



System i
i5/OS Memo to Users

Version 6 Release 1 (July 2012 update)





System i
i5/OS Memo to Users

Version 6 Release 1 (July 2012 update)

Note

Before using this information and the products it supports, be sure to read the information in "Notices" on page 81.

This edition applies to version 6, release 1, modification 0 of the IBM i5/OS (Program 5761-SS1) and to all subsequent releases and modifications until otherwise indicated in new editions. This version does not run on all reduced instruction set computer (RISC) models nor does it run on CISC models.

© **Copyright IBM Corporation 2008.**

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

Contents

About i5/OS Memo to Users vii

Who should read this memorandum	vii
Additional incompatibility information	vii
Accessibility features for i5/OS Memo to Users	vii
What's new	viii
Installing V6R1 over V5R3.	x
Discontinued support for certain software and hardware	x
PTF numbers in this memorandum	x
Memorandums for previous releases	x
Prerequisite and related information	x
How to send your comments	xi

Chapter 1. Read this first 1

Current customers only — before you install i5/OS V6R1	1
Server firmware requirements for V6R1	2
Process for ordering cumulative PTF packages	2
Program conversion	3
Optical media for V6R1.	4
Ensure that the cluster version is current before installing V6R1	4
Plan for your Operations Console installation or upgrade	4
Information APARs for software installation	5
Information APAR II14310	5
Upgrade planning	5
Prerequisite information	6
Support for I/O expansion units.	6
Setting up Ethernet connectivity from i5/OS logical partitions to the Hardware Management Console	6
Key messages in this document	6

Chapter 2. i5/OS 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
Changes to system printer files and other IBM-supplied objects	7
Adopt authority for Java programs requires new PRPQ.	7
ALCOBJ for data queues	8
API changes	9
API binding directory updates	9
Changes to flight recorder APIs (QWTDMPFR and QWTDMPLF)	10
Create User Space (QUSCRTUS) API	10
Get File Information - stat() API	10
List Objects Secured by Authorization List (QSYLATLO) API returns path names using Unicode CCSID	10
Perform Shared Memory Control Operations (shmctl and shmctl64) APIs	10

Retrieve Command Information (QCDRCMDI) API not returning help bookshelf	10
Serial number format changes for hardware resource APIs and commands	10
Backup and recovery changes	11
Omitting objects on the SAV or RST commands	11
Saving and restoring journal receivers to a previous release	11
Changes to save and restore limits.	11
Virtual tape changes	12
GO SAVE, GO RESTORE, GO LICPGM changes	12
Bidirectional data changes	12
C and C++ changes.	12
C language runtime changes.	12
New 64-bit time type definitions and routines	13
Changes affecting static C++ objects	13
Clustering changes	13
Installation requirements for the cluster version	13
Change Cluster Resource Group (QcstChangeClusterResourceGroup) API.	13
Cluster command changes	14
Clusters GUI changes	14
Clustering jobs	15
Device cluster resource group changes	15
i5/OS command changes.	15
Change Network Server Desc (CHGNWSD) command	15
Convert SQL C++ Source (CVTSQLCPP) command	15
Copy from Import File (CPYFRMIMPF) and Copy To Import File (CPYTOIMPF) commands	16
Copy To Stream File (CPYTOSTMF) and Copy From Stream File (CPYFRMSTMF) commands.	17
CPYF and CPYSRCF will not update 'Last source update date/time'	17
Create Network Server Desc (CRTNWSD) command changes	17
Create NWS Configuration command parameter default changed	17
Delete User-Defined FS command	18
Display Program References (DSPPGMREF) command changes	18
*DOC option removed from VFYSRVCFG command	18
Extended Serial Number field for hardware resource commands	18
IBM commands no longer exist in secondary language libraries	18
New records for the outfile for the Display Authorization List Objects command	19
Print Performance Explorer (PRTPEXRPT) command positional parameter limit	19
Start Database Monitor (STRDBMON) command	19
Start Object Conversion (STROBJCVN) command	19
Text description TEXT(*CMDPMT) retrieved dynamically	19

Teraspace-enabled parameter changes for CHGPGM and CHGSRVPGM commands	19	Exit point changes	34
*VRT option removed from the Change Device Desc (Optical) (CHGDEVOPT) command	20	QIBM_QPWFS_FILE_SERV exit point changes	34
Work with Network Server Status (WRKNWSSTS) command (change to SVRTYPE parameter).	20	QIBM_QWC_PRERESTRICT and QIBM_QWC_PWRDWN SYS exit point changes	34
Submit Remote Command (SBMRMTCMD) command	20	File-level backup through QNTC file systems	34
Conversion of authorization format	20	FTP changes	34
Database changes	20	Handling inquiry messages	35
Change in cursor positioning after a ROLLBACK HOLD statement for cursors.	20	Host Ethernet Adapter changes.	35
Changes to LEFT and RIGHT Scalar functions.	21	Host print transform changes	35
DB2 DECFLOAT changes.	21	Host server job changes	35
Open data path (ODP) I/O feedback changes	21	Integrated file system names	35
SQL materialized query table changes	21	Integrated Language Environment (ILE) Compiler changes.	36
Start Database Monitor (STRDBMON) changes	21	IPv6 changes	36
QAQQINI option changes	22	New host table file	36
International Components for Unicode sort sequence tables	22	QUSRSYS/QATOCTCPIP (TCP/IP domain information) file	36
SQL and DB2 compatibility	22	TCP/IP configuration changes	37
User-defined functions for common table expressions invoked according to SQL standards	22	IPv6 socket changes	37
SQL Call Level Interface changes	23	Java Database Connectivity (JDBC) changes	38
Object management authority changes	23	java.sql.DatabaseMetadata changes	38
Additional AF (Authority Failure) journal entries created by host server jobs.	24	New QueryCloseImplicit JDBC connection property	38
Substitution character warning changes	24	Java programs created by CRTJVAPGM have no DE code.	38
SQL function statement changes	24	Job queue object changes	39
DBCLOB CCSID changes	24	Job termination changes	39
RRN scalar function changes	24	Display Program Messages screen no longer shown during job termination.	39
Resolving an SQL0338 failure	24	Message waiting light not processed at job end.	39
DESCRIBE OUTPUT changes	25	Journaling enhancements	40
Trimming blanks for C NULL terminated	25	LICOPT usage	43
SQE caching of temporary result sets when UDFs and UDTFs are involved	25	Machine interface (MI) instruction changes	43
SQL programming changes	26	DESMTX MI instruction	43
SQL Query changes in collation of errors might result in a SQL0802.	28	MATAL MI instruction	43
Badly formed data can trigger data mapping error.	28	MATINXAT MI instruction	43
Implicit numeric conversions may now return an SQL0302, SQL0678, CPD4363 or CPD3131	29	MATMATR MI instruction	43
Queries that call a user-defined function changes	29	MATPRATR MI instruction	43
Queries converted to INTEGER type	30	MATRMD MI instruction option hex 12 and 20	43
Queries converted to REAL type	30	POWER6 systems with EnergyScale	44
JOIN with a USING clause change	30	Teraspace location locks	44
Date and time function changes	30	NetServer User Profile changes	45
Database Create File program changes	31	Network File System changes	45
Changes for user-defined external functions that call a service program	31	Network server host (NWSH) changes	46
Create Table Like changes for Implicitly Hidden columns	31	Network storage space changes.	46
SQLTABLEPRIVILEGES now returns Index privileges	32	New message queue for electronic customer support messages	46
DRDA package management changes.	32	Optical media changes.	46
New RUNSQL command.	32	Performance changes	46
Device files may experience performance problems	32	Collection Services changes	46
Domain name system (DNS) changes.	32	Performance database file (QAPM*) field definition changes	47
ENDTCP before ENDSBS *ALL *IMMED	34	Performance data collector changes	47
		Performance Explorer changes	47
		Performance Management Agent	48
		Performance Management Agent command changes.	48
		Power-down system confirmation panel	48
		Program conversion	48
		Conversion of module, program, and service program objects	48

XML Toolkits 5733-XT1 and 5733-XT2	79	Programming Interface Information	82
IBM Cryptographic Support for AS/400 (5722-CR1)	80	Trademarks	82
5798FAX - IBM Facsimile Support for IBM i	80	Terms and conditions	83
Notices	81		

About i5/OS Memo to Users

This information describes the changes in version 6, release 1, modification 0 (V6R1) 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 i5/OS Memo 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 V6R1. 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, and changes that could affect the way things operate or appear in the new release. This chapter is intended for all users of the IBM System i computers.
- **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 System i computers.
- **Licensed programs** contains new release changes that might affect existing applications. These changes might also affect applications saved on a V6R1 system to be restored on a previous release server. This chapter is intended for application programmers and system programmers who use the System i computers 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 i5/OS Memo to Users, updates to this document will be available in the English Internet version of the i5/OS Information Center at this Web site:

<http://www.ibm.com/systems/i/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, PTF cover letters, see preventive service planning (PSP), and authorized problem analysis reports (APAR) information, at the Support for IBM System i Web site:

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

Click **Technical Databases**.

Accessibility features for i5/OS Memo to Users

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

The V6R1 i5/OS Memo to Users also is available as a text file on the Support for IBM System i Web site. Click **Technical Databases > Preventive Service Planning (PSP) > All Preventive Service Planning Documents by Release > R610**.

What's new

| The following revisions or additions have been made to the Memo to Users since the first V6R1 publication:

- | • **July 2012 update:**
 - | – “Change in cursor positioning after a ROLLBACK HOLD statement for cursors” on page 20 was updated.
 - | – “Additional AF (Authority Failure) journal entries created by host server jobs.” on page 24 was updated.
 - | – “Resolving an SQL0338 failure” on page 24 was updated.
 - | – “SQLCODE and SQLSTATE changes” on page 26 was updated.
 - | – “Database Create File program changes” on page 31 was updated.
 - | – “SQLTABLEPRIVILEGES now returns Index privileges” on page 32 was added.
 - | – “New RUNSQL command” on page 32 was added.
 - | – “Message waiting light not processed at job end.” on page 39 was added.
 - | – “NetServer User Profile changes” on page 45 was added.
- | • **May 2011 update:**
 - | – “Implicit numeric conversions may now return an SQL0302, SQL0678, CPD4363 or CPD3131” on page 29 was updated.
 - | – “Watch for event support changes” on page 58 was added.
 - | – “5798FAX - IBM Facsimile Support for IBM i” on page 80 was added.
- | • **March 2011 update:**
 - | – “Job termination changes” on page 39 was added.
 - | – “Database Create File program changes” on page 31 was updated.
 - | – “Trimming blanks for C NULL terminated” on page 25 was added.
- | • **September 2010 update:**
 - | – “SQE caching of temporary result sets when UDFs and UDTFs are involved” on page 25 was added.
 - | – “Additional AF (Authority Failure) journal entries created by host server jobs.” on page 24 was added.
 - | – “IBM Cryptographic Support for AS/400 (5722-CR1)” on page 80 was added.
- | • **June 2010 update:**
 - | – “Implicit numeric conversions may now return an SQL0302, SQL0678, CPD4363 or CPD3131” on page 29 was added.
 - | – “XML Toolkits 5733-XT1 and 5733-XT2” on page 79 was added.
- | • **April 2010 update:**
 - | – “BRMS networking changes:” on page 63 was added.
- | • **February 2010 update:**
 - | – “GO SAVE, GO RESTORE, GO LICPGM changes” on page 12 was added
 - | – “Embedded SQL with WHENEVER SQLWARNING handlers” on page 27 was added
 - | – “SQL0445 and SQLWARN1 for stored procedure output parameters” on page 27 was added
 - | – “SQL Query changes in collation of errors might result in a SQL0802” on page 28 was added
 - | – “Device files may experience performance problems” on page 32 was added
 - | – “ENDTCP before ENDSBS *ALL *IMMED” on page 34 was added
 - | – “Performance database file (QAPM*) field definition changes” on page 47 was added
- | • **October 2009 update:**
 - | – “Domain name system (DNS) changes” on page 32 was updated.

- | – “Performance Management Agent command changes” on page 48 was added.
- | – “Create Table Like changes for Implicitly Hidden columns” on page 31 was added.
- | – “Install Linux Server (INSLNXSVR) command” on page 60 was updated.
- | – “IBM Extended Integrated Server Support for i5/OS (5761-LSV)” on page 64 was updated.
- | – “Trigger cache manager” on page 67 was added.
- | – “Web server search engine and Web crawler support” on page 67 was added.
- | • **May 2009 update:**
 - | – “DRDA package management changes” on page 32 was added.
 - | – “Submit Remote Command (SBMRMTCMD) command” on page 20 was added.
 - | – “Changes for user-defined external functions that call a service program” on page 31 was updated.
 - | – “Badly formed data can trigger data mapping error” on page 28 was added.
 - | – “ILE CL Call Procedure (CALLPRC) advice for call to RPG now fails” on page 72 was added.
- | • **March 2009 update:**
 - | – “System i Navigator removal of support for memory-resident database monitors next release” on page 77 was added.
- | • **January 2009 update:**
 - | – “Server firmware requirements for V6R1” on page 2 was updated.
 - | – “Changes for user-defined external functions that call a service program” on page 31 was added.
 - | – “Special situations for program conversion” on page 49 was updated to add program examples.
- | • **November 2008 update:**
 - | – “Resolving an SQL0338 failure” on page 24 was updated.
 - | – “Convert SQL C++ Source (CVTSQLCPP) command” on page 15 was added.
 - | – “Change in cursor positioning after a ROLLBACK HOLD statement for cursors” on page 20 was added.
- | • **July 2008 update:**
 - | – “IBM WebSphere Application Server V6.0 (5733-W60) and IBM WebSphere Application Server Version 6.1 (5733-W61)” on page 70 was updated.
- | • **May 2008 update:**
 - | – “SQLCODE and SQLSTATE changes” on page 26 was updated.
 - | – “Plug-ins and LoadModule directives” on page 66 was updated.
 - | – “System i Navigator wireless support” on page 76 was updated.
 - | – “Host print transform changes” on page 35 was added.
- | • **March 2008 update:**
 - | – “List Objects Secured by Authorization List (QSYLATLO) API returns path names using Unicode CCSID” on page 10 was added.
 - | – “Pre-V6R1 PTFs for restoring modules, programs, or service programs” on page 51 was added.
 - | – “System i Navigator wireless support” on page 76 was added.
- | • **February 2008 update:**
 - | – “Database Create File program changes” on page 31 was added.
 - | – “File-level backup through QNTC file systems” on page 34 was added.
 - | – “FTP changes” on page 34 was added.
 - | – “Host Ethernet Adapter changes” on page 35 was added.
 - | – “Journaling enhancements” on page 40 was updated.
 - | – “Starting and ending TCP/IP” on page 54 was updated.
 - | – Other minor technical corrections have been made.

Installing V6R1 over V5R3

If you are installing i5/OS V6R1 over V5R3, you should also read the *iSeries Memorandum to Users* for V5R4. This document contains incompatibility-related information about the new functions and enhancements incorporated into V5R4. You can order this document by typing the following command:

```
SNDPTFORD SF98016
```

The Memo to Users is also available in the PSP information at the Support for IBM System i Web site:
<http://www.ibm.com/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 System i announcement material. To get the most current information on discontinued products or features along with suggested replacements, go to the System i Planning Web site:

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

Select the **Migration & upgrades** link for more information.

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 *Memorandum to Users* using the SNDPTFORD command, you can also view these documents at the Support for IBM System i Web site:

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

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

Prerequisite and related information

Use the i5/OS Information Center as your starting point for looking up System i technical information.

You can access the Information Center two ways:

- From the following Web site:
<http://www.ibm.com/systems/i/infocenter/>
- From CD-ROMs that ship with your operating system order:
i5/OS Information Center CD, SK3T-4091.

The i5/OS Information Center contains advisors and important topics such as Java, TCP/IP, Web serving, secured networks, logical partitions, high availability, control language (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 *System i Access for Windows DVD*, SK3T-4098. This DVD
 - | contains the IBM System i Access for Windows licensed program. System i Access Family offers
 - | client/server capabilities for connecting personal computers to System i computers.
-

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 i5/OS 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 i5/OS Information Center:
RCHINFOC@us.ibm.com

Be sure to include the following:

- The name of the book or i5/OS 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 i5/OS V6R1

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 *i5/OS Information Center CD*, SK3T-4091, CD-ROM or on the Internet at the following Web site:

<http://www.ibm.com/systems/i/infocenter/>

Note: After the release of the i5/OS Information Center, updates will be available in the English Internet version. To review these updates, click **Updates since V6R1 announcement** on the Information Center home page.

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 **Installing, upgrading, or deleting i5/OS and related software** topic in the i5/OS Information Center contains software preinstallation information and information about 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 preventive service planning (PSP) Information provides information about software problems you might encounter as you install the new release. The PSP number for the online V6R1 i5/OS Memo to users is SF98026. You can download PSP by using electronic customer support from the Support for IBM System i Web site:

<http://www.ibm.com/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 V6R1 is SF98020. Information within the PSP describes late-breaking installation information, and is grouped by product area. To receive this PSP information by using electronic customer support, type the following command on an i5/OS command line:

```
SNDPTFORD SF98020
```

- The PSP identifier for information that pertains to problems discovered since the availability of the current cumulative PTF package is SF99610. Information within this PSP describes all PTFs released since the current cumulative PTF package began shipping. It also contains information about 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 i5/OS command line:

```
SNDPTFORD SF99610
```

- The PSP identifier for information that pertains to installing V6R1 hardware is MF98610. Review this PSP information before you install either new System i models or hardware devices. To receive this PSP information by using electronic customer support, use this command:

```
SNDPTFORD MF98610
```

- The PSP identifier for information that pertains to system upgrades and data migrations is SF98169. Information within this PSP describes upgrade and migration corrections. Review this PSP information before you upgrade your system model or migrate your data between systems. To receive this PSP information by using electronic customer support, type the following command on an i5/OS command line:

- The System i PTF maintenance strategy. A PTF maintenance strategy is recommended for all System i customers. This might reduce the impact to System i operations that result from unplanned outages or program failures. For more information about the System i maintenance strategy, do the following:
 1. Go to: <http://www.ibm.com/systems/support/i>.
 2. Under the Popular links category, click **Fixes**.
 3. Click **Guide to fixes**.
 4. Click the **Server maintenance** tab.
 5. Click **Create a maintenance strategy**.

Server firmware requirements for V6R1

Before you upgrade to the next release, verify the server firmware level that is supported. Your server firmware must be at a certain minimal level to be able to support V6R1. (This does not include System i 8xx models that support V6R1.)

For information about viewing existing firmware levels, go to the Viewing existing firmware levels topic in the IBM Systems Hardware Information Center. Click **Customer service, support, and troubleshooting > Updates > Managing updates > Viewing existing firmware levels**.

- | For information about verifying firmware requirements for the new release, do either of the following:
 - | • Go to the Verifying firmware requirements for the new release topic in the i5/OS Information Center. Click **i5/OS and related software > Installing, upgrading, or deleting i5/OS and related software > Upgrading or replacing i5/OS and related software > Preparing to upgrade or replace software > Preparing the system for i5/OS[®] software upgrade or replacement > Verifying firmware requirements for the new release**.
 - | • Go directly to the System i Prerequisite site:
 - | <http://www.ibm.com/systems/support/i/planning/migrationupgrade.html>
 - | On this site, select the IBM prerequisite under the Hardware upgrade information heading, and then select the Software tab and specify product 5761-SS1 to find the data.
- | In previous releases, fixes for the server firmware were delivered as PTFs for the Licensed Internal Code. In V5R4M5 and later, the fixes for the server firmware are delivered as PTFs for an i5/OS[®] licensed program. The product identifier that is used depends on the system model that you are running; for example, 5733-905 for POWER5 models. You can determine the product identifier that is used for the server firmware fixes by using the Display Firmware Status (DSPFMWSTS) command.
- | **Note:** The firmware of the tape drives and controllers does not need to be updated for V6R1.

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 V6R1 (SF99610) 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.

If an extended period of 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 Send Program Temporary Fix Order (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 IBM System i Web site:

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

Click **Technical Databases > Preventive Service Planning - PSP**.

Refer to the software installation preventive service planning (PSP) document, SF98020, 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 IBM 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.

Program conversion

For V6R1, conversion is required for all programs created under previous releases that use the i5/OS machine interface (MI). This conversion upgrades and refreshes programs to provide enhanced system integrity, improved performance and new function, by taking advantage of many new operating system and processor capabilities. For a program to be converted, its creation data must be available. Programs created for V5R1 or later releases have creation data available for use during conversion. Creation data is always available for use during conversion when a program and its associated modules were created for V5R1 or later.

Preparing for these conversions of programs in libraries can have a significant effect on the amount of planning time before the upgrade. After the upgrade, some conversions might have a significant performance impact when program objects are not converted until accessing the program for the first time.

Other types of conversions also will occur for V6R1, including Java programs in directories, spooled files, and integrated file system names in file systems that are not case-sensitive.

IBM provides the Analyze Object Conversion (ANZOBJCVN) command for i5/OS V5R4 and V5R3 to help you plan these conversions. This tool helps you identify potential conversion difficulties, if any, and estimates the conversion times.

To begin preparing your systems for V6R1 conversions listed above, follow these directions:

1. Review Information APAR II14306 for the commands and function included with the conversion PTFs. This information APAR is in the preventive service planning information about the Support for IBM System i Web site. Click **Technical Databases > Authorized Problem Analysis Reports (APARs)**.
2. Read the IBM Redpaper publication *i5/OS Program Conversion: Getting ready for i5/OS V6R1*, REDP-4293, which you can obtain from the following Web site:

<http://www.redbooks.ibm.com>

The Redpaper also includes the PTF numbers to use for loading the ANZOBJCVN command.

- | For more information about program conversion, also see "Program conversion" on page 48, the **Installing, upgrading, or deleting i5/OS and related software** information center topic, and the System i Upgrade planning Web site:

<http://www.ibm.com/systems/support/i/planning/upgrade/index.html>

For more information about integrated file system name conversions, see "Integrated file system names" on page 35.

Optical media for V6R1

The optical media for the i5/OS V6R1 operating system is DVD. If you use physical media to install your V6R1 release, you must have a DVD device. Previous typical releases required approximately 13 CD-ROMs. The V6R1 release is available on three optical media. Some individual licensed programs are on CD-ROM. For a description of the V6R1 media, see **Media labels and their contents**.

Ensure that the cluster version is current before installing V6R1

If you are using clusters or high availability, ensure that the cluster version is current before installing i5/OS V6R1 on your system. Clustering supports only one cluster version difference. A V5R4 cluster must have a cluster version of 5 before you install V6R1.

Plan for your Operations Console installation or upgrade

If you are upgrading to V6R1 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 > Installing, upgrading, or deleting i5/OS and related software > Upgrading or replacing i5/OS and related software**.

Prerequisite information for Operations Console users upgrading to, or installing, V6R1:

You must comply with the following before upgrading or installing your software (operating system, Licensed Internal Code) to V6R1:

1. For all upgrades and installations, you need to establish a connection between the system 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 system. 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 system, 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.
2. It is recommended that you update System i Access for Windows to V6R1 before you upgrade the operating system. For more information, see the **Installing System i Access for Windows** topic in the i5/OS Information Center.

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

Important: During a manual IPL of the system, and if no console has been specified before, you will receive two additional screens to confirm the setting of the console type. 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 type 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 might happen on your first manual IPL of V6R1; 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 system model upgrade

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

New models

If you are using a System i 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 **Managing consoles, interfaces, and terminals > Changing consoles, interfaces, and terminals** topic in the **IBM Systems Hardware Information Center**.

Embedded Ethernet ports on POWER5 models 5xx (except 595) will now be the default port to be used for Operations Console (LAN) on systems that are not HMC managed. The System i 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 recovery and console can 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 recovery and console can 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 i5/OS Information Center.

Information APARs for software installation

Information APAR II14310

As you plan for your i5/OS installation or upgrade, ensure that you review Information APAR II14310 for V6R1 in the preventive service planning information.

If you are upgrading from V5R3, for example, you will need to load and apply certain PTFs to be able to accept online software agreements. If you are using image catalogs to upgrade from either V5R3 or V5R4, certain PTFs are required. Both of these preparation steps are included in the Installing, upgrading, or deleting i5/OS and related software information center topic.

To review Information APAR II14310, go to the Support for IBM System i Web site:

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

Click **Technical Databases > Authorized Program Analysis Reports (APARs)**.

Upgrade planning

The IBM System i Support: Planning (<http://www.ibm.com/systems/support/i/planning/>) Web site provides links to a variety of planning tools and information. The page on **Upgrade planning - i5/OS and OS/400® level mapping** shows which releases of i5/OS or OS/400 are supported by each of the System i, iSeries, and AS/400® RISC models.

The IBM System i Support: Planning - Migration and upgrades Web site (<http://www.ibm.com/systems/support/i/planning/migrationupgrade.html>) includes product previews, statements of directions, and information about products or features that might not be supported in future releases of the i5/OS operating system or on future System i models.

Prerequisite information

For information about the prerequisite information for features that you currently have or plan to add to your system, see the IBM Systems Web site (<http://www.ibm.com/systems/>). Search for **IBM Prerequisite**.

Support for I/O expansion units

You can obtain the latest information about software prerequisites for new I/O features and new I/O expansion units that are available at the Support for IBM System i Web site:

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

Setting up Ethernet connectivity from i5/OS logical partitions to the Hardware Management Console

Follow the installation steps for setting up Ethernet connectivity from i5/OS logical partitions (LPARs) to the Hardware Management Console (HMC) for HMC-managed systems. You can set up Ethernet connectivity in the following ways:

- Use an Ethernet adapter that the LPAR owns.
- Use the virtual LPAR-to-LPAR LAN that the Hypervisor provides. The virtual LAN must have a route to a real customer LAN using a bridge or routing partition that owns a LAN adapter.

| You must set up Ethernet connectivity between the HMC and i5/OS on POWER6® processor-based
| systems. If i5/OS forwards the problem report to HMC over the Ethernet connection, the HMC service
| function can work on the hardware problem. HMC must be aware of the problem before it can help fix
| the problem.

For more information, click **Managing the Hardware Management Console (HMC) > Installing the HMC** in the IBM Systems Hardware Information Center.

Key messages in this document

Each of the topics in this document provides important compatibility information that you should become familiar with before you install the new release. Some of these topics have a larger impact on more customers than others, and you should review these topics carefully. In V6R1, the following topics are particularly important:

- “ALCOBJ for data queues” on page 8
- “Saving and restoring journal receivers to a previous release” on page 11
- “Copy from Import File (CPYFRMIMPF) and Copy To Import File (CPYTOIMPF) commands” on page 16
- “IBM commands no longer exist in secondary language libraries” on page 18
- “Install Windows Server (INSWNTSVR) command changes” on page 60
- “Domain name system (DNS) changes” on page 32
- “Integrated file system names” on page 35
- “IPv6 changes” on page 36
- “Journaling enhancements” on page 40
- “Program conversion” on page 48
- “Spooled file changes” on page 54
- “System i Access for Windows (5761-XE1)” on page 73

Chapter 2. i5/OS operating system

This chapter describes changes to the i5/OS 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) might be affected by changes to IBM-supplied system output files in this release. IBM commands and APIs that generate database output files add new fields to the end of record formats for additional information returned each release. Adding new fields to the record format, even to the end of the record, changes 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, review the application to determine which system file it uses. New fields have been added to IBM-supplied database files in each release of i5/OS and System i.

Security audit record changes

Changes made to security auditing for this release might affect applications that read those audit records. Actions that were not audited in previous releases might now be audited. Existing audit records might 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 might 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 i5/OS functions that use IBM-supplied control language (CL) commands that are not library-qualified in this release might be changed in a future release to specify a specific library, *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) command that allow your exit program to get control and possibly change the command that is used.

Changes to system printer files and other IBM-supplied objects

The MAXRCDS parameter for QSYSPRT and QPSAVOBJ printer device files has always been 100 000. During an upgrade, the default value has not changed for system printer files. Customization of IBM-supplied printer files is lost on a release upgrade. To preserve those changes, you need to rerun your changes to printer system files each release.

Changes to many types of IBM-supplied objects are lost during the upgrade, because the copy of the object in the IBM product library is replaced by the new copy of the object.

Adopt authority for Java programs requires new PRPQ

This only applies to Java applications that are run using the Classic Java Virtual Machine. The IBM Technology for Java Virtual Machine does not support adopting authority from Java programs.

By default, Java adopt authority will be disabled in V6R1. Java applications which have native methods that depend on adopting authority from Java programs will no longer work in V6R1. If you have Java applications with native methods which rely on adopting authority from Java programs, there are two options:

1. Change the application so that it no longer depends on adopting authority from Java programs. This is the option that IBM strongly encourages customers to choose. Documentation on the preferred methods for changing your Java applications is available in the Java topic collection in the Programming category in the i5/OS Information Center at <http://www.ibm.com/systems/i5infocenter>.
2. You can order PRPQ 5799-AAJ that enables Java adopt authority in V6R1. IBM does not recommend this option because the PRPQ will be withdrawn in a future release. If PRPQ 5799-AAJ is installed, Java adopt authority will be enabled only for Java applications that are run with the Classic Java Virtual Machine. The IBM Technology for Java Virtual Machine does not support adopting authority from Java programs.

If you are unsure whether any of your Java applications rely on adopting authority from Java programs, you can look in the i5/OS Information Center for more information about the resources available to assist with the transition away from adopting authority from Java programs.

For Java programs created on releases before V6R1 using the Create Java Program (CRTJVAPGM) command and specifying USRPRF(*OWNER), the programs can be saved and restored to V6R1. However, the adopted authority attributes of the program will not be used unless PRPQ 5799-AAJ is installed. If the PRPQ is installed, the adopt authority attributes in the Java program will be used. Java programs created on V6R1 with a target release of V5R3 or V5R4 will have the adopt attribute stored in the Java program. When that Java program is saved on a V6R1 or later system and restored to a V5R3 or V5R4 system, those adopt attributes will be used as they are today on those releases.

ALCOBJ for data queues

For V6R1, you can change a data queue so that data queue operations can enforce the lock obtained with the Allocate Object (ALCOBJ) CL command. In previous releases, a lock with ALCOBJ was ignored by data queue operations.

- If a data queue is created and changed with the Change Data Queue (QMHQCDQ) API to enforce locks, it is not permitted to be saved to a target release previous to V6R1.
- Some applications might have used the Allocate Object (ALCOBJ) command in the past but some programs using the data queues from those same applications might have been using the queues without specifying the ALCOBJ command. If the queue is changed to use the new lockable attribute, the program that was not using ALCOBJ, and even those using ALCOBJ, can now see locking errors that they did not see in the past.

Potential deadlock and new failure when using the new ALCOBJ function:

A permanent deadlock might occur when using the Change Data Queue (QMHQCDQ) API to change data queue attributes to enforce data queue locking. If the enforce data queue locking option is turned on, and a thread gets an exclusive thread-scoped lock on the queue with the ALCOBJ CL command, and then it performs a receive (or dequeue) with a wait, and no messages are available that satisfy the receive (or dequeue), the thread enters a permanent deadlock. This situation occurs because no other thread or job can get in to perform a send (or enqueue) because the receive (or dequeue) thread that is waiting is holding the thread-scoped exclusive lock (that was obtained using ALCOBJ). To recover from the deadlock, cancel the job.

A new failure might occur when using the Change Data Queue (QMHQCDQ) API to change the data queue attributes to enforce data queue locking. This situation can occur when the following is true in thread 1:

- The Enforce data queue locking option is turned on,
- the data queue is journaled, and
- thread 1 performs a send (or enqueue) operation that results in an extend of the data queue (and an internal seize during the extend operation).

Meanwhile, thread 2 tries to perform a receive (or dequeue) operation.

With this scenario, thread 2 will fail because the data queue lock will conflict with the seize held by thread 1. This failure is a correctable problem. Thread 2 would need to try again the receive (or dequeue) which will no longer conflict with thread 1 when the extend operation completes.

API changes

API binding directory updates

The service programs, which export APIs associated with C and C++ runtime applications, have been added to the system binding directory QUSAPIBD. This might result in the diagnostic message CPD5D03 (definition supplied multiple times) and a binding failure if you have developed your own procedure exports with the same name as these APIs. If this error is encountered, you should rename your exported procedures. As a temporary circumvention, you can also remove the service program, which exports the duplicate name, from the QUSAPIBD binding directory using the Remove Binding Directory Entry command (RMVBNDDIRE) or Work with Binding Directory (WRKBNDDIR) command.

Use the Display Service Program (DSPSRVPGM) command with `DETAIL(*PROCEXP *DTAEXP)` to see the list of exports for that service program. You might get message CPD5D03 for Data Exports and Procedure Exports. The following are the new service programs:

- QC2SYS
- QC2IO
- QC2NIO
- QC2UTIL1
- QC2UTIL2
- QC2UTIL3
- QC2IFS
- QC2NIFS
- QC2POSIX
- QC2NPOSX
- QC2ULOCL
- QC2VLOCC
- QC2VLOCL
- QC2ASYSIG
- QC2TS1
- QC2TSI
- QC2NTSI
- QC2TSU
- QC2TSV
- QC2TOOLS
- QC2TSO
- QC2NTSO
- QC2TSP
- QC2NTSP

Changes to flight recorder APIs (QWTDMPFR and QWTDMPFLF)

To use the Dump Flight Recorder (QWTDMPFR) and Dump Lock Flight Recorder (QWTDMPFLF) APIs, a caller is now required to have service (*SERVICE) special authority. In addition, the default public authority for the QWTDMPFLF API has been changed from *EXCLUDE to *USE.

Create User Space (QUSCRTUS) API

In i5/OS V6R1, the default value for the Optimum Alignment parameter of the Create User Space (QUSCRTUS) API has changed from 0 to 1, so optimum alignment has become the default. This change reduces the maximum amount of space available for user spaces created with the default value. You can programmatically retrieve the maximum size of an optimally-aligned space using option Hex 0003 of the Materialize Machine Data (MATMDATA) MI instruction.

Recommendation: Because the performance benefits of using optimally-aligned spaces might vary over time, IBM recommends creating a user space with optimum alignment, even if an application does not experience a performance benefit from using an optimally-aligned user space on one particular release.

Get File Information - stat() API

The `stat()` API, when issued against an object in the QfileSvr.400 file system now returns the primary group ID on the remote system for the object. Because the QFileSvr.400 file system is used to access objects on a remote system, the primary group ID might not exist on the system where the `stat()` API is being called. Additionally, the primary group ID might exist, but might not reference the same user profile as on the remote system. Previously, this value was always set to zero in the information returned from the `stat()` API.

List Objects Secured by Authorization List (QSYLATLO) API returns path names using Unicode CCSID

The List Objects Secured by Authorization List (QSYLATLO) API now returns path names in a Unicode CCSID instead of the job CCSID for formats that support path names. The CCSID of the path name is returned as part of the path name structure. Use this CCSID value to determine the CCSID of the returned path. You can use the Code Conversion (`iconv`) API to convert the Unicode path name back to the job CCSID.

Perform Shared Memory Control Operations (shmctl and shmctl64) APIs

The `shmctl()` API might return a new error number, `E_OVERFLOW`, when the `IPC_STAT` command is performed. This occurs when the shared memory segment size cannot be represented correctly in the `shm_segsz` field of the `shmid_ds` structure. In particular, sizes larger than 4 294 967 295 bytes (4 GB - 1) cannot be represented correctly. To obtain the size in such cases, you must use the new `shmctl64()` API.

Retrieve Command Information (QCDRCMDI) API not returning help bookshelf

The Retrieve Command Information (QCDRCMDI) API will no longer return any help bookshelf information. The offset to the help bookshelf information and length of bookshelf information in return formats CMDI0100 and CMDI0200 will always be zero.

Serial number format changes for hardware resource APIs and commands

API changes:

To account for the longer serial number, an Extended Serial Number field has been added to the end of the data structure for the following APIs:

- Retrieve Hardware Resource Information (QGYRHRI, QgyRtvHdwRscInfo)
- Retrieve Resource Information (QRZRRSI)
- Search Hardware Resource Entry (QRZSCHE)

When running on hardware that supports the new format, the old serial number field is set to an extracted version of the 11S-format serial number for the APIs listed above. In all cases, the Extended Serial Number field contains either the older 10-character serial number or the 11S-format serial number padded to the right with blanks. It is recommended that users of these APIs use the new Extended Serial Number field on the current API or *OUTFILE interfaces.

Command changes:

The following commands, for both display and print, have also been changed to use the longer serial number format:

- Display Hardware Resource (DSPHDWRSC) command (display, *PRINT, and *OUTFILE outputs)
- Work with Hardware Resources (WRKHDWRSC) command
- Work with Hardware Products (WRKHDWPRD) command

For the commands listed above, the serial number field contains the 11S-format serial number, when running on hardware that supports the new format. For hardware that does not support the new format, the older 10-character serial number format is used.

Backup and recovery changes

Omitting objects on the SAV or RST commands

In the previous release, if you specified to omit objects using the OBJ parameter on the SAV or RST commands, objects from any unmounted user-defined file systems (UDFSs) would *not* have been omitted.

Likewise, if you had previously specified to omit objects using the PATTERN parameter on the SAV or RST commands, objects from any unmounted UDFSs that matched the PATTERN would not have been omitted.

In V6R1, these objects will now be omitted from the save or restore operation

Saving and restoring journal receivers to a previous release

In V6R1, to save and restore journal receivers in a network between a V6R1 system and systems that are running earlier releases, you must apply the following PTFs to the earlier release systems:

- SI27618 for V5R3
- SI27617 for V5R4

You must also apply these PTFs if you intend to use remote journals between a V6R1 system and a system that is running an earlier release.

Changes to save and restore limits

You cannot use save files with *NOMAX specified for the MAXRCDS parameter to save to a previous release if the number of records in the save file is larger than 2 146 762 800.

The QSRLSAVF API can return a value larger than 2 146 762 800 in the records field in SAVF0100 format. This field has been changed from an INT to a UNSIGNED field to accommodate the larger value. If you use this API, you should recompile to use an UNSIGNED field.

Virtual tape changes

If you run the Save System (SAVSYS) command using a virtual tape volume that is too small to install the system, inquiry message CPA370B is sent. For unattended backup operations, you should add a reply list entry (ADDRPYLE command) to automatically reply to this message.

GO SAVE, GO RESTORE, GO LICPGM changes

The GO SAVE (options 21-23, 40), GO RESTORE (options 21-23, 40), and GO LICPGM (option 40) menus were recently enhanced to end TCP before ending the subsystems. Commands that are issued as part of these options are:

- End TCP/IP Server (ENDTCPSVR)
- End Host Server (ENDHOSTSVR)
- End TCP/IP (ENDTCP)

To allow time for these commands to process, two Delay Job (DLYJOB) commands are issued. This adds a total of 10 minutes delay before the End Subsystem (ENDSBS) command is issued.

Bidirectional data changes

The system bidirectional layout services has been updated to an enhanced version. This update might cause different conversion results if bidirectional data is being processed.

C and C++ changes

C language runtime changes

The following C language runtime changes conform better with ANSI and POSIX specifications.

The output from the **printf()** family of routines has changed when a floating-point value of INFINITY or Not-A-Number (NaN) results. In earlier releases, HUGE_VAL and *NaN were used to represent these special values. In V6R1, INFINITY and NAN are used for these special values to match the ANSI and POSIX specifications.

The output from **strfmon()** and **wcsfmon()** has changed when a floating-point value of INFINITY results. In earlier releases, HUGE_VAL was used to represent this special value. In V6R1, INFINITY is used.

The output from the **printf()** family of routines has changed when a negative zero (floating-point value of zero with the sign bit set) results. In earlier releases, the sign was ignored with a value of zero and 0 was output. In V6R1, the output is -0 to match the ANSI and POSIX specifications.

The **printf()** and **scanf()** family of routines no longer accepts and ignores F and N as size modifiers. In earlier releases, the F and N characters were allowed as size modifiers and were ignored. In V6R1, these modifiers are no longer allowed.

The **printf()** and **scanf()** family of routines no longer accepts and ignores a size modifier of ll for non-integer conversion specifiers. In earlier releases, the ll size modifier was accepted and ignored for non-integer conversion specifiers. In V6R1, the ll size modifier is only allowed for d, i, o, u, x, X, and n conversion specifiers.

New type definitions and macros have been added to the <math.h> include file. The type definitions **float_t** and **double_t** and the macro definitions INFINITY, NAN, and HUGE_VALL have been added to match ANSI and POSIX specifications. Application code which contains definitions of the above items might require changes to remove the local definitions and use the standard definitions in the <math.h> include file.

New macros have been added to the <limits.h> include file. The macro definitions LONGLONG_MIN, LONGLONG_MAX, and ULLONGLONG_MAX have been added to match ANSI and POSIX specifications. Application code which contains definitions of the above items might require changes to remove the local definitions and use the standard definitions in the <limits.h> include file.

New 64-bit time type definitions and routines

New 64-bit time type definitions and 64-bit time routines have been added to the <time.h> and <sys/types.h> include files. The type definition **time64_t** and routines **ctime64()**, **ctime64_r()**, **difftime64()**, **gmtime64()**, **gmtime64_r()**, **localtime64()**, **localtime64_r()**, **mktime64()**, and **time64()** have been added to allow the C runtime to support dates past 2038. Application code which contains definitions of the above items might require changes to remove the local definitions and use the standard definitions within these include files.

Changes affecting static C++ objects

Applications that contain static C++ objects or bind to service programs that contain static C++ objects might experience different behavior in V6R1 if the constructor for a static C++ object calls the C runtime **exit()** function, or if it fails in an abnormal manner. Possible sources of abnormal failure include a call to the C Runtime **abort()** function, an unhandled exception message, an unmonitored signal, or a thrown C++ object with no corresponding catch clause.

The behavior has changed for the following two situations:

- The activation group, in which a constructor is run, matches the activation group of the **main()** function of the application and the constructor calls the C Runtime **exit()** function. In earlier releases, the application silently exited without running any portion of the **main()** function of the application and with no messages in the job log. In V6R1 and later releases, the application receives message MCH3203 f/AiEagerActivator with an associated vlog entry.
It is recommended that you do not call the C runtime **exit()** function from the constructor of a static C++ object.
- The activation group, in which a constructor is run, does not match the activation group of the **main()** function of the application and the constructor fails in an abnormal manner. In earlier releases, all constructors attempted to run. The activation group associated with constructors that had failed abnormally were destroyed, as well as the activation group associated with the **main()** function of the application. In V6R1 and later releases, as soon as one constructor fails in an abnormal manner, no more constructors attempt to run, and all activation groups created as a result of the call to the **main()** function of the application are destroyed. In V6R1 and later releases, the activation group for the **main()** function of the application is not destroyed if the activation group already exists in the job.

Clustering changes

Installation requirements for the cluster version

For more information about the installation requirements for cluster versions, see “Ensure that the cluster version is current before installing V6R1” on page 4.

Change Cluster Resource Group (QcstChangeClusterResourceGroup) API

A new restriction has been placed on the existing Change Cluster Resource Group (QcstChangeClusterResourceGroup) API. A cluster administrative domain cluster resource group (CRG) cannot be changed with this API. An attempt to change a cluster administrative domain CRG will result in a CPFBB00 message. Use the QcstChangeClusterAdminDomain API to change a cluster administrative domain instead.

Cluster command changes

The Change Cluster Configuration (CHGCLUCFG) command in the i5/OS operating system was removed. The replacement command in the 5761-HAS LP is Change Cluster (CHGCLU), which includes some additional functions.

The following cluster commands have moved from QSYS to the new IBM System i High Availability Solutions Manager (5761-HAS) product. The V5R4 source level of these commands exists in the QUSRTOOL library.

Cluster commands:

- Add Cluster Node Entry (ADDCLUNODE)
- Change Cluster (CHGCLU) (renamed from CHGCLUCFG)
- Change Cluster Node Entry (CHGCLUNODE)
- Change Cluster Version (CHGCLUVER)
- Create Cluster (CRTCLU)
- Delete Cluster (DLTCLU)
- Display Cluster Information (DSPCLUINF)
- End Cluster Node (ENDCLUNOD)
- Remove Cluster Node Entry (RMVCLUNODE)
- Start Cluster Node (STRCLUNOD)
- Work with Cluster (WRKCLU)

Cluster resource group (CRG) commands:

- Add Cluster Resource Group Node Entry (ADDCRGNODE)
- Change Cluster Resource Group (CHGCRG)
- Change Cluster Resource Group Primary (CHGCRGPRI)
- Create Cluster Resource Group (CRTCRG)
- Delete Cluster Resource Group from Cluster (DLTCRGCLU)
- Display Cluster Resource Group Information (DSPCRGINF)
- End Cluster Resource Group (ENDCRG)
- Remove Cluster Resource Group Node Entry (RMVCRGNODE)
- Start Cluster Resource Group (STRCRG)

Device domain entry commands:

- Add Device Domain Entry (ADDDEVDMNE)
- Remove Device Domain Entry (RMVDEVDMNE)

Cluster resource group device entry commands:

- Add Cluster Resource Group Device Entry (ADDCRGDEVE)
- Change Cluster Resource Group Device Entry (CHGCRGDEVE)
- Remove Cluster Resource Group Device Entry (RMVCRGDEVE)

Cluster administrative domain commands:

- Create Cluster Admin Domain (CRTCAD) (renamed from CRTADMDMN)
- Delete Cluster Admin Domain (DLTCAD) (renamed from DLTADMDMN)

Clusters GUI changes

In V6R1, all the new functions added to the Clusters GUI are delivered through the 5761-HAS LP, which runs on IBM Systems Director Navigator for i5/OS. The name is the Cluster Resource Services GUI. The V5R4 Clusters GUI is still available in System i Navigator, but “as-is” at the V5R4 level.

Clustering jobs

Before V6R1, cluster (QCSTLCTL, QCSTCRGM, and *crg-name*) jobs were in the QSYSWRK subsystem. In V6R1, cluster jobs are system jobs. When using the Work with Active Jobs (WRKACTJOB) command, find the list of system jobs rather than the list of jobs in the QSYSWRK subsystem.

Device cluster resource group changes

Changes to switchover:

A minor change to the device Cluster Resource Group (CRG) switchover behavior simplifies user actions if a failure occurs during a vary on operation for a configuration object. If all the vary on operations are successful, the switchover behavior is still the same. Most users will benefit from the change and require no additional action. You can still obtain the old behavior with a programming change.

On a switchover, if a vary on operation on the new primary node fails, then with the old behavior, a switch back to the original primary node occurs. With the new behavior, a switch back to the original primary does not occur. Instead, a new exit program action code dependent data value of VaryFailed is passed into the exit program indicating that any vary on operation failed. Additionally, the device CRG is ended.

To preserve the old behavior, the exit program should return Failure if the exit program action code dependent data is VaryFailed. This causes a switch back to the old primary node.

Support for additional device types:

Beginning in V6R1, a device CRG supports devices other than independent auxiliary storage pool (IASP) devices, including the following devices:

- Tape
- Optical
- Communication adapters
- IXS servers

In V5R4 and earlier releases, a device CRG supports only independent ASP devices.

i5/OS command changes

Change Network Server Desc (CHGNWSD) command

Change Network Server Desc (CHGNWSD) command has parameter changes that might require you to recompile any existing CL programs using these parameters.

- The Shutdown TCP port (SHUTDPORT) is an obsolete parameter that was never used. It has been removed from the command.
- The *MLTPHGRP special value has been removed from the Removable media path (RMVMEDPTH) parameter.

Convert SQL C++ Source (CVTSQLCPP) command

| The CVTSQLCPP command will no longer be supported in future releases. The CVTSQLCPP command is
| part of Licensed Program DB2 Query Mgr and SQL DevKit (57xxST1). It is suggested to use the Create
| SQL ILE C++ Object (CRTSQLCPPI) command, which is also included as part of Licensed Program
| 57xxST1.

Copy from Import File (CPYFRMIMPF) and Copy To Import File (CPYTOIMPF) commands

CPYTOIMPF now implements the new parameter, Stream File Authority (STMFAUT). This parameter enhances the authority settings that this command applies when you use it to export database files to a newly created stream file. This parameter is ignored if the object already exists, as the command does not change the authorities of an existing object even if MBROPT(*REPLACE) is specified; only the data is replaced, not the object.

This parameter supports four options:

***DFT** The owner of the stream file is granted *RWX data authority to the stream file. The primary group and *PUBLIC have *NONE data authority to the stream file. Object authorities are based on the object authorities for the directory where the stream file is to be created.

***INDIR**

The authority settings of the output stream file are based on the authority for the directory where the stream file is to be created. The stream file is assigned the same public authority, private authorities, primary group, primary group authority, authorization list, and auditing value as the directory in which it is created. If the destination file system does not support the *INDIR special value, the command will fail.

***FILE** The authority settings of the output stream file are based on the authority for the object specified on the From file (FROMFILE) parameter. The stream file is assigned with the same public authority, private authorities, primary group, primary group authority, authorization list, and auditing value as the from-file object being copied. If the destination file system does not support one or more of these values, the unsupported values will be ignored.

If the from-file object is a multisystem file, the default value *DFT is used instead of the *FILE value.

***INDIRFILE**

The resulting authority information is similar to that produced by copying and pasting a stream file using the System i Navigator graphical user interface. The authority information for the stream file is initially based on the directory where the stream file is created. Then, authority information from the object specified on the From file (FROMFILE) parameter is copied to the object. This action might replace some of the initial authority information obtained from the directory.

If the from-file object is a multisystem file, the default value *INDIR is used instead of the *INDIRFILE value.

The CPYFRMIMPF command has enhanced support for the Mixed PC-ASCII CCSID (Coded Character Set Identifier) data. Previously when the command tried to process stream files or physical files (PF) containing data encoded in Mixed PC-ASCII CCSID, it could not interpret the difference between Single Byte Character Set (SBCS) from the Double Byte Character Set (DBCS) in the data stream. To solve this problem, the code for Mixed PC-ASCII CCSID examines the data stream and identifies DBCS data from the SBCS data.

Previously, when you used the CPYFRMIMPF command, the message CPF2973 "Data from file &1 in &2 truncated to &6 characters" was sent when the record length of the to-file was not big enough to hold the from-file data. When the from-file is a stream file, this message did not display the name of the from-file neither the library name. Instead empty spaces displayed for each. In V6R1, the message CPIA083 with text "Stream file copied to object with truncated records" is displayed when a truncation operation occurs and the from-file is a stream file.

The CPYFRMIMPF command now displays the message CPF2845 reason code 13, where in previous releases, the expected message was CPF2845 reason code 98. Also, CPYFRMIMPF now displays the message CPF2846 reason code 1, where in previous releases, the expected message was CPF2845 reason code 99.

Copy To Stream File (CPYTOSTMF) and Copy From Stream File (CPYFRMSTMF) commands

The authority requirements have changed for the Copy To Stream File (CPYTOSTMF) and Copy From Stream File (CPYFRMSTMF) commands. Both commands now require read (*R) data authority to the conversion table, if one is specified. The CPYFRMSTMF command now requires write (*W) data authority to the target database file, and *ADD authority to the library if the specified member does not exist. The command is now more consistent with the authority requirements for the Copy Object (CPY) command. Users who have the minimum authority documented for releases earlier than V6R1 for the database file or conversion table will experience command failures. Such users will need more authority to these objects. You might need to change any applications that created the objects or set their authorities.

In addition, the CPYTOSTMF command might now fail when the specified database file member is in use by another process. In previous releases, it was possible that the data in the file could be updated by another process while that data was being copied by the CPYTOSTMF command. This behavior was also inconsistent with the behavior of the CPY command and could produce an output stream file containing incorrect data. In V6R1, the command will fail when the source member is locked with *EXCLRD or *SHRUPD locks. These locks cannot be acquired while the command is accessing any of the data.

These commands have also been enhanced to accept a coded character set identifier (CCSID) for the stream file, rather than a code page. This allows greater use of these commands in circumstances that require data conversion between dissimilar encoding schemes. However, the command changes might cause created stream files to be assigned different CCSIDs than in previous releases. This might affect later operations that assume or require the stream file to have a particular CCSID. Because of these changes, you should use the new Stream file CCSID (STMFCSSID) parameter on these commands rather than the Stream file code page (STMFCODPAG) parameter. The STMFCODPAG parameter might be removed in a future release.

CPYF and CPYSRCF will not update 'Last source update date/time'

In earlier releases (V5R3M0 SI25796 and V5R4M0 SI25844), the Copy File (CPYF) and Copy Source File (CPYSRCF) commands set the **Last source update date/time** for the target source member to the date and time the command was run.

Starting in V6R1, the **Last source update date/time** of the target source member will retain the value of the source member being copied, unless MBROPT(*ADD) was used. If you want the **Last source update date/time** for the target source member to be the date and time the copy command was run, you will need to use the CPYSRCF command and specify SRCCHGDATE(*NEW).

Create Network Server Desc (CRTNWSD) command changes

The Create Network Server Desc (CRTNWSD) command has parameter changes that might require you to recompile any existing CL programs using these parameters.

- The Shutdown TCP port (SHUTDPORT) is an obsolete parameter that was never used. It has been removed from the command.
- The *MLTPHGRP special value has been removed from the Removable media path (RMVMEDPTH) parameter.

Create NWS Configuration command parameter default changed

The default for the Enable unicast (ENBUNICAST) parameter of the Create NWS Configuration (CRTNWSCFG) command has changed to *YES to reflect the preferred method to discover and

communicate with the remote server's service processor. IBM Director discovery of the remote server is more reliable in most environments using this method.

Delete User-Defined FS command

The CPCA089 (Link removed) message is no longer sent by the Delete User-Defined FS (DLTUDFS) command when it removes a link for an object.

Display Program References (DSPPGMREF) command changes

File-level information is now returned for query definition (*QRYDFN) objects when you use the Display Program References (DSPPGMREF) command. Where-used Object Information Repository (OIR) data is now added to the *QRYDFN objects that are created in V6R1. Where-used OIR data keeps information about the files used in the *QRYDFN objects.

There is a performance impact to using the DSPPGMREF command because *QRYDFN object references are also returned when *ALL is specified for the PGM and OBJTYPE keywords.

If you create and save a query definition from the previous release, it works the same way when you restore it to V6R1. However, if you restore a V5R4 query definition (*QRYDFN) onto V6R1, you will not get the where-used OIR information unless you specifically convert the query definition in V6R1. The query definition is not automatically converted when the system is upgraded to V6R1. To convert the query definition, you will need to apply a special PTF. Contact Fix Central, Database category for the PTF number at the following Web site:

<http://www.ibm.com/eserver/support/fixes/>

The objects created and saved in V6R1 can be restored to the previous release and they should work the same way.

*DOC option removed from VFYSRVCFG command

The Information Center update is no longer supported by Universal Connection Manager to connect to IBM. Therefore, the *DOC option for the SERVICE parameter on the Verify Service Configuration (VFYSRVCFG) has been removed. Any CL programs that contain a VFYSRVCFG command with SERVICE(*DOC) will need to be changed.

Extended Serial Number field for hardware resource commands

An Extended Serial Number field has been added for the DSPHDWRSC, WRKHDWRSC, and WRKHDWPRD commands. For details, see "Serial number format changes for hardware resource APIs and commands" on page 10.

IBM commands no longer exist in secondary language libraries

In prior releases, CL command (*CMD) objects were part of the language product load for the operating system and other IBM licensed programs. This meant that *CMD objects were installed with other language-type objects, like message files, in each installed secondary language library. These libraries are named QSYS29 nn where nn is a two-digit number that identifies the national language version. Starting in V6R1, *CMD objects are part of the code product load for the operating system and for new or refreshed IBM licensed programs. The *CMD objects are installed with other code-type objects, like programs, in the product library. The QSYS29 xx libraries still exist; however they no longer contain IBM command (*CMD) objects for i5/OS and for new or refreshed licensed programs. If you have any CL programs that directly refer to commands in secondary-language libraries by a library that qualifies the command name with 'QSYS29 nn /' , those CL commands will fail because the command will not be found.

New records for the outfile for the Display Authorization List Objects command

Starting in V6R1, if OUTPUT(*OUTFILE) is specified on the Display Authorization List Objects (DSPAUTLOBJ) command, the generated output file contains records for directory objects and objects in directories that are secured using the specified authorization list. If you have programs that process DSPAUTLOBJ outfile records, you will need to change the code to handle the new type of records. A record for a directory or an object in a directory has blanks in the existing record format fields for object name (AONAME) and object library (AOLIB). You can find the new fields for directory objects at the end of record format QSYDALO of model outfile QADALO in library QSYS.

Print Performance Explorer (PRTPEXRPT) command positional parameter limit

The maximum number of positional parameters allowed by the Print PEX Report (PRTPEXRPT) command is three instead of *NOMAX (No maximum positional limit specified).

Start Database Monitor (STRDBMON) command

The Start Database Monitor (STRDBMON) command has been enhanced with the Filter Query Governor (FTRQRYGOVR) parameter. The FTRQRYGOVR parameter replaces the earlier use of the STRDBMON COMMENT parameter to specify a filter query governor value. You can use the COMMENT parameter to specify FTRQRYGOVR(*COND) or FTRQRYGOVR(*ALL) only if you do not specify the FTRQRYGOVR parameter, or if you specify *NONE for the FTRQRYGOVR parameter.

Start Object Conversion (STROBJCVN) command

The Priority (PTY) parameter has been removed from the Start Object Conversion (STROBJCVN) command.

Text description TEXT(*CMDPMT) retrieved dynamically

When creating a regular command with the Create Command (CRTCMD) command or a proxy command with the Create Proxy Command (CRTPRXCMD) command and *CMDPMT is specified, the default text description stored in the *CMD object will be *CMDPMT. When the text description is later retrieved or displayed for the command, the text is the same as it would appear for the command title if the command was prompted. In prior releases, *CMDPMT caused the text description to be a static copy of the command title prompt text. Changing the text description for an existing command by running Change Object Description (CHGOBJD) or Change Command (CHGCMD) or Change Proxy Command (CHGPRXCMD) and specifying TEXT(*CMDPMT) will cause the text description for the command to be retrieved dynamically.

Teraspace-enabled parameter changes for CHGPGM and CHGSRVPGM commands

All modules, Integrated Language Environment[®] (ILE) programs and service programs, and original program model (OPM) programs are automatically made teraspace-enabled on V6R1M0 and later releases. Any program running on V6R1 can process teraspace addresses. All of these objects created on releases earlier than V6R1 are converted before they are run on V6R1 or later releases of i5/OS. The information panels produced by the Display Module (DSPMOD), Display Program (DSPPGM), and Display Service Program (DPSRVPGM) commands will show that the object is teraspace-enabled unless the object was created on a release earlier than V6R1 and has not yet been converted.

Before V6R1, specifying a value other than *SAME or the current value for the TERASPACE parameter on the Change Program (CHGPGM) or Change Service Program (CHGSRVPGM) commands caused the program or service program to be re-created.

Starting in V6R1, for objects created for V6R1 or later releases, the TERASPACE parameter value on the CHGPGM and CHGSRVPGM commands is ignored and does not re-create the program or service program. If the program or service program has a target release (TGTRLS) value earlier than V6R1M0, specifying a value different than the current TERASPACE attribute re-creates the object as teraspace-enabled, and stores the specified value in the object template information. The value in the object template is used if the object is saved for a release earlier than V6R1.

***VRT option removed from the Change Device Desc (Optical) (CHGDEVOPT) command**

The *VRT option for the RSRCTYPE parameter on the Change Device Desc (Optical) (CHGDEVOPT) command has been removed. Any CL programs that contain a CHGDEVOPT command with RSRCTYPE(*VRT) need to be changed and recreated from the source.

Work with Network Server Status (WRKNWSSTS) command (change to SVRTYPE parameter)

A new special value, *ALL, has become the default value for the SVRTYPE (Server type) parameter of the Work with Network Server Status (WRKNWSSTS) command. SVRTYPE(*ALL) displays information for all types of network servers.

Submit Remote Command (SBMRMTCMD) command

- | SBMRMTCMD has changed in V6R1 to return all of the server-side messages for the executing command, with the exception of those messages issued by programs called on the server. In releases prior to V6R1, a maximum of ten messages were returned. Client applications will have more message information returned in V6R1 than in previous releases.

Conversion of authorization format

The current format limits the number of objects that can be put on an authorization list. The new format requires a conversion, which occurs on the first touch of the object. The old format of authorization list is no longer supported.

Space requirements for the new format, which includes a resilient index, is approximately six times the current space that is used by the object. The resilient index is more complex and requires more space than a flat array.

Database changes

Change in cursor positioning after a ROLLBACK HOLD statement for cursors

- | Beginning with V6R1, the cursor positioning after a ROLLBACK HOLD statement has changed for some UPDATE or DELETE WHERE CURRENT OF SQL statements. To observe the change in behavior, the UPDATE OR DELETE WHERE CURRENT OF SQL statements must include ALL of the following conditions:

- | • Contains an ORDER BY clause
- | • Does not contain a FOR UPDATE clause

- | With V6R1, the cursor position is reset the same as it was at the start of the unit of work. The cursor position was not reset before V6R1.

- | To retain the cursor positioning across ROLLBACK HOLD statements, complete the following steps:
 - | 1. Add a SAVEPOINT statement after the FETCH statement:

```
| SAVEPOINT SVPT_NAME ON ROLLBACK RETAIN CURSORS
| 2. Modify the ROLLBACK HOLD statement:
| ROLLBACK TO SAVEPOINT SVPT_NAME
```

| This results in a rollback operation to savepoint SVPT_NAME but retains the position of the cursor.

Changes to LEFT and RIGHT Scalar functions

LEFT and RIGHT SQL Scalar functions are now character based instead of byte based. The second argument now indicates the number of characters instead of the number of bytes. This change does not affect LEFT and RIGHT functions where the first argument is a single-byte CCSID (for example, 37 or 500). This change only affects the result of LEFT and RIGHT functions where the first argument is a mixed-byte CCSID, UTF-8, or UTF-16.

In the SELECT LEFT statement below, assume that FIRSTNAME is a VARCHAR(12) column, encoded in Unicode UTF-8, in T1. One of its values is the 6-character string Jürgen:

```
SELECT LEFT(FIRSTNAME, 2) FROM T1
```

Before V6R1, the above statement returns the value Jâ (x'4AC3') (because 2 means 2 bytes). In V6R1, the above statement returns the value Jü (x'4AC3BC') (because 2 means 2 characters).

DB2 DECFLOAT changes

Programs created before V6R1 can create a user-defined type named DECFLOAT. When these programs are run in V6R1, unqualified references to the type DECFLOAT might resolve to the new system data type in QSYS2, if QSYS2 appears in the SQL path before the library with the user-defined type, whereas in previous releases they resolved to the user-defined type.

Open data path (ODP) I/O feedback changes

The join default value area of the Open data path (ODP) I/O feedback has been reduced from 32 bits to 24 bits. This change accommodates additional feedback flags for decimal floating point. Applications that depended on bits 25-32 for join default value information will have to be modified because this information is no longer available.

SQL materialized query table changes

SQL materialized query tables (MQT) created before V6R1 can be based on protected system files. An MQT over a system-protected file, such as a cross-reference file, causes errors when trying to change the coded character set identifier (CCSID) of the system file. Starting in V6R1, users will not be able to create MQTs over protected system files.

Start Database Monitor (STRDBMON) changes

Database monitor QAQQDBMN model file changes

The record format for the database monitor file has changed. The model file QAQQDBMN reflects the changes to existing fields QQC181, QQC182, and QQC183 and several new fields. Because of these changes, monitor files from previous releases are not compatible. The STRDBMON command will fail with error message CPF436A if the OUTFILE parameter refers to a monitor file created on a previous release. A new database monitor file will need to be created using the STRDBMON command. Because the format of the database monitor file has changed, any programs that reference the file record format for QAQQDBMN should be recompiled.

Database Monitor table format changes

In V6R1, the format of the table used by the Start Database Monitor (STRDBMON) command has changed. If you reuse a Database Monitor output file from a previous release that has not been updated

to match the V5R4 format, the STRDBMON command fails with the message CPF436A. To recover, you can specify a different output file, delete the output file, or use certain functions on System i[®] Navigator to update the file to the new format.

The Database Monitor 1000 record has been incompatibly changed. When QQRID=1000, the SQL_Path is found in the QQCLOB2 column instead of being spread over a combination of QVC1000, QWC1000, QVC5001, QVC5002, QVC3001, QVC3002, and QVC3003.

The following System i Navigator functions update existing pre-V6R1 monitor files to the new format:

- Analyze or Show Statements for a detailed SQL Performance Monitor
- Compare of two detailed SQL Performance Monitors

QAQQINI option changes

In V6R1, the default value for the QAQQINI option for IGNORE_DERIVED_INDEX has changed from *NO to *YES. The default behavior now runs supported queries through the SQE even if a select/omit logical file index is created over any of the tables in the query. In V6R1, the SQE optimizer supports many types of derived indexes. The usage of the QAQQINI option for IGNORE_DERIVED_INDEX applies to only select/omit logical file indexes.

International Components for Unicode sort sequence tables

SQL for DB2[®] i5/OS might remove support for International Components for Unicode (ICU) sort sequence tables that are based on ICU version 2.6.1 in a future release. In V5R3 and V5R4, all ICU sort sequence tables are based on ICU version 2.6.1. In V6R1, support is being added for ICU sort sequence tables based on version 3.4. The naming convention for ICU sort sequence tables is based on the locale. For example, QSYS/FR specifies to use ICU sort version 2.6.1 for French. To use ICU sort version 3.4 for French, specify the new V6R1 sort sequence table QSYS/I34FR.

It is recommended that applications be changed to specify version 3.4 of the ICU sort sequence table instead of ICU version 2.6.1. This includes, but not limited to, recompiling SQL applications, changing JDBC connection properties, and recreating indexes. In V6R1, the SQL Query Engine (SQE) is adding support for sort sequence. SQE supports ICU version 3.4 and does not support ICU version 2.6.1. Also, you cannot use indexes created specifying ICU sort sequence table of version 2.6.1 to implement queries that specified an ICU sort sequence table of version 3.4 or vice-versa.

SQL and DB2 compatibility

On an SQL call statement, the SQLERRD(1) and DB2_RETURN_STATUS return the procedure return value if the RETURN statement is used in the SQL procedure. It was unclear in previous versions of the *SQL Reference* if the SQLERRD(1) and DB2_RETURN_STATUS applied to external procedures. For external procedures, if the procedure returns with an SQLCODE that is greater than or equal to 0, the specified target for DB2_RETURN_STATUS in a GET DIAGNOSTICS statement is set to a value of zero. If the external procedure returns with an SQLCODE that is less than zero, the specified target for DB2_RETURN_STATUS in a GET DIAGNOSTICS statement is set to a value of -1. See the SQL RETURN statement in the *DB2 for i5/OS SQL Reference* for additional information about how the DB2_RETURN_STATUS value is handled for SQL procedures.

User-defined functions for common table expressions invoked according to SQL standards

In V6R1, user-defined functions (UDFs) for common table expressions (CTEs) are invoked consistent with SQL architecture standards. This means that a UDF of this type within a CTE might be invoked more often or less often than it did before V6R1, but now it is consistent in its operation and therefore its (the CTEs) resulting contribution to the main query.

SQL Call Level Interface changes

SQL Call Level Interface API changes

SQL Call Level Interface (CLI) has had a number of changes in V6R1 to align the API with the DB2 Call Level Interface, the CLI API most commonly used for accessing other DB2 platforms. These changes were made to improve compatibility and portability of applications written to interface with the CLI APIs and to conform to the ISO CLI standard.

The most notable change was made to the constant values that represent BINARY and VARBINARY data types in the API. The values of the constants were changed as follows:

Table 1. Changes to the value of constants

Constant name	Old value	New value
SQL_BINARY	96	-2
SQL_VARBINARY	97	-3

Important: This change will require that applications that use CLI, which reference these constants, to recompile their programs that contain the reference before running the program on the V6R1 release. In addition, any application that has a hard-coded value of 96 to represent SQL_BINARY needs to change to use the new value. Any application that has a hard-coded value of 97 to represent SQL_VARBINARY needs to change to use the new value.

Applications that are not recompiled before running on V6R1 will see SQL0804 if an incorrect data type is used in the application.

SQL Call Level Interface metadata changes

Some SQL Call Level Interface (CLI) Metadata functions have changed to improve compatibility and portability of applications written to interface to conform to the ISO CLI standard. Refer to the Compatibility section in the V6R1 SQL Call Level Interface book for details.

SQLColAttribute and SQLColAttributes API changes

The SQL Call Level Interface (CLI) `SQLColAttribute()` and `SQLColAttributes()` APIs have changed in V6R1. The `SQL_DESC_DISPLAY_SIZE` option returns an extra byte in the output length if null terminated strings are being used for the connection. In previous releases, the `SQL_DESC_DISPLAY_SIZE` output length is one byte short, in some cases.

Object management authority changes

The *DB2 for i5/OS SQL Reference* states that when granting or revoking table or view privileges, the statement should fail if the user does not have object management authority (*OBJMGT) for the table or view. The *DB2 for i5/OS SQL Reference* also says that the authorities required to prepare a statement are the same as the authorities required to execute the statement.

Before V6R1, the database did not check for *OBJMGT authority for the table or view during a PREPARE statement. Now that the correct checking is done, an SQL0551 failure might be seen on V6R1, where it had not been seen in earlier releases.

Running the prepared statement has always correctly enforced the *OBJMGT authority level. In earlier releases, the SQL0551 message was seen during the EXECUTE statement, whereas now this message will be seen during the PREPARE statement.

Additional AF (Authority Failure) journal entries created by host server jobs.

Some QAQQINI query options can only be changed if the user has *JOBCTL user special authority or QIBM_DB_SQLADM function usage authorization. QUERY_TIME_LIMIT and STORAGE_LIMIT are examples of query options that require this higher level of authorization, because the change in query behavior may impact system performance.

In V6R1, an attempt to increase either of these values over the current job setting value without *JOBCTL authority will result in the request not being honored and an AF audit journal entry created. In previous releases, the request was not honored but no AF audit journal entry was created.

Substitution character warning changes

Before V6R1, substitution character warnings were not being returned even though the *DB2 for i5/OS SQL Reference* stated that they would be returned. With V6R1, these warnings are now returned.

When a character string is translated from one CCSID to another and one or more character substitutions occurs, the SQLWARN8 value in the SQLCA should contain W, the SQLSTATE should be '01517', and the SQLCODE should be +335. In some cases, such as translations involving UCS-2, neither SQLWARN8 nor SQLSTATE contain the expected values.

Logic involving character translation has been changed to set SQLWARN8, SQLSTATE, and SQLCODE.

SQL function statement changes

The SQL statements that allow you to specify SPECIFIC FUNCTION or SPECIFIC PROCEDURE all run as if you specified SPECIFIC ROUTINE. The result is that you can drop a FUNCTION by using the DROP SPECIFIC PROCEDURE statement. In V6R1, if the type of the object found does not match the type specified in the statement, an SQL0440 failure is returned for the DROP, COMMENT ON, GRANT, and REVOKE SQL statements.

DBCLOB CCSID changes

Before V6R1, if a CCSID was not specified for a DBCLOB and no associated CCSID existed, the DBCLOB column was given a CCSID value of 65535. In V6R1, this was changed so that the DBCLOB column uses a CCSID value of 1200.

Before V6R1, if a DBCLOB host variable did not have a CCSID specified and no associated CCSID existed, an SQL0330 message was returned. In V6R1, this was changed so that the DBCLOB host variable uses a default CCSID value of 1200.

RRN scalar function changes

The RRN() scalar function has been incompatibly changed in V6R1 to fail when used against a view based on a user-defined table function (UDTF). For example:

```
Select RRN(x) from view x when the view x is based on a UDTF.
```

Before V6R1, this usage might return results with no error, In V6R1, it fails with errors CPD43AD and SQL0391.

Resolving an SQL0338 failure

In V6R1, a query might fail with SQL0338 for queries with a JOIN. In earlier releases, this usage was allowed. The error needs to be issued because of the implied order of the join condition.

For example, the following query will begin to fail on V6R1 because the reference to a.c1 is at an incorrect scope for the join predicate.

```
| Select a.c1
| From BASE1 a LEFT OUTER JOIN BASE3 c
| LEFT OUTER JOIN BASE2 b on a.c1 = b.c1
| on b.c1n = c.c1n
```

| The SQL0338 can also be issued when a table is specified followed by a comma and a joined table.

```
| Select a.c1
| From BASE1 a, BASE2 b INNER JOIN BASE3 c on a.c1 = c.c1
```

| To resolve the SQL0338 failure, restructure the query.

| The Start Database Monitor (STRDBMON) command can be used to identify queries that need to change.

| Example:

```
| STRDBMON OUTFILE(QGPL/FIND338S) JOB(*ALL/*ALL/*ALL) FTRSQC(-338)
| ENDDBMON JOB(*ALL/*ALL/*ALL)
| SELECT
| QQC21 AS "operation",
| QVC102 AS "user",
| QQI8 AS "SQLCODE",
| QQC104 AS "Application library",
| QQC103 AS "Application program",
| QQ1000 AS "SQL statement text",
| A.*
| FROM QGPL.FIND338S A
| WHERE QQRID=1000
| ORDER BY QQSTIM
```

DESCRIBE OUTPUT changes

In V6R1, the DESCRIBE OUTPUT SQL statement processing has been changed to return SQL0237 and SQL0239 warnings indicating that not enough SQLVAR entries are supplied in the SQLDA. In previous releases, warnings were not always returned. In some cases, the value returned for SQLD and in byte 7 of the SQLDAID has changed, in addition to the warning changes. In all cases, SQLD multiplied by the value in byte 7 of the SQLDAID (when not blank) matches the required number of SQLVAR entries.

Trimming blanks for C NULL terminated

| Prior to V6R1, retrieval assignment was not always trimming blanks in order to add a null terminator to the end of the data when assigning character data to the C null terminated data type. In V6R1, retrieval assignment will now trim trailing blanks in order to null terminate the data when assigning to the C null terminated data type.

SQE caching of temporary result sets when UDFs and UDTFs are involved

| Prior to V6R1, the SQE optimizer did not cache temporary result sets across queries when the query referenced one or more User Defined Functions (UDFs) or User Defined Table Functions (UDTFs) which were created as DETERMINISTIC and NO EXTERNAL ACTION. In V6R1, the behavior changed so temporary result sets may now be cached for these types of queries. Some customers have experienced issues where the DETERMINISTIC UDF is running queries and re-use of the cached temporary is not seeing updates to the underlying tables referenced by the queries in the UDF since the UDF is not re-run.

| To allow compatibility with previous behavior, a new QAQQINI option, DETERMINISTIC_UDF_SCOPE, is available by PTF.

| DETERMINISTIC_UDF_SCOPE - Specifies the scope or lifetime of the deterministic setting for User Defined Functions (UDFs) and User Defined Table Functions (UDTFs).

| Supported Values:

- | • *DEFAULT - The default value is *ALWAYS.
- | • *ALWAYS - The UDF is always considered deterministic. Query temporary objects might be shared across query opens and the UDF might not run for a particular query open.
- | • *OPEN - The UDF is considered deterministic only for a single instance of a query open. Query temporary objects are not shared across query open. The UDF is run at least once in the query for a given set of input parameters. This is the pre-V6R1 behavior.

| V6R1 PTFs for the DETERMINISTIC_UDF_SCOPE QAQQINI option are:

- | • XPF - SI36907 and co-reqs
- | • SLIC - MF48244(V6R1M0) MF48245(V6R1M1)

SQL programming changes

Close SQL cursor SET OPTION changes

In V6R1, the Close SQL Cursor (CLOSQLCSR) SET OPTION default has changed from *ENDACTGRP to *ENDMOD for SQL routine bodies.

When CLOSQLCSR is not specified, SQL procedures, scalar functions, and triggers are created with CLOSQLCSR=*ENDMOD. In previous releases, CLOSQLCSR=*ENDACTGRP was used. For SQL UDTFs, the CLOSQLCSR option is ignored and the UDTF is always created with *ENDACTGRP.

Because of this change, the default behavior is different in the following case.

Statements prepared in an SQL routine are closed when the routine exits. A statement cannot be prepared on one call of an SQL code body built with CLOSQLCSR=*ENDMOD and executed on a subsequent call.

CLOSQLCSR option change

The default for the Close SQL Cursor (CLOSQLCSR) option on the Run SQL Statements (RUNSQLSTM) command was changed from *ENDACTGRP to *ENDMOD.

SQLCODE and SQLSTATE changes

SQL routines that contain references to SQLCODE and SQLSTATE within an IF-THEN-ELSE statement will not work the same. SQLCODE and SQLSTATE will be set to 0 and the diagnostic area cleared before the first statement is executed in the THEN or ELSE clause. For example:

```
IF SQLCODE < 0 THEN
  SET SAVE_CODE = SQLCODE;
END IF;
```

| SQL routines can be changed to avoid this problem by using the GET DIAGNOSTICS SQL statement.

```
| DECLARE LOCAL_SQLCODE INTEGER;
| DECLARE LOCAL_SQLSTATE CHAR(5);
| ...
| GET DIAGNOSTICS condition 1
|   LOCAL_SQLCODE=DB2_RETURNED_SQLCODE;
|   LOCAL_SQLSTATE=RETURNED_SQLSTATE;
| IF LOCAL_SQLCODE < 0 THEN
|   SET SAVE_CODE = LOCAL_SQLCODE;
| END IF;
```

If SQLCODE is negative when the IF is executed, SAVE_CODE is set to 0. In previous releases, SAVE_CODE was set to the negative value contained in SQLCODE.

| These changes apply only to functions and procedures that are built using the LANGUAGE SQL keyword on the CREATE PROCEDURE, CREATE FUNCTION, CREATE TRIGGER, ALTER PROCEDURE, or ALTER FUNCTION SQL statements. They do not apply to higher-level language programs.

Data mapping errors

When truncation occurs during the mapping of strings for the SET and VALUES statement in an SQL routine, a warning is returned instead of a hard error. In V5R4, the SET statement in procedure **proc** returns SQLCODE =-303, SQLSTATE='22001'. In V6R1, the SET statement returns SQLCODE=0, SQLSTATE='01004'. To get an error on the truncation, specify SET OPTION SQLCURRULE=*STD for the SQL routine:

```
create procedure proc ()
LANGUAGE SQL
begin
  declare v1 char(5) default 'aaaaa';
  set v1 = 'xxxxxxx';
end
```

Data mapping errors, such as truncation and divide by zero, that occur on the RETURN statement of a UDTF will be hard errors. On previous releases, an invocation of a user-defined table function that resulted in data mapping errors usually returned a warning.

```
create function func_divide (p1 int, p2 int)
returns table (x1 int) language sql disallow parallel
  return (select p1 / p2 from qsys2/qsqptab1)

select * from table(func_divide(1,0)) dt
```

The previous invocation of **func_divide** on V5R4 returned SQLCODE=+802, SQLSTATE='01564'. On V6R1, the function returns SQLCODE=-802, SQLSTATE='22012'.

```
create function func_concat (p1 varchar(4), p2 varchar(4))
returns table (x1 varchar(4)) language sql disallow parallel
  return (select p1 || p2 from qsys2/qsqptab1)

select * from table(func_concat('aaa', 'bb')) dt
```

The previous invocation of **func_concat** on V6R1 resulted in SQLCODE =-303, SQLSTATE = '22001'. On V5R4, the function returned the truncated result aaab.

Embedded SQL with WHENEVER SQLWARNING handlers

PTF SI36142 changed the behavior of DELETE and UPDATE SQL statements for some Embedded SQL applications. When DELETE or UPDATE is used without a WHERE clause and the target table contains no rows, a warning condition is indicated in the SQLCA, along with SQLCODE=+100 and SQLSTATE='02000'. Prior to this PTF, the warning condition was not specified. If the application had declared a handler using WHENEVER SQLWARNING, the handler did not receive control. After application of this PTF, the handler will receive control. This behavior change will only impact an application which is executing a DELETE or UPDATE statement against an empty table, with no WHERE clause and has a WHENEVER SQLWARNING handler defined and that handler treats the warning condition as a failure.

SQL0445 and SQLWARN1 for stored procedure output parameters

In releases prior to V6R1 the mapping of stored procedure output parameters was inconsistent.

The inconsistencies included:

- SQLWARN1 was never being set to 'N' for C null terminated strings (NTS)
- The parameter's indicator variable was not set for some cases with the original length of the truncated string
- C Null terminated strings were not being padded for *CNULRQD

In V6R1 these inconsistencies were corrected.

The following is a summary of the changes where SLEN and TLEN are the source and destination string length attribute:

- | • Padding of C NTS for *CNULRQD when SLEN < TLEN. The NULL terminator is added at position TLEN.
- | • No padding of C NTS for *NOCNULRQD when SLEN < TLEN. The NULL terminator will be added at position SLEN + 1.
- | • For *NOCNULRQD, issue SQLWARN1 'N' warnings when only the NULL terminator is truncated (SLEN = TLEN). When SLEN > TLEN issue SQLWARN1 'W'.
- | • For *CNULRQD, only issue SQLWARN1 'W' warnings when SLEN >= TLEN. A NULL terminator will always be added at position TLEN.
- | • Send SQL0445 warning when there is truncation, and set the indicator with the length of the source string that was truncated.

| For fixed length data types (CHAR(N) and GRAPHIC(N)), the source string length attribute to consider is the length attribute N, even though the string might look like: 'a '. That is, when doing retrieval assignment (output parameter mapping), the trailing blanks are significant and generate an SQL0445 when truncated. That is not the case for storage assignment, such as when mapping into the procedure parameter, which allow the silent trimming of blanks.

| A that used to not send truncation warnings in a prior release might send truncation warnings in 6.1. The SQL0445 warning appears as SQLCODE = +445. This will affect applications that are coded to use the EXEC SQL WHENEVER SQLWARNING pre-compiler statement and applications that use SQLCODE <> 0 logic.

| **SQL Query changes in collation of errors might result in a SQL0802**

| In PTF MF45590, a change was made to not collate errors for SQL queries with distinct or aggregation. This was done to avoid providing wrong results when a result in error can not be included in aggregation or distinct result sets. A selection error is now signalled in this scenario, using the SQL0802 message with SQLCODE=-802 and SQLSTATE='22003'. The SQL0802 message will be preceded by CPF5035 and CPD4019 escape messages.

| To allow compatibility with previous behavior, a new QAQQINI option, COLLATE_ERRORS, is available by PTF.

| COLLATE_ERRORS - Specifies how data errors are handled on the GROUP BY and ORDER BY expression during hash or sort processing within queries.

| Supported values:

- | • *DEFAULT - The default is *NO.
- | • *NO - A value of *NO causes the query to be ended with an error when a grouping or ordering expression results in an error.
- | • *YES - A value of *YES indicates that the grouping or sort continues.

| The PTFs needed for the new QAQQINI option are:

- | • V6R1M0 - MF46798, SI35322 and SI35323
- | • V6R1M1 - MF46947, SI35322 and SI35323

| **Badly formed data can trigger data mapping error**

| Prior to V6R1, if there was badly formed data in decimal or zoned fields within physical files, it was dependent on the query implementation and the interface as to whether an SQL0406 or an SQL0802 data mapping error would be signalled. In V6R1, there is better detection for data mapping errors caused by badly formed decimal and zoned data in non-SQL tables. It will be more likely that an SQL0406 or SQL0802 will be seen on V6R1. It will be necessary to correct the badly formed data within the physical file to avoid the data mapping error.

| **Implicit numeric conversions may now return an SQL0302, SQL0678, CPD4363 or CPD3131**

| Implicit numeric conversions between numeric types and a ' ' (blank) or an " (empty string), will now return one of the following:

- | • SQL0302 with error type 6
- | • CPD4363 with reason code 1
- | • SQL0678
- | • CPD3131 (if using OPNQRYF)

| The (' - blank, and " - empty string) are no longer considered character zero.

| Example query:

```
| select *  
| from qsys2.qsqptab1  
| where 0 = ''
```

| **Queries that call a user-defined function changes**

Queries that call a user-defined function (UDF) that is defined with a parameter style of GENERAL or GENERAL WITH NULLS might return different results than in releases earlier than V6R1. If the return value is a character type that is 16 bytes or less, the result is now right-justified. In the following example, an 'A' gets dropped from the result:

The CREATE FUNCTION statement:

```
CREATE FUNCTION NEWSTR(CHAR(5)) RETURNS CHAR(5)  
LANGUAGE C EXTERNAL NAME 'MYLIB/STRFUNCS(NEWSTR)'  
NO EXTERNAL ACTION  
PARAMETER STYLE GENERAL
```

The code:

```
#include <stdio.h>  
#include <string.h>  
#include <stdlib.h>  
  
EXEC SQL INCLUDE SQLCA;  
EXEC SQL INCLUDE SQLDA;  
  
EXEC SQL BEGIN DECLARE SECTION;  
char hvdb[41]; /* input parm */  
typedef struct {  
    char output[6]; /* result parm */  
} rdb;  
EXEC SQL END DECLARE SECTION;  
rdb USCG01VC01(char *hvdb)  
{  
    rdb rdb1;  
    strcpy(rdb1.output,hvdb);  
    return (rdb1);  
}
```

The query:

```
select mylib.newstr(char('ABCDE')) from qsys2.qsqptab1
```

Previous to V6R1, the query returned:

```
'ABCDE'
```

On V6R1, the query returns:

```
'BCDE '
```

Queries converted to INTEGER type

For query operations involving SMALLINT (2-byte integer) and character string types, both operands are now converted to INTEGER (4-byte integer) before performing the operation. Before V6R1, the character string were converted to SMALLINT and the SMALLINT operand was not converted. This affects the resulting data type for the operation.

For example, the resulting data type for `SELECT MAX('123', SMALLINT(1.23E3)) FROM QSYS2/QSQPTABL` will now be INTEGER instead of SMALLINT. As another example, `SELECT * FROM TBL1 A, TBL2 B WHERE A.CHAR1 = B.SMALLINT1`, an index over B.SMALLINT1 is no longer considered for optimization.

Queries converted to REAL type

For query operations involving REAL (single precision floating point) and character string types, both operands are now converted to DOUBLE (double precision floating point) before performing the operation. Before V6R1, the character string was converted to REAL and the REAL operand was not converted. This affects the resulting data type for the operation.

For example, the resulting data type for `SELECT MAX('123', REAL(1.23E3)) FROM QSYS2/QSQPTABL` will now be DOUBLE instead of REAL. As another example, `SELECT * FROM TBL1 A, TBL2 B WHERE A.CHAR1 = B.REAL1`, an index over B.REAL1 is no longer considered for optimization.

JOIN with a USING clause change

The implementation of JOIN with a USING clause specified has been changed to match the standard.

The implicit column list of a select list with a USING clause consists of the common columns, followed by the remaining columns of the first file, and then followed by the remaining columns of the second file.

```
create table t1 (a1 int, a2 int, a3 int, a4 int)
create table t2 (a2 int, a3 int, a4 int, a5 int)
```

For the select statement

```
select * from t1 join t2 using (a2,a3)
```

The result column list is

```
a2, a3, t1.a1, t1.a4, t2.a4, t2.a5
```

The implicit column list of single table does not include the USING columns.

For the select statement

```
select t1.* from t1 join t2 using (a2,a3)
```

The result column list is

```
t1.a1, t1.a4
```

The following statement is *not* valid. Because a2 is in the USING clause, it cannot be qualified in the statement.

```
select t1.a2 from t1 join t2 using (a2,a3)
```

The following is a valid statement. Because a2 is in the USING clause, a single a2 is in the result, so it should not be qualified. On previous releases, the SQL0203 message is issued.

```
select a2 from t1 join t2 using (a2,a3)
```

Date and time function changes

The ADD_MONTHS and LAST_DAY scalar functions have been changed in V6R1 to return a timestamp when the input argument is a timestamp; otherwise a date is returned. In V5R4, a date was always returned.

The NEXT_DAY scalar function has been changed in V6R1 to return a date when the input argument is date; otherwise a timestamp is returned. In V5R4, a timestamp was always returned.

Database Create File program changes

A change has been made to allow the format-level identifiers of SQL tables, indexes and views to be the same as the format-level identifiers of physical and logical files that are generated with data description specifications (DDS). However, this change might cause different format-level identifiers to be generated for any SQL table, SQL index, SQL view or database file that was originated from an SQL created file.

Another change has been made with the Database Create File program which will now generate the correct format-level identifiers for record formats in the following situations:

- The SQL tables are created again with an SQL Create Table statement, and the table contains at least one variable-length, date, time, or timestamp field.
- An alter table request or Change Physical File (CHGPF) command which specifies a SRCFILE, CCSID, LANGID, or SRTSEQ keyword, is performed against a specific SQL table with a date, time, or timestamp field.
- A Copy File (CPYF) request that requires a new format to be generated, such as a CPYF with CRTFILE(*YES) where the new file created spans either basic System ASPs and/or IASP boundaries.

The format-level identifiers are *not* corrected for record formats for the SQL tables with a date, time, or timestamp field in the following situations:

- The files are created again by doing a restore request.
- The files are created again by applying the journal entries.
- The files are created again by replaying the create operation using journal entries.
- New files are created using the Create Duplicate Object command using a specific file as the From Object command. The Copy Library (CPYLIB) command uses the Create Duplicate Object command to create the database files in the new library.

Changes for user-defined external functions that call a service program

In V6R1, a user-defined external function that calls a service program that returns values must be created with the GENERAL parameter style, or the call will return incorrect results. In V5R4, the system allowed parameter styles other than GENERAL to work for user-defined external functions that were defined in a service program. In V6R1, styles other than GENERAL are no longer allowed. A user-defined function not created with the GENERAL parameter style that calls a service program will either return incorrect results, or will in some cases issue message CPF426A with a reason code value of 8.

Create Table Like changes for Implicitly Hidden columns

Application of V6R1 5761SS1 PTF SI36162 changes the way Create Table Like handles columns that have the Implicitly Hidden attribute. Without the PTF applied, if an SQL table with an implicitly hidden column is the source for a new table created using Create Table Like, the column in the new table is not hidden. After application of the PTF, this behavior is changed and the column in the new table is created with the Implicitly Hidden attribute. In the example that follows, the behavior before the PTF would be to return two columns (C1 and C2) with the select. After application of the PTF, the select will return only the non-hidden column (C1).

```
CREATE TABLE QTEMP.T2 (C1 INT, C2 INT IMPLICITLY HIDDEN);
CREATE TABLE QTEMP.T2B LIKE QTEMP.T2;
INSERT INTO QTEMP.T2B (C1) VALUES (5);
SELECT * FROM QTEMP.T2B;
```

SQLTABLEPRIVILEGES now returns Index privileges

Application of V6R1 5761SS1 PTFs SI46751 and SI46765 will change the behavior of SYSIBM/SQLTABLEPRIVILEGES. Prior to these PTFs, we did not return the INDEX privilege for SYSIBM/SQLTABLEPRIVILEGES. That was a conscious choice because INDEX is a synonym for the ALTER privilege in DB2 for i. With these PTFs, both ALTER and INDEX privileges are now returned. The main reason for this change is that if an attempt was made to generate the SQL to port a CREATE TABLE and the privileges to another DB2, in order to get the same authority, both the ALTER and INDEX privileges would be needed.

DRDA package management changes

In V6R1, DRDA packages that are accessed in a DRDA connection will be locked shared read during the life of the DRDA connection. This will prevent other users from dropping the DRDA package while it is in use. A connection can replace or drop a DRDA package only if it can obtain an exclusive lock on the package.

New RUNSQL command

DB2 PTF Group SF99601 level 25 (or higher) includes support for a new CL command named RUNSQL (Run SQL).

If you have your own version of a command named RUNSQL and that command is called without library qualification, the new QSYS/RUNSQL command will be the command called.

To determine whether RUNSQL commands exist on a machine, use WRKCMD CMD(*ALL/RUNSQL).

To overcome the command conflict problem, consider the following options:

1. Change applications to library qualify use of non-QSYS RUNSQL commands.
2. Rename the IBM supplied RUNSQL command to RUNSQLIBM.

Note: this option should be considered as a short term fix.

Device files may experience performance problems

Application programs that use printer, display, and communications device file (*FILE) objects might run slower after upgrading a system to the IBM i 6.1 release. Enhancements were made in 6.1 to further protect system data at the beginning of the *FILE objects from accidental or intentional tampering. Device files created on 6.1 will not experience a performance degradation. Existing device file objects will be automatically converted the first time the file objects are restored from a save file or save media and will not experience a performance degradation after the files are converted.

Domain name system (DNS) changes

Conversion of customer configuration information

Conversion of customer configuration information located in /QIBM/UserData/OS400/DNS to BIND 9 format is required. This conversion is automatic and occurs on first attempt to start a DNS server or first attempt to change a DNS configuration using the Change DNS Server Attributes (CHGDNSA) CL command or System i Navigator.

The contents of the file /QIBM/UserData/OS400/DNS/ATTRIBUTES will identify whether conversion needs to occur. If this file shows V4R2M0 as the first line, it indicates you are running BIND 4 and conversion is required. If this file shows V5R1M0 as the first line, it indicates you are running BIND 8 and conversion is required. If this file shows V6R1M0, it indicates you are running BIND 9 and no conversion is required.

When conversion is complete, you can add or change the DNS configuration information.

Care should be taken when restoring saved configuration information, as data backed up before any conversion will cause all updates to be lost and the automatic conversion to occur again.

| Domain name system (DNS) security related changes

| Two security related changes for DNS running BIND 9 may require configuration updates.

| • **The default behavior has changed for the "allow-query-cache" and "allow-recursion" ACLs.**

| If the "allow-query-cache" ACL is not explicitly set in named.conf, the default for host cache access is determined by other ACL configurations as follows:

- | 1. The "allow-recursion" ACL, if specified.
- | 2. The "allow-query" ACL, if specified.
- | 3. If neither "allow-recursion" nor "allow-query" are specified, the default of (localnets; localhost;) is used.

| If the "allow-recursion" ACL is not explicitly set in named.conf, the default for recursive access is determined by other ACL configurations as follows:

- | 1. The "allow-query-cache" ACL, if specified.
- | 2. The "allow-query" ACL, if specified.
- | 3. If neither "allow-query-cache" nor "allow-query" are specified, the default of (localnets; localhost;) is used.

| For example, do the following if you want all clients to be able to query the record recursively:

- | 1. Stop the DNS server
- | 2. Open the /QIBM/UserData/OS400/DNS/<DNS_NAME>/named.conf and add the following statements:

```
| allow-query  
| {  
| any  
| };
```

- | 3. Start the DNS server

| • **If the server on IBM i is a slave DNS server which receives dynamic DNS updates from DHCP servers or DNS clients, you need to make this change to your configuration:**

| Specify security-related option "allow-update-forwarding" and add the IP addresses of the DHCP servers or clients that are allowed to make dynamic DNS updates. Note that the special value any can be used to allow updates from any IP address.

| For example, if the IP address of your DHCP server is 10.1.1.1, and the IP address of one client is 10.2.2.2, then you would do the following:

- | 1. Stop the DNS server
- | 2. Open the /QIBM/UserData/OS400/DNS/<DNS_NAME>/named.conf and add the following statements:

```
| allow-update-forwarding  
| {  
| 10.1.1.1;  
| 10.2.2.2;  
| };
```

- | 3. Start the DNS server

ENDTCP before ENDSBS *ALL *IMMED

If the Telnet server is not ended prior to the ending of subsystems, it might not automatically restart. To circumvent this and other potential problems it is suggested that prior to ENDSBS *ALL *IMMED these steps be taken:

1. ENDTCPSVR *ALL
2. ENHOSTSVR *ALL
3. Allow the commands to complete and the servers to end. If issuing the commands from a CL program include a DLYJOB DLY(300) to allow the commands to complete.
4. ENDTCP

Exit point changes

QIBM_QPWFS_FILE_SERV exit point changes

The definition for exit point QIBM_QPWFS_FILE_SERV will be different when V6R1 is installed over an earlier release. When an exit program is added or removed for the exit point, a reprocessing program will require the user doing the add or remove operation to have *ALLOBJ and *SECADM special authorities. The Allow change parameter will be set to *NO which means the user will not be able to change the exit point definition. If a user had been changing the exit point definition in earlier releases, they will not be able to do so in V6R1. In V6R1, users who could not retrieve the path name to pass to the exit program registered for exit point QIBM_QPWFS_FILE_SERV will be able to retrieve the path name.

QIBM_QWC_PRERESTRICT and QIBM_QWC_PWRDWNSYS exit point changes

Batch immediate jobs which call the exit programs for the Pre-restricted State Exit Point (QIBM_QWC_PRERESTRICT) and the PWRD0200 format of the Power Down System Exit Point (QIBM_QWC_PWRDWNSYS) are no longer running under the QSYS user profile when the exit programs are called. These jobs will be running under the user profile that was used to initiate the job, which is the current user profile of the job which issued the End System (ENDSYS), End Subsystem (ENDSBS *ALL), or Power Down System (PWRDWNSYS) command. If an exit program requires additional authority, it will need to adopt authority or switch to a different user profile.

File-level backup through QNTC file systems

If your PC has Windows 2000 installed and you upgrade to V6R1, file-level backup through QNTC might no longer work. The QNTC file system does not correctly return all of the objects in QNTC. This problem is related to the number of objects in a directory and the length of the names of those objects. This known problem in Windows 2000 has been fixed in Windows 2003.

Perform one of the following actions to avoid this problem:

- To *permanently* resolve this problem, upgrade the Microsoft server to an operating system that Microsoft supports, and install the most recent Microsoft service pack.
- To *temporarily* resolve this problem, divide directories that cannot be read into several smaller directories or rename long file names to shorter file names.

FTP changes

In V6R1, the FTP client attempts to use Extended Passive Mode for data connections. This is a change from previous releases. Not all FTP servers support this new function and users might receive the FTP550 error before receiving the data. You can issue the SENDEPSV FTP client subcommand to shut off this function for the current session before issuing any data command.

Handling inquiry messages

The function to call user exit programs for exit point QIBM_QMH_HDL_INQEXT, to intercept an inquiry message sent to *EXT in an interactive job, is available with PTF SI29311. If an inquiry handling exit program is defined for this exit point, you might see a behavior change for inquiry messages sent to *EXT that were shown on the Display Program Messages screen. Exit programs for this exit point can send a reply to the inquiry message so that you are no longer interrupted to obtain a reply to an inquiry message. In addition, depending what is done in an exit program, you might see longer processing times or even additional screens shown that were not shown before.

Host Ethernet Adapter changes

- | In V6R1, the default link status behavior for the Host Ethernet Adapter (HEA) has changed.
- | At V5R4M5, the configuration status of the HEA Ethernet line description stays active regardless of the status of the link for the physical HEA to the network. This behavior allows interpartition communication to continue regardless of the status of the external physical link.
- | With V6R1 and PTF MF43768 applied, the configuration status for the Ethernet line description always shows the link state of the physical HEA port. Knowing the link state is important if you are using the HEA in a proxy address resolution protocol (ARP) for virtual IP configuration. This change allows the TCP/IP stack to detect a physical link failure, and allows the proxy agent to switch over to the active line description. However, with this new default behavior, interpartition communications using the HEA might be interrupted if the physical link is interrupted. You can change this function back to the way it worked at V5R4M5 by running the Advanced Analysis macro:
| `heaignorephysical -on`
- | Search for document **477927749** in the Software Support Knowledgebase for further instructions on using the Advanced Analysis macro. To view this document, go to the following Web site:
| <http://www-912.ibm.com/ImprovedSearch/searchoptions.jsp>

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.

Host server job changes

Audit, journal, and accounting information now show more work being done under the host server job's actual user profile, and less work under the QUSER user profile. Additionally, single-use jobs (prestart jobs and batch-immediate jobs) finish under the user's profile, rather than QUSER's profile.

Integrated file system names

File systems store names in Unicode. Those file systems that are not case sensitive are affected by changes in characters and the casing rules for a specific Unicode standard. File systems that are not case sensitive, such as the "root" (/) and user-defined file systems (UDFS) that are created with CASE(*MONO), support Unicode Standard 4.0 as of i5/OS V6R1.

The conversion for names in "root" (/) and UDFSs in basic user auxiliary storage pools (ASPs) begins automatically shortly after V6R1 is installed. This conversion does not significantly impact your system activity, as it runs in a low-priority background job. The conversion for UDFSs in independent ASPs starts automatically after the independent ASP is varied on. The conversion runs until all directories have been converted to use Unicode Standard 4.0. It is restarted after every IPL or independent ASP vary on

until all eligible file systems have been converted. The order of the conversion is the “root” (/) file system, user-defined file systems in ASPs 1 through 32, and user-defined file systems in independent ASPs, as they are varied on. This might cause additional CPU usage and an increased number of input and output (I/O) operations in the System i environment. This is normal. The CPU utilization and the I/O rate returns to normal after the directory conversion function has completed. To check the status of the conversion, see the messages that are sent to the QSYSOPR message queue. Error messages are sent to the QHST history log.

For more information, see the following materials:

- **Converting names to support additional characters** in the Integrated file system topic in the i5/OS Information Center.
- Info APAR II14306, which describes the Analyze Object Conversion (ANZOBJCVN) command.

Integrated Language Environment (ILE) Compiler changes

Activation group termination routines run during job termination:

In earlier releases, activation group termination routines were not run for either the user or system state activation groups during job termination. In V5R2, a change was added to call activation group termination routines for the user state default activation group before resetting the activation group when the job was being reused. In V6R1, activation group termination routines for the default activation group are run during job termination. Activation group termination routines include high-level language termination routines (such as C++ destructors), routines registered using the Register Activation Group Exit Procedure (CEE4RAGE and CEE4RAGE2) APIs, or routines registered using the ILE C runtime atexit API. Programs and service programs that activate into the caller's activation group can be activated into the default activation group.

IPv6 changes

New host table file

The current host table file, QUSRSYS/QATOCHOST, is documented in the **Database and device files used by CL commands** topic in the i5/OS Information Center as a file that can be used by a user program. The contents of the old version of the host table file will be preserved indefinitely so as not to break any customer applications that might be referencing it. This file will be maintained with the same contents as in the new host table file with the following restrictions:

- IPv6 addresses are not supported; only IPv4 addresses are supported.
- Up to four host names per IPv4 address. If more than four host names are specified for an IP address, those host names are not copied to the old file. The first four host names specified on the ADDTCPHTE and CHGTCPHTE commands are copied.

To access any of the new functions provided by the new host table file, use sockets APIs **gethostent()**, **sethostent()**, and **endhostent()**; or their thread-safe equivalents: **gethostent_r()**, **sethostent_r()**, and **endhostent_r()** to obtain host table information.

QUSRSYS/QATOCTCPIP (TCP/IP domain information) file

The current TCP/IP domain information file, QUSRSYS/QATOCTCPIP, is documented in the **Database and device files used by CL commands** topic in the i5/OS Information Center as a file that can be used by a user program. The contents of the following record will be preserved indefinitely so as not to break any customer applications that might be referencing it:

- Domain name server record. Currently, the RMTNMESRV record contains the IPv4 addresses of the domain name servers. It will continue to contain this information as long as IPv4 addresses are specified for the domain name servers.

To access the data contained in the QATOCTCPIP file, use the Retrieve TCP/IP Attributes (QtocRtvTCPA) API.

TCP/IP configuration changes

TCP/IP configuration physical files QUSRSYS/QATOCIFC and QUSRSYS/QATOCRTE will be converted to add a text description field to the end of each of their record formats. References to the new text description field will be added to logical files QUSRSYS/QATOCLIFC, QUSRSYS/QATOCLRTE, and QUSRSYS/QATOCLRT2. This conversion is performed when i5/OS first uses the physical files. These files should not be used directly by customer application programs. To obtain TCP/IP interface information programmatically, use the List Network Interfaces (QtocLstNetIfc) API. To obtain TCP/IP route information, use the QtocLstNetRte command.

Also, the default values for the TCP receive buffer size (TCPRCVBUF) and TCP send buffer size (TCPSNDBUF) parameters on the Change TCP/IP Attributes command (CHGTCPA) command have changed for V6R1 from 8192 to 65 536. This change has no effect on the values for those parameters that are stored in the TCP/IP configuration, but if a customer has a CL program that contains TCPRCVBUF(*DFT) or TCPSNDBUF(*DFT) specified and that program is run, the new default value takes effect. In most cases, the new default values result in better performance when sending and receiving data using TCP.

The message TCP2617 (TCP/IP connection to remote system &2 closed, reason code &5) is no longer sent to the QSYS/QSYSOPR message queue. This message continues to be sent to the QUSRSYS/QTCP message queue and the history log. It is sent to the QTCPWRK job log. In addition, the * value on the LCLINTNETA and RMTINTNETA parameters on the ENDTCPCNN command have a slightly different meaning for V6R1. It used to mean "IPv4 unspecified address". Now, because the command supports IPv6, it means "both IPv4 and IPv6 unspecified addresses". This command will function the same way as in V5R4 and previous releases: the IPv4 connection identified by * is ended. But for V6R1, if an associated IPv6 connection exists, it also is ended.

To end the IPv4 connection only, specify the null IPv4 address (0.0.0.0). To end the IPv6 connection only, specify the null IPv6 address (::).

IPv6 socket changes

In V6R1, the header file **netinet/ip6.h** will no longer contain the following constants and structures:

- IP6OPT_BINDING_UPDATE
- IP6OPT_BINDING_ACK
- IP6OPT_BINDING_REQ
- IP6_BUF_ACK
- IP6_BUF_HOME
- IP6_BUF_COA
- IP6_BUF_ROUTER
- **struct ip6_opt_binding_update**
- **struct ip6_opt_binding_ack**
- **struct ip6_opt_binding_request**

These IPv6 binding options and constants are obsolete; they are inconsistent with the binding message structures described in RFC 3775. In addition, the field names of struct **ip6_opt_home_address** were updated in **netinet/ip6.h** to comply with RFC 4584. In previous releases, **getnameinfo()** truncated a node name or service location if the buffer provided was not large enough. In V6R1, **getnameinfo()** will fail with a error return value of **EAI_OVERFLOW** if the buffer for a node name or service location is not large enough.

Java Database Connectivity (JDBC) changes

java.sql.DatabaseMetadata changes

The following changes might affect users of the Java Database Connectivity (JDBC) interface: `java.sql.DatabaseMetadata`. JDBC is an application programming interface (API) included in the Java platform that enables Java programs to connect to a wide range of databases.

The implementation of the `DatabaseMetaData` functions has been changed to match the JDBC 4.0 specification and to match the information returned by the same functions on other DB2 platforms. Consequently, the following behavior changes might be noticeable.

- The native JDBC driver previously permitted the use of “localhost” as the catalog name for most of the metadata methods. The JDBC specification states that the catalog name must match the catalog name as stored in the database. In V6R1, no information is returned if “localhost” is specified as the catalog name.
- The native JDBC driver always returned an empty `ResultSet` when the nullable parameter for the `getBestRowIdentifier` method was set to false. This parameter now returns the correct result.
- The values returned by the `getColumn` method for the `BUFFER_LENGTH`, `SQL_DATA_TYPE`, and `SQL_DATETIME_SUB` columns might be different. These values should not be used in a JDBC application because the JDBC specification defines these columns as unused.
- The native JDBC driver used to treat the table and schema parameters of the `getCrossReference`, `getExportedKeys`, `getImportedKeys`, and `getPrimaryKeys` methods as a pattern. This behavior has been fixed. The JDBC specification states that the name must match the name as stored in the database.
- Some values returned by the `getProcedureColumns` for various data types have been changed. The values now match the values returned by other platforms in the DB2 family.
- Some column names returned by the `getProcedure` method are different. These columns are columns 4, 5, and 6. The JDBC specification does not define these columns. The column names now match the names used by other platforms in the DB2 family.
- The Remarks columns for `getProcedures` used to return an empty string if no information was available. In this case, the column now contains null. This behavior is consistent with the DB2 family behavior.

New `QueryCloseImplicit` JDBC connection property

This release introduces performance improvements, including the new native JDBC driver connection property, `QueryCloseImplicit`. By default, this new property will be enabled.

Using this new `QueryCloseImplicit` property provides significant performance improvement for many applications that use native JDBC to query the System i database.

Although it is unlikely, with this new connection attribute, applications might see slightly different semantics in certain error scenarios. These error scenarios are APARable situations with or without the new function. With the new function, the error on the `CLOSE` call, which should never happen except in APARable situations, is reported on the `ResultSet next()` method call in addition to the `ResultSet close()` method to the JDBC application. Without the new function, the error is only returned on the `ResultSet close()` method.

Java programs created by CRTJVAPGM have no DE code

This section applies to only Java applications that are run using the Classic Java Virtual Machine.

Beginning with V6R1, the System i Java Virtual Machine no longer uses statically-compiled machine instructions in Java programs. These machine instructions were previously referred to as direct execution (DE) code in the Java programs. If you have a Java program from a previous release that contains DE

code, that DE code is removed either when you use the Create Java Program (CRTJVAPGM) command on that file in V6R1, or when the existing Java program is run on V6R1 or later.

You can still create a Java program with the CRTJVAPGM CL command. The resulting Java program contains class information which is used by the Java runtime interpreter or the Java JIT (Just In Time) compiler. Use of the CRTJVAPGM CL command on all .jar and .zip files containing Java classes is optional. However, use of the CRTJVAPGM command is recommended because it helps consolidate and reduce system storage utilization.

Job queue object changes

Conversion of job queue objects

The attributes of a job queue are being moved from an index entry to the associated space of the job queue object to promote object resiliency. This requires a conversion of existing job queue objects, which occurs during the IPL after installation of the base operating system. This conversion results in an increase in storage of 4 kilobytes for each job queue object.

Job notification exit point changes

To register or deregister a data queue for the Job Notification exit point, QIBM_QWT_JOBNOTIFY now requires job control (*JOBCTL) special authority, *CHANGE authority to the data queue, and *EXECUTE authority to the data queue library. If the data queue or the data queue library does not exist, a user with all object (*ALLOBJ) and job control (*JOBCTL) special authorities is allowed to register or deregister a data queue for the Job Notification exit point.

The Job Notification exit point QIBM_QWT_JOBNOTIFY uses the authority of the QSYS user profile when sending to the identified data queues. This change makes the job queue notifications consistent with the job start and job end notifications.

Job termination changes

Display Program Messages screen no longer shown during job termination.

After applying PTF SI43173, the Display Program Messages screen will no longer be displayed when an interactive job has been requested to be ended or while a job is in termination phase. The Display Program Messages screen is shown for inquiry or informational messages sent to *EXT. If an inquiry message is sent, the default reply will be sent to the inquiry message and the job termination will continue, without allowing an end user to reply to the message. The default reply will not be allowed to be rejected by a reply handling exit program. Before this change, a message sent to *EXT during job termination would be displayed and wait for a response from the end user. While waiting for a response from the screen, an End Job (ENDJOB) command or End Subsystem (ENDSBS) command would have no effect on the job, and the End Job Abnormal (ENDJOBABN) command was required to end the job. With the PTF, messages sent to *EXT during job termination will not be displayed resulting in ENDJOB and ENDSBS ending the job as expected.

Message waiting light not processed at job end.

PTF SI46730 includes a performance improvement: the message waiting light will no longer be processed when a job is ending. This means the message waiting light could behave differently after this fix is applied. For example, if group jobs are being used with a group message queue and the light is on, when the job ends, the message waiting light may be turned off.

Journaling enhancements

Reduced lock levels required for journal commands and APIs:

The lock levels have been reduced for several journal CL commands and APIs. In earlier releases, the lock level was O_RDONLY and O_SHARE_NONE for integrated file system objects and *EXCL for all other object types. The lock levels are now O_RDONLY and O_SHARE_RDWR for integrated file system objects and *EXCLRD for all other object types. The following are the commands and the APIs that have changed:

- End Journal Physical File (ENDJRNPf) - file. (**Note:** You cannot end journaling of a physical file that has uncommitted changes.)
- End Journal Access Path (ENDJRnAP) - file and access paths of members.
- End Journal Object (ENDJRnOBJ) - data areas and data queues.
- End Journal (ENDJRn) and QjoEndJournal API - data areas, data queues, and integrated file system objects.
- Start Journal Access Path (STRJRnAP) - file and access paths of members.
- Start Journal Object (STRJRnOBJ) - data areas and data queues.
- Start Journal (STRJRn) and QjoStartJournal API - data areas and data queues.

Messages for journaling commands:

Message handling has changed for the following journaling commands:

- End Journal (ENDJRn)
- End Journal Access Path (ENDJRnAP)
- End Journal Object (ENDJRnOBJ)
- End Journal Physical File (ENDJRnPF)
- Start Journal (STRJRn)
- Start Journal Access Path (STRJRnAP)
- Start Journal Object (STRJRnOBJ)
- Start Journal Physical File (STRJRnPF)

Individual completion messages for each object that successfully starts or ends journaling are no longer sent by default during the processing of the above commands. The message behavior has changed for existing users of these commands. To get the old messaging behavior for these commands, specify LOGLVL(*ALL).

Processing changes when starting and ending journaling for physical files:

| Before V6R1, the Start Journal Physical File (STRJRnPF) and End Journal Physical File (ENDJRnPF)
| commands stopped at the first file that failed to start or end journaling. Starting with V6R1, if more than
| one file is specified on the command, the STRJRnPF and ENDJRnPF commands process the entire list of
| file names passed rather than stop at the first file that fails to start or end journaling. A diagnostic
| message is issued for each failure and the command ends with the CPF700A (from STRJRnPF) or
| CPF700B (from ENDJRnPF) escape message. These two messages summarize the total number of objects
| that are selected for processing and the total number of objects that are successfully processed.

| Because of this change, when more than one file is specified, the commands no longer return any
| exception other than CPF700A (from STRJRnPF) or CPF700B (from ENDJRnPF). If you specify only one
| file, then the commands return the actual failing message as an escape message, as they have in past
| releases. If you have existing programs that specify a list of files for either of these commands, and
| monitor for a specific message from one of these commands (an example would be CPF7030 - Object
| already journalled), your programs will have to change. The programs no longer receive any escape
| message other than CPF700A or CPF700B.

Changes required when handling the CPF7030 escape message:

Existing applications that monitor for the CPF7030 escape message will need to change. The CPF7030 is now sent as a diagnostic message if the object is already journaled *and* the Start Journal Physical File (STRJRNPf) or Start Journal Object (STRJRNOBJ) command specified different journaling attributes than the attributes of the currently journaled object. For example, you could create a file named **myfile** that is journaled with attributes IMAGES(*BOTH) and OMTJRNE(*OPNCLO). If you were to issue a STRJRNPf command to journal **myfile** with attributes IMAGES(*AFTER) OMTJRNE(*OPNCLO), the command fails with the following messages:

- Diagnostic message CPF7030.
- Diagnostic message CPD7002.
- Escape message CPF700A.

Instead, if you issue a STRJRNPf command for **myfile** and specify IMAGES(*BOTH) OMTJRNE(*OPNCLO) (which match the current journaling attributes of the file), the command issues the completion message CPC7031.

Journaled libraries (SQL schema)

A newly created SQL schema is an automatically journaled library. Any file created into the schema is automatically journaled. You can now use the Display Library Description (DSPLIBD) command to view the inheritance rules of the schema.

When a library (schema) is journaled, the journal entries for logical files that are built over SQL table functions are recorded in the journal for the library. In earlier releases, journal entries for these objects were recorded in the QSYS2/QSQJRN journal.

Journaled libraries (STRJRNLIB command):

The Start Journal Library (STRJRNLIB) command holds a lock level of *EXCLRD on the library until the library object is journaled. For customers who used the QDFTJRN data area for automatic journaling, the Create Data Area (CRTDTAARA) and Change Data Area (CHGDTAARA) commands held a lock level of *SHRUPD on the library.

Authority to the journal:

No authority to the journal is required when a newly created object is automatically journaled. The automatic journaling includes the following scenarios:

- An integrated file system object can be automatically journaled when created in a directory with the inherit journaling attribute on.
- A file, data area, or data queue can be automatically journaled when created in a journaled library or a library that contains the QDFTJRN data area.
- A SQL table can be automatically journaled when created in a schema.

Before V6R1, the system required you to have *OBJOPR and *OBJMGT authority to the journal for automatic journaling to complete successfully. The one exception was that SQL automatic journaling did not require you to have any authority to the journal.

Note: If you are using a System i web client that is communicating with a system that is at an earlier release than V6R1, the system still performs the authority checking. For example, creating a stream file in an integrated file system directory that has the inherit journaling attribute on still requires the user on the client system to have *OBJOPR and *OBJMGT authority to the journal.

APYJRNCHGX command:

The support for the Apply Journaled Change Extended (APYJRNCHGX) command will be withdrawn in a future release. The APYJRNCHGX command was first delivered in i5/OS V5R2 and provided

additional functions that the Apply Journalled Changes (APYJRNCHG) command did not have. In V6R1, the APYJRNCHG command and other CL commands provide the functions that were provided by the APYJRNCHGX command.

Automatic journaling using the QDFTJRN data area:

Automatic journaling support using the QDFTJRN data area was introduced in V5R3. In V6R1, this function has been replaced by journaling the library. Support of the QDFTJRN data area will be withdrawn in a future release of i5/OS.

In V6R1, if a library is journalled and it contains a QDFTJRN data area, the information in the QDFTJRN data area is used to automatically start journaling. The inherit rules for the journalled library are ignored in this situation.

Message CPC7031 changes:

For message CPC7031, if the number of objects that are journalled using the Start Journal Physical File (STRJRNPF) command exceeds 32 767, the first field in the message data fields returns -1. Additional message data fields were added to the message to accurately report how many objects started journaling even when the number of objects exceeds 32 767.

QJOCHRVC API changes:

The Change Journal Recovery Count (QJOCHRVC) API required authority has changed from *ALLOBJ to *JOBCTL. The QJOCHRVC API had an incorrect *ALLOBJ authority requirement as no external objects are accessed by the API. The most appropriate authority to require is *JOBCTL as jobs on the system may be affected by the API. Therefore message CPF802 is no longer issued, and audit entry T AF with subtype A is no longer recorded. Instead, message CPF69AF is issued and audit entry T AF with subtype K is recorded.

DSPJRN command changes:

By default, the Display Journal (DSPJRN) command no longer returns entries that are marked to be ignored during an Apply Journalled Changes or Remove Journalled Changes operation (where the Ignore APY/RMV value is Yes) unless the journal entry has a journal code of B. These entries are not returned for any output operation (displaying to the screen, writing to an output file, or printing) unless INCHIDENT(*YES) is specified on the command.

Change to DSPJRN output file processing:

The DSPJRN command no longer fails with diagnostic message CPF7063 (Cannot use file &1 as OUTFILE) and exception CPF9860 (Error occurred during output file processing) when the specified output file is journalled to the same journal specified on the JRN parameter. This results in a journal entry being written to the journal for each record that is written to the output file.

Difference in journal receiver size reporting:

In V6R1, a journal receiver's size might change after it has been detached from the journal if the Remove internal entries (*RMVINTENT) value had been specified for the receiver size options (RCVSIZOPT) for the journal. Be aware this will result in the Work with Journal Attributes (WRKJRNA) command or the QjoRetrieveJournalInformation API returning a larger size for the journal receiver than other system interfaces (Display Object Description for example). This size change is the result of the system freeing up the storage consumed by the internal journal entries when they are no longer needed for recovery. This size change happens after the journal receiver is detached from the journal.

Change in handling authority failures for applying and removing journalled changes:

Before V6R1, when using the Apply Journalled Changes (APYJRNCHG), Apply Journalled Changes Extended (APYJRNCHGX), and Remove Journalled Changes (RMVJRNCHG) commands, if you did not have authority to a journalled file or object in the library and you specified to work with all objects in a library, no changes for the object would be applied or removed and no error message would be issued. Starting in V6R1, these commands fail with an authority failure message (CPF9822 or CPF9802) if the object is journalled to the specified journal. If you do not have authority to the object and it is either not journalled, journalled to a different journal, or not eligible to be journalled, no authority failure will result.

LICOPT usage

The `MinimizeTeraspaceFalseEAOs` value on the LICOPT parameter, used on the CHGPGM and CHGSRVPGM commands, for example, has a different meaning now that teraspace is hardware-based instead of software-based. Effective Address Overflow exceptions are now much less likely to occur when working with teraspace. If you use the `MinimizeTeraspaceFalseEAOs` value, it might improve performance to stop using this option for V6R1.

Machine interface (MI) instruction changes

DESMTX MI instruction

The DESMTX MI instruction is being changed to allow mutexes to be destroyed using a mutex copy. When a mutex is successfully destroyed using a mutex copy, the mutex copy is set to binary zero, and the caller receives the normal (0) return value instead of the ENOTSUP error that was returned previously. When a mutex is destroyed using a mutex copy, attempts to use the originally created mutex or another copy of that mutex results in the EDESTROYED or EINVAL error return value.

MATAL MI instruction

The **number of bytes available for materialization** field in the MATAL instruction returns values up to 2 147 483 647. If more bytes are available, the number of bytes available is returned in the new **materialize size value** field and **number of bytes available for materialization** is set to -1. If you need to materialize more than 2 147 483 647 bytes, you can use either the **information requirements** option 72 to materialize entries into an independent index or you can use the **array of type and subtype codes** option to subset your MATAL output.

MATINXAT MI instruction

In the MATINXAT MI instruction, the **maximum entry length** field has been renamed **maximum entry length attribute** and the return value of 1 now indicates that the maximum entry length is between 2000 and 32 000 bytes.

MATMATR MI instruction

The Materialize Machine Attributes (MATMATR) selection value Hex 013C for timed power on has been changed. The time and date at which the system should automatically power on has been changed from local system time to Coordinated Universal Time (UTC).

MATPRATR MI instruction

i5/OS no longer provides an access group when initiating a process. The Materialize Process Attributes (MATPRATR) MI instruction now returns a null pointer for the process access group. Any customer program that uses the MATPRATR MI instruction and depends on receiving a non-null pointer will need to be changed.

MATRMD MI instruction option hex 12 and 20

To account for the longer serial number, an Extended Serial Number field has been added to the end of the template. When running on hardware that supports the new format, the old serial number field is set

to an extracted version of the 11S-format serial number. In all cases, the Extended Serial Number field contains either the older 10-character serial number or the 11S-format serial number padded to the right with blanks.

It is recommended that current users of MATRMD options 12 and 20 use the new Extended Serial Number field.

POWER6 systems with EnergyScale

IBM POWER6-based systems with EnergyScale™ are able to reduce power consumption by varying processor speed. The POWER6 processor provides timekeeping facilities to track absolute and relative processor time to enable software to account for dynamic processor speed variations. The *processor time* is processing time independent of the processor speed. The *scaled processor time* is the processor time at full speed. If the processor is running at half speed, the scaled processor time accumulates at half the rate of the processor time. The i5/OS machine interface has been updated to support variable-speed processors. Interfaces that report processor time have been changed to report both processor time and scaled processor time, and to indicate whether the two times can differ during the IPL.

The updated MI instructions include:

MATMATR option Hex 20 (Processor attributes)

The scaled processor time attribute is materialized.

MATRMD options Hex 26 (Processor utilization data) and Hex 28 (Multiprocessor utilizations)

The scaled processor time attribute is materialized. The scaled processor utilized time, scaled processor stolen time, scaled processor donated time, scaled processor idle time, and scaled processor interrupt time are materialized.

MATRMD option Hex 27 (Shared processor pools utilization information)

The scaled processor utilized time is materialized.

MATPRATR options Hex 21 (Process resource usage attributes), Hex 23 (Thread Performance Attributes), Hex 24 (Execution Status Attributes), Hex 32 (Threads execution status attributes), Hex 37 (Performance statistics and execution attributes), and Hex 38 (Workload management information)

The scaled processor time attribute is materialized.

MATMIF option Hex 0001

The scaled processor time attribute is materialized.

MATMIF option Hex 0002

The scaled processor time attribute is materialized. The scaled processor utilized time is materialized.

Teraspace location locks

LOCKSL, LOCKTSL: Mapping a teraspace location to a single-level storage address is no longer permitted. The **Type of teraspace storage location lock** field is no longer supported. Any value specified for this field is ignored.

Starting with V6R1, the LOCKSL instruction is the preferred method for locking space locations, including teraspace locations. The UNLOCKSL instruction also is the preferred method for unlocking space locations. If future enhancements are necessary, only LOCKSL and UNLOCKSL will be changed. The LOCKTSL and UNLOCKTSL instructions will not be updated for any future enhancements.

LOCKTSL: Locking a teraspace location using a mapped single-level storage address is no longer allowed.

MATAOL, MATOJLK, MATSELLK: Materialization instructions interpret teraspace addresses from the context of the thread in which they are running.

MATPRK: If the current thread does not belong to the process whose locks are being materialized, a null pointer is returned for any teraspace locations.

NetServer User Profile changes

Application of PTFs V6R10: MF55659 or V6R1.1: MF55660 will result in behavior changes for NetServer user-profiles.

The CHGUSRPRF command will no longer re-enable a NetServer user profile that has become disabled.

The following options are available for re-enabling a disabled NetServer user profile:

1. Using System i Navigator from the PC

System i Navigator can be used to display a list of user profiles that are disabled for NetServer use. Access the Disabled User IDs screen by taking the following steps in System i Navigator:

Note: The Network option in System i Navigator is not installed by default. This method of displaying disabled user IDs is only available in System i Navigator if the Network function is installed on the PC.

- a. Expand My Connections.
- b. Expand the connection name for the system.
- c. Expand Network.
- d. Expand Servers.
- e. Select TCP/IP.
- f. A list of TCP/IP Servers will be displayed on the right side of the System i Navigator window. Right-click on i5/OS NetServer and select the option for Disabled User IDs.

Profiles may be re-enabled for NetServer use by clicking the Enable User ID button on this window.

2. Using the GO NETS Tool

GO NETS is shipped with IBM i operating system and is located in the QUSRTOOL library.

See the following URL for details on how to create the tool:

<http://www.ibm.com/systems/i/software/netserver/qusrtool.html>

The NETS menu (on current versions of the tool) will include Option 12, Work with Netserver Users (WRKNSVUSR), which displays a list of profiles which have been disabled for NetServer use (similar to the list in System i Navigator). Then use Option 7 to re-enable any profiles that are displayed.

3. Call the Change Server Information (QZLSCHSI) API directly

Create a user-defined program that calls the QZLSCHSI API interface using the ZLSL0200 request format.

Details on the API parameters and usage can be found in the IBM i Information Center.

<http://publib.boulder.ibm.com/infocenter/iseries/v6r1m0/index.jsp>

Network File System changes

The Network File System (NFS) server now runs as a threaded job. When the NFS server is started, two QNFSNFSD jobs are submitted to the QSYSWRK subsystem. One job is the threaded server job, which is used for access to all file systems except the Document library services file system (QDLS). The second job runs a single thread to service access to the QDLS file system.

The Start Network File System Server (STRNFSSVR *SVR) command automatically starts both jobs. The NBRSVR and NBRBIO parameters are ignored on the STRNFSSVR command.

NFS version 4 is not supported with the base i5/OS operating system. It will be enabled later in a PTF.

Network server host (NWSH) changes

You can now configure the hardware resource name in V6R1 by specifying the Network Server Host Port resource name, which is in the form CMNxx by default.

Note: This change is new. In V5R4, the hardware resource name was configured by specifying the Network Server Host Adapter resource name, which was in the form LINxx by default.

Existing Network Server Host Adapters device descriptions are *not* automatically reconfigured. You must manually reconfigure your NWSHs to point to the new resource name before you can use them.

Eight 6B01 Ethernet port resources will show up as “Not detected” under each Network Server Host Port. You can delete these resources by using the Hardware Service Manager. Object conversion is done automatically. You should be able to use existing objects from the previous release.

Network storage space changes

When storage spaces exist in user auxiliary storage pools (ASPs) or independent ASPs, the process to save the storage space in releases earlier than V6R1 requires two save operations. The first save operation is for the storage space directory that exists in the system ASP: '/QFPNWSSTG/MYDISK'. The second save operation is for the user-defined file system (UDFS) that exists in the user ASP or independent ASP: '/dev/QASP25/MYDISK.udfs'. In V6R1, a storage space is now a single complex object. Therefore, you would save a storage space in a single step by specifying only the storage space directory to be saved: '/QFPNWSSTG/MYDISK'. This save operation includes the data in the user-defined file system without user intervention.

To make a storage space a complex object that must be saved and restored as a single entity requires existing storage spaces to be altered. This step is done the first time the storage space is used by the system. The storage space directory is marked as a savable object, but all of the other objects that make up the storage space are modified to be nonsavable objects. Therefore, saving object '/QFPNWSSTG/MYDISK/QFPCONTROL' does not save any objects, because the file is marked as nonsavable. When saving object '/QFPNWSSTG/MYDISK' the directory and all the objects that make up that storage space are saved as a single complex object.

New message queue for electronic customer support messages

In releases before V6R1, messages sent from electronic customer support (ECS) programs were sent to the message queue QSYSOPR. In V6R1, the default message queue used by ECS programs is the message queue QSERVICE instead of QSYSOPR. In addition, you can use a new parameter (ECSMSGQ) on the Change Service Attributes (CHGSRVA) command to select the message queue to which ECS programs should send messages.

Optical media changes

In earlier releases, recordable CD optical media was reported as type *CD-ROM. In V6R1, recordable CD optical media is reported differently; for example, *CD-R or *CD-RW. Applications that use an outfile created by the Display Optical (DSPOPT) command might be affected, because new values might be returned for the Media Type field.

Performance changes

Collection Services changes

Maximum interval time for the Disk, IPCS, and IOP categories:

Collection Services has changed the maximum interval time for the Disk, Interprocess communication status (IPCS), and Input/output processor (IOP) categories to 5 minutes. Five minutes is already the fixed default, although you can set the maximum interval time to a larger value in *CUSTOM profile definitions. This limit enables you to use the default collection interval while still maintaining a maximum five-minute interval time. When you set the default interval to an interval less than 5 minutes, you get this data collected consistent with other categories.

- Collection Services starts automatically at IPL and database files are created concurrent with data collection. Previously, this was the typical system behavior if Performance Management (PM) Agent was allowed to run. Now Collection Services will start regardless of what the PM Agent setting is.
- The automatic deletion of management collection objects when expired is now only done for those objects located in the configured collection library. This change solves customer problems related to having these objects deleted when the intent was to save them by moving to another library and system.
- For new system installs, there are default configuration changes that increases the amount of storage used:
 - The Management Collection Object retention is changed to 5 days.
 - Database files are created in the background.

Collection of disk performance data:

The collection of disk performance data was enhanced in i5/OS V6R1, which allows the performance reports to show more accurate data. Specifically, disk service time and disk wait (or queue) time will be much more accurate, and the new values might be considerably different than the values reported in previous releases. In releases before V6R1, these numbers were estimated based on a simplified statistical model. In V6R1, these values are based on directly-measured data.

Performance database file (QAPM*) field definition changes

- | The QAPM* performance database files have some fields that contain hexadecimal data and are defined with data types of either character or hexadecimal. SQL processing of these files, depending on settings, can result in undesired CCSID transformations on this data.
- | To prevent the CCSID transformation, the affected fields have been redefined as binary character fields.

Performance data collector changes

Low-level system support for program call tracing, instruction tracing, and performance data collecting is consolidated under a single point of control. Program creation on V6R1 enables this control, by default, in most procedures of a program object. Users benefit because they can collect Performance Data Collector (PDC) entry and exit events from existing applications and new programs, which run on V6R1, without needing to re-create special versions, as is required on earlier releases.

Because performance measurement code is present in most procedures, by default, the Enable performance collection (ENBPFCOL) option is ignored when creating modules and programs on V6R1 or later. The ENBPFCOL option is found on compiler commands and the CHGMOD, CHGPGM, and CHGSRVPGM commands.

To enable performance collection for all procedures in a program, which was the effect of option ENBPFCOL(*ALLPRC), you can create your program with optimization level 30 or less.

Beginning with V6R1, the **CallTracingAtHighOpt** Licensed Internal Code option is ignored during module creation, because this is now the default.

Performance Explorer changes

In V6R1, the following special values for the Program events (PGMEVT) parameter of the Add Performance Explorer Definition (ADDPEXDFN) command are not supported:

- *MIPRECALL
- *MIPOSTCALL
- *JVAPRECALL
- *JVAPOSTCALL

The Convert Performance Collection (CVTPFRCOL) command now converts PEX data from N-2, N-1 to release N. For example, you can convert PEX data from V5R3 or V5R4 to the V6R1 format.

Object keys and statement numbers for 16 levels of the stack are saved for some events, increasing the required database file storage size by 50% in some cases.

Performance Management Agent

Performance Management (PM) Agent has a significant change regarding the contact information. You must provide the mandatory customer information before being able to transmit your data back to IBM.

Performance Management Agent command changes

- | The Configure PM Agent Line commands (CFGPMLIN) and (PMLINMON) will not be shipped with the next release of IBM i. Any CL program that includes these commands might need to be altered for any release after V6R1.

Power-down system confirmation panel

By default, the Power Down System (PWRDWN SYS) command will display a confirmation screen when the command is executed in an interactive job. If you want your system to behave as it did before V6R1 and not display this confirmation screen, use this command:

```
ADDENVVAR ENVVAR(QIBM_PWRDWN SYS_CONFIRM) VALUE(*NO) LEVEL(*SYS)
```

*ENVVAR is the default value for the CONFIRM parameter. If the environment variable is not defined (or does not contain *YES, *NO, or *INTERACT), *INTERACT will be used in V6R1, and a confirmation will be displayed in an interactive job. Before V6R1, *NO was used as the CONFIRM default value and no confirmation was displayed.

Program conversion

Conversion of module, program, and service program objects

The conversion from software-based process local addressing to hardware-based process local addressing requires that all modules, programs, and service programs created on releases before V6R1 be converted to be used on V6R1 and later releases.

The conversion will take some time. There are three different options for when conversions can occur:

- You can convert objects as they are restored onto a V6R1 system. This type of conversion slows down the restore operation, but eliminates the need for any later conversions.
- You can use the Start Object Conversion (STROBJCVN) command to explicitly convert objects at some point after the restore operation, but before calling the objects for the first time.
- You can convert an object the first time that it is called, which can cause some delay until all objects have been converted. This option does not require any explicit conversion actions.

For any of the conversion methods, the length of time it will take for a particular object to be converted depends on several factors. These factors include the size of the object to be converted, and the processor speed and memory characteristics of the system on which the conversion occurs.

Modules, programs, and service programs created for releases before V5R1 might have their creation data completely removed. These pre-V5R1 objects with no creation data can be restored on V6R1, but cannot be converted, so they cannot be used on V6R1.

For more information, see “Program conversion” on page 3.

Conversion impacts on digital signatures

Programs, service programs and module objects created on releases before V6R1 that are digitally signed and then restored onto a V6R1 or later system will be treated as unsigned. Pre-V6R1 signatures are not recognized on V6R1 and later releases for these object types. Similarly, programs, service programs and module objects created and digitally signed on V6R1 and later releases will be treated as unsigned if restored on a release before V6R1.

Moving programs, service programs and modules created on releases earlier than V6R1 to V6R1 and later releases will convert these objects to a format compatible with V6R1 before the objects can be used. Programs, service programs and modules created on V6R1 and later releases that are restored to a release before V6R1 will be converted to pre-V6R1 format. The conversion process will cause all digital signatures on these objects to be removed.

Adaptive code generation

Beginning with V6R1, the MI (Machine Interface) architecture creates module, program, and service program objects that take advantage of the latest hardware features of the system on which they are created or converted. This policy allows unique features available on new POWER6 processors, for example, to be used immediately. Before this policy change, the system did not use new processor features until all machine models supported by the current release of the operating system had the new feature.

Programs that use the latest hardware features might require conversion when taken to a different system at the same release level, but running on an older generation of the hardware. Even though such conversions are performed automatically as needed, the conversion process removes digital signatures and uses machine resources similar to those required when creating the object.

New options are available in V6R1 to help you override the default policy if you want to avoid conversion when programs are deployed to different machine models. For example, to create module objects that do not use unique features, you can use `LICOPT("CodeGenTarget=Common")` on the module creation command, or set the environment variable `QIBM_BN_CREATE_WITH_COMMON_CODEGEN` to a value of 2 before creating module objects. For details, refer to “Adaptive Code Generation” in the “Advanced Optimization Techniques” chapter of the ILE Concepts publication.

Special situations for program conversion

| Null pointer values might no longer be produced as the result of explicit or implicit pointer arithmetic on
| a null pointer value. *Explicit pointer arithmetic* is when a value is added to a pointer in a program
| statement. An example of *implicit pointer arithmetic* is when a pointer references a field that is not the first
| field defined in a data structure. When arithmetic is performed on a pointer that contains a null pointer
| value, an exception might not be signaled, and the pointer result might no longer compare equal to null.
| Subsequent attempts to use the modified pointer value might cause an MCH3601 (pointer does not exist)
| exception or an MCH0601 (space addressing violation) exception. However, if a very large value,
| requiring an 8-byte binary field to hold it, is explicitly added to a null pointer value, then a subsequent
| attempt to use the resulting pointer might not cause an exception. In that case, the pointer reference is
| within the job's teraspace. Refer to “Example: RPG program with implicit pointer arithmetic on a null
| pointer value” on page 50 and “Example: COBOL program with implicit pointer arithmetic on a null
| pointer value” on page 51 for program examples.

An operation that copies data from one location to another under conditions that cause an MCH0601 (space addressing violation) exception might not copy any of the data in situations where a partial copy might have occurred on a previous release.

It is now possible for bound programs that monitor for exception MCH3601 (pointer does not exist) while using the SCANX, TESTPTR or SETSPFP built-in instructions to receive the exception twice for one exception condition, if the exception handling action passes control to the next hardware instruction.

Programs that erroneously attempt to use an integer value-preserving pointer (a space pointer assigned an integer value) to reference storage might not receive an MCH3601 (pointer does not exist) as would likely have occurred on previous releases. Instead, a variety of exceptions might be signaled, most commonly MCH0601 (space addressing violation), or no exception may be signaled. If no exception is signaled, the reference will appear to succeed, even though references using these pointers are not defined. Any such reference will be within the job's teraspace.

Example: RPG program with implicit pointer arithmetic on a null pointer value

The following is an example of an RPG program that implements implicit pointer arithmetic on a null pointer value.

```
|      D a                ds                based(p)
|      D b                10a
|      D c                10a
|      D p2               s                *
|      D basedp2         s                10a  based(p2)
|      /free
|          dsply ('Subfield is at position 1, offset zero');
|          p2 = %addr(b);
|          exsr chkPointer;
|          dsply ('Subfield is at position 11, offset 10');
|          p2 = %addr(c);
|          exsr chkPointer;
|          return;
|
|          begsr chkPointer;
|          if p2 <> *null;
|              dsply 'pointer is not null';
|              basedp2 = 'abc';
|          else;
|              dsply 'pointer is null';
|          endif;
|      endsr;
```

When compiled into a program and called on a V5R4 system, the source code produces the following output:

```
| DSPLY Subfield is at position 1, offset zero
| DSPLY pointer is null
| DSPLY Subfield is at position 11, offset 10
| DSPLY pointer is null
```

When compiled into a program and called on a V6R1 system, the source code produces the following output:

```
| DSPLY Subfield is at position 1, offset zero
| DSPLY pointer is null
| DSPLY Subfield is at position 11, offset 10
| DSPLY pointer is not null
| Pointer or parameter error (C G D F).
```

Example: COBOL program with implicit pointer arithmetic on a null pointer value

The following is an example of a COBOL program that implements implicit pointer arithmetic on a null pointer value.

```
IDENTIFICATION DIVISION.
PROGRAM-ID. VERIFYNULL.
DATA DIVISION.
LINKAGE SECTION.
01 A.
   05 B PIC X(10).
   05 C PIC X(10).
01 basedp2 PIC x(10).
WORKING-STORAGE SECTION.
01 P usage pointer.
01 P2 usage pointer.
*
PROCEDURE DIVISION USING A.
TEST1-INIT.
   display "Subfield is at position 1, offset zero".
   set P2 to address of B.
   perform chkPointer.
   display "Subfield is at position 11, offset 10".
   set P2 to address of C.
   perform chkPointer.
   goback.
chkPointer.
   IF P2 NOT = NULL
      DISPLAY "POINTER IS NOT NULL"
      set address of basedp2 to P2
      MOVE "abc" to basedp2
   ELSE
      DISPLAY "POINTER IS NULL"
END-IF.
```

When compiled into a program named "verifynull" and called on a V5R4 system, the source code produces the following output:

```
Subfield is at position 1, offset zero
POINTER IS NULL
Subfield is at position 11, offset 10
POINTER IS NULL
```

When compiled into a program named "verifynull" and called on a V6R1 system, the source code produces the following output:

```
Subfield is at position 1, offset zero
POINTER IS NULL
Subfield is at position 11, offset 10
POINTER IS NOT NULL
Pointer not set for location referenced.
Function check. MCH3601 unmonitored by VERIFYNULL
at statement 0000000036, instruction X'0000'.
Message 'MCH3601' in program object 'VERIFYNULL'
in library 'YOURLIB' (C D F G).
```

Pre-V6R1 PTFs for restoring modules, programs, or service programs

Information APAR II14370 contains information about PTFs needed for releases before V6R1 to restore modules, programs or service programs that were created or converted on V6R1.

PTF group changes

The system can now keep more than one level of a PTF group on the system. The default number of levels of each PTF group that is kept on the system is 2. You can change this value by using the Change Service Attributes (CHGSRVA) command. All PTF group commands and APIs have new parameters for managing which level of the PTF group to use.

The Work with PTF Groups (WRKPTFGRP) print panel has been changed. The PTF group name was reduced from displaying the entire PTF group name (30 characters) to displaying the first 18 characters of the name. This format now matches the interactive display format.

New values for the status of a PTF group are returned. The new status values are shown by the Work with PTF Groups (WRKPTFGRP) command as well as the List PTF Groups (QpzListPtfGroups) and List PTF Group Details (QpzListPtfGroupDetails) APIs.

PTF groups for previous operating system releases with a status of Not Applicable are deleted during i5/OS upgrade and during Operational Assist (OA) cleanup.

PTF media changes

A new value (*DVDROM) has been added in the Contact Information for getting PTFs in physical media. The *AUTOMATIC selection on distribution media for PTFs has changed to default *DVDROM if the system does not detect any device.

Remote Command and Distributed Program Call server exit program changes

The Remote Command and Distributed Program Call server exit point, QIBM_QZRC_RMT, is used for both remote command requests and distributed program call requests. For remote command requests, in V5R4 and earlier releases, there are two consecutive reserved fields. These fields total 20 bytes of reserved space. In V6R1, these reserved fields are replaced by a 4-byte CCSID value, followed by 16 bytes of reserved space, for a total of 20 bytes. The CCSID value represents the CCSID of the command string. The following are valid values:

- 0 (job CCSID)
- 1200 (UTF-16)
- 1208 (UTF-8)

If you use a common user exit program for multiple releases, you might need to make changes to handle the CCSID value.

Secure Sockets Layer changes

The Secure Sockets Layer version 2.0 protocol (SSLv2) is now disabled by default for system Secure Sockets Layer (SSL).

System SSL is the i5/OS Licensed Internal Code implementation of SSL functionality. It is tightly coupled with the operating system and the sockets code specifically providing extra performance and security.

System SSL is available to application developers using two different programming interfaces and one JSSE implementation:

- Global Secure Toolkit (GSKit) APIs, including ILE C APIs accessible from other ILE languages.
- Native i5/OS SSL APIs, including ILE C APIs accessible from other ILE languages. (This API set is not recommended; use GSKit instead.)
- Native i5/OS JSSE implementation.
- The shipped default JSSE implementation for JDK 1.4.

SSL applications created by IBM, IBM business partners, independent software vendors (ISV), or customers that use one of these interfaces to System SSL are affected. FTP and Telnet are examples of IBM applications that use system SSL.

You can re-enable SSLv2 by changing the new system value QSSLPCL. For additional information, see the SSL topic in the i5/OS Information Center.

The SSL default cipher specification list has changed for system SSL. The default cipher specification list for system SSL no longer contains any Transport Layer Security (TLS) Protocol Version 1 or SSLv3 ciphers that use a key with less than 128 secret bits. As the SSL Version 2 (SSLv2) protocol is disabled, the default cipher specification list also no longer contains any SSLv2-only ciphers. The Advanced Encryption Standard (AES) 128-bit cipher is first in the ordered default cipher specification list; previously it was third in the list.

The ciphers that are no longer in the default cipher specification list are still supported by System SSL for applications that code to specifically use those ciphers.

You can control the ciphers that system SSL supports by using the new system values QSSLCSL and QSSLCSLCTL. Although you cannot control the default cipher specification list, you can alter the order of ciphers indirectly by changing the order of the ciphers contained in the QSSLCSL system value. For additional information, see the SSL topic in the i5/OS Information Center.

The default cipher specification list for system SSL has been updated as follows:

- *RSA_AES_128_CBC_SHA
- *RSA_RC4_128_SHA
- *RSA_RC4_128_MD5
- *RSA_AES_256_CBC_SHA
- *RSA_3DES_EDE_CBC_SHA

Server firmware PTFs on systems running V5R4M5 or later releases

- | The Server IPL Required field on the QPZRTVFX and QpzListPTF APIs is not applicable to server firmware PTFs on systems running Licensed Internal Code V5R4M5 or later.

For POWER5 and POWER6 hardware running V5R4M5 or later releases, server firmware PTFs will no longer be associated with Product ID 5761-999 (Licensed Internal Code). The PTFs will be associated with a new Product ID, 5733-9xx, where xx depends on the hardware platform. The operational changes to specify a different Product ID, 5733-9xx, for installing and displaying server firmware PTFs will be primarily for individual PTF processing when the server firmware is managed by the operating system. The procedures to install a cumulative PTF package containing server firmware PTFs will not change. A new DSPFMWSTS (Display Firmware Status) command is provided to display the status of the server firmware.

The process for distributing server firmware products and PTFs to systems of a different hardware platform (different firmware) will change. A new CRTFMWPRD (Create Firmware Product) command is provided to create a server firmware product on a system for a different hardware platform. Instead of saving and restoring the server firmware with the Licensed Internal Code, the server firmware PTFs will be saved and restored with the server firmware product (5733-9xx).

Shared data pool changes

A V5R4 Add function PTF SI23027 was created that enabled you to manually configure a subsystem and allocate a private storage pool to use for iSCSI network servers.

If you are using shared data pools, you will need to change your configuration to allocate a shared data pool using the WRKSHRPOOL command. You will need to change network server descriptions (NWSDs) to use the created shared data pool.

Because the QGPL/QFPHIS subsystem description is no longer used, you should remove it.

Existing NWSDs will use the default *BASE memory pool until these changes are made.

Spooled file changes

Start Spool Reclaim (STRSPLRCL) command:

The Start Spool Reclaim (STRSPLRCL) CL command has replaced the QSYS/QSPFIXUP program. This program will remain on the system. Calling this program results in error message CPF338E, which points you to the new command.

Output queue conversion:

During the V6R1 installation IPL, all output queues are deleted and recreated. The object associated space will remain the same, but the creation date and time on the output queue object will change. This function reclaims the storage used by old output queue entries. This conversion will not result in any compatibility issues.

Work with Document Print Queue (WRKDOCPRTQ) command:

The Work with Document Print Queue (WRKDOCPRTQ) command no longer lists jobs on job queues. It still lists printer output on output queues. This was first announced in the Memo to Users in V5R4.

Updating the Last Changed Date/Time of output queue objects:

The Last Changed Date/Time of the output queue object is no longer updated when a spooled file is added or removed from an output queue, because the output queue object is no longer physically accessed during these operations. The Last Used Date and Days Used Count will continue to be updated for these operations.

Delete Expired Spooled files (DLTEXPSPLF) command:

When upgrading to V6R1, when you use the Delete Expired Spooled files (DLTEXPSPLF) CL command, you must have *USE authority to all auxiliary storage pool (ASP) devices in an ASP group. This change is being done because requiring only *EXECUTE to the ASP group is inconsistent with the rest of the i5/OS security requirements for use of an ASP group.

Define or Change the System at IPL:

The display commands used during IPL from the Define or Change the System at IPL menu no longer produce printed output.

Starting and ending TCP/IP

With this release, a new system job named QTCPCTL is started during IPL of i5/OS. This job coordinates and processes requests to start or end TCP/IP interfaces, and is persistent from IPL to power down. A second system job QTCPWRK is started as part of IPL processing of i5/OS. This job ends when TCP/IP is ended and is started when TCP/IP is started.

In previous releases, these functions were performed by the QTCPPIP and QTCPMONITR jobs in the QSYSWRK subsystem. These two jobs started and ended along with TCP/IP. Taking i5/OS to restricted

state automatically ends all active TCP/IP interfaces and the TCP/IP protocol stack (as before). With this release, ending the QSYSWRK subsystem does not end the active TCP/IP interfaces. Messages, which were sent to the job logs of the QTCPIP and QTCPMONITR jobs, are now sent to the job logs of the QTCPCTL and QTCPWRK system job logs.

| TCP/IP messages, which are sent to the history log, might now appear after CPF0968, System ended to
| restricted condition.

With V6R1, the STRTCP command now allows the command to be run while TCP/IP is active. In this case, the system performs the processing associated with any STRTCP parameter specified as *YES. For example, if TCP/IP is started with the command STRTCP STRSVR(*NO) STRIFC(*NO) STRTPPRF(*NO) STRIP6(*NO), the IPv4 protocol stack is started, but no TCP/IP servers, interfaces, point-to-point profiles, or the IPv6 protocol stack are started. Subsequently running the command STRTCP STRSVR(*YES) STRIFC(*YES) STRTPPRF(*NO) STRIP6(*YES) starts the IPv6 protocol stack and all TCP/IP servers and interfaces configured as AUTOSTART(*YES).

Note: If the TCP/IP protocol stack was not started when TCP/IP was started, you can start the IPv6 part of the TCP/IP protocol stack by issuing another STRTCP CL command and specifying the STRIP6 parameter as (*YES). This action applies to both normal operating mode and restricted state.

Subsystem monitor job changes

If you use the QSYS/QWTSBSIASP data area to run a subsystem monitor job with an auxiliary storage pool (ASP) group in its library name space, you must use the ASPGRP parameter on the Create Subsystem Description (CRTSBSD) or Change Subsystem Description (CHGSBSD) command instead. In addition, you must ensure that the job descriptions, which the subsystem uses when creating user jobs, specify an Initial ASP Group (INLASPGRP) parameter that matches the ASPGRP parameter of the subsystem description. The data area is ignored when starting the subsystem. The ASP group for the user job overrides the ASP group of the subsystem.

System value changes

QLMTDEVSSN system value:

When the QLMTDEVSSN system value is changed, message CPF1806 is sent to the history log (QHST) instead of message CPF1815.

QLOCALE system value:

When the QLOCALE system value is changed, message CPF18C9 is sent to QHST instead of message CPF1806.

Changes to default values:

The shipped defaults for system values QACTJOB, QADLACTJ, QADLTOTJ, QTOTJOB, and QUPSDLYTIM have changed. This change affects existing systems only if the system value object needs to be rebuilt or if the system value object on the install media is used.

Telnet changes

New parameters have been added to the user exit interface for the Telnet server to hold the longer internet addresses required by IPv6 clients and servers. The current documentation for the Telnet server user exit (INIT0100 format), Format of Connection Description Information, states that an IP address is provided in the client and server internet address fields.

However, if the remote (connecting) client is an IPv6 client, the internet address no longer fits in the IPv4 Client internet address field. Therefore, the field is filled with nulls (hex '00') instead. In addition, the server (local) internet address no longer fits in the IPv4 server (local) internet address field and also is filled with nulls, when the remote client is an IPv6 client. You will need to access these longer internet addresses using the new IPv6 Client internet address, and IPv6 Server (local) internet address fields.

Teraspace changes

Storage for a teraspace-shared memory segment

The storage for a teraspace-shared memory segment is now allocated when it is accessed. As a result, the storage is charged against the temporary storage limit for the job that accesses the storage. If the job does not have enough temporary storage to satisfy the request, the job is ended. Previously, the storage for a teraspace-shared memory segment was allocated when it was first attached to a job or when it was resized after being attached to a job. As a result, the job performing the operation was charged for all of the storage. If the job didn't have enough temporary storage to satisfy the attach/resize request, the `shmat()` and `shmctl()` APIs failed with an ENOMEM error.

The `shmat()` API attaches to teraspace-shared memory segments using a new internal alignment factor. This change might result in a different address being returned by the API than would have been returned in a previous release. As a result, callers of the API should not make any assumptions about the address at which a shared memory segment gets attached. In particular, the `shmaddr` parameter on the API is the address at which the calling thread wants the shared memory segment attached, but has never been guaranteed to be granted by the system.

Teraspace mutex changes

Unless they are explicitly destroyed, pointer-based mutexes that are created in teraspace exist for the duration of their lifetime scope. In earlier releases, the lifetime scope of these mutexes was determined by the lifetime of the creating process or the lifetime of the teraspace in which they were created, depending on the attributes of the teraspace. The lifetime scope of these mutexes is now determined only by the lifetime of the teraspace in which they are created.

Time management changes

Changes for file and volume time values recorded on UDF optical media

Media formatted with Universal Disk Format (UDF) supports time zones. Before V6R1, i5/OS stored the file and volume dates on optical media using machine time, with no reference to the time zone. Beginning in V6R1, Coordinated Universal Time (UTC) will be recorded on optical media with the UDF media format. Times stored in releases before V6R1 will be interpreted as UTC. As a result, those times stored in previous releases on UDF media might be off by up to 23 hours. Optical media with the High Performance Optical File System (HPOFS) media format is not affected by this change.

Time zone changes

Operational procedures:

Non-Gregorian users using the time zone object year offset value will notice that the calculations for day-of-week in a non-Gregorian year will now match the correct day-of-week, as determined in the current Gregorian year.

Based on the Time zone (QTIMZON) system value, the Leap year adjustment (QLEAPADJ) system value is managed by the system and is now read-only. Any attempt to change this system value will result in a CPF1030 diagnostic message.

Internal time values for jobs and spooled files are now in UTC. If you change the Time zone (QTIMZON) system value or change the time zone object specified by the QTIMZON system value, the stored UTC values are not affected. These values are displayed or retrieved in whatever local system time is in effect when the values are displayed or retrieved. Before V6R1, these values were stored in local system time and the values displayed or retrieved showed the local system time in effect when the value was stored.

Internal timers used by the operating system are now in UTC. A change to the Time zone (QTIMZON) system value or a change to the time zone object specified by the QTIMZON system value does not affect a UTC-timed wait. This change affects the Resume job time (RSMTIME) parameter of the Delay Job (DLYJOB) command, the Schedule date (SCDDAT) and Schedule time (SCDTIM) parameters on the Submit Job (SBMJOB) command, and the next job submission performed by the QJOBSCD system job.

Times specified on the Delay Job (DLYJOB) command, the Submit Job (SBMJOB) command, the Add Job Schedule Entry (ADDJOBSCDE) command, and the Change Job Schedule Entry (CHGJOBSCDE) are still in local system time. However, if you specify a time that is skipped by a transition into daylight savings time, the result might not exactly match previous behavior and the order for multiple operations is not preserved.

Semantics of use:

The Timed power-on value that is returned using the Materialize Machine Attributes (MATMATR) MI instruction option 013C has changed from local system time (LST) to Coordinated Universal Time (UTC).

The TZ string that is returned for the Retrieve Time Zone (qwcrvtz) API for format RMTZ0200 is changing because of evolving industry standards. The previous interpretation of the industry standard allowed non-valid characters to be substituted in the Standard Time and Daylight Saving Time abbreviated names with “*”. The new standard disallows the “*” and the API now returns an “x” for non-valid characters. Also, the abbreviated names are quoted between < and >.

The text messages (CPX092E and CPX092F) that currently have AST (Alaska Standard Time) and ADT (Alaska Daylight Time) have changed to have AKST and AKDT for their abbreviated names. The IBM-supplied QN0900AST time zone uses these messages.

Setting a year offset during installation:

If you want to define the year offset for a system before upgrading from one release to another, create a data area named QWCYEAROFS in the QSYS library. The QWCYEAROFS data area must be a character data area with a minimum length of 14 characters. The format is defined as follows:

Byte 1 Sign for current year offset (either '+' or '-').

Bytes 2 - 4

Year offset from Gregorian (000 to 140)

Bytes 5 - 14

Time zone description name, left justified and padded on the right with blanks. This is the name of the time zone that is to be changed with the specified year offset during the installation IPL. You can use a special value of *SYSVAL to indicate that the current time zone object specified in the Time zone (QTIMZON) system value is to be changed. If a name is specified, a check is made to make sure it exists. The time zone description is changed with the year offset, and QTIMZON is changed to this specified time zone. To prevent future release upgrade problems, it is recommended that you specify a user-defined time zone object, and not one of the IBM-supplied time zone objects.

All IBM-supplied time zone objects have a year offset of 0.

Messages with blank dates and times

You can use new message data types (*UTC, *UTCD, *UTCT) in message descriptions in V6R1. Some IBM-supplied messages have changed to use the new UTC message data types. For example, CPF1124 (Job started) and CPF1164 (Job ended) have changed. When the history log is viewed after installing V6R1, the dates and times in these messages will have blank dates and times for messages that were sent before the installation. If you view messages with the Display Log (DSPLOG) CL command, pressing F1 with the cursor under a message shows the date and time that the message was sent. Similarly, after the installation, if a job log from before the installation contains one of the changed IBM messages, message text for dates and times might show blanks. Blanks appearing in the message text is a temporary annoyance around the installation time, but messages and job logs generated after the install no longer show the blanks.

| **Watch for event support changes**

| **Watch for event changes for CPF1124 and CPF1164**

| With PTF SI42126 for V6R1, when either CPF1124 (Job started) or CPF1164 (Job ended) are specified as a message id to watch, if those messages occur in job QSCWCHPS, which is a job that is used to process watch events, they will not trigger the user exit program. CPF1124 and CPF1164 will be ignored in jobs that process watch events.

| **Watch for event user exit program job changes**

| In V6R1, when a user exit program is called because a watched event occurs, the exit program will now run in a prestart job named QSCWCHPS in subsystem QUSRWRK. In V5R4, a unique job was setup for each watch that was started and the user exit program would run in the job named the same as the watch session id. For example, when Service Monitor started, it would start over 70 watches so this would result in over 70 jobs being created, (for example. job name SRVMON0001), and the exit program for each watch would be called in the appropriate job. In V5R4, the watch processing jobs would exist until the watch was ended or the job was ended, compared with V6R1 where the QSCWCHPS job ends after the exit program is called and a new QSCWCHPS job is created for subsequent watch processing. Starting in V6R1, there will no longer be a job started for each watch session id.

Chapter 3. Options

This chapter describes changes to options of the i5/OS operating system.

User Tools (QUSRTOOL) (Option 7)

Most of the cluster resource services commands that had been in the QSYS library in V5R4 are now in User Tools (QUSRTOOL), Option 7. QUSRTOOL now contains the V5R4 source code for these commands. For information about how to create the commands, see the member TCSTINFO in the file QUSRTOOL/QATTINFO, and the **High availability function in the base operating system** topic in the information center.

NetWare Enhanced Integration (Option 25)

Support for option 25 of i5/OS removed

In V6R1, support for NetWare Enhanced Integration (i5/OS option 25) is removed. The following CL commands and APIs that were part of option 25 are no longer supported.

Commands:

- ADDNTWAUTE
- CHGNDSCCTX
- CHGNTWAUTE
- CHGNTWVOL
- CRTNTWVOL
- DLTNTWVOL
- DSPNDSCCTX
- DSPNTWAUTE
- DSPNTWCNN
- DSPNTWVOL
- ENDNTWCNN
- RMVNTWAUTE
- STRNTWCNN
- VFYNTWAUTE
- WRKNTWAUTE
- WRKNTWCNN
- WRKNTWVOL

APIs:

- QfpzAddNtwAutE
- QfpzChgNtwAutE
- QfpzEndNtwCnn
- QfpzListNtwAutE
- QfpzRmvNtwAutE
- QfpzStrNtwCnn
- QfpzVfyNtwAutE

The following i5/OS commands no longer support *NETWARE for the TYPE parameter:

- Add Mounted FS (ADDMFS or MOUNT)
- Remove Mounted FS (RMVMFS or UNMOUNT)

The following API no longer supports QNetWare for the file-system type information as input. It also has been removed from the current constant which represents all dynamically-mounted file systems (which includes QNetWare):

- Perform File System Operation (QP0LFLOP) API

Note: Because the dynamically-mounted file system constant is being updated, anyone who uses that constant should recompile in V6R1; otherwise they will receive an error.

APIs that supported QNetWare file system on releases before V6R1 no longer support QNetWare.

Other CL commands that supported values specific to NetWare have been changed to remove all NetWare support, such as the following commands:

- Change NWS Attributes (CHGNWSA)
- Change NWS User Attributes (CHGNWSUSRA)
- Display NWS Attributes (DSPNWSA)
- Display NWS User Attributes (DSPNWSA)
- Submit Network Server Command (SBMNWSCMD)
- Work with Network Server Status (WRKNWSSTS)
- Work with NWS User Enrollment (WRKNWSENR)

Integrated Server Support (Option 29)

Install Windows Server (INSWNTSVR) command changes

The Install Windows Server (INSWNTSVR) command has a number of parameter changes that might require you to recompile any existing CL programs using these parameters.

- The default value for the Synchronize date and time (SYNCTIME) parameter has changed to *NONE. All currently supported integrated servers have battery-backed CMOS date and time support built into the hardware. Thus, using this parameter is virtually unnecessary or potentially confusing to applications when alternative time synchronization mechanism are also used, such as synchronization with an Active Domain controller.
- The Shutdown TCP port (SHUTDPORT) is an obsolete parameter that was never used. It has been removed from the command.
- The default value for the Initialize service processor (INZSP) parameter has changed to *NONE to be consistent with currently supported service processor hardware configurations.
- Enable unicast (ENBUNICAST) parameter has changed its default to *YES to match supported hardware configurations.
- Enclosure identifier (EID) parameter has changed its default to *AUTO to be consistent with the default ENBUNICAST parameter and recommended usage.
- The Default IP security rule (DFTSECRULE), IP security rule (IPSECRULE) and SP certificate identifier (SPCERTID) parameters were removed. These parameters were restricted to only allowing the default value to be specified in V5R4.

Install Linux Server (INSLNXSVR) command

- | In the next release of IBM i, the Install Linux Server (INSLNXSVR) command will be removed. Any
- | existing CL programs that use the INSLNXSVR CL command may need to be modified.

In V6R1 the Install Linux Server (INSLNXSVR) command has a number of parameter changes that might require you to recompile any existing CL programs using these parameters. These changes include iSCSI implementation changes, the removal of support for all Linux operating systems on the IXS/IXA platforms, and the removal of support for SLES9 on the iSCSI platform. For further details of the removal of support for Linux operating systems versions, please refer to the text in the 5761-LSV product section.

- Linux server distribution (LNXSVRDST) parameter values *RHEL3, *RHEL4, *SLES8, *SLES9 have been removed.
- Only iSCSI-attached integrated servers can be installed using the Install Linux Server (INSLNXSVR) command in this release. The Resource name (RSRCNAME) parameter was removed.
- No supported Linux distributions can be installed using the Linux source directory (LNXSRCDIR) parameter. This parameter was removed.
- The default value for the Synchronize date and time (SYNCTIME) parameter was changed to *NONE. All currently supported integrated servers have battery-backed CMOS date and time support built into the hardware making this support unnecessary.
- The Shutdown TCP port (SHUTDPORT) is no longer needed and has been removed from the command.
- The default value for the Initialize service processor (INZSP) parameter was changed to *NONE to be consistent with currently-supported hardware configurations.
- The default value for the Enable unicast (ENBUNICAST) parameter was changed to *YES to match currently-supported hardware configurations.
- The default value for the Enclosure identifier (EID) parameter was changed to *AUTO to be consistent with the default ENBUNICAST parameter and recommended usage.
- The Default IP security rule (DFTSECRULE), IP security rule (IPSECRULE), and SP certificate identifier (SPCERTID) parameters were removed to simplify the command. These parameters had been restricted to only allowing the default value to be specified in V5R4. They have been removed because they do not allow any other value to be specified.

Linux file-level backup

Using Linux file-level backup requires Linux iSCSI and Linux logical partition (LPAR) systems to require the following i5/OS products to be installed:

- 5761-SS1 Option 12 - Host Servers
- 5761-SS1 Option 29 - Integrated Server Support
- 5761-SS1 Option 34 - Digital Certificate Manager
- 5761-SS1 Option 35 - CCA Cryptographic Service Provider
- 5761-LSV - Extended Integrated Server Support

Qshell (Option 30)

ipcs utility

The Qshell **ipcs** (Report interprocess communication status) utility uses more blank spaces to pad the SEGSZ extended description output data. This is required to display the size of 64-bit shared memory segments correctly.

Qshell PATH changes

Before V6R1, if the PATH variable was not set, a default Qshell PATH of **/usr/bin:** was used. Starting in V6R1, the default Qshell PATH is **/usr/bin.:/QOpenSys/usr/bin.**

i5/OS Portable Application Solutions Environment (Option 33)

i5/OS PASE support for ioctl changed

The i5/OS PASE **ioctl** function is changed to include IPv6 socket addresses in results when called with the **SIOCGIFCONF** command. IPv6 interfaces on lines that do not have an associated IPv4 interface

return the line-description name in field **ifr_name** (instead of an address in dotted-name form). **ioctl** is also enhanced to support the CSIOCGIFCONF, OSIOCGIFCONF, and SIOCGSIZIFCONF commands, with results equivalent to AIX®.

For more information about these commands, see the “**ioctl** Socket Control Operations” topic in the “Technical Reference: Communications, Volume 2,” in the AIX information center. The i5/OS PASE commands work in the same way as on AIX.

i5/OS PASE support for real-time semaphore interfaces

i5/OS PASE is enhanced to support runtime functions **sem_open**, **sem_close**, **sem_unlink**, **sem_init**, **sem_destroy**, **sem_getvalue**, **sem_post**, **sem_wait**, **sem_trywait**, and **sem_timedwait** (using the same system support as like-named ILE functions).

i5/OS PASE support for execution priority changed

The mapping between “nice” values and i5/OS RUNPTY is changed so a single unit of “nice” always maps to two units of RUNPTY (instead of a variable mapping to 1-4 units of RUNPTY). This causes the defaults for interactive job use of the i5/OS PASE “nice” command to run at RUNPTY(40) instead of RUNPTY(60). Clients that want “nice” to run work at batch RUNPTY(50) must specify a “nice” value of 15.

Execution priority values returned by i5/OS PASE runtime functions (such as field **pi_pri** in structures returned by **getprocs**) are now RUNPTY values. This causes the i5/OS PASE **ps** command to show RUNPTY values under the heading “PRI”.

i5/OS PASE applications can now control thread execution priority using **pthread_setschedparam** or an attributes object passed to **pthread_create**. i5/OS restricts thread priority so applications can only set priority values that make a thread equal to or less favored than RUNPTY for the job

Chapter 4. Licensed programs

This chapter contains information about individual licensed programs that have been modified or changed for V6R1.

Backup Recovery and Media Services (5761-BR1)

WRKMEDBRM changes:

A change to the WRKMEDBRM panel shows a column that once was **Expire** but is now **Status**. The items that can show under that column now are *INZ, *EXP, *ACT, and *ERR.

Printer file changes:

BRMS no longer deletes printer file attributes when upgrading. Before V6R1, when you installed 5761-BR1, the printer files that existed on the system were deleted and replaced by the new shipped printer files. Now when you upgrade, BRMS saves the printer file attributes, installs the new files, and applies the saved attributes to the new printer files. BRMS preserves any changes made to the attributes before upgrading.

BRMS networking changes:

In V6R1, BRMS uses the SQL CLI function to find scratch media from other systems to use on a save. This requires that relational database directory (RDB) entries be created for each system in the network on each system in the network. For information on how to set up RDB entries on BRMS networked systems see :

<http://www.ibm.com/systems/i/support/brms/brmstcpip.html>

IBM DB2 Content Manager OnDemand for i5/OS (5761-RD1)

Print Report from OnDemand (PRTRPTOND) command

The Print Report from OnDemand (PRTRPTOND) command has had two keywords removed. The removed keywords are COPIES and PAGERANGE. These keywords had defaults and they would almost never have been specified. Because of how the printing was done, these keywords never functioned and they are being removed to eliminate confusion.

Start Disk Storage Management (STRDSMOND) command

The Start Disk Storage Management (STRDSMOND) command has had one keyword removed. The removed keyword is VALIDATE. This keyword had a default and would almost never have been specified. The functions run by the command are common code that runs on several platforms, the validate function enabled by the VALIDATE keyword is not needed on the i5/OS system. The keyword is being removed to eliminate confusion.

IBM DB2 Query Manager and SQL Development Kit for i5/OS (5761-ST1)

Create SQL Fortran Program command

In V6R1, the SQL precompiler support for the Create SQL Fortran Program (CRTSQLFTN) CL command has been removed. The command is deleted from the QSQL and QSYS libraries when the V6R1 release is installed.

IBM DB2 Extenders Version 9.1 for i5/OS (5761-DE1)

Perform the following steps after installing IBM DB2 XML Extender (5761-DE1, option 2):

1. You must install IBM XML Toolkit for System i (5733-XT2, option 1) before you can use DB2 XML Extender.
2. After you complete installing DB2 XML Extender over an earlier release of XML Extender, enter the following command to run the migration:

```
CALL QDBXM/QZXMMIGV
```

IBM Developer Kit for Java (5761-JV1)

IBM Technology for Java versions

The System i platform supports multiple versions of the Java Development Kits (JDKs) and the Java 2 Platform, Standard Edition. If neither the `JAVA_HOME` environment variable nor the Java property `java.version` is specified, the default JDK chosen in this multiple JDK environment depends on which 5761-JV1 Options are installed.

Note: If you install only one JDK, the default JDK is the one you installed. If you install more than one JDK, the following order of precedence determines the default JDK:

1. Option 8 - IBM Technology for Java 5.0 32-bit
2. Option 9 - IBM Technology for Java 5.0 64-bit
3. Option 7 - Classic 5.0
4. Option 11- IBM Technology for Java 6 32-bit
5. Option 12 - IBM Technology for Java 6 64-bit
6. Option 10 - Classic 6
7. Option 6 - Classic 1.4

Java Runtime.exec changes

In V6R1, the IBM Technology for Java Virtual Machine has changed the default for output from the Java `Runtime.exec()` function. When this function runs as an i5/OS PASE executable program, the output of the child process defaults to the character coding that is specified by the `file.encoding` property. This property is normally an ASCII-character encoding. In the previous release, the output of the child process used an EBCDIC-character encoding that matched the job CCSID. This changed behavior enables your Java application to read output from the child process without accounting for the CCSID differences. This change does not affect cases where you use `Runtime.exec()` to call an ILE program, whose default output might still be EBCDIC.

If you want the `Runtime.exec()` of an i5/OS PASE executable program to output EBCDIC character encoding, set the environment variable `QIBM_JAVA_EXEC_CONV=T`.

IBM Extended Integrated Server Support for i5/OS (5761-LSV)

| IBM plans for IBM i 6.1 to be the final release to support Extended Integrated Server Support for i5/OS
| (5761-LSV). In i 6.1, 5761-LSV contains integrated server support for VMware ESX support on iSCSI
| attached BladeCenter and System x servers, Linux on iSCSI attached BladeCenter and System x servers as
| well as Linux running in IBM i hosted logical partitions.

| VMware ESX on iSCSI attached BladeCenter and System x server will continue to be supported in the
| next release of IBM i in Option 29 - Integrated Server Support.

| For Linux on iSCSI attached BladeCenter and System x servers, 5761-LSV provides Linux operating
| system installation, save while active support for virtual disk (storage spaces), File Level Backup of files

| on the Linux server, dynamic linking and unlinking of virtual disk, and virtual optical and tape support.
| These functions are not planned to be available with the next IBM i release.

| Suggested Replacement: - Use IBM i support for VMware ESX running on iSCSI attached BladeCenter
| and System x servers to host Linux servers. However, the VMware ESX hosted Linux servers would not
| support save while active, file level backup, or virtual optical and tape support.

| For Linux running in IBM i hosted partitions, 5761-LSV provides save while active support for virtual
| disk (storage spaces) and File Level Backup of files on the Linux server. These functions are not planned
| to be available with the next IBM i release. Note that IBM i will continue to support hosting of Linux
| partitions, but that save while active and file level backup is not planned to be supported with the next
| IBM i release.

Linux support for IXS and IXA platforms:

In V6R1 support for all Linux operating systems running on the Integrated xSeries Server (IXS) and xSeries Adapter (IXA) platforms for integrated servers has been removed. The support for Linux on the IXS/IXA will be maintained in previous i5/OS releases for as long as the i5/OS release is supported and the Linux distribution supports the operating system.

| An IXS/IXA server can be migrated from a previous release to a new system, but you will not receive
| any support or fixes for the previous release. The function that is expected to work (but not maintained)
| for an IXS/IXA migrated server is as follows:

- | • Starting and shutting down the server.
- | • Storage space linking.
- | • The optical and tape locking utility (**ixsdev**).
- | • Statistics reported with the Work with Network Server Status (WRKNWSSTS) command.

The function that will *not* work in V6R1 is, as follows:

- Installations on IXS/IXA hardware.
- The post-installation utility (**ixssetup.sh**).
- The update utility (**ixsupdt**).

SuSE Enterprise Linux 9 support for iSCSI platforms:

Support for SuSE Enterprise Linux 9 (SLES9) running on an iSCSI platform for integrated servers will be removed. The support for SLES9 on the iSCSI platform will be maintained in previous i5/OS releases for as long as the i5/OS release is supported and the Linux distribution supports the operating system. An iSCSI-based integrated server running SLES9 can be migrated but support is “as is.” The function that is expected to work (but not maintained) for a migrated server is, as follows:

- Starting and shutting down the server.
- Storage space linking.
- The optical and tape locking utility (**ixsdev**).
- Statistics reported with the WRKNWSSTS command.

The function that will *not* work in V6R1 is, as follows:

- Installations on iSCSI hardware for *SLESe9.
- The post-installation utility (**ixssetup.sh**).
- The update utility (**ixsupdt**).

IBM HTTP Server for i5/OS (5761-DG1)

Third-party modules

The APR major version is changing from 0 to 1. Binary compatibility is not possible across a change in the major version. There are significant API changes, such as structure size changes and function name changes, that could cause compatibility problems for older programs.

All third-party modules must be recompiled for Apache HTTP server version 2.2 before being loaded.

Few customers have third-party modules, but one group that is affected are any customers who have installed an open source Tomcat connector. The connector is considered a third-party module.

Plug-ins and LoadModule directives

- | These instructions apply to all IBM HTTP servers that have been associated with a WebSphere
- | Application Server Version 6.0 or Version 6.1 when upgrading to V6R1. Get the latest version of the
- | WebSphere or Domino plug-ins for IBM HTTP Server through the PTF.

The LoadModule directives used by external HTTP servers that are associated with WebSphere Application Servers Versions 6.0 and 6.1 have changed.

You can update the LoadModule directive from the **IBM Web Administration for i5/OS** Web page: Select the HTTP server, then choose **Edit Configuration file**.

The following directions apply to HTTP servers that were already associated with a WebSphere Application Server Version 6.0 before upgrading to V6R1. You must do these steps before starting the servers.

- | **Note:** The server start code has been modified to correct the following directive changes, but the code
- | might not be able to use the new values in all cases. If errors occur on start of the servers, follow
- | these manual directions:

1. Apply PTF 5733-W60 SI29611.
2. Change the LoadModule directive from `LoadModule was_ap20_module /QSYS.LIB/QWAS6.LIB/QSVTAP20.SRVPGM` to `LoadModule was_ap20_module /QSYS.LIB/QWAS6.LIB/QSVTAPGT.SRVPGM`.

The following directions apply to HTTP servers that were already associated with a WebSphere Application Server Version 6.1 before upgrading to V6R1. You must do these steps before starting the servers.

1. Change the LoadModule directive from `LoadModule was_ap20_module /QSYS.LIB/product_installation_library/QSVTAP20.SRVPGM` to `LoadModule was_ap20_module /QSYS.LIB/QHTTPSVR.LIB/QSVTAP20.SRVPGM`. The product installation library is a name such as QWAS61A or QWAS61B.
2. In all cases, change the library name to QHTTPSVR.

Tomcat

Tomcat servers will no longer be shipped as part of the 5761-DG1 product in V6R1. Remove the following directive from customer configurations:

```
LoadModule jk_module /QSYS.LIB/QHTTPSVR.LIB/QZTCJK.SRVPGM
```

In addition, unless an open source version of Tomcat with the Tomcat connector is used instead, remove the following directives from customer configurations:

JkAsfTomcat
JkLogFile
JkLogLevel
JkMount
JkMountCopy
JkWorkersFile

If an open source version of the Tomcat connector is used, the information above regarding third-party modules applies to the module that is loaded into the HTTP server.

| **Trigger cache manager**

| **Last release for Trigger cache manager support**

| IBM i V6R1 will be the final release that will include the Trigger cache manager (5761-DG1) product (Option 1).

| **Web server search engine and Web crawler support**

| **Last release for Web server search engine and Web crawler support**

| IBM i V6R1 will be the final release that will include and support a Web server search engine and a Web crawler in the (5761-DG1) product.

| The following commands will not be shipped with the next release of IBM i:

| CFGHTTPSCH
| STRHTTPCRL
| ENDHTTPCRL
| RSMHTTPCRL

| Any existing CL programs that use these commands may need to be modified for future releases.

IBM System Manager for i5/OS (5761-SM1)

RSTAPARDTA might be deleted when the 5722-SM1 product is deleted:

The Restore APAR Data (RSTAPARDTA) command has moved from the IBM System Manager product to the i5/OS product. This change keeps both the Save APAR Data (SAVAPARDTA) and RSTAPARDTA commands within the same product. If you have V5R3 or V5R4 of System Manager (5722-SM1) installed on a V6R1 system, install the following 5722-SM1 PTFs to prevent problems when deleting or migrating System Manager:

- V5R3 PTF S128255
- V5R4 PTF S129132

If the RSTAPARDTA command has been accidentally deleted from V6R1, you can recover it by running the following command:

```
CPY OBJ('/QSYS.LIB/RADBKP.COMD') TOOBJ('/QSYS.LIB/RSTAPARDTA.COMD')
```

IBM TCP/IP Connectivity Utilities for i5/OS (5761-TC1)

SMTP configuration data changes:

The simple Mail Transfer Protocol (SMTP) configuration data that is stored in member CONFIG of physical file QATMSMTP in library QUSRSYS is automatically converted when the data is first referenced. This conversion is done when the SMTP server is started by the STRICPSVR or STRICP command, when the SMTP attributes are changed by running the CHGSMTPA command, or an update to

SMTP attributes is made through System i Navigator. This conversion is necessary to add new fields to the configuration file and to migrate information which was maintained in data areas in earlier releases to new fields in the configuration data file member.

The data areas that are migrated are QTMSIDVRFY, QTMSNOBRLF, and QTMSPOPOVR in the QUSRSYS library. All of these data areas are converted to their corresponding configuration fields. After the conversion, the data areas are deleted.

If you perform a save operation of the SMTP configuration file before doing the conversion, you should also save data areas QUSRSYS/QTMSIDVRFY, QUSRSYS/QTMSNOBRLF, and QUSRSYS/QTMSPOPOVR, if they exist.

IBM Universal Manageability Enablement for i5/OS (5722-UME)

CIM function moves from the base operating system to Licensed Program:

IBM System i Common Information Model Object Manager (CIMOM) server and providers are moved from the base operating system into IBM Universal Manageability Enablement for i5/OS, Licensed Program (LP) 5722-UME, V1R2M0. The CIM server (QYCMCIMOM) that shipped with i5/OS in the previous release is replaced by 5722-UME in V6R1. 5722-UME will install on V5R4 and V6R1 versions of i5/OS. 5722-UME is installed by default with i5/OS V6R1.

Dependency for CIM server startup:

Users must install these LPs and options before the CIM server can be started without error:

- 5761-SS1 option 33 (Portable Application Solutions Environment).
- 5733-SC1 option 1 (OpenSSH, OpenSSL, zlib)

CIM runs in i5/OS PASE:

CIM server and providers can run in Portable Application Solutions Environment (i5/OS PASE). Commands (**cimconfig**, **cimmof** and **cimprovider**) are moved to i5/OS PASE. You need to run the `call qp2term` command before running i5/OS PASE commands.

IBM Director compatibility:

The 5722-UME LP might not work with IBM Director earlier than version 5.20.2. IBM Director 5.20.2, or later, is needed.

Development interfaces:

The CIM server is based on OpenPegasus V2.5.1. Any external products using the restricted i5/OS V5R3 or V5R3 CIM provider interfaces must adapt to the new OpenPegasus Software Development Kit (SDK) V2.5.1 and obtain updated i5/OS-specific interface documentation from IBM. The development interfaces remain restricted with 5722-UME, and a limited availability agreement is required for their use. If external products have installed CIM providers in V5R4 or V5R3 in i5/OS, those CIM providers will not work with 5722-UME until they are updated to run in i5/OS PASE and to use updated interfaces. The CIM server supports “out-of-process” providers.

CIM schema:

5722-UME, V1R2M0, includes the Distributed Management Task Force (DMTF) CIM Schema V2.14.

CIMOM TCP/IP server entry in System i Navigator:

You can start and end the CIM server as a TCP/IP server. Using System i Navigator, expand **Network->Servers->User-Defined** to find the “CIMOM” entry.

You also can use Systems Director Navigator for i5/OS by selecting **Management->Network->User-Defined**, then click **OK**. Then you can use the web page to start or stop CIMON.

Repository migration:

If V6R1 is upgraded over V5R3 or V5R4, the CIM Server repository is migrated from DMTF CIM Schema V2.7 (or V2.9) to DMTF CIM Schema V2.14 during the first startup of CIM server in 5722-UME LP. This migration will take a while, depending on the size of the repository, processor speed, and system utilization.

The CIM server will be unavailable to process CIM requests until the repository migration completes. Stopping the server job during migration might result in a loss of data.

The following objects in the repository are not migrated:

- CIM Provider Registration is not migrated.
- Static instances of Metric Definition are not migrated. With 5722-UME V1R2M0, CIM providers dynamically collect information and implement the same functionality.

The CIM server log might have the following messages related to the repository migration:

- When migration starts, message PGS10100 is written in the CIM server log. The default location is `/QOpenSys/QIBM/UserData/UME/Pegasus/logs`.

PGS10100: The CIM server is starting to restore repository and then migrate the repository from an earlier version. This will take several minutes, during which the server will not be available. Stopping the server job might result in the loss of data.

- When migration ends without any error, message PGS10101 is written in the CIM server log.

PGS10101: The Common Information Model (CIM) repository has been migrated successfully.

Authentication methods:

The CIM server in 5722-UME LP removes Kerberos authentication support, and adds Secure Sockets Layer (SSL) support as authentication methods. New configuration properties are added, such as **sslTrustStoreUserName**, **exportSSLTrustStore**, **sslTrustStore**, **sslCertificateFilePath**, **enableSSLExportClientVerification**, **crlStore**, **sslKeyFilePath**, and **sslClientVerificationMode**.

Configuration properties:

The CIM server in 5722-UME LP changed some configuration properties. These properties are changed to fixed properties: **home**, **daemon**, **slp**, and **repositoryDir**. These properties are obsolete: **tempLocalAuthDir** and **kerberosServiceName**.

You can use the new properties **httpBindAddress** and **httpsBindAddress** to set the addresses that the remote client can access by HTTP/HTTPS protocols. The default values are 127.0.0.1 and ALL. The **enableHttpLocalConnection** configuration is obsolete.

If V6R1 is upgraded over V5R4 or V5R3, the first time the CIMOM is started, the CIM server migrates data from the i5/OS CIM server configuration from the previous releases to the i5/OS PASE CIM server configuration. Most of the configuration property values are migrated intact. The following are exceptions, and they are not migrated:

- The **logdir**, **home**, **daemon**, **slp**, **repositoryDir**, **tempLocalAuthDir** and **kerberosServiceName** properties are not migrated.

- If the **httpAuthType** property is Kerberos in the CIM server from the previous releases, **enableHttpsConnection**, **sslClientVerificationMode**, and **httpAuthType** are not migrated. **enableHttpConnection** is set to false.
- The **enableHttpLocalConnection** property is not migrated.

The property values being migrated are not validated. If the CIM server configuration properties from the previous releases are not set to function correctly, this situation might prevent the 5722-UME LP CIM server from starting and working correctly.

Provider properties:

Compared with V5R4 CIM providers in the previous releases, some provider properties are changed:

Table 2. Provider properties

CIM Classes	Property	V5R4 CIMOM provider in previous releases	5722-UME LP provider
IBMOS400_NetworkPort	RequestedState	7 (test)	5 (no change)
IBMPSG_PhysicalDisk	PowerOnHours	Implemented	Removed
IBMPSG_PhysicalMemory	HotSwappable	Implemented	Removed
IBMPSG_Chassis	HotSwappable, IsLocked	Implemented	Removed
IBMPSG_Processor	CPUStatus	Property type: Array of uint16	Property type: uint16
IBMPSG_PhysicalnetworkAdapter	HotSwappable	Implemented	Removed
IBMPSG_UserAccount	InstallDate	Implemented	Removed

Metrics provider access control:

Metrics provider access control is changed. A new Application Administration Group, **System Management Operations** is under the **CIMOM Server** product. The entry **Access to the CIM Performance Provider** can be in this group. The Metric value provider (including its instance and association) uses this entry to control user access.

Query language of indication:

With the 5722-UME LP, the indication filter query language supports **WQL** and **DMTF:CQL**. In V5R4 and V5R3, i5/OS CIM providers in the previous releases only supported **CIM:CQL**.

Provider parameter validity checking:

The CIM providers use DMTF-standard parameter validity checks on the CIM interface. If non-valid parameters are passed to the CIM interface, the CIM providers will notify users with the error and the reason.

IBM WebSphere Application Server V6.0 (5733-W60) and IBM WebSphere Application Server Version 6.1 (5733-W61)

Installation of WebSphere Application Server:

In V5R4, IBM Web Enablement for i5/OS was shipped with i5/OS, and included IBM WebSphere Application Server Express V5.1 for iSeries and IBM WebSphere Application Server Express V6 for OS/400.

- | For V6R1, IBM Web Enablement for i5/OS is shipped with i5/OS, and it includes IBM WebSphere Application Server Express V6 for OS/400 and IBM WebSphere Application Server Express V6.1 for

| i5/OS. IBM WebSphere Application Server Express V5.1 for iSeries is not included or supported with IBM
| Web Enablement for i5/OS for V6R1. Additionally, Application Server V5.1 is not supported on i5/OS
| V6R1 and is not operational on i5/OS V6R1.

| If you are currently using WebSphere Application Server V5.1 and upgrading to V6R1, you need to
| migrate to the WebSphere Application Server V6.0.2.23 or V6.1.0.11 products. Because migration does not
| require Application Server V5.1 to be operational, migration can occur before or after upgrading to V6R1.

The minimum required level of WebSphere Application Server V6.1 on V6R1 is 6.1.0.11. Installing WebSphere Application Server V6.1 on V6R1 requires a refreshed installation version of the product. The refreshed installation version of WebSphere Application Server V6.1 is based on level 6.1.0.11. If WebSphere Application Server V6.1 is on the system before upgrading to V6R1, then apply fix pack 11 (6.1.0.11) or later to WebSphere Application Server V6.1 before or after upgrading to V6R1.

The minimum required level of WebSphere Application Server V6.0 on V6R1 is 6.0.2.23. Installing WebSphere Application Server V6.0 on V6R1 requires a refreshed installation version of the product. The refreshed installation version of WebSphere Application Server V6.0 is based on level 6.0.2.23. If WebSphere Application Server V6.0 is on the system before upgrading to V6R1, then apply fix pack 23 (6.0.2.23) or later to WebSphere Application Server V6.0 before or after upgrading to V6R1. See the WebSphere Application Server for OS/400 V6 and WebSphere Application Server for i5/OS V6.1 documentation at the following Web sites for more information:

<http://publib.boulder.ibm.com/infocenter/wsdoc400/index.jsp>

<http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/index.jsp>

| Also, for additional information about changes to WebSphere Application Server for i5/OS V6R1, see
| <http://www.ibm.com/support/docview.wss?rs=180&context=SSEQTP&q1=enablejvm&uid=swg21305405>
| http://www.ibm.com/support/docview.wss?rs=180&context=SSEQTP&q1=enablejvm&uid=swg21305405&loc=en_US&cs=utf-8&lang=en.

LoadModule directives:

LoadModule directives are used by external HTTP servers that are associated with WebSphere Application Server Versions 6.0 and 6.1. For information about updating the LoadModule directives, see “Plug-ins and LoadModule directives” on page 66.

IBM WebSphere Development Studio for System i (5761-WDS)

Teraspaces-enabled parameter changes for CRTCMOD, CRTBNDC, CRTCPPMOD, and CRTBNDCPP commands

Starting in V6R1, all module or program objects are created as teraspaces-enabled even if the *NO special value is specified for the TERASPACE parameter on the CRTCMOD, CRTBNDC, CRTCPPMOD, and CRTBNDCPP commands. If the module or program object is being created for a target release (TGTRLS parameter) of V6R1M0 or later, the TERASPACE parameter is ignored. If the module or program object is being created for a target release before V6R1M0, the specified TERASPACE parameter value is stored in the object template.

To allow objects to be created on V6R1 with a TERASPACE value of *NO for an earlier TGTRLS, the CL commands will allow TERASPACE to be specified and will store the specified value in the object template if the TGTRLS value is earlier than V6R1M0. Even though the object displays on V6R1 with *YES for the TERASPACE attribute, it uses the value stored in the object template if the object is saved for an earlier release. The object shows the specified TERASPACE value when restored on a release earlier than V6R1.

ENBPFCOL parameter ignored on compiler commands

Starting in V6R1, the Enable performance collection (ENBPFCOL) parameter value will be ignored. Programs and ILE modules created by the WebSphere Development Studio and CL compilers will always have performance measurement code generated into the objects.

ILE C and C++ Compiler INLINE option changes

For the Create Bound C Program (CRTBNDC) and Create C Module (CRTCMOD) commands, the default mode for the INLINE option has changed from *NOAUTO to *AUTO. If the INLINE option is set to *ON and no inline mode is specified, the inline mode is *AUTO instead of *NOAUTO. The use of *AUTO permits additional functions to be inlined that were not previously inlined. Because the INLINE option is set to *OFF by default, only objects created with the INLINE option set to *ON and with no inline mode specified are affected by this change.

For the Create Bound C++ Program (CRTBNDCPP) and Create C++ Module (CRTCPPMOD) commands, the default mode for the INLINE option has changed from *NOAUTO to *AUTO. If the INLINE option is set to *ON and no inline mode is specified, the inline mode is *AUTO instead of *NOAUTO. The use of *AUTO permits additional functions to be inlined that were not previously inlined. Because the INLINE option is set to *OFF by default, only objects created with the INLINE option set to *ON and with no inline mode specified are affected by this change.

ILE CL Call Procedure (CALLPRC) advice for call to RPG now fails

| An ILE CL program that uses a 2-byte character variable to receive a 1-byte character value from an RPG procedure call may require a change in the CL source in V6R1 and later releases.

| Some versions of the RPG Programmer's Reference prior to V5R1 advised users to declare a 2-byte character variable and use a substring operation to receive the 1-byte return value. This method was incorrectly based on undefined implementation details that have changed in V6R1.

| To determine whether you have code that should be changed, you can compile a list of potentially affected CL source physical file members, examine the CL source, and make changes as needed. To compile a list of ILE CL source physical file members and line numbers of Return Value (RTNVAL) uses, the following command can be used.

| **Note:** This command assumes the ILE CL source is stored in files named QCLSRC within libraries.

```
| QSH CMD('find /QSYS.LIB -name "*.MBR" | grep "/QCLSRC.FILE/"  
| | xargs grep -E in "rtnval(\+|\()" > myoutfile.txt')
```

| To display the results file: **WRKLNK myoutfile.txt**

| To determine if the ILE CL source requires a change, examine the listed source file members. If a 2-byte character value is returned from a call to an RPG procedure, and only the first character is used by the CL procedure, the code will no longer run as expected on V6R1 and later releases.

| To correct the ILE CL source, change the declare of the character value being returned to use TYPE(*CHAR) LEN(1). The substring operation is no longer needed and can be removed.

| For CL programs that will only be run on V6R1 and later releases, specify V6R1M0 or later for the TGTRLS parameter for the compile of the CL module or program. No change is needed in the RPG source.

| If a single version of the CL program must be able to run on multiple releases that include V6R1 and a previous release, then the source of the RPG procedure being called also requires a change. Change the RPG source by adding the EXTPROC(*CL) keyword to the prototype (PR) and procedure interface (PI). Any RPG programs or modules that call the changed RPG procedure must be recompiled. An alternative

| correction that would not involve recompiling the RPG modules is to create a new "wrapper" procedure
| with EXTPROC(*CL). The CL module would be changed to call the new procedure, which would call
| through to the original, unchanged RPG procedure.

Lotus Domino for i5/OS (5733-LD7 and 5733-LD8)

Lotus Domino 7.0.3 for i5/OS (5733-LD7) and Lotus Domino 8.0.1 for System i (5733-LD8) work with V6R1. Some licensed programs, such as Lotus Domino and related products, (Lotus Sametime, Lotus Enterprise Integrator (LEI), Lotus Quickr, and others), require additional time for conversion and might require the installation of certain PTFs or hot fixes. For additional information, see the IBM Redpaper publication, *i5/OS Program Conversion: Getting ready for i5/OS V6R1*, from the IBM Redbooks Web site (<http://www.redbooks.ibm.com>).

For additional information about Lotus products, see the Lotus Domino for i5/OS Web site (<http://www.ibm.com/systems/i/software/domino/support>).

System i Access for Windows (5761-XE1)

Windows Installer for System i Access for Windows

New installation DVDs:

The iSeries Setup and Operations CD has been replaced by the System i Access for Windows DVDs. These DVDs contain all languages.

Note: The i5/OS Information Center assumes that there is only one DVD for the installation. For additional information about installing from DVD, visit the following Web site:

<http://www.ibm.com/systems/i/software/access>

Installing multiple languages for System i Access for Windows

Installing multiple languages for System i Access for Windows on your PC is not supported this release. However, you can install any available language as the primary language for System i Access for Windows on your PC during the initial installation. If you are installing from the product DVD, all languages are available and you can select to install any one of these languages as the primary language on your PC. If you are installing from the Licensed Program 5761-XE1, you can install multiple secondary languages for this product onto your i5/OS operating system. During the initial installation to your PC, you can select to install any of these secondary languages, in addition to the primary language installed for 5761-XE1, as the primary language on your PC.

Windows Installer wizard dialogs for Hebrew and Arabic languages:

The installation technology for System i Access for Windows was updated to the Windows Installer technology this release. The authoring tool used to create the installation does not contain Arabic and Hebrew. When you are installing this product, the installation wizard dialogs will not be available in Hebrew or Arabic. However, the rest of the product still supports Arabic and Hebrew, as in previous releases.

Installation image size for System i Access for Windows:

System i Access for Windows, 5761-XE1, takes up more space in the integrated file system than previous releases. This product contains three installation images; one for 32-bit PCs, one for 64-bit AMD/eMT PCs, and one for 64-bit Windows Itanium. In the previous release, the size of the image directory was 180 megabytes (MB); now the size of the image directories is 450 MB.

Changes to Windows installer:

The installation technology for System i Access for Windows was updated to the Windows Installer technology this release. This update results in the following changes:

- Administrators no longer can use tailored installations to control the features that their users can install. The tailored installation has been replaced by Windows Installer transforms.
- The Windows Installer technology does not use response files for silent installations. Instead the administrator can control the behavior of the installation by using transforms, public properties, setting the user-interface level, and changing the **setup.ini** file.
- The i5/OS Information Center topic for System i Access for Windows contains information for administrators on how to deploy and control the installations for their users.

Location of PC5250 data files:

Beginning with V6R1, the two predefined folders in which the workstation profile (.ws) files, and all other PC5250 data files are stored, changed locations. The old location that was based on the System i Access for Windows install path is now based on the user-specific Application Data folder. At the base location, the folder IBM\Client Access\Emulator\private is created to store the PC5250 files.

You can configure these paths for all users in two ways:

- By setting the /pc5250path parameter and running the **cwbcfg.exe** tool.
- By using the **System i Access Properties** under the **PC5250** tab.

At the first logon after an install, for each user who has configured one of the predefined folders that are identified above, settings are automatically changed and files are automatically copied to the new location, however, shortcut icons are not reconfigured. For example, a shortcut icon that refers, by full path, to a .ws file to launch a PC5250 session is not changed. The old folder icon is still usable to start a PC5250 session, however, configuration changes from that session are not saved to the new folder. It is strongly recommended that you delete and re-create, or change, such shortcut icons to specify the new folder location.

The full paths of these new locations are not the same for all versions of the Windows operating systems, however, all versions do provide two environment variables that you can use to refer to or find these locations. The user-specific Application Data folder name is stored in the environment variable APPDATA, and the common Application Data folder name in the environment variable ALLUSERSPROFILE. Environment variable values are obtained by enclosing them with percent signs(%). You can change PC5250 shortcut icons by replacing the part of the path that refers to your My Documents folder with %APPDATA%, and by replacing the part of the path that refers to the System i Access for Windows install path with %ALLUSERSAPPDATA%. For example, change the shortcut icon that refers to "C:\Documents and Settings\user5\My Documents\IBM\Client Access\Emulator\private\System1.ws", to "%APPDATA%\IBM\Client Access\Emulator\private\System1.ws", and it should access the .ws file in the new folder location. Consider making backup copies of shortcut icons before changing them.

Upgrading 5722-XE1 licensed program on V5R4 system

You must apply PTF SI28716 to your V5R4 system before upgrading to the 5761-XE1 licensed program. If you do not have this PTF applied, you will need to remove the 5722-XE1 licensed program before installing 5761-XE1.

Latest System i Access for Windows Service Pack required before upgrading client PCs

Before upgrading to V6R1 using the Check Version program at logon time, you must apply the latest service pack to your V5R3 or V5R4 System i Access for Windows client installation.

Management Central changes

Collection services changes

Starting in V6R1, you do not need any special authority to create database files using Collection Services. Earlier, you were required to have *ALLOBJ authority to create database files using Collection Services.

Management Central inventories with Blades hardware

Because Blades hardware is virtual, most of the APIs to access hardware information are not supported. Therefore, Management Central Hardware Inventory, Search in Hardware Inventory, and Export Hardware Inventory might not work. Processes and tasks, which display system information, such as Export Software Inventory, might display limited or no information.

Managing fixes

Send Function: Send Requisite Fixes option changes

When you select the Send Function: Send Requisite Fixes option, distribution requisites are sent along with prerequisite and corequisite fixes. You must have a V6R1 Central System selected. Requisite fixes are distributed to any endpoint while processing either Send and Install or Send functions.

Compare and Update changes

The Compare and Update functions can now compare fix groups when you compare fixes. When the results of the Compare and Update are distributed, the fix group information is sent from the model system to the target systems. Fixes distribution remains the same as in previous releases.

Important: When working with fix groups, it is strongly recommended to set the Central System with the highest release available to avoid incompatibility issues with older releases.

Restart of failed monitors

In releases before V6R1, failed monitors were automatically restarted only if the QYPS_AUTORETRY property was set on and &RESTART was specified as part of the name of the monitor. In V6R1, failed monitors are automatically restarted if the QYPS_AUTORETRY property is set on, regardless of whether &RESTART is part of the monitor name.

System i Access for Windows .NET Provider

The DB2 for i5/OS .NET Data Provider has made some incompatible changes, requiring the provider's Assembly version to change from 10.0.0.0 to 12.0.0.0. This change requires existing applications to either recompile using Visual Studio 2.0, or to create or change a configuration file, so that the .NET runtime can redirect requests for the 10.0.0.0 provider version to the 12.0.0.0 provider version.

For more details, see “Incompatible changes from previous releases” in the System i Access for Windows .NET Provider technical reference. For other late-breaking updates, refer to the Access for Windows .NET Provider web site at:

<http://www.ibm.com/systems/i/software/access/windows/dotnet/>

System i Navigator

Remote Access Services GUI (Properties and New Profiles panels)

In V5R4, the Remote Access Services Graphical User Interface (GUI) had **New** and **Open** buttons under the **Connection** tab for originator and receiver connection profiles. In V6R1, the interface has changed. The **Connection** tab does not have the **New** button anymore, but it still has the **Open** button. The **Open** button does both functions. If you enter a line name that is not in the list, the button works as the **New**

button. If you select or enter a line name that is in the list, the button works as an **Open** button. Because of API changes, the V5R4 System i Navigator Remote Access Services GUI is not compatible with a V6R1 system. Users will have unexpected errors trying to use the V5R4 System i Navigator Remote Access Services GUI with a V6R1 system.

System i Navigator on Windows Itanium

System i Navigator, GUI Builder, and Resource Script Converter are not supported on Windows iA64 (Itanium systems).

Copying files using System i Navigator

In releases before V6R1 using System i Navigator to copy a file might fail unexpectedly, even though the same function works when using the Copy Object (CPY) command. Copying some files with System i Navigator required *ALLOBJ special authority. In V6R1, you do not need *ALLOBJ special authority to copy files using System i Navigator.

In releases before V6R1, if you encountered an error while using System i Navigator to copy a complete directory, the copy processing would end and the remaining files would not be copied. In V6R1, the copy process attempts to copy the remaining files after a file copy has failed.

Location of System i Navigator plug-ins:

Plug-ins to System i Navigator must be located in the integrated file system in the QIBM/ProdData/OpNavPlugins or QIBM/UserData/OpNavPlugins directory. Plug-ins that are located in the old QIBM/ProdData/GUIPlugin or QIBM/UserData/GUIPlugin directories are not detected.

| System i Navigator wireless support

| Beginning in V6R1, System i Navigator for Wireless can run using only the Lightweight Infrastructure Web application server as the servlet engine. Previously, you could choose from an Apache Software Foundation (ASF) Jakarta Tomcat servlet engine or a WebSphere Application Server to run the servlet.

| To use this wireless support, ensure that you apply the following corequisite PTFs:

- | • PTF SI31126
- | • PTF SI30251

| For additional prerequisites, default port information, and instructions for configuring the Lightweight Infrastructure Web application server, see the Configuring a Web application server topic in the i5/OS Information Center. Click on Connecting to System i > System i Navigator > System i Navigator for Wireless > Setting up System i Navigator for Wireless > Configuring your wireless environment > Configuring a Web application server.

| Starting and Stopping System i Navigator for Wireless

| The command line methods for starting and stopping System i Navigator for Wireless have changed in V6R1.

| Use this command to start System i Navigator for Wireless:

| STRTCPSVR SERVER(*HTTP) HTTPSVR(QIWIRELESS)

| Use this command to stop System i Navigator for Wireless:

| ENDTCPSPV SERVER(*HTTP) HTTPSVR(QIWIRELESS)

| For additional instructions, see the Starting and stopping System i Navigator for Wireless topic in the
| i5/OS Information Center. Click on Connecting to System i > System i Navigator > System i Navigator
| for Wireless > Starting and stopping System i Navigator for Wireless.

| **System i Navigator removal of support for memory-resident database monitors next release**

| V6R1 will be the last release for System i Navigator and Systems Director to provide support for
| memory-resident database monitors. Memory-resident database monitors are identified as type Summary
| in the SQL Performance Monitors folder.

| **Note:** The memory-resident database monitor APIs will still be available.

| In prior releases, the main advantage of the memory-resident database monitor over the Detailed monitor
| (STRDBMON) was its smaller impact on system resources. This was due to the fact that it maintains the
| collected information in memory and writes to files only if and when the user requests that it do so. The
| result was the ability to gather database performance statistics with a minimal impact to the performance
| of the system as whole, or to the performance of individual SQL statements. However, significant
| enhancements introduced in V5R4 and V6R1 allow additional filtering on detailed SQL performance
| monitors, which greatly reduces the system impact of this type of monitor. SQL Plan Cache viewing and
| snapshot capabilities have also been added to provide access to the "always on" statement information
| cache. These new and redesigned tools provide great detail while maintaining low system impact. These
| improvements have eliminated the advantage the memory-resident database monitor once provided.

| If you plan to continue working with the memory-resident database monitors, you can directly call the
| memory-resident database monitor API interfaces.

Web enablement of System i Access

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

The changes made to web enable these functions caused several minor changes to the content and flow of
graphical panels for System i Access.

You can find more information in the **System i Navigator** information center topic and at the following
Web site:

<http://www.ibm.com/systems/i/solutions/management/products.html>

Click the **IBM System i Navigator** link.

IBM Systems Director Navigator for i5/OS:

Beginning in V6R1, you also can access System i Navigator tasks on the Web using the new IBM Systems
Director Navigator for i5/OS web console. This console displays the main task categories in a task
navigation area, and the selected tasks are displayed as portlets in the console work space. Ensure that
the HTTP Administration server is running, then enter the following URL to go to the console:

`http://{system}:2001`

For more information, see the **Systems Director Navigator for i5/OS** information center topic.

Note: The **i5/OS Tasks** page, which used to be shown when you entered the following URL:
`http://{system}:2001` now must be accessed from the **IBM Systems Director Navigator for i5/OS Welcome** page. Click the **i5/OS Tasks Page** link at the bottom of the **Welcome** page to view the tasks from the i5/OS Tasks page.

System i Navigator tasks on the web:

Most System i Navigator tasks are available from a Web browser by typing a specific URL. These tasks can be used without installing System i Access for Windows or System i Navigator. As of V6R1, over 300 tasks are now available on the web. Ensure that the HTTP Administration server is running, then enter the following URL to go to the tasks on the web home page:

```
http://{system}:2001/webnav/WnServlet?task=home
```

From this home page, you can view all web tasks, create a customized task URL, or save favorites of all web tasks for inclusion into your browser's favorites list.

Unicode support for System i Access for Windows and Linux clients

The System i Access for Windows and Linux clients also have a switch for overriding the new Unicode support, thus going back to using the job CCSID for all requests. You can set the `CWB_RMTCMD_V5R4_COMPATIBILITY` environment variable on the client to enable the override. The following values are valid for this variable:

***ALL** All processes use the old message format.

XXX.EXE

This process name (in 8.3 format) uses the old message format.

Visual Basic Wizards changes

The Visual Basic Wizards feature has been removed from the System i Access for Windows licensed program. The purpose of the Visual Basic Wizard was to help an application developer in writing ActiveX Data Objects (ADO). The usage and need for these wizards has diminished with the arrival of the next generation of database tools included in industry-standard development tools.

System i Access for Web (5761-XH2)

Before you install V6R1 System i Access for Web:

If you currently have a configured, running version of System i Access for Web on your i5/OS system, you must run the System i Access for Web configuration commands again when you install the new version before you can use the product. Running the command enables the new functionality provided with the new release. When running `CFGACCWEB2` for the WebSphere Application Server environments, be aware that the value input to the `WASPRF` parameter is case sensitive and should be entered exactly as shown in the `/QIBM/UserData/Access/Web2/config/instances.properties` file. For detailed instructions on installing and upgrading the product and running the configuration commands, see the V6R1 i5/OS Information Center: Click **Connecting to System i > System i Access > System i Access for Web**.

Web application servers and WebSphere Portal environments:

If you currently have a configured, running version of System i Access for Web on your i5/OS system, that web application server environment might no longer be supported by System i Access for Web. You will need to migrate your System i Access for Web configuration to a supported web application server. For additional information, see the V6R1 i5/OS Information Center: Click **Connecting to System i > System i Access > System i Access for Web**.

V6R1 System i Access for Web does not support the following web application servers:

- WebSphere Application Server V5.0 for iSeries
- WebSphere Application Server V5.0 - Express for iSeries
- WebSphere Application Server V5.1 for iSeries
- WebSphere Application Server V5.1 - Express for iSeries
- Apache Software Foundation (ASF) Tomcat

V6R1 System i Access for Web does not support the following WebSphere Portal environments:

- WebSphere Portal - Express for iSeries V5.0
- Workplace Services Express V2.5

V6R1 System i Access for Web can migrate user-generated data from an unsupported web application server environment to a supported web application server environment when new web application servers are configured. You can migrate user-generated data by using the CFGACCWEB2 command and specifying input for the SRCSVRTYPE/SRCSVRINST/SHRUSRDTA parameters.

V6R1 System i Access for Web also lets you to remove the System i Access for Web configuration from the unsupported web application server environment using the RMVACCWEB2 command. Be sure you have loaded and applied the latest V6R1 System i Access for Web PTFs before attempting to migrate from or remove unsupported web application server environment configurations.

Apache Software Foundation (ASF) Tomcat:

V6R1 System i Access for Web does not support Apache Software Foundation (AFS) Tomcat. It is recommended that you migrate ASF Tomcat configurations to the integrated Web application server environment.

Configuration commands:

In previous releases of System i Access for Web, only CL commands were provided to manage System i Access for Web configurations. QShell script commands were later added to those previous releases using PTFs to support newer web application servers. In V6R1 System i Access for Web, the CL commands and QShell script commands are provided as part of the base release. You can use either CL commands or QShell script commands to manage your System i Access for Web configuration. For information about how to use the commands, see the V6R1 i5/OS Information Center: Click **Connecting to System i > System i Access > System i Access for Web**.

Digital signatures:

V6R1 System i Access for Web contains and installs objects to the i5/OS system that are digitally signed. These objects are built with version 6 information. If V6R1 System i Access for Web is installed to a V5R4 i5/OS system, the digital signatures will be lost. To restore the digital signatures to the objects, follow these steps:

1. Upgrade the system to V6R1 i5/OS.
2. Reinstall V6R1 System i Access for Web.
3. Reload and apply V6R1 System i Access for Web PTFs.

XML Toolkits 5733-XT1 and 5733-XT2

The XML Toolkit 5733-XT1 has been withdrawn in V6R1. The XML Toolkit 5733-XT2 is the replacement for 5733-XT1. An application that uses the APIs included in 5733-XT1 will not run on V6R1. You must change the source code of the application to use the APIs provided in the XML Toolkit 5733-XT2. Note that the differences between the two versions include differences in both syntax and run-time behavior. For more information on converting your application, refer to the documentation included in the XML Toolkit 5733-XT2.

IBM Cryptographic Support for AS/400 (5722-CR1)

The IBM Cryptographic Support for AS/400 (5722-CR1) has been withdrawn in V6R1. Customers looking for a general cryptographic API set should consider the following alternatives:

- The IBM i Cryptographic Services API
- The Common Cryptographic Architecture (CCA) API in IBM i Option 35 along with the IBM i Cryptographic Coprocessor feature
- Java Cryptography Extension (JCE)

Customers looking for financial PIN processing support should consider using the CCA API in IBM i Option 35 along with the IBM i Cryptographic Coprocessor feature.

5798FAX - IBM Facsimile Support for IBM i

There are two versions of Facsimile Support for IBM i that are supported on IBM i 6.1.

- Facsimile Support for iSeries(5798FAX)V5R2M0 - This version is supported to run on IBM i 5.4 through IBM i 6.1. When this version is installed on IBM i 6.1, object conversion is needed. This version will not be supported in future releases.
- Facsimile Support for IBM i(5798FAX)V5R7M0 - This version is supported to run on IBM i 6.1 and later releases. There are enhancements included in this release, and no object conversion is needed. IBM recommends this version of Facsimile Support on IBM i 6.1 and later releases.

For the recommended fixes for each version of Fax support on IBM i 6.1, refer this document:

http://www.ibm.com/s_dir/s1kbase.NSF/fe291e47c3bd573c86256a64006793fc/4dbb96f3d430a0e286256dd50058fffe?OpenDocument

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 grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
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 106-0032, 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 YBWA
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, IBM License Agreement for Machine Code, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Programming Interface Information

This i5/OS Memo to Users publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of i5/OS.

Trademarks

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

AFS
AIX
AS/400
DB2
Domino
EnergyScale

i5/OS
IBM
IBM (logo)
Integrated Language Environment
iSeries
Lotus
Lotus Enterprise Integrator
OS/400
POWER5
POWER6
Quickr
Redbooks
Sametime
System i
Tivoli
WebSphere
Workplace
xSeries

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

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

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

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

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

Terms and conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

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, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

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

System i
i5/OS Memo to Users
Version 6 Release 1 (July 2012 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 email 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.

Email 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