behavior when using mark values that cannot be accurately represented in 32-bit values.

- ILE programs should replace the QleActBndPgm and QleGetExp APIs with QleActBndPgmLong and QleGetExpLong, respectively.
- OS/400 PASE programs should replace runtime functions _ILELOAD and _ILESVM with _ILELOADX and _ILESVMX, respectively.
- All programs that use the following MI instructions with 32-bit activation or invocation marks should change to use 64-bit marks as documented in the following MI instructions:
 - ACTBPGM
 - FNDRINVN
 - MATACTAT
 - MATACTEX
 - MATAGPAT
 - MATHSAT
 - MATINVAT
 - MATINVE

- MATINVS
- MATPRAGP
- MATPRMSG
- MATPTR
- RINZSTAT

Universal Connection replaces IGN on the Work with PM eServer iSeries customization display

References to IBM Global Network[®] (IGN) will be removed from the Work with PM eServer iSeries customization display. Universal Connection is the replacement for the removed IGN support.

The product previously referred to as PM/400 is now referred to as PM eServer iSeries, or PM iSeries. The Configure PM/400 (CFGPM400) and GO PM400 commands, however, have not changed.

New security auditing for special authority violations

A new violation type of "K" has been defined for the Authority Failure (AF) audit record. The "K" violation type indicates that a special authority violation was detected. As a result of this new violation type, the CPF2220, CPF4AAE and CPF2246 messages will no longer be sent to the history log (QHST); instead, look for violation type "K" AF audit records to determine if a special authority violation has occurred.

Some violation type "A" AF audit records, which were generated for special authority violations, will be changed to the new violation type "K" AF audit records.

Restricting access to audit settings

The auditing values returned by some APIs, in some output files, and on some screens and GUI panels will not be the current auditing values if you do not have either all object (*ALLOBJ) or audit (*AUDIT) special authority. The special value *NOTAVL (not available), or appropriate replacement value, will be returned or displayed instead.

Java Database Connectivity (JDBC) coded character set identifier (CCSID) requirement

In V5R3, the native JDBC driver has changed to require that the CCSID of the job running JDBC be set to something other than 65535 in order to start the JDBC driver. If the job CCSID is set to 65535, the JDBC driver will encounter the following SQL exception:

```
SQLState: HY000
Message:  INTERNAL ERROR:  Creation of DB2Driver object for
           registering with DriverManager failed.
Vendor:    -99999
```

This SQL exception has the following chained SQL exception.

```
SQLState: 22522
Message:  CCSID value is not valid.
Vendor:    -99999
```

When using the native JDBC driver, DB2[®] UDB for iSeries converts data using the job CCSID. When the job CCSID is set to 65535, no translation takes place, which results in unexpected errors when running SQL statements. In previous releases the JDBC driver would start but errors would occur later that would require you to change the job CCSID to something other than 65535.

Remote command and distributed program call server

In V5R3, the remote command and distributed program call server can support commands and parameter lists that are greater than 6000 bytes in length. This may impact exit programs registered to exit point QIBM_QZRC_RMT format CZRC0100 that are examining or logging the command or parameter list.

Saving UIM panels to a previous release

In V5R3, the User Interface Manager (UIM) defaults to providing the abbreviated time zone name when using the PRTHEAD tag without the PRTTIME attribute. If UIM panel groups compiled on V5R3 are to be saved to a prior release, then the PRTTIME and PRTPDATE attributes must be specified. As a result, the application must then provide the time and date values that are to be displayed on the panel.

Reported utilization for shared partitions

For V5R3M5 LIC POWER5 processor systems, the accuracy of CPU utilization reported for shared processor partitions has been improved. The method used to quantify CPU utilization now correlates more closely to how it is quantified for dedicated processor partitions. With this improvement, shared and dedicated processor partitions with similar capacity will show similar CPU utilizations when running equivalent workloads.

On POWER5 shared processor partitions, CPU utilization reported in performance reports as well as existing tools may appear lower with this change.

High availability clustering considerations

Before you upgrade your system to V5R3, if your system uses clustering and is one of the cluster nodes, you should first check the current cluster version using the Display Cluster Information (DSPCLUINF) command to ensure that it is at a level that is compatible with V5R3. At V5R3, the potential cluster version is 4. For a V5R3 node to be able to join the cluster, the current cluster version level must be at least 3. If your V5R2 cluster version is 2, you must increase the version to 3 before you upgrade.

See the topic Cluster version for details about this requirement. To learn more about clusters, click **Systems management > Clusters**.

Command and API changes

Authority change for CHGDDMTCPA command

In prior releases, the Change Distributed Data Management TCP/IP Attributes (CHGDDMTCPA) CL command required only *IOSYSCFG special authority in order to change or set the minimum level of password security required when a client system connects to this system as a server (PWDRWD parameter). In V5R3, the CHGDDMTCPA CL command requires *IOSYSCFG, *ALLOBJ, and *SECADM special authorities to change or initially set the PWDRQD value. The only special authority required to use the CHGDDMTCPA command to change the AUTOSTART value remains *IOSYSCFG.

Authority change to remove *LOCAL RDB directory entry

In prior releases, only *IOSYSCFG special authority was required in order to remove the *LOCAL RDB directory entry (either through the Remove Relational Database Directory Entry (RMVRDBDIRE) CL command, specifying *ALL for the RDB parameter; or by using option 4 from the Work with Relational Database Directory Entry (WRKRDBDIRE) panel). In V5R3, *IOSYSCFG, *ALLOBJ, and *SECADM special authorities are required in order to remove the *LOCAL RDB directory entry.

CPYFRMIMPF and CPYTOIMPF command changes

No conversion performed if TOFILE field has no CCSID or a 65535 CCSID

In prior releases, the CPYFRMIMPF and CPYTOIMPF commands would create a temporary database file when copying a stream file to a database file. The temporary file was created with CCSID(*JOB). The temporary file was used as an intermediate file that the stream file was copied to before the import file text was processed and copied to the target database file. Any input stream file text was converted to the job's CCSID when the stream file was copied to the temporary database file. If the target database file (TOFILE parameter) field did not have an explicit CCSID or had a CCSID of 65535, the copied import file text would remain in the job's CCSID because it was already converted when the import file was copied to the intermediate, temporary database file.

In V5R3, when the CPYFRMIMPF and CPYTOIMPF commands copy an import file from a stream file to a target database file, an intermediate database file is not used. If the TOFILE field does not have an explicit CCSID or has a CCSID of 65535, no text conversion will occur and the database file will contain the same hexadecimal code points that existed in the original stream file.

If you need the text in the import file to be converted from the original CCSID, make sure that the TOFILE field CCSIDs are specified and are not 65535.

For additional information on ways to handle CPYFRMIMPF and CPYTOIMPF command changes in V5R3, please refer to Informational APAR II13784.

Record delimiter (RCDDL) parameter change

RCDDL(*ALL) will find the first occurrence of a single or double character combination of carriage-return and line-feed and use that value as the record delimiter.

Character strings containing a null character (X'00') allowed only in binary fields

In previous releases, the CPYFRMIMPF and CPYTOIMPF commands treated character strings containing a null character (X'00') as acceptable values for any field. In V5R3, character strings containing a null character (X'00') are allowed only for binary type fields. Binary type fields are BINARY, VARBINARY, BLOB, CHAR FOR BIT DATA, or CHAR CCSID(65535). A null character (X'00') encountered in a non-binary field will produce a data mapping error.

Create Validation List (CRTVLDL) command changes

In V5R3, validation list (*VLDL) objects will be created with a maximum size of 1 terabyte instead of a maximum size of 4 gigabytes. Creating *VLDL objects with a maximum size of 1 terabyte allows for more entries to be stored in the *VLDL object. *VLDL objects with a maximum size of 1 terabyte cannot be saved to a release prior to V5R2. If you need to be able to create *VLDL objects that can be saved to a release prior to V5R2, you can create a data area (*DTAARA) object with the name of QSYVLDV5R1 in QUSRSYS library using the Create Data Area (CRTDTAARA) command. If the specified *DTAARA object exists, then the CRTVLDL command will create *VLDL objects with a maximum size of 4 gigabytes.

Digital certificate management APIs require DCM option

The Digital Certificate Management APIs that accept a certificate as input or return certificate information now require that Digital Certificate Manager (DCM), option 34 of the operating system (5722-SS1) be installed.

Also, the Retrieve Object Signatures (QYDORTVO, QydoRetrieveDigitalSignatures) API now requires that DCM be installed since it uses the Digital Certificate Management APIs.

Journaling command and API changes

Change Journal (CHGJRN) command and the QjoChangeJournalState API changes

In V5R3, the Change Journal (CHGJRN) command and the QjoChangeJournalState API ignore a request to change the state of a local journal to *INACTIVE. The value is allowed, but an informational message CPI7014 is issued that says the request was ignored. Any local journals found to be inactive during the install of the release will be made active.

Create Journal (CRTJRN) command changes

The default value for MNGRCV on the CRTJRN command will change from *USER to *SYSTEM. During an IPL the system identifies the *SYSTEM managed journals and performs a CHGJRN command to create and attach a new journal receiver and reset the journal sequence number of journals that are not needed for commitment control recovery for that IPL. Since new journal receivers are attached during the IPL, specify *CURCHAIN rather than *CURRENT for the RCVRNG parameter associated with the DSPJRN, RCVJRNE, and RTVJRNE commands and the QjoRetrieveJournalEntries API. This change to *SYSTEM as the default value does not affect any journals created before you install V5R3.

Create Journal (CRTJRN) and Change Journal (CHGJRN) command changes: QSYS/QAUDJRN audit journal

The required value for the RCVSIZOPT on the CRTJRN and CHGJRN commands for the audit journal QSYS/QAUDJRN will be *MAXOPT1. Previously the default value was *NONE. The *MAXOPT1 value allows the journal receiver associated with the audit journal to grow larger than 2 gigabytes. This limits situations where the journal receiver hits its maximum size and auditing is disabled. In that situation, you must restart auditing manually.

Create Journal Receiver (CRTJRNRCV) command changes

The default value for the THRESHOLD on the CRTJRNRCV command will change from *NONE to 1.5GB. This will cause a message CPF7099 to be sent to the journal MSGQ when the journal receiver reaches 1.5GB for a USER managed journal. No message was issued for journal receivers with a THRESHOLD value of *NONE resulting in the receiver growing to its maximum size. At that point attempts to insert additional journal entries would result in exception CPF7003. For a SYSTEM managed journal, a new journal receiver will automatically be attached. Message CPF7020 is sent when the current journal receiver is detached.

Apply Journalized Changes (APYJRNCHG) command changes

The Apply Journalized Changes (APYJRNCHG) command now requires object existence (*OBJEXIST) authority for an object that exists in a library for which changes are being applied.

The APYJRNCHG command now requires Add authority for the library of the object for which changes are being applied.

The default value for the Commitment boundary (CMTBDY) parameter of the APYJRNCHG command changed from *NO to *YES.

The default value for the Member specification on the Journalized file identification (FILE) parameter of the APYJRNCHG command changed from *FIRST to *ALL.

The APYJRNCHG command parameter of Objects (OBJ) default value changed for Member from *FIRST to *ALL.

The APYJRNCHG command no longer issues message CPC7040 for a successful apply. In the past, up to 8,192 messages could be sent out of the possible 65,535 objects that had entries applied. A new parameter, File to receive output (OUTFILE), was added to which the apply results can be recorded. Only failure messages are now sent to the job log.

Remove Journalled Changes (RMVJRNCHG) command changes

The default value for the Member specification on the Journalled file identification (FILE) parameter of the Remove Journalled Changes (RMVJRNCHG) command changed from *FIRST to *ALL.

The default value for the Member specification on the Objects (OBJ) parameter of the RMVJRNCHG command changed from *FIRST to *ALL.

The default value for the Commitment boundary (CMTBDY) parameter of the RMVJRNCHG command changed from *NO to *YES.

The RMVJRNCHG command will no longer issue message CPC7044 for a successful remove. In the past, up to 8,192 messages could be sent out of the possible 65,535 objects that had entries removed. A new parameter, File to receive output (OUTFILE), was added to which the remove results can be recorded. Only failure messages will now be sent to the job log.

Apply Journalled Changes Extended (APYJRNCHGX) command changes

The Apply Journalled Changes Extended (APYJRNCHGX) command will be withdrawn in a future release. When this command is withdrawn, the APYJRNCHG command will have all the functionality of the withdrawn command.

The APYJRNCHGX command no longer issues message CPC7040 for a successful apply. In the past, up to 8,192 messages could be sent out of the possible 65,535 objects that had entries applied. A new parameter, File to receive output (OUTFILE), was added to which the apply results can be recorded. Only failure messages are now sent to the job log.

Remove Journalled Changes (RMVJRNCHG), Apply Journalled Changes (APYJRNCHG) and Apply Journalled Changes Extended (APYJRNCHGX) command changes

The Remove Journalled Changes (RMVJRNCHG), Apply Journalled Changes (APYJRNCHG) and Apply Journalled Changes Extended (APYJRNCHGX) commands will no longer immediately end when an entry cannot be processed. Instead, no more entries will be processed for the object that had the error, but the commands will continue for the other objects.

For integrated file system objects, there are some operations (for example, rename, move, unlink, and link) that are actually the object's parent directory's operations. Therefore, when processing stops early for an integrated file system object and continues for other objects, any following directory operations that involve the parent directory of that integrated file system object will still be processed. In previous releases when the first error was encountered, the command ended. If you still want the behavior of the command ending on the first error, you can specify a new parameter, Object error option (OBJERROPT), on the RMVJRNCHG, APYJRNCHG, and APYJRNCHGX commands; valid values of this parameter are *CONTINUE and *END. *CONTINUE is the default.

The RMVJRNCHG, APYJRNCHG and APYJRNCHGX commands will now have a limit of 300,000 objects for which entries can be processed. The previous limit on these commands was 65,535. However, a given journal still can actively journal only 250,000 objects at one time.

If an object was saved with save-while-active, then the RMVJRNCHG, APYJRNCHG and APYJRNCHGX commands no longer check to be sure that the version restored onto the system was the last saved version. Instead, these commands start at the start of the save journal entry corresponding to the version of the object that was restored. This could cause more entries to be applied than if the latest version of the object were restored. You will not see the inquiry message (CPA7050) in this case. The inquiry message would have warned you that the latest version of the object had not been restored. There is also the potential that some changes to the object will be lost. There are some operations that cause changes to the object that would be saved with the object, but do not cause a journal entry to be deposited, so those changes cannot be replayed by the RMVJRNCHG, APYJRNCHG, and APYJRNCHGX commands. If the

more recent version of the object is restored, then the changes will be there. By restoring an older version of the object and applying journal changes, the changes are not made (since there is no journal entry), so they are effectively lost.

System managed access path (SMAPP) related changes

System default access path recovery time changed from 70 minutes to 60 minutes: The semantics of the special value *SYSDFT, on the EDTRCYAP, CHGRCYAP, and DSPRCYAP commands will be changed from 70 minutes to 60 minutes. Therefore, the system default access path recovery time will be changed from 70 minutes to 60 minutes (that is, the value that is used if the user takes no action to change the SMAPP times). If any user has explicitly changed this value, then they are not affected by this update.

Specified recovery times no longer rounded to 10 minute increments: Also, the semantics of a specified recovery time on the CHGRCYAP command or EDTRCYAP panel will be changed to no longer round the input value to the next 10 minute increment. The recovery time has always been specified in minutes, but it was rounded to the next 10 minute increment. For example, specifying 13 minutes for the recovery time actually resulted in the system using 20 minutes for the recovery time. Specifying 13 minutes will now result in the system using 13 minutes for the recovery time.

Change Journal Recovery Count (QJOCHRVC) now a formal API

In previous releases, you might have used the undocumented program QJOCHRVC to control how often changes to journaled objects were forced to auxiliary storage.

In V5R3, the QJOCHRVC (Change Journal Recovery Count) program has been changed to become a formal API. As part of this change, the authorities and parameters associated with QJOCHRVC have changed. QJOCHRVC is now shipped with a public authority of *EXCLUDE and requires *ALLOBJ special authority to run. QJOCHRVC also now requires two parameters: a 4-byte integer that specifies the journal recovery count and the standard error code structure.

You can find additional information concerning the Change Journal Recovery Count (QJOCHRVC) API in the iSeries Information Center.

adjtime(), gettimeofday(), and settimeofday() APIs updated to access the system clock

In V5R3, the adjtime(), gettimeofday(), and settimeofday() APIs access the system clock. In previous releases, these APIs accessed the software clock. Existing applications that use these APIs continue to access the software clock. Applications recompiled for V5R3 access the system clock.

The software clock is a system facility that can be set to a different time than the system clock and uses the QUTCOFFSET system value to calculate the UTC time. Since the software clock is not integrated with the rest of the system, it will be removed in a future release.

| For V5R3, you can access the software clock with the Qp0zAdjTime(), Qp0zGetTimeofday(), and
| Qp0zSetTimeofday() APIs. Any C based module using adjtime(), gettimeofday() or settimeofday() which
| is compiled for a release prior to V5R3 will automatically reference the appropriate software clock
| function (Qp0zAdjTime(), Qp0zGetTimeofday() and Qp0zSetTimeofday() respectively). Programs created
| on V5R3 that target a prior release and that contain C modules using the software clock APIs will need to
| have the imbedded C modules recompiled targeting a release prior to V5R3. For additional information
| please contact software support.

Open List of ASPs (QYASPOL) API changes

Record format YASP0600 of the Open List of ASPs (QYASPOL) API no longer distinguishes between integrated file system and library users of an independent auxiliary storage pool. Refer to the "Use identification" field description for additional details.

The QYASPOL API returns a YASP0300 record format for all disk drives in all independent ASPs. Previously, it returned this record format only for disk drives in an ACTIVE or AVAILABLE independent ASP.

Changes to profile handle and token APIs might require source code changes

Changes to the profile handle and token APIs might require source code changes to your applications. These changes apply to all security levels.

Get Profile Handle (QSYGETPH) API

When specifying a password for the password parameter, the password length and CCSID parameters are required.

When specifying a special value for the password parameter, password length and CCSID parameters are not allowed.

Get Profile Handle (QsyGetProfileHandle) API

A value of *CURRENT is not allowed for the user ID parameter on this API. Use QSYGETPH API or QsyGetProfileHandleNoPwd API instead.

Special values are not allowed for the password parameter on this API. Use QSYGETPH API or QsyGetProfileHandleNoPwd API instead.

Generate Profile Token (QSYGENPT) API

When specifying a password for the password parameter, the password length and CCSID parameters are required.

When specifying a special value for the password parameter, password length and CCSID parameters are not allowed.

Generate Profile Token Extended(QsyGenPrfTknE) API

Special values are not allowed for the password parameter on this API. Use QsyGenPrfTkn API instead.

Generate Profile Token (QsyGenPrfTkn) API

Password values are not allowed for the password parameter on this API. Use QsyGenPrfTknE API instead.

Reply handling exit point changes

An additional screen can be shown to indicate screens are being copied with the Start Copy Screen (STRCPYSCN) command. After the user has responded to the inquiry message to indicate whether they want their screens to be copied, another screen will be shown to inform a user if their screens are being copied. Even if a user indicates they do not want screens copied, a user exit program could exist for the QIBM_QMH_REPLY_INQ exit point in the system registration facility, that could replace a reply value provided by a user. When screens are copied, the new message informing the user their screens are being copied will be the first screen copied.

The send reply function (SNDRPY command and QMHSNDRM API) will issue a new status message if a reply is replaced by a reply handling exit program. The new status message is CPF2458. A status message can be monitored by a program but it is not placed on a program message queue or in the job log, so this may cause some unexpected errors depending what the program does. If an error handler tries to resend

the CPF2458, an error could occur. If a receive is done, blanks or a different message could be returned. So a program using the send reply function needs to ensure it handles the new status message appropriately. If an API error code is used instead of allowing messages to be issued, a status message is not returned in the error code information.

Activation group default changing on Create Program (CRTPGM) command

Prior to V5R3, the default for the ACTGRP (Activation group) parameter on the Create Program (CRTPGM) command was *NEW. In V5R3, the default is changed to *ENTMOD, which can impact applications where the entry module for the ILE program is written in ILE CL, ILE RPG, or ILE COBOL. With this change, recreating applications having an entry module written in these languages cause the application to run in the QILE or QILETS activation group (rather than a new activation group).

Integrated file system byte stream file APIs now use the 1200 CCSID for storing path names

The integrated file system byte stream file APIs now use the 1200 CCSID for storing path names. Prior to V5R3, path names were stored using CCSID 13488. CCSID 1200 is a superset of 13488; therefore, no change to customer programs is required, although the CCSID 1200 will be returned as the encoding.

Add Mounted File System (MOUNT, ADDMFS) command changes

The default mount option for a UDFS mount changed from 'rw' to 'rw,suid' on the Add Mounted File System (MOUNT, ADDMFS) command. This change affects you only if you look at the M_NO_SUID mount flag on a UDFS mount, since this flag is now turned off by default.

Change to path name format for Qlg (NLS-enabled) APIs

The path name format is common across application programming interfaces that work with objects that are supported across the integrated file systems. These APIs require a path name to identify the object with which the API will work. For any interfaces that use the Qlg path name format, and which result in the creation of an integrated file system object, a value of null characters is no longer valid if the path name delimiter is not null characters. Instead, errno EBADNAME is returned. Examples of such interfaces are the QlgOpen(), QlgCreate and QlgRenamexx APIs. If the open interfaces specify the O_CREAT flag, this error checking will occur, even if the object already exists.

Display DLO Auditing Level (DSPDLOAUD) command changes

You must have all object (*ALLOBJ) or audit (*AUDIT) special authority to use the DSPDLOAUD command. Previously, the command would also succeed if you had at least use (*USE) authority to the document or folder.

Work with Object Links (WRKLNK) and Display Object Links (DSPLNK) command changes

You must have all object (*ALLOBJ) or audit (*AUDIT) special authority to see the auditing value or create object auditing value for an object. If you do not, the value *NOTAVL will be returned in those fields in the display or printed report.

Get Attributes (Qp0IGetAttr) API changes

You must have all object (*ALLOBJ) or audit (*AUDIT) special authority to retrieve the auditing value (QP0L_ATTR_AUDIT) for an object. If you do not, the value *NOTAVL will be returned in the auditing value field.

Additionally, if only the auditing value is requested, then execute (*X) authority will not be required for each directory, preceding the last component, in the path name of the specified object. Because of this, if

the object whose audit attribute is being requested is in QFileSvr.400, then the QSECOFR user profiles on the source and target system must be enabled, and their passwords must match for the operation to succeed.

FRCACCPH(*YES) parameter on the CHGLF command

The FRCACCPH(*YES) parameter on the CHGLF command no longer starts the key-by-key disk write of each changed access path. Instead, such access paths are identified as candidates for System Managed Access Path Protection (SMAPP), which is a more efficient and effective means of reducing IPL recovery duration than the FRCACCPH(*YES) approach.

You do not need to take any action; you do not need to convert any objects.

In V5R3, the FRCACCPH(*YES) parameter is ignored. In some future release, IBM may remove the FRCACCPH parameter.

Access paths that are not selected by SMAPP for such protection will continue to experience the high overhead associated with the FRCACCPH(*YES) option. Thus, you might want to examine such access paths, and change their settings to FRCACCPH(*NO) if you want to reduce this high performance overhead. You can run CL and SQL queries to identify such access paths, as shown in the following examples:

```
DSPFD FILE(*ALL/*ALL) TYPE(*ATR) OUTPUT(*OUTFILE) FILEATR(*LF)
      OUTFILE(QGPL/LF)
```

```
SELECT LGFILE, LGLIB, LGKAP FROM QGPL/LF
      WHERE LGFKAP = 'Y'
AND LGFILE NOT LIKE 'Q%'
```

Alternatively, you can view the current list of access paths that are open and exhibiting this behavior by running the Edit Recovery for Access Paths (EDTRCYAP) or Display Recovery for Access Paths (DSPRCYAP) commands and then selecting option F13.

Removal of support for ATM and ISDN

The ATM and ISDN network interfaces have been removed in V5R3. This has resulted in the following changes to commands.

Removed CL commands

The following CL commands have been removed in V5R3:

- CRTNWIISDN
- CHGNWIISDN
- CRTNWIATM
- CHGNWIATM
- CRTLINIDLC
- CHGLINIDLC
- CRTLINNET
- CHGLINNET
- ADDCNNLE
- CHGCNNLE
- RMVCNNLE
- RNMCNNLE
- WRKCNNLE
- CRTCNL
- CHGCNNL

Note: For objects that have been configured using the following parameters, operating system functions are no longer supported.

Create Line PPP (CRTLINPPP) and Change Line PPP (CHGLINPPP) command changes

The Create Line PPP (CRTLINPPP) and Change Line PPP (CHGLINPPP) commands have changed:

- The special value *NWID for the RSRCNAME parameter is no longer a valid choice.
- The following parameters are no longer supported and might be removed in a future release:
 - Outgoing connection list (CNNLSTOUT)
 - Connection list entry (CNNLSTOUTE)
 - Incoming connection list (CNNLSTIN)
 - Attached nonswitched NWI (NWI)
 - NWI channel number (NWICHLNBR)
 - Switched NWI list (SWTNWILST)
 - Information transfer type (INFTRFTYPE)
 - Switched NWI selection (SWTNWISLCT)

Create Line X25 (CRTLINX25) and Change Line X25 (CHGLINX25) command changes

The Create Line X25 (CRTLINX25) and Change Line X25 (CHGLINX25) commands have changed:

- The special value *NWID for the RSRCNAME parameter is no longer a valid choice.
- The special value *X31 for the INTERFACE parameter is no longer a valid choice.
- The following parameters are no longer supported and might be removed in a future release.
 - Outgoing connection list (CNNLSTOUT)
 - Connection list entry (CNNLSTOUTE)
 - Incoming connection list (CNNLSTIN)
 - Packet mode (PKTMODE)
 - Information transfer type (INFTRFTYPE)
 - Attached nonswitched NWI (NWI)
 - NWI channel type (NWICHLTYPE)
 - NWI channel number (NWICHLNBR)
 - Switched NWI list (SWTNWILST)
 - Switched NWI selection (SWTNWISLCT)

Create Line Ethernet (CRTLINETH) and Change Line Ethernet (CHGLINETH) command changes

The Create Line Ethernet (CRTLINETH) and Change Line Ethernet (CHGLINETH) commands have changed:

- The special value *ATM for the NWITYPE parameter is no longer a valid choice.
- The following parameters are no longer supported and might be removed in a future release.
 - ATM access type (ACCTYPE)
 - PVC identifiers (PVCID)
 - Use LECS address (USELECSADR)
 - LES ATM address (LESATMADR)
 - Emulated LAN name (EMLLANNAME)
 - LEC disconnect time out (LECDSTIMO)

Create Line Token Ring (CRTLINTRN) and Change Line Token Ring (CHGLINTRN) command changes

The Create Line Token Ring (CRTLINTRN) and Change Line Token Ring (CHGLINTRN) commands have changed:

- The special value *ATM for the NWITYPE parameter is no longer a valid choice.
- The following parameters are no longer supported and will be removed in a future release.
 - ATM access type (ACCTYPE)
 - LEC frame size (LECFRAME)
 - PVC identifiers (PVCID)
 - Use LECS address (USELECSADR)
 - LES ATM address (LESATMADR)
 - Emulated LAN name (EMLLANNAME)
 - LEC disconnect time out (LECDSTIMO)

Create Controller APPC (CRTCTLAPPC) and Change Controller APPC (CHGCTLAPPC) command changes

The Create Controller APPC (CRTCTLAPPC) and Change Controller APPC (CHGCTLAPPC) commands have changed:

- The special value *IDLC for the LINKTYPE parameter is no longer a valid choice.
- The following parameters are no longer valid and might be removed in a future release.
 - IDLC window size (IDLCWDWSIZ)
 - IDLC frame retry (IDLCFRMRTY)
 - IDLC response timer (IDLCRSPTMR)
 - IDLC connect retry (IDLCCNNRTY)
 - Outgoing connection list (CNNLSTOUT)
 - Connection list entry (CNNLSTOUTE)

Create Controller Host (CRTCTLHOST) and Change Controller Host (CHGCTLHOST) command changes

The Create Controller Host (CRTCTLHOST) and Change Controller Host (CHGCTLHOST) commands have changed:

- The special value *IDLC for the LINKTYPE parameter is no longer a valid choice.
- The following parameters are no longer valid and might be removed in a future release.
 - IDLC window size (IDLCWDWSIZ)
 - IDLC frame retry (IDLCFRMRTY)
 - IDLC response timer (IDLCRSPTMR)
 - IDLC connect retry (IDLCCNNRTY)
 - Outgoing connection list (CNNLSTOUT)
 - Connection list entry (CNNLSTOUTE)

Create Controller RWS (CRTCLRWS) and Change Controller RWS (CHGCLRWS) command changes

The Create Controller RWS (CRTCLRWS) and Change Controller RWS (CHGCLRWS) commands have changed:

- The special value *IDLC for the LINKTYPE parameter is no longer a valid choice.
- The following parameters are no longer valid and might be removed in a future release.
 - IDLC window size (IDLCWDWSIZ)
 - IDLC frame retry (IDLCFRMRTY)
 - IDLC response timer (IDLCRSPTMR)

- IDLC connect retry (IDLCCNNRTY)
- Outgoing connection list (CNLSTOUT)
- Connection list entry (CNLSTOUTE)

Create Device Net (CRTDEVNET) command change

The Create Device Net (CRTDEVNET) command no longer allows the special value *TEL for the TYPE parameter as a valid choice.

Default value on AUT parameter changed for some commands

In V5R3, the default value for the AUT parameter has changed from *LIBCRTAUT to *CHANGE on the following commands:

- CRTCFGL (Create Configuration List)
- CRTCOSD (Create Class-of-Service Description)
- CRTCTLAPPC (Create Controller Description (APPC))
- CRTCTLASC (Create Controller Description (Async))
- CRTCTLBSC (Create Controller Description (BSC))
- CRTTLFNC (Create Controller Description (Finance))
- CRTCTLHOST (Create Controller Description (SNA Host))
- CRTCTLWS (Create Controller Description (Local WS))
- CRTCTLNET (Create Controller Description (Network))
- CRTTLRTL (Create Controller Description (Retail))
- CRTTLRWS (Create Controller Description (Remote WS))
- CRTTLTAP (Create Controller Description (Tape))
- CRTTLVWS (Create Controller Description (Virtual WS))
- CRTDEVAPPC (Create Device Description (APPC))
- CRTDEVASC (Create Device Description (Async))
- CRTDEVASP (Create Device Description (ASP))
- CRTDEVBSC (Create Device Description (BSC))
- CRTDEVCRP (Create Device Description (Crypto))
- CRTDEVDKT (Create Device Description (Diskette))
- CRTDEVDSP (Create Device Description (Display))
- CRTDEVFNC (Create Device Description (Finance))
- CRTDEVHOST (Create Device Description (SNA Host))
- CRTDEVINTR (Create Device Description (Intra))
- CRTDEVMLB (Create Device Description (Media Lib))
- CRTDEVNET (Create Device Description (Network))
- CRTDEVOPT (Create Device Description (Optical))
- CRTDEVPRN (Create Device Description (Printer))
- CRTDEVRTL (Create Device Description (Retail))
- CRTDEVSNPT (Create Device Description (SNPT))
- CRTDEVSNUF (Create Device Description (SNUF))
- CRTDEVTP (Create Device Description (Tape))
- CRTLINASC (Create Line Description (Async))
- CRTLINBSC (Create Line Description (BSC))
- CRTLINDDI (Create Line Description (DDI))
- CRTLINETH (Create Line Description (Ethernet))

- CRTLINFAX (Create Line Description (Fax))
- CRTLINFR (Create Line Description (Frame Relay))
- CRTLINPPP (Create Line Description (PPP))
- CRTLINS DLC (Create Line Description (SDLC))
- CRTLINTDLC (Create Line Description (TDLC))
- CRTLINTRN (Create Line Description (Token-Ring))
- CRTLINWLS (Create Line Description (Wireless))
- CRTLINX25 (Create Line Description (X.25))
- CRTMODD (Create Mode Description)
- CRTNTBD (Create NetBIOS Description)
- CRTNWIFR (Create Network Interface (FR))
- CRTNWS D (Create Network Server Description)

Change TCP/IP Attributes (CHGTCPA) command changes

TCP keep alive (TCPKEEPALV) parameter changed

In V5R3, the TCPKEEPALV attribute on the CHGTCPA command, which controls the length of time a connection can remain idle before the transport begins to probe for the other end of the connection, has been changed.

In previous releases, a timer was used for each connection that had the SO_KEEPALIVE option turned on. The timer was set for the interval specified by the TCPKEEPALV attribute. When the interval expired, if the connection had been idle for the entire interval, probing for the other end of the connection began. Otherwise, the keep alive timer was reset. The tolerated idle interval could be up to twice the interval specified by the TCPKEEPALV attribute.

In V5R3, probing begins after an eligible connection has been idle for an interval exactly as long as the TCPKEEPALV attribute.

IPQOSBCH and IPQOSTMR parameters ignored in V5R3

Two parameters on the CHGTCPA command, IP QoS datagram batching (IPQOSBCH) and IP QoS timer resolution (IPQOSTMR), are ignored in V5R3. In a future release, these parameters might be removed from the command.

Perform Hardware Configuration Operation (QYHCHCOP) API restriction

In V5R3, the Perform Hardware Configuration Operation (QYHCHCOP) API will return a "not supported" response when called with a Request type of LPARRequest (1) on an @server i5. The DiskRequest (2) Request type will continue to be supported. The alternative interface for performing logical partition (LPAR) configurations is to use the Hardware Management Console support.

WAITFILE parameter on CRTPF and CRTLF commands changed

In September of 2002, PTFs were created for releases V5R1 and V5R2 of the operating system that changed the behavior of the *IMMED value for the WAITFILE (Maximum file wait time) parameter for database physical files and logical files. The problem fixed by these PTFs was to treat WAITFILE(*IMMED) as a wait time of zero seconds (instead of a wait time of 30 seconds) when the file is opened.

If you created physical files or logical files with the WAITFILE(*IMMED) option, but had applications that needed to have a non-zero wait time, you were instructed to use the Change Physical File (CHGPF) or Change Logical File (CHGLF) command for the existing file to specify a value (for example, 30) for the WAITFILE parameter.

- | For V5R3, the default value for the WAITFILE parameter for the Create Physical File (CRTPF) and Create Logical File (CRTLFL) commands was changed from *IMMED to 30. If you want to have a wait time of zero, you will need to specify WAITFILE(*IMMED) on the CRTPF or CRTLFL command. If you have multiple jobs concurrently referencing the file, you should have a file wait time other than *IMMED to allow for proper synchronization of file access between jobs.
- | The SQL CREATE TABLE statement uses the same default for file wait time as the WAITFILE parameter on the CRTPF command. If you have SQL tables created in releases prior to V5R3, you should consider using the CHGPF command with WAITFILE(30) to set the table's wait time to 30 seconds.

Backup and recovery changes

Save command and API changes

The default value on the Save access path (ACCPTH) parameter has been changed from the previous default of *NO to a new default of *SYSVAL for the following commands and APIs:

- SAVLIB command
- SAVOBJ command
- SAVCHGOBJ command
- SAVRSTLIB command
- SAVRSTOBJ command
- SAVRSTCHG command
- QSRSAVO API

The new system value, QSAVACCPTH, specifies whether or not access paths will be saved. The default value for the new system value will be '1' (Save access paths).

This change means that access paths are saved by default, which avoids lengthy access path rebuilds when files are restored. However, save operations might take longer to complete, and might use more media. If you want to avoid this new default behavior, you can either explicitly specify ACCPTH(*NO) or change the QSAVACCPTH system value to be '0' (Do not save access paths).

For OUTPUT(*PRINT), the value shown for 'Size' in the library summary information is changed from the size of the library object to the total size of objects saved for the library.

SAVLIB LIB(*NONSYS), SAVLIB LIB(*ALLUSR), and SAVCHGOBJ LIB(*ALLUSR) commands save libraries in a different order

The order in which libraries are saved by SAVLIB LIB(*NONSYS), SAVLIB LIB(*ALLUSR), and SAVCHGOBJ LIB(*ALLUSR) has changed. Libraries QSYS2, QGPL, QUSRSYS, and QSYS2xxxxx (where xxxxx is an independent ASP number) are saved first, if they are located on the ASPs specified by the ASPDEV parameter. Specifying one of these libraries on the STRLIB parameter causes more libraries to be saved than in previous releases.

If you use RSTLIB SAVLIB(*NONSYS) or RSTLIB SAVLIB(*ALLUSR), or if you use a backup and recovery solution that uses these commands, you are not affected by these changes.

If you have built your own customized backup procedures, you might need to change your procedures to ensure that these libraries are restored before restoring other user libraries. The IBM supplied recovery procedures have been updated to document this requirement.

Output file will contain additional records when INFTYPE(*MBR) is specified

When OUTPUT(*OUTFILE) INFTYPE(*MBR) is specified on the Save Library (SAVLIB), Save Object (SAVOBJ), Save Changed Objects (SAVCHGOBJ), Restore Library (RSTLIB), Restore Object (RSTOBJ) commands, an additional output file record that contains a blank member name will be created for each database file. This record provides the journal and journal receiver information needed when recovering the entire file, and matches the data provided by OUTPUT(*OUTFILE) INFTYPE(*OBJ).

When OUTPUT(*OUTFILE) INFTYPE(*MBR) is specified on the Save Object List (QSRSAVO) API, an additional output file record that contains a blank member name will be created for each database file. This record provides the journal and journal receiver information needed when recovering the entire file, and matches the data provided by OUTPUT(*OUTFILE) INFTYPE(*OBJ).

Cannot save a file restored with partial transactions to an earlier release

A database file that has any member that has been restored with partial transactions will not be allowed to be saved for a release prior to V5R3. The save operation will skip that file and an error message (CPI3215) will be sent. See the Save Active Wait Time (SAVACTWAIT) parameter topic in the iSeries Information Center for more information about partial transactions.

Save-while-active processing changes

In V5R3, save-while-active processing for values SAVACT(*SYSDFN) and SAVACT(*LIB) converged, making the process more efficient. The operational characteristics of *SYSDFN and *LIB are now identical to each other. This might result in faster times to reach a library checkpoint, the amount of time that specific objects are locked during checkpoint processing might change (*SYSDFN times may increase), and the overall duration of the save operation might change.

SAVACT(*SYNCLIB) remains the value recommended by IBM in most cases.

If you specify SAVACT(*LIB) or SAVACT(*SYSDFN) when using a media definition, SAVACT(*SYNCLIB) is no longer substituted for the value specified, and message CPI3823 is no longer issued.

PTF adds additional error checking to the restore function

- | In V5R3, the program temporary fix (PTF) SI22948 adds additional error checking to the restore function.
- | If you try to restore a device file in a job that already has the file open, the file will not be restored, and the following messages are sent:
 - | • CPF0675 – Device file is in use.
 - | • CPD3737 – Object not restored.
- | This additional error checking has also been added to the restore function in V5R2 via PTF SI24174 and in V5R4 via PTF SI23983.

DB2 UDB for iSeries changes

Software prerequisite for DB2 UDB for iSeries inter-operation with DB2 Connect and DB2 UDB for Linux/UNIX/Windows Version 8

In V5R3, DB2 UDB for iSeries requires the following service levels for inter-operation with the DB2 UDB products:

- | • DB2 UDB Linux/UNIX/Windows Version 7 (however, Version 8 FP4 is recommended)
- | • DB2 Connect™ Version 7 (however, Version 8 FP4 is recommended)

You can get the FixPak from the DB2 Technical Support Web site at the following URL:
<http://www.ibm.com/cgi-bin/db2www/data/db2/udb/winos2unix/support/index.d2w/report>

New DB2 tables and views added to the QSYS2 library

Several new DB2 tables and views were added to the QSYS2 library in V5R3. This was done to more closely match the ISO and ANSI SQL standards. Since QSYS2 is in the system portion of the library list, users who have files with the same name in other libraries in the library list might need to make changes to their applications.

If you have files with the same name, you can take either of two recommended circumventions:

- Change all applications that use any of the duplicate-named files to use the qualified name for all references to them.
- Rename any of the duplicate-named files that are not in QSYS2 and change all applications to refer to the new name.

The following alternate circumventions are possible, but not recommended, as they might cause problems in applications that refer to the objects in QSYS2. Since these circumventions do not require application changes, however, you might consider trying one of them as a short-term fix until you can implement one of the recommended circumventions:

- Remove QSYS2 from the system portion of the library list and add it to the end of the user portion of the library list.
- Rename the corresponding duplicate file in QSYS2; however, this will cause problems for any third-party applications that attempt to refer to the new files.

The following new tables and views were added to the QSYS2 library:

- SYSSEQOBJ
- CHECK_CSTS
- COLUMNS
- COLUMNS_S
- PARAMETERS
- PARM_S
- REF_CST1
- REF_CST2
- ROUTINES
- ROUTINES_S
- SCHEMATA
- SCHEMATA_S
- SQL_LANG_S
- SYSCAT1
- SYSCAT2
- SYSCHARSET
- SYSCHRSET1
- SYSCHRSET2
- SYSSEQ
- SYSTABDEP
- TABLE_CSTS
- TABLES
- TABLES_S

- UDT_S
- UDTs
- VIEWS

Note: In V5R2, the following new objects were added in QSYS2:

- QSYS2.SYSCATALOGS
- QSYS2.SYSROUTINEDEP

In V5R2, the following new objects were added in SYSIBM:

- sysibm.SQLTables
- sysibm.SQLSchemas
- sysibm.SQLTableTypes
- sysibm.SQLColumns
- sysibm.SQLPrimaryKeys
- sysibm.SQLForeignKeys
- sysibm.SQLSpecialColumns
- sysibm.SQLStatistics
- sysibm.SQLTablePrivileges
- sysibm.SQLColPrivileges
- sysibm.SQLProcedures
- sysibm.SQLProcedureCols
- sysibm.SQlUDTs
- sysibm.Tables
- sysibm.Tables_S
- sysibm.Views
- sysibm.Schemata
- sysibm.Schemata_S
- sysibm.Information_Schema_Catalog_Name
- sysibm.Catalog_name
- sysibm.Character_Sets
- sysibm.Character_Sets_S
- sysibm.Columns
- sysibm.Columns_S
- sysibm.Table_Constraints
- sysibm.Referential_Constraints
- sysibm.Ref_constraints
- sysibm.Check_Constraints
- sysibm.User_Defined_Types
- sysibm.UDT_S
- sysibm.Routines
- sysibm.Routines_S
- sysibm.Parameters
- sysibm.Parameters_S
- sysibm.SQL_Languages_S
- sysibm.SysjarContents
- sysibm.SysjarObjects

DYNAMIC keyword handled as if SENSITIVE DYNAMIC has been specified

For V5R3, the sensitivity and scrollability attributes that can be specified in DECLARE CURSOR and in the ATTRIBUTES clause are being expanded to more closely match what other DB2 UDB platforms support. This change exposed an incompatibility between what is done on the iSeries and what is specified in the SQL standard.

In releases prior to V5R3, the DYNAMIC keyword modified a scrollability option, and indicated that a scrollable cursor can be used for updates. In the SQL standard the DYNAMIC keyword modifies a sensitivity option.

For V5R3, the DYNAMIC keyword will be handled as if SENSITIVE DYNAMIC had been specified. For existing applications that specify DYNAMIC SCROLL, this will have the desired effect of specifying a scrollable cursor that can be used for updates.

These changes will cause the following compile-time and runtime differences in V5R3:

- The syntax INSENSITIVE DYNAMIC SCROLL will not be allowed in V5R3. The SQL precompiler will signal an error at compile time.
- A cursor declared as DYNAMIC SCROLL in V5R2 is ASENSITIVE. In V5R3, it will be SENSITIVE. The difference between the two is that an error is signaled for a SENSITIVE cursor if a sensitive cursor cannot be implemented. For an ASENSITIVE cursor, an insensitive cursor is implemented and no error is signaled.
- If a cursor was created in a prior release declared as DYNAMIC SCROLL and the query implementation would require the creation of a temporary result table, a runtime error will be signaled. Conditions that require the creation of a temporary result table are listed in messages CPI4324 and CPI4325 in message file QCPFMSG.

New EXTRACT keyword

EXTRACT was added as a new SQL keyword. If you have an existing query that has a function or column named EXTRACT, the function or column must be referred to as "EXTRACT" (that is, enclosed within quotation marks).

BINARY data type has a release-to-release compatibility issue

The BINARY data type has a release-to-release compatibility issue in that it is possible for an application to create a user-defined type with a name of BINARY or VARBINARY. Unqualified references to BINARY or VARBINARY data types will resolve to the appropriate system data type in library QSYS2. The user must qualify the data type name if the intention is to use a user-defined type with the same name.

Stricter enforcement of ROWID compatibility rules

The *DB2 UDB for iSeries SQL Reference* states that a ROWID operand can be assigned only to another ROWID operand. However, in some situations in prior releases, the ROWID data type was treated as equivalent to VARCHAR. Specifically, the CALL, FETCH, SELECT INTO, VALUES INTO, and SET statements did not distinguish between ROWID and VARCHAR. Also, for the INSERT and UPDATE statements, a VARCHAR field could be assigned to a ROWID field, but the reverse was not permitted.

In V5R3, the compatibility rules for ROWID will be enforced as stated in the SQL Reference. The most common scenario where users may see a change in behavior is fetching a ROWID column into a character host variable. This was permitted in V5R2 but will be prohibited in V5R3. In V5R3, a ROWID host variable must be used instead. CLI and ODBC do not support a separate ROWID host variable type, so fetches into character host variables will still be permitted in CLI and ODBC.

SQL DELETE and journal entries

An SQL DELETE statement that does not contain a WHERE clause will delete all rows of a table. Prior to V5R3, each row was always individually locked, deleted, and if the file was journaled, a DL entry was written to the journal for each row. Starting in V5R3, if all rows are being deleted from a table using such an SQL DELETE statement, the rows might be deleted using either a clear operation (if not running under commitment control) or a change file operation (if running under commitment control). This implementation will be much faster than individually deleting each row. If running under commitment control, the deletes can still be committed or rolled back.

The only difference between this implementation and prior releases is that the journal will contain only a single entry that indicates that all rows were deleted (either a CR or CG journal entry) as opposed to an entry for each deleted row. This new technique will be used only if all of the following conditions are true:

- A significant number of rows are being deleted
- The job issuing the DELETE statement does not have an open cursor on the file (not including pseudo-closed SQL cursors)
- No other job has a lock on the file
- The file does not have a delete trigger
- The file is not the parent to a referential constraint with a CASCADE, SET NULL, or SET DEFAULT delete rule
- The user issuing the DELETE statement has *OBJMGT or *OBJALTER authority in addition to *OBJOPR and *DLT

Reduction of scale no longer performed on SQL query comparisons involving numeric literals

In V5R3, reduction of scale is no longer performed on SQL query comparisons that involve numeric literals. In previous releases, reduction of scale could have been performed on literals.

Comparisons involving numeric literals now follow the rules contained in the Numeric Comparisons topic in the DB2 UDB for iSeries SQL Reference.

Enable concurrent write (ECW) default behavior on record inserts

In V5R3, the default behavior for record inserts changed to Enable Concurrent Write (ECW) mode; ECW mode can offer significant performance improvement when new records are added concurrently to a file that uses the REUSEDLT(*YES) attribute.

Prior to V5R3, attempts to perform concurrent adds to the same file often turned into sequential rather than concurrent behavior as the system validated relative record numbers (RRNs). For example, if two jobs were concurrently adding 10 records each, and the second job completed first, the validation of the RRNs would enforce a sequential validation process that held the second job until the first job was validated.

In V5R3, with ECW mode, this validation step completes concurrently. One difference in this mode is when the file is journaled. Since the validations occur in parallel, and since one job might complete before another, then it is possible that the journal entry for a record in the second job will reach the journal before the journal entry for a record in the first job. In this event, deleted records will be assigned RRNs so that the RRN of the record in the second job can be validated. In the journal, you will simply see an F/IZ journal entry for the RRNs that are filled with deleted rows. When the deleted records are subsequently reused, you will see the normal R/PX journal entry.

Prior to V5R3, you needed to explicitly enable the ECW mode to gain the performance advantages it offered. For V5R3, all physical files that use the REUSEDLT(*YES) attribute (which includes all SQL tables

created using default attributes) automatically inherit this mode. In the unlikely event that you want to revert to the pre-V5R3 mode of operation, in which validation of RRNs is done sequentially, you must call your service provider.

New CACHE_RESULTS Query QAQQINI option

A new option, CACHE_RESULTS is supported for the QAQQINI control file in V5R3. For queries involving temporary results (for example, sorts and hashes), the database often saves the results across query pseudo-close/ pseudo-open, as long as the result set is not large, hoping to reuse the results for the next run of the query. Beginning in V5R3, the database also begins saving these temporary results even when a job is finished with them, assuming another job can later reuse the results.

The database automatically controls the caching of these results, removing cache results as storage usage becomes large. However, the amount of temporary storage used by the database can be noticeably more than in previous releases. This option provides a way for the user to control the cache results processing.

The following options are valid:

***DEFAULT**

The default value is the same as *SYSTEM.

***SYSTEM**

The database manager may cache a query result set. A subsequent run of the query by that job or (if the ODP for the query has been deleted) by any job, will consider reusing the cached result set.

***JOB**

The database manager may cache a query result set from one run to the next for a job, as long as the query uses a reusable ODP. When the reusable ODP is deleted, the cached result set is destroyed. This value mimics V5R2 processing.

***NONE**

The database does not cache any query results.

Rules for specifying SQLSTATE values enforced in SQL routines

SQLSTATE values are required to be uppercase characters (A-Z) or numeric values (0-9). Prior to V5R3, an error was not issued if lowercase characters (a-z) were specified in an SQLSTATE value in SQL routines. On V5R3, an error will be issued for SQLSTATE values that are not valid.

You might have specified these SQLSTATE values in SIGNAL and RESIGNAL statements and in condition-declarations and handler-declarations in an SQL procedure, function, or trigger.

Authority to create a view

Prior to V5R3, you needed object operational (*OBJOPR) authority and any one of the data rights (*READ, *ADD, *UPD, *DLT, or *EXECUTE) to create a view. On V5R3, you need *OBJOPR authority and read (*READ) data authority (or SELECT privilege). These authorities are consistent with DB2 UDB requirements on other servers.

Error message change when comparing numeric and string fields

In V5R3, DB2 UDB for iSeries added support to allow comparisons between numeric and string fields. In prior releases, such comparisons would result in an SQL0401 error message (Operands not compatible). For V5R3 and later releases, such comparisons will either run with no error message (if the string field contains a valid numeric value, like '123'), or will result in error message SQL0302 or SQL0678 (if the string field contains a value which is not a valid numeric value, like 'abc'). If you have a program that is expecting error message SQL0401 for these types of field comparisons, the program will need to be changed.

| **ILE RPG SQL precompiler enforces unique variable names**

| In releases prior to V5R3 (beginning with V3R2), the ILE RPG SQL precompiler has not supported variable scoping. The SQL programming with embedded languages topic states that "the names of host variables must be unique within the program, even if the host variables are in different procedures." In these prior releases, this restriction was not enforced.

| In V5R3, this restriction is enforced in the ILE RPG SQL precompiler. As a result, some programs that successfully compiled on previous releases might fail in V5R3.

| The program temporary fixes (PTFs) SI17120, SI17121, SI17122 were created to mitigate the impact. The ILE RPG SQL precompiler now issues SQL0314 as a severity 11 error instead of severity 35 error whenever the duplicated variable names have the same type. Instead of having to change the source, you can change the GENLVL parameter on the CRTSQLRPGI command to be 11 or higher.

| **Updating rows with NULL value fails with SQL0407 error when NULL values are not allowed**

| On V5R2 and earlier releases, an SQL0407 failure was not returned in all cases specified by the SQL standard.

| In V5R3, the failure will be seen by applications attempting to update a row with a NULL value when the SQL table has been defined to not accept a NULL value.

| For example, consider the following example:

```
| CREATE TABLE QTEMP.TT1 (F1 INT NOT NULL, F2 INT NOT NULL)
| CREATE TABLE QTEMP.TT2 (F1 INT NOT NULL, F2 INT NOT NULL)
| insert into qtemp.tt1 values(1, 0)
| insert into qtemp.tt1 values(2, 0)
| insert into qtemp.tt2 values(1, 123)
|
| UPDATE QTEMP.TT1 A
|   SET F2 = (SELECT IFNULL(B.F2, A.F2)
|             FROM QTEMP.TT2 B
|            WHERE A.F1 = B.F1 )
```

| Before the update operation, the records for table TT1 and TT2 are the following:

```
| F1  F2  F1  F2
| 1   0   1  123
| 2   0
```

| After the update operation completes, the records for table TT1 and TT2 are the following:

```
| F1  F2  F1  F2
| 1  123  1  123
| 2   0
```

| In both V5R2 and V5R3, the data in the target table (QTEMP.TT1) is correct. In V5R3, however, the UPDATE statement will fail with an SQL0407 after updating the first row.

| **Format level identifiers changed for record formats of new SQL tables that have one or more TIMESTAMP columns**

| The way in which SQL tables with TIMESTAMP data types generate the record format level identifier has changed. See Information APAR II14134 for details.

Machine Interface (MI) instruction changes

Materialize Machine Attributes (MATMATR) MI instruction

In prior releases, if the MATMATR (Materialize Machine Attributes) MI instruction was used to retrieve the amount of main memory (option X'12C') and the system either had more than 32,767 megabytes of main memory or required more than 32,767 megabytes of memory to run at optimum performance, a negative number was returned in the respective field.

In V5R3, two additional fields have been added to return the correct amount for usable memory and minimum required memory. The preexisting fields (labeled "Usable memory installed" and "Minimum memory required") now contain 32,767 in the situations that would have previously resulted in negative numbers in previous releases.

You can use these two new fields (labeled "Total usable memory installed" and "Minimum required memory") to ensure that you always have the correct amounts.

Materialize Resource Management Data (MATRMD) MI instruction

The Materialize Resource Management Data (MATRMD) MI instruction option X'13' returns information on a maximum of 32 processors even if there are more installed. If information on more than 32 processors is required, then the MATRMD MI instruction option X'28' must be used.

Materialize Journal Port Attributes (MATJPAT) MI instruction no longer reports a local journal as inactive

In V5R3, the journal state of a local journal can no longer be changed to *INACTIVE. During the install of V5R3, any local journal that is found in an inactive state will automatically become active. As a result, the Materialize Journal Port Attributes (MATJPAT) MI instruction no longer supports the *Prevent entries from being deposited* field.

Materialize System Object (MATSOBJ) MI instruction changed

In V5R3, the Materialize System Object (MATSOBJ) MI instruction will return X'FF' for the *Object audit attributes* field to programs that run without all object (*ALLOBJ) or audit (*AUDIT) special authority. Programs that have either *ALLOBJ or *AUDIT special authority will return the same values as in previous releases for the *Object audit attributes* field.

Process access group to be eliminated

In a future release, jobs will be initiated without an associated access group. The process access group no longer provides a performance benefit and is no longer used by the machine when creating objects. In the future, the access group will no longer be created. Code that uses the Materialize Process Attributes (MATPRATR) MI instruction to obtain the address of the process access group will receive a null pointer.

Optical support changes

Optical device support changes

Default file CCSID changed for optical files on UDF volumes

A different default file CCSID is returned for optical files on UDF (Universal Disk Format) volumes.

In previous releases, a CCSID value of 65535 was returned as the default for optical files on UDF that were not created with a CCSID.

In V5R3, the default file CCSID for optical files on UDF volumes is the default job CCSID of the requester for optical files that were created through the operating system, or (for optical files that were created on another operating system) the ASCII equivalent of the default job CCSID.

Default file CCSID for optical files created through Hierarchical File System programming interfaces or through operating system save commands

Optical files that are created on volumes formatted in UDF (Universal Disk Format) will be created with the default job CCSID when they have been created through the Hierarchical File System programming interfaces or through operating system save commands. In prior releases, optical files created through these interfaces were not created with a CCSID, and a CCSID value of 65535 was returned through the integrated file system interfaces.

Some UDF optical volumes recognized by a different volume name

Some UDF (Universal Disk Format) optical volumes that are formatted on other operating systems are recognized by a different volume name in V5R3M0 than they were in a previous release of OS/400. UDF allows for a *logical volume identifier* and a *physical volume identifier*. In V5R3, the logical volume identifier is displayed instead of the physical volume identifier. On most UDF volumes, these two volume identifiers are the same, but on volumes formatted on other platforms, they can be different. It is on these volumes that a change in volume names will be visible in V5R3. Note that UDF is the logical format the operating system creates on DVD-RAM.

Optical command changes

Add Optical Cartridge command confirmation display removed

In prior releases, the Add Optical Cartridge (ADDOPTCTG) command displayed a confirmation screen that instructed the user to put a cartridge in the input/output station and press enter and also gave the user the option of canceling the operation. In V5R3, this confirmation screen is no longer shown.

Change Optical Attributes command now used to specify media format

An optical CD or DVD can have two media formats. In V5R2, a user could create a data area (*DTAARA) object and set a value that indicated whether to use the basic or extended media format. In V5R3 the data area will no longer be used. The user will need to use the Change Optical Attributes (CHGOPTA) command to specify which media format to use (EXTMEDFMT parameter). If the data area exists and the value in the data area says to use the extended format this value will not be brought forward to V5R3. The CHGOPTA command will have to be used to set this value.

Reclaim Optical command directory index (DIR) parameter changed

The default value for the directory index (DIR) parameter of the Reclaim Optical (RCLOPT) command has changed from *YES to *NO. This change improves the performance when OPTION(*RESET) or OPTION(*UPDATE) is selected when the default for DIR is maintained. The effect of DIR(*NO) is that the directory index for the optical volume(s) is not rebuilt, but instead is built on demand when needed. If you want to rebuild the directory index, you should explicitly specify DIR(*YES).

System value changes

Retrieve System Values (RTVSYVAL) against QAUDLVL system value might not return the complete list of audit values

In V5R3, if you run Retrieve System Values (RTVSYVAL) command against the QAUDLVL system value, it might not return the complete list of audit values if you are also using the new QAUDLVL2 system value.

If you make no changes to the way you are doing auditing, then everything is compatible. If you use QAUDLVL and QAUDLVL2, then you will need to retrieve both system values to get the complete list.

New system value (QTHDRSCAFN) and a new job attribute for changing and displaying the main storage affinity

In V5R2, the program QWCMEMAF (provided by PTF) allowed you to change and display the main storage affinity.

In V5R3, this program is replaced by a new system value (QTHDRSCAFN) and a new job attribute. The QWCMEMAF program object has been removed in V5R3.

When you are upgrading to V5R3, the QTHDRSCAFN system value defaults to *NOGROUP, *NORMAL which is equivalent to option('2') on the QWCMEMAF program object. The default previously was set to '0' (which meant 'No affinity') in V5R2. 'No affinity' is no longer an option in V5R3.

New system value QENDJOB LMT

The QENDJOB LMT system value replaces the QENDJOB LMT data area (*DTAARA) in the QSYS library that was supported in previous releases (via PTF). The QENDJOB LMT data area is ignored and can be deleted.

Change System Value (CHGSYSVAL) command requires *ALLOBJ authority to change the date and time system values

In prior releases, you needed only *USE authority to the Change System Value (CHGSYSVAL) command to change the following date and time system values:

- QCENTURY
- QDATE
- QDAY
- QHOUR
- QMINUTE
- QMONTH
- QSECOND
- QTIME
- QYEAR

In V5R3, you must have all object (*ALLOBJ) special authority in order to change these date and time system values, as well as the new system values QDATETIME and QTIMZON.

Changes to system QIGC value

Prior to this release, the system value QIGC (DBCS version installed indicator), was set on if a DBCS version of OS/400 was installed. In V5R3, any NLV can support DBCS, so this value will always be set to on. If you have applications that check this value, you should update them to use the job level DBCS indicator. You can find more information about this indicator in the Retrieve Job Information (QUSRJOBI) API.

QPRCMLTTSK system value changed

In V5R3, the QPRCMLTTSK (Process multitasking) system value can have character values of 0, 1, or 2. Programs that assume that the values for this system value are only 0 or 1 will need to be changed.

Default for QPWRDWN LMT system value changed

The default for system value QPWRDWN LMT will be changed from 600 seconds (10 minutes) to 900 seconds (15 minutes). This will change the value for existing customers who have QPWRDWN LMT set to 600 seconds.

Because more cleanup is done during the processing of a Power Down System (PWRDWN SYS) command, customers may need to adjust their QPWRDWN LMT system value to allow more time for the processing to complete.

QUTCOFFSET and QTIMZON system value changes

In V5R3, changes to the QUTCOFFSET system value will not occur if the specified offset conflicts with the offset calculated by the server using the time zone description set in the QTIMZON system value. In a future release, the QUTCOFFSET system value will become a read-only system value.

If you want to define the time zone for a server before you install V5R3, you should create a data area named QWCTIMZON in the QSYS library. The QWCTIMZON data area should be a character data area with a minimum length of 16 characters. The format is defined in the following way:

- Byte 1: Sign for UTC offset (either '+' or '-')
- Bytes 2 and 3: Hours for UTC offset (00 to 12)
- Byte 4: Time separator (ignored; included to improve readability of information in data area)
- Bytes 5 and 6: Minutes for UTC offset (00 to 59)
- Bytes 7 through 16: Time zone description name, left justified and padded on the right with blanks.

Because the server automatically adjusts the local system clock for Daylight Saving Time support in V5R3, you should remove other mechanisms for changing the time, such as job schedule entries and system startup program logic, that you might have created in previous releases.

For detailed information about time management, see the Time Management topic in the iSeries Information Center.

Work management changes

New default configuration

A new default job description, QDFTSVR in QGPL, is provided on the system. This job description is used by some server jobs that used to use the QDFTJOB in QGPL. The following server jobs use the QDFTSVR job description:

- QPWFSERVSO
- QPWFSERVSS
- QPWFSERVS2
- QZDAINIT
- QZLSFILE
- QANEAGNT
- QSQSRVR
- QSRRATBL
- Q5BWHSRV
- QRWTSRVR
- QZDASOINIT
- QZDASSINIT

If you have previously changed QDFTJOB to change job attributes for these server jobs, you should now change QDFTSVR or the job description used by the server.

In addition, job descriptions QZBSJOB and QPWFSERV now specify JOBMSGQMX(8) and JOBMSGQFL(*WRAP). In V5R2, QZBSJOB used JOBMSGQMX(*SYSVAL) and JOBMSGQFL(*WRAP). QPWFSERV used JOBMSGQMX(*SYSVAL) and JOBMSGQFL(*SYSVAL).

Fewer "Job started" and "Job ended" messages sent to the history log

The Job started (CPF1124) and Job ended (CPF1164) messages are no longer sent to the history log and the job log for batch immediate jobs that are started for QShell and OS/400 PASE applications. Because QShell and OS/400 PASE applications can start many jobs that are short-lived, the number of Job started and Job ended messages can become very large. The Job ended message will continue to be sent if the job ends abnormally (with an end severity of 20 or greater).

This is the first release where the system will reduce sending the Job started and Job ended messages. If you rely on these messages, you need to move to alternative ways of getting similar information, because it is likely that future releases will further reduce the sending of these messages.

Alternative ways of getting information about jobs started and ended on the system include:

- Job notification exit point. This exit point has been enhanced in V5R3 to include additional information when jobs start and end
- Security auditing (the JS audit record)
- Job accounting

Test request key no longer recognized by operating system

The test request key is no longer recognized by the operating system. Pressing the test request key on a signon panel has no effect.

Run priority changed for new batch immediate jobs

Batch immediate server jobs that are started by the system to perform work use the run priority of the thread that issued the request rather than the run priority of the job that issued the request. Batch immediate jobs started by the spawn() API continue to use the run priority of the job that issued the request. A new option on spawn() allows the caller to use the run priority of the thread that issued the request.

Jobs that provide system services for independent ASPs are now system jobs

The jobs that provide system services for independent auxiliary storage pools (ASPs) are now system jobs. On the Work with Active Jobs (WRKACTJOB) display and in the iSeries Navigator hierarchy of jobs, these jobs are no longer found in the QSYSWRK subsystem. They are now found with other system jobs. With this change, taking the system to a restricted state no longer causes independent ASPs to become unavailable.

Work management enforcement of object authorities

In V5R3, the output queue and sort sequence table for the following commands are now found in the **library name space** of the submitted job:

- Submit Job (SBMJOB)
- Batch Job (BCHJOB)

In previous releases, these were found in the library name space of the job processing the command. See the Set ASP Group (SETASPGRP) command for more information about library name space.

The user of the submitted job is the user that must be authorized to the sort sequence table. In previous releases, the user issuing the SBJJOB, SBMDBJOB, or STRDBRDR command had to be authorized to the sort sequence table. Also, the user issuing the SBJJOB, SBMDBJOB, or STRDBRDR command no longer needs to be authorized to objects touched when the command processing verifies the authorities of the user of the submitted job.

The following commands now require use (*USE) authority to the user profile specified on the command, and the user profile must exist at the time the command is issued; previously, you needed only object operational (*OBJOPR) and read (*READ) authority and you needed to be authorized to the user profile only if the object existed.

- Add Communications Entry (ADDCMNE)
- Add Prestart Job Entry (ADDPJE)
- Change Communications Entry (CHGCMNE)
- Change Prestart Job Entry (CHGPJE)

If the user who issues the command does not have all object (*ALLOBJ) special authority, the job description must exist at the time the following commands are run; previously, you needed to be authorized to the job description only if the object existed:

- Add Autostart Job Entry (ADDAJE)
- Add Communications Entry (ADDCMNE)
- Add Prestart Job Entry (ADDPJE)
- Add Work Station Entry (ADDWSE)
- Change Autostart Job Entry (CHGAJE)
- Change Communications Entry (CHGCMNE)
- Change Prestart Job Entry (CHGPJE)
- Change Work Station Entry (CHGWSE)

You now need use (*USE) authority to the job queue specified on the following commands; previously, you needed only read (*READ) authority:

- Change Job (CHGJOB)
- Transfer Job (TFRJOB)
- Transfer Batch Job (TFRBCHJOB)
- Add Job Schedule Entry (ADDJOBSCDE)
- Change Job Schedule Entry (CHGJOBSCDE)

End Job (ENDJOB) command change

In some cases, a second ENDJOB OPTION(*IMMED) is no longer ignored. For these cases, the CPF1361 message is no longer sent.

Host print transform changes

- | The host print transform for MFRTYPMDL *NECPCPR201 and the underlying QWPNEC201 WSCST object are no longer supported. You can continue to use this transform, but it will not be updated or maintained.

Chapter 3. Options

This chapter describes changes to specific options of the operating system.

S/36 and S/38 Migration (Option 4) and S/36 Migration Assistant (Option 11)

S/36 and S/38 Migration (Option 4) and S/36 Migration Assistant (Option 11) have been withdrawn effective with V5R3.

All displays, files, programs and messages related to these two options have been removed.

Support for the following commands is withdrawn in V5R3:

- ANZS34OCL (Analyze S/34 OCL)
- ANZS36OCL (System/36™ OCL Analysis)
- CHKS36SRCA (Check S/36 Source Attributes)
- CHGS34LIBM (Change S/34 Library Member)
- CVTBASSTR (Convert S/36 Stream File)
- CVTBASUNF (Convert S/36 Unformatted File)
- CVTBGUDTA (Convert BGU Data)
- CVTS36CFG (Convert S/36 Configuration)
- CVTS36FCT (Convert S/36 Forms Control Table)
- CVTS36JOB (Convert S/36 Migration Job)
- CVTS36QRY (Convert S/36 Query)
- CVTS38JOB (Convert S/38 Migration Job)
- GENS36RPT (Generate S/36 Report)
- GENS38RPT (Generate S/38 Report)
- MGRS36 (Complete System/36 Migration)
- MGRS36APF (System/36 APF Migration)
- MGRS36CBL (System/36 Cobol Migration)
- MGRS36DFU (System/36 DFU Migration)
- MGRS36DSP (System/36 Display File Migrate)
- MGRS36ITM (Migrate S/36 Item)
- MGRS36LIB (System/36 Library Migration)
- MGRS36MNU (System/36 Menu Migration)
- MGRS36MSGF (System/36 Message File Migrate)
- MGRS38OBJ (Migrate S/38 Objects)
- MGRS36QRY (System/36 Query Migration)
- MGRS36RPG (System/36 RPG II Migration)
- MGRS36SEC (System/36 User ID Migration)
- MIGRATE (Migration Menu)
- PRTMUDOC (Print Migration Documentation)
- QMUS36 (System/36 Command Selection)
- RCVMGRDTA (Receive Migration Data)

- RESMGRNAM (Resolve Migrated Name)
- RSTS36FLR (Restore S/36 Folder)
- RSTS38AUT (Restore S/38 Authorities)
- SNDMGRDTA (Send Migration Data)
- STRS36MGR (Start S/36 Migration)
- STRS38MGR (Start S/38 Migration)

Host Servers (Option 12)

Open List APIs moved to library QSYS

Prior to V5R3, the Open List APIs were part of the Host Servers option (option 12).

In V5R3, these APIs are part of base operating system (option 0) and are located in library QSYS instead of library QGY.

For V5R3, Open List router programs that are in library QGY will transfer control to the corresponding API in QSYS.

These router programs and the library QGY will be removed in a future release. You should change any applications that make library-qualified calls to the Open List APIs in library QGY to call the APIs in library QSYS. The parameters for these APIs did not change in V5R3.

The following commands will be removed from library QGY when the QGY library is removed in a future release:

- EDTWSOAUT
- GRTWSOAUT
- RVKWSOAUT
- SETCSTDTA

Ultimedia System Facilities (Option 16)

All commands and APIs related to Ultimedia System Facilities (Option 16) have been removed in Version 5 Release 3.

This includes the following commands:

- ADDUSFCNNE
- ADDUSFDEVE
- ADDUSFSVRE
- CHGUSFDEVE
- ENDUSF
- RMVUSFCNNE
- RMVUSFDEVE
- RMVUSFSVRE
- RSTUSFCNR
- SAVUSFCNR
- STRUSF
- WRKUSFCNNE
- WRKUSFDEVE

- WRKUSRSVRE

This also includes the following Ultimeia System Facilities APIs:

- QUMBCOBJ, QumbCaptureObject
- QUMBCLRD, QumbClearData
- QUMBCTLP, QumbControlProcess
- QUMBCPYO, QumbCopyObject
- QUMBCPYD, QumbCopyObjectData
- QUMBCRTO, QumbCreateObject
- QUMBCTBN, QumbCreateThumbnail
- QUMBDLTO, QumbDeleteObject
- QUMBDLTD, QumbDeleteObjectData
- QUMBOSEL, QumbObjectSelection
- QUMBMOVO, QumbMoveObject
- QUMBMOVD, QumbMoveObjectData
- QUMBPFIL, QumbPresentFile
- QUMBPOBJ, QumbPresentObject
- QUMBPRQS, QumbProcessRequest
- QUMBQCON, QumbQryConnectors
- QUMBQDEV, QumbQryDevices
- QUMBQOBJ, QumbQryObjects
- QUMBQATT, QumbQryObjectsByAttribute
- QUMBQLOC, QumbQryOriginatingLocation
- QUMBQUSR, QumbQryUser
- QUMBQWPS, QumbQryWinPositionAndSize
- QUMBRSVK, QumbResolveKeyObject
- QUMBRAUT, QumbRtvObjectAuthority
- QUMBRTVD, QumbRtvData
- QUMBRMMM, QumbRtvMultimediaMessage
- QUMBRATT, QumbRtvObjectAttribute
- QUMBRTBN, QumbRtvObjectThumbnail
- QUMBROWN, QumbRtvObjectOwner
- QUMBRUID, QumbRtvUniqueID
- QUMBRUAU, QumbRtvUserAuthorityToObject
- QUMBRTND, QumbReturnData
- QUMBSAID, QumbSendNamedAsyncID
- QUMBSATT, QumbSetObjectAttribute
- QUMBSAUT, QumbsetObjectAuthority
- QUMBSOWN, QumbSetObjectOwner
- QUMBSTBN, QumbSetObjectThumbnail
- QUMBSWPS, QumbSetWinPositionAndSize
- QUMBVCLS, mciCloseAll
- QUMBVDLG, mciDeleteGroup
- QUMBVGID, mciGetDeviceID
- QUMBVGES, mciGetErrorString

- QUMBVGNM, mciGetNotifyMsg
- QUMBVMKG, mciMakeGroup
- QUMBVSND, mciSendCommand

OS/400 Portable Application Solutions Environment (Option 33)

Qp2RunPase API change

The Qp2RunPase API changed to send an escape message if the `symbolName` argument is *not* a null pointer. OS/400 PASE programs can call ILE or OPM code using runtime functions `_ILECALL` and `_PGMCALL` to obtain data (such as MI pointers) that cannot be expressed as null-terminated character string arguments or environment variables.

OS/400 PASE program name on job displays

The name of the OS/400 PASE program currently running in a job is now shown in the **Function** column on job list displays such as those displayed with the Work With Active Jobs (WRKACTJOB) command. If the name is too long, only the rightmost part is shown.

International Components for Unicode (Option 39)

IBM intends to withdraw the 1.6, 1.8.1, and 2.0 versions of ICU in a future release. In anticipation of this, you should begin using ICU version 2.6.

Chapter 4. Licensed programs

This chapter contains information about individual licensed programs that have been modified or changed for V5R3.

Backup, Recovery and Media Services (5722-BR1)

BRMS product initialization required

The Backup, Recovery and Media Services (BRMS) for iSeries (5722-BR1) product must be initialized following installation of the Licensed Program Product or any subsequent re-installation of the BRMS primary language. A product initialization job is automatically submitted to the batch queue upon successful installation of the BRMS primary language. You need take no additional action other than to monitor for the successful job completion message in the QSYSOPR message queue. The qualified job name is:

```
job-number/install-userid/BRMSPRDINZ
```

You cannot use the BRMS menus, commands, APIs and the BRMS iSeries Navigator client until product initialization is complete. You will get the exception, BRM40A2 - BRMS product initialization required, if you attempt to use BRMS functions before the product is initialized. If you continue to receive this exception and you cannot find the initialization job on the system, you can perform the product initialization by running the following command:

```
INZBRM OPTION(*RUNPRDINZ)
```

BRMS product initialization can be a long running process. You can optimize its performance by running it when there is little or no other activity on the server. To find out the approximate length of time the process will require, you can run the following command:

```
INZBRM OPTION(*ESTPRDINZ)
```

For more precise estimates, always run this command in the same system environment as the initialization job.

Special values for ASPDEV parameter changed

In V5R3, the allowed special values for ASPDEV parameter have changed. The Auxiliary storage pool (ASPDEV) parameter on the following commands has been changed to disallow "*" as a possible value. For some commands, the default value was also changed as indicated.

Command	Default value	Changed
Save Library using BRM (SAVLIBBRM)	ASPDEV(*SYSBAS)	Yes
Save Object using BRM (SAVOBJBRM)	ASPDEV(*SYSBAS)	Yes
Save Object List using BRM (SAVOBJLBRM)	ASPDEV(*SYSBAS)	Yes
Save System using BRM (SAVSYSBRM)	ASPDEV(*ALLAVL)	No

DB2 UDB XML Extender (5722-DE1, Option 2)

- | In prior releases, the DB2 UDB XML Extender shipped the XML Parser technology. In V5R3, XML Extender will use the XML Parser which is delivered in the XML Toolkit for iSeries (5733-XT1 option 7).
- | This product and option must be installed before you use DB2 UDB XML Extender for iSeries on V5R3.

If you are installing DB2 UDB XML Extender on V5R3 over a prior release of XML Extender, and your database has already been enabled for XML Extender using the `enable_db` Administration command, you must complete the following migration step; otherwise, the `disable_db` Administration command might fail. From the command line, enter the following command:

```
CALL QDBXM/QZXMIGV
```

IBM HTTP Server (5722-DG1)

Java CGI no longer supported

Java CGI is no longer supported with IBM HTTP Server for iSeries. Java CGI was supported on past releases with Java Developer Toolkit (JDK) Version 1.1 only. There is no support for JDK levels other than 1.1. Since JDK 1.1 is no longer shipped with the 5722-JV1 product in V5R3, Java CGI is no longer supported with IBM HTTP Server (5722-DG1) in this release.

Configuration update for Tomcat 3.2.4

If you use Apache Software Foundation's Jakarta Tomcat 3.2.4 that is shipped with IBM HTTP Server, you need to update your Tomcat configuration to use version 1.3 of the Developer Kit for Java (5722-JV1). This change is necessary because Version 1.2 of the Developer Kit for Java has been removed in V5R3 (see "Developer Kit for Java (5722-JV1)" on page 57).

Error log file expanded

The DDS format for the HTTP server Error Log file has been expanded to include more information. The QHTTPSVR/QAZHBERR file has been updated to contain two additional fields. The ERRLVL and ERRTXT fields have been appended to the ERRSRC record format. Programs accessing this file should have LVLCHK(*NO) specified.

HTTP Server (original) withdrawn

HTTP Server (original) has been withdrawn from the IBM HTTP Server for iSeries (5722-DG1) product in V5R3. IBM first announced plans to withdraw the original server in the V5R2 iSeries Upgrade Planning Web site at the following URL:

- | <http://www/systems/support/i/planning/>

If you use HTTP server, you must migrate your HTTP Server (original) configurations to HTTP Server (powered by Apache). Web applications that include static (HTML) pages, CGI programs, or Server Side Includes (SSI) do not require migration and are fully supported on HTTP Server (powered by Apache).

A migration wizard is provided with the IBM Web Administration for iSeries interface to assist in the migration. You should migrate from the HTTP Server (original) to HTTP Server (powered by Apache) before you install V5R3. The migration wizard is available with the IBM Web Administration for iSeries interface in V5R1, V5R2, and V5R3. The migration is recommended prior to V5R3 so that you can test the migration to HTTP Server (powered by Apache) while continuing to serve production Web sites using HTTP Server (original) during the migration. Original server configurations stored in QUSRSYS/QATMHTTTPC will not be altered or deleted during a V5R3 install. This will allow you to migrate original server configurations to HTTP Server (powered by Apache) after installing V5R3, should you choose to do so. Original server instances stored in QUSRSYS/QATMHINSTC will be removed

during a V5R3 install so that they no longer appear on the IBM Web Administration for iSeries interface after installation of V5R3. Original server instances will be deleted, but original server configurations will be left to allow for migration.

Developer Kit for Java (5722–JV1)

Support for option 3 (JDK 1.2), option 4 (JDK 1.1.8), and JPDC removed

Support for option 3 (JDK 1.2) and option 4 (JDK 1.1.8) in the Developer Kit for Java (5722–JV1) product has been removed in V5R3. If you have applications that require these JDK levels, you must upgrade to the newer JDK level. JDK levels 1.3 or 1.4 are fully compatible with JDK level 1.2, and an upgrade should pose no problems. However, Sun did make interface changes between SDK 1.1.x and J2SDK. If your application requires JDK 1.1.8, you might need to make code changes in order to run with J2SDK (or a later JDK) on V5R3.

Support for the Java PEX Data Converter (JPDC) has been removed in V5R3. Prior to this release, JPDC was shipped as an optional extension with the Developer Kit for Java product. The functions provided by JPDC are now provided by the Performance Trace Data Visualizer (PTDV) tool, a PC-based tool for viewing and analyzing data generated by Performance Explorer. You can learn more about this tool at the following URL:

<http://www.alphaworks.ibm.com/tech/ptdv>

Support for Remote AWT withdrawn

The Abstract Windowing Toolkit (AWT) classes provide graphics functions to Java programmers. The iSeries server originally provided Remote AWT (RAWT) to delegate graphics services and display to a remote device with native graphics capability running Java. As Java and the capabilities available on the iSeries evolved, Native AWT (NAWT) was provided so the generation and manipulation of graphics (but not display) could be carried out without having a remote device involved. The NAWT support has matured to the point where it provides a superior solution for graphics operations and display (through an attached X-server); therefore, support for RAWT is withdrawn in V5R3. For more information about NAWT, see the Native Abstract Windowing Toolkit topic in the iSeries Information Center.

Performance Tools (5722-PT1)

Changes to interactive and batch workload reporting by Performance Tools in V5R3

In previous releases, the Performance Tools product categorized DDM and iSeries Access jobs as interactive work. Starting with V5R3, DDM jobs are categorized as batch work. iSeries Access jobs will be categorized as interactive or batch, depending on their job type. If an iSeries Access job has a job type of "I", it will be considered interactive; otherwise, it will be considered batch. This means that iSeries Access jobs could appear in either the report sections for interactive workload or in report sections for non-interactive workload, or both. DDM jobs will appear in the non-interactive workload sections.

Additionally, the formulas to calculate overall Interactive CPU utilization and batch CPU utilization changed. In previous releases, Performance Tools categorized workload as interactive or batch based on the following criteria:

- Interactive: Interactive, DDM, iSeries Access, Passthrough, MRT, and System/36
- Batch: Batch, Spool, Evoke, and Autostart

Starting with V5R3, overall interactive CPU utilization will report the total CPU consumed by all jobs that the machine considers to be 5250 OLTP (online transaction processing) capable jobs. Overall batch CPU utilization will include all other jobs and system work not reported as interactive CPU utilization.

Performance report percentage values

In V5R3, a performance report that is based on the same data as a report from an earlier release (converted database files or database files generated from the same *MGTCOL) will show different percentage values for most CPU reporting. Percentages for total system utilization, including interactive feature utilization, database utilization and secondary workload utilization, are reported the same. However, percentages for individual jobs and groups of jobs (for example, batch, interactive, and so on) will change. This is because the basis for the calculated percent changed from a virtual processor perspective to one based on actual available CPU time. This change provides a more realistic or expected view of the data.

IBM Content Manager OnDemand for iSeries (5722-RD1)

STRASMOND command changed

Three keywords have been removed from the STRASMOND command in V5R3: LOGSTS, SNDFAILMSG, and USER. Previously, these parameters were ignored.

DB2 Query Manager and SQL Development Kit for iSeries (5722-ST1)

CRTSQLC command removed

The CRTSQLC command is no longer shipped because there has not been a supported compiler for it to use for many releases. The supported C precompiler command is CRTSQLCI.

Precompiler commands changed

In V5R3, the CRTSQLRPG, CRTSQLRPGI, CRTSQLCI, CRTSQLCBLI, CRTSQLCBL, CRTSQLCPPI, CRTSQLPLI precompiler commands have changed.

If you specify INCFIL(*LIBL/*SRCFILE) in V5R3, the command will fail. In V5R3, you must specify this as INCFIL(*SRCFILE).

WebSphere Development ToolSet (5722-WDS)

Application Development Manager (Option 22) and Application Dictionary Services (Option 23) licensed programs withdrawn

Application Development Manager (Option 22) and Application Dictionary Services (Option 23) licensed programs have been withdrawn effective with V5R3.

You can use numerous products as an alternative to these discontinued features. Among these alternatives are the following products:

- Softlanding TurnOver Change Management. See the information at the following Web site for details:
<http://www-1.ibm.com/support/docview.wss?uid=nas1084b0c4f0f30547486256c5a00778940>
- Aldon Lifecycle Manager. See the information at the following Web site for details:
http://www-304.ibm.com/jct09002c/gsdod/showcase.do??solution_id=soq96494123171205002|30

All displays, files, programs, and messages related to these options have been removed.

In addition, support for the following commands is withdrawn effective with V5R3:

- ADDADMLANG (Add ADM Language)
- ADDADMTYPE (Add ADM Type)
- ADDPRJLIBL (Add Project Library List)

- ADDPRJUSR (Add Project User)
- BLDPART (Build Part)
- CHGADMACN (Change ADM Action)
- CHGADMLANG (Change ADM Language)
- CHGGRP (Change Group)
- CHGPART (Change Part)
- CHGPARTINF (Change Part Information)
- CHGPRJ (Change Project)
- CHGPRJUSR (Change Project User)
- CHKINPART (Check Part In)
- CHKOUTPART (Check Part Out)
- CMPPART (Compare Part)
- CPYPART (Copy Part)
- CRTGRP (Create Group)
- CRTPART (Create Part)
- CRTPRJ (Create Project)
- CVTPART (Convert Part)
- DLTGRP (Delete Group)
- DLTPART (Delete Part)
- DLTPRJ (Delete Project)
- DSPPART (Display Part)
- EXPPART (Export Part)
- FNDSTRPART (Find String in Parts with PDM)
- FNDSTRPAR2 (Find String in Parts with list)
- IMPPART (Import Part)
- MRGPART (Merge Part)
- PRMPART (Promote Part)
- PRTADMLANG (Print ADM Language)
- PRTADMTYPE (Print ADM Type)
- PRTPARTINF (Print Part Information)
- PRTPRJ (Print Project)
- PRTPRJLOG (Print Project Log)
- PRTPRJUSR (Print Project User)
- QRYPART (Query Part)
- QRYPRJ (Query Project)
- RCLPRJ (Reclaim Project)
- RCVPART (Receive Part)
- RMVADMLANG (Remove ADM Language)
- RMVADMTYPE (Remove ADM Type)
- RMVPRJLIBL (Remove Project Library List)
- RMVPRJUSR (Remove Project User)
- RNMPART (Rename Part)
- RTVPARTINF (Retrieve Part Information)
- STRADS (Start AppDict Services/400)
- WRKAPPDCT (Work with Dictionaries)

- WRKPRJPDM (Work with Projects using PDM)
- WRKGRPPDM (Work with Groups using PDM)
- WRKPARTPDM (Work with Parts using PDM)

New reserved words added to ILE COBOL

In V5R3, the following new reserved words were added to ILE COBOL:

- AS
- CONSTANT
- END-XML
- FILE-STREAM
- PARSE
- PROCESSING
- XML
- XML-CODE
- XML-EVENT
- XML-NTEXT
- XML-TEXT

ILE COBOL programs that use any of these words as a data name will not compile successfully in V5R3 (although they would have compiled successfully in previous releases). To compile successfully in V5R3, you must rename any data names that exactly match any of the above new reserved words.

Existing ILE COBOL programs that use these reserved words are not affected unless they are recompiled using V5R3 ILE COBOL with TGTRLS *CURRENT.

IBM Open Class - source and samples (Option 55) withdrawn

IBM Open Class[®] - source and samples (Option 55) has been withdrawn effective with V5R3M0.

ILE C and ILE C++ *PRV Compilers not separate options

In V5R3, the ILE C compiler (Option 51) and ILE C++ compiler (Option 52) include support for creating ILE modules for V5R3M0 as well as previous releases, V5R2M0 and V5R1M0. Therefore, Option 53 (Compiler - ILE C *PRV) and Option 54 (Compiler - ILE C++ *PRV) are not needed and are withdrawn effective with V5R3.

Independent ASP support enhancements

In V5R3, SEU, SDA and RLU were enhanced to include the auxiliary storage pool (ASP) name in the recovery entries; this prevents the recovery function from being performed on the wrong member with the same name and library in different ASPs.

Complete support for independent ASPs on the Start Program Development Manager (STRPDM) and Work with Libraries using Program Development Manager (WRKLIBPDM) commands provides the same parameters as the Work with Libraries (WRKLIB) command. This support allows PDM to access libraries in different ASPs with selection criteria on the ASP device or ASP number.

iSeries Integration for Windows Server (5722-WSV)

For V5R3, the Install Windows Server (INSWNTSVR) command has changed several parameters which might require CL programs that use the command to have source changes and be recompiled:

- The Windows server version (WNTVER) parameter removed the special values *NT40 and *WIN2002, and added special value *WIN2003 when installing a Windows Server 2003.

- The Domain role (DMNROLE) parameter has removed the special value *BKUCTL which was only allowed when WNTSVR(*NT40) was specified.
- The Server storage space sizes (SVRSTGSIZE) minimum value for the first element (Install source size) has increased to 500 MB. The minimum value for the second element (System size) has increased to 1024 MB.
- The License mode (LICMODE) parameter removed the special value *PERSEAT and *PERSESSION for the third parameter element (Terminal services), and added special values *PERDEVICE and *PERUSER when installing a Windows Server 2003 server with the Terminal Services optional component. The *NT40 number of terminal services license numbers has also been removed.
- The Server domain name (SVRDMN) parameter has been removed.
- The Cluster configuration (CLUCFG) parameter added two new parameter elements (quorum resource size and ASP device name) and removed two existing parameter elements (account name and account password).

iSeries Access for Windows (5722-XE1)

Support for Application Development folder and Interprocess Communication function in iSeries Navigator

The Application Development function and the Interprocess Communication subfunction have been removed from the iSeries Access for Windows installation media in V5R3. If you require this function, you have two alternatives:

- If the iSeries Navigator Network component is already installed on the PC, install any iSeries Access for Windows service pack, which will create the Application Development folder and the Interprocess Communication sub-function in the iSeries Navigator tree. If the Network component is not installed, run Selective Setup from an installation image that contains any iSeries Access for Windows service pack; this installs the iSeries Navigator Network component and creates the Application Development folder automatically.
- You can call two programs to replace the functionality provided by the Interprocess Communication subfunction. You can get status about Interprocess Communication objects by using the `ipcs` utility. The output produced by `ipcs` when the `-E` option is specified has been changed to provide the same type of information as was available with the Interprocess Communication sub-function. You can delete Interprocess Communication objects by using the `ipcrm` utility. Documentation for `ipcs` and `ipcrm` is available in the iSeries Information Center.

Migration of PC5250 workstation profiles

In V5R3, the new version of PC5250 that is included in iSeries Access for Windows requires that certain supporting files (such as `.ws` files) be migrated to a new version level. This migration occurs automatically the first time a file is used, unless you take certain steps to prevent it. Migrating files can result in users with earlier versions of PC5250 (which are in earlier versions of iSeries Access for Windows) not being able to use the migrated files. This can happen, for example, if multiple users use the same files from a shared location. Not migrating a file may result in that file being unusable to PC5250. Click the following paths to control the migration of supporting files using iSeries Navigator:

- **iSeries Access for Windows Properties > PC5250**
- **Application Administration > Central Settings > Client Applications > Advanced Settings > PC5250**

Note: When files are automatically migrated, the original file is saved with a slightly modified extension (for example `abc.ws` is saved as `abc.ows`).

iSeries Navigator changes

Offset from Greenwich Mean Time (GMT) system value

In V5R3, you can no longer change the Offset from Greenwich Mean Time (GMT) system value with iSeries Navigator. The offset is now determined by the time zone specified in the QTIMZON system value.

Messages sent with iSeries Navigator

In V5R3, when you send messages with iSeries Navigator, your user name is no longer appended to the end of the message text. You can view the sender information on the properties page or in the *From User* column on the message list. The user name will continue to be appended for an iSeries server running on releases earlier than V5R3.

iSeries Access POP server options adjusted

The POP server options on the iSeries have been adjusted to be consistent with what is now available through iSeries Access.

Special MAPI interfaces through iSeries Access (SNA, IPX, IP, Address Book) were removed in prior releases:

- In V4R3, SNA and IPX through Client Access was removed.
- In V5R1, the MAPI interfaces for IP and Address Book were removed.
- In V5R2, the POP servers that supported MAPI interfaces were removed.

The POP server that communicates with any POP client through the standard TCP/IP POP interface is still available.

Windows Operating Systems support changes

In V5R3, you can no longer install iSeries Access for Windows on Windows 95, Windows 98, or Windows ME operating systems. You can install it on Windows NT®, Windows 2000, Windows XP, and Windows Server 2003. Future releases will not support installing on Windows NT. If you want to upgrade your PC to one of the supported Windows operating system, you will need to do the following steps:

1. Uninstall iSeries Access for Windows.
2. Upgrade the Windows operating system.
3. Install iSeries Access for Windows.

AFP and SCS Printer Driver considerations

Beginning with Windows 2000 and Windows XP, when the Driver Signing option is set to Warn or Block, and Windows encounters a file that has not been digitally signed, the Windows operating system interrupts install programs and issues an Windows Logo Testing warning/error message. In previous versions of Windows operating systems, you were not required to take any action because Windows allowed install programs to bypass the setting for Driver Signing Options.

The AFP™ and SCS printer drivers components are not digitally signed by Microsoft®. For this reason, in V5R3, printer drivers have been removed from a Typical install.

If you select to install printer drivers through one of the other installation methods and you do not change the Driver Signing option to Ignore, you will need to add or update your printer drivers using Microsoft's directions in their help text, specifying the c:\Program Files\IBM\Client Access\CWBAFP directory location for the AFP printer driver and specifying the c:\Program Files\IBM\Client Access\CWBSCS directory location for the SCS printer driver after the iSeries Access for Windows installation program completes.

Upgrading from a previous product or release

In V5R3, you cannot migrate from AS/400 Client Access for Windows 95/NT to iSeries Access for Windows. You will need to uninstall AS/400 Client Access for Windows 95/NT and then install iSeries Access for Windows.

The supported releases that you can upgrade to V5R3 iSeries Access for Windows are V5R1 AS/400 Client Access Express and V5R2 iSeries Access for Windows.

Distributing iSeries Navigator plug-ins and add-ins

Future releases will support installing iSeries Navigator plug-ins only from the iSeries server integrated file system QIBM/USERDATA/OpNavPlugin directory. If you currently use the QIBM/USERDATA/GUIPlugin directory, you should plan to migrate your plug-in to the QIBM/USERDATA/OpNavPlugin directory. IBM iSeries Navigator plug-ins will be migrating to the QIBM/PRODDATA/OpNavPlugin directory. (These directories are not automatically created by the system; you might need to create them manually.) Also, future releases will not support installation of add-ins.

Data Transfer considerations

Starting with V5R3, Data Transfer takes advantage, by default, of the new UTF-8 (CCSID 1208) DB2 UDB for iSeries support when creating files on the host when the client file contains Unicode data. V5R3 clients accessing pre-V5R3 host systems will, by default, use UCS-2 (CCSID 13488) DB2 UDB support when creating these files. You should consider not using these defaults when the client is in a network accessing both pre-V5R3 and V5R3 host systems in order to avoid possible compatibility concerns when later upgrading the pre-V5R3 system to V5R3.

If you use the Microsoft Excel Data Transfer Add-in, an additional change has been made. In V5R2, when scanning a Microsoft Excel spreadsheet, cells containing character data would result in the cell being tagged, by default, as CHAR. This would result in the host file field being created as CHAR. In V5R3, these cells will be tagged by default as UNICODE. This would result in the host file field being created as GRAPHIC. This may cause problems with host applications that read the data from the host file. If you think this will cause a problem, change the field type from UNICODE to CHAR after the spreadsheet is scanned.

Ultimedia System Facilities

All APIs related to Ultimedia System Facilities have been removed in V5R3.

iSeries Access for Web (5722-XH2)

Before you install V5R3 iSeries Access for Web:

If you currently have a configured, running version of iSeries Access for Web on your iSeries server, you must run the iSeries Access for Web configuration commands again when you install the new version before you can use the product. For detailed instructions on installing and upgrading the product and running the configuration commands, see the V5R3 iSeries Information Center; click **Connecting to iSeries > Choose your connection > iSeries Access > iSeries Access for Web**.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
500 Columbus Avenue
Thornwood, NY 10594-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation

Software Interoperability Coordinator, Department 49XA
3605 Highway 52 N
Rochester, MN 55901
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, or other countries, or both:

AFP
AS/400
DB2
DB2 Connect
Enterprise Storage Server
eServer
i5/OS
IBM
IBM Global Network
iSeries
Open Class
Operating System/400
OS/400
POWER5
Redbooks
System/36
System p5
TotalStorage
WebSphere

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.

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

Terms and conditions for downloading and printing publications

Permissions for the use of the publications you have selected for download are granted subject to the following terms and conditions and your indication of acceptance thereof.

Personal Use: You may reproduce these Publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these Publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these Publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these Publications, or reproduce, distribute or display these Publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the Publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the Publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations. IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE

All material copyrighted by IBM Corporation.

By downloading or printing a publication from this site, you have indicated your agreement with these terms and conditions.

Readers' Comments — We'd Like to Hear from You

iSeries
Memorandum to users
Version 5 Release 3 (May 2008 update)

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Comments:

Thank you for your support.

Submit your comments using one of these channels:

- Send your comments to the address on the reverse side of this form.
- Send a fax to the following number: United States and Canada: 1-800-937-3430
Other countries: 1-507-253-5192
- Send your comments via e-mail to: RCHCLERK@us.ibm.com

If you would like a response from IBM, please fill in the following information:

Name

Address

Company or Organization

Phone No.

E-mail address



Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM CORPORATION
ATTN DEPT 542 IDCLERK
3605 HWY 52 N
ROCHESTER MN 55901-7829



Fold and Tape

Please do not staple

Fold and Tape



Printed in USA