



IBM Systems - iSeries
Memorandum to Users

Version 5 Release 4 (May 2008 update)





IBM Systems - iSeries
Memorandum to Users

Version 5 Release 4 (May 2008 update)

Note

Before using this information and the products it supports, be sure to read the information in "Notices" on page 53.

Eighth Edition (May 2008)

This edition applies to version 5, release 4, modification 0 of the IBM i5/OS (Program 5722-SS1) and iSeries licensed programs. This version does not run on all reduced instruction set computer (RISC) models nor does it run on CISC models.

© Copyright International Business Machines Corporation 2006, 2008. All rights reserved.

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

Contents

About iSeries Memorandum to Users vii

Who should read this memorandum	vii
Additional incompatibility information	vii
What's new	vii
Installing V5R4 over V5R2	viii
Discontinued support for certain software and hardware	viii
Installation considerations	viii
PTF numbers in this memorandum	ix
Memorandums for previous releases	ix
Prerequisite and related information	ix
iSeries Navigator	ix
How to send your comments	x

Chapter 1. Read this first 1

Current customers only — before you install V5R4	1
V5R4 installation requirements	2
Plan for your Operations Console installation or upgrade	3
Upgrading to V5R4 using image catalog support	5
I/O software prerequisites	5
Support for I/O expansion units	5
Key messages in this document	5

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
Prepare for Install authority change	7
Load source disk unit must be at least 17 GB	7
Booting i5/OS from load source disk unit attached through 2847 I/O processor	7
Removal of System Object Model (SOM) related files, parameter values, and MI instructions	8
New IPL attribute for recovery of spooled files	8
Program QSPGETF removed	8
QAITMON output file format has changed	9
Conversion of the CIM repository	9
Job log pending status more common in V5R4	10
Constant data items now exported from read-only storage	10
Data queues and user queues converted on first use	10
Object conversions will have marginal effect on system performance after installation	11
File delete requests on Microsoft Windows XP SP2	11
First Failure Data Capture changes	12
New printer file DDS keyword RELPOS for relative field positioning	12
Changes for Electronic Service Agent	12
QLOCALE system value set to default locale during installation	12
Locales updated to use Euro	13
Special authority failure procedure changes	13

Dump Tape (DMPTAP) command authority change	13
Password authentication failure changes	13
Restricting access to audit settings	13
Command menu removed	14
System now ignores signals sent to or within system jobs	14
Default job description changing for some server jobs	14
Smaller QJOBMSGQMX values for some server jobs	15
Message handler authority level changes	15
Display call stack function enhanced	15
Create object auditing attribute restrictions	15
API procedures exported by service program QP2USER can run in any activation group	16
Support for NLV 2950 removed	16
Process access group to be removed in a future release	16
Changes for files created by Performance Explorer	16
Troubleshooting Point-to-Point Protocol (PPP)	16
Point-to-Point Protocol (PPP) job list display changes	17
Hits per second in HTTP Summary in Performance Tools system report	17
Individual CPU utilization in Performance Tools system report	17
Larger job accounting fields	17
Changes to system job accounting codes	17
Work with Active Jobs (WRKACTJOB) display changes	17
IBM commands copied to QSYS library will be proxy commands	18
Journaling-related changes	19
Image catalog byte stream file CCSID changes	19
TCP/IP server configuration information logged in QTCPSTSVRS job log	19
Changed behavior for /QNTC directory	20
Change to journaling of files created with Copy File (CPYF) command	21
iSeries NetServer changes affecting thread behavior	21
System Request authority level changes	22
SSL default cipher specification list will change in the future	22
High availability clustering considerations	22
Command and API changes	22
Removal of Display Access Group (DSPACCGRP) and Analyze Access Group (ANZACCGRP) commands	22
Send PTF Order (SNDPTFORD) command and Universal Connection configuration changes	22
Vary Configuration (VRYCFG) command changes	23
Remove Server Storage Link (RMVNWSSSTGL) command changes	23
Create Network Server Description (CRTNWSD) command changes	24
Service Configuration commands changed	24

Authority changes for the image catalog commands	24	Parameter marker attributes	35
Create Image Catalog (CRTIMGCLG) command changes.	25	Database host server and DRDA/DDM server jobs set the current directory	35
Changes to Change Contact Information (CHGCNTINF) command	25	Temporary storage	36
Add Contact Information (ADDCNTINF) command removed.	25	Format level identifiers changed for record formats of new SQL tables that have one or more	
Return format for Retrieve Service Attributes (QESRSRVA) API changed	25	Timestamp columns	36
Open List of ASPs (QYASPOL) API with mirrored disks	25	Changes to the SQL LOCATE scalar function	36
Work with Spooled Files (WRKSPLF) command supports generic user data for SELECT	25	Machine interface (MI) instruction changes	37
Changes to CHGAUD, CHGAUT, CHGOWN, and CHGPGP command processing	26	Header file for the DESMTX MI instruction changed	37
Reclaim Activation Group (RCLACTGRP) command does not end some IBM activation groups	26	Maximum number of outstanding allocations field removed from the output of the MATHSAT and MATHSAT2 MI instructions	37
*WINDOWS value replaces *WINDOWSNT value on network server commands	26	The MATSOBJ field for returning information on the storage protection of encapsulated objects made obsolete	37
CMDTYPE parameter ignored on Submit Network Server Command (SBMNWSCMD)	26	LOCKOL, UNLOCKOL, and SETOBFPF MI instructions not supported	37
Install Windows Server (INSWNTSVR) command changes.	26	New option added to the MATRMD MI instruction.	37
Changes to Trace commands.	27	MATMIF MI instruction changed	38
New TCP/IP API formats to support IPv6	27	Header file for the MATMTX MI instruction changed	38
List Database File Members (QUSLMBR) API change	28	Restriction on LOCKTSL MI instruction	38
Default will change for QUSCRTUS parameter	28	Capacity donation and processor utilization	38
Changes to the Copy From Import File (CPYFRMIMPF) and Copy To Import File (CPYTOIMPF) commands	28	Server firmware PTFs on systems running Licensed Internal Code V5R4M5 or later releases	39
Changes to Change Job (CHGJOB) and End Job (ENDJOB) commands	28	File-level backup through QNTC file systems	39
Changes to Create Job Description (CRTJOB) and Change Job Description (CHGJOB) commands.	29	Host print transform changes	39
Backup and recovery changes	29	Chapter 3. Options	41
Changes to Display Save and Restore media panels	29	GDDM (Option 14).	41
Changes for parallel save operations to tape	29	NetWare Enhanced Integration (Option 25).	41
PTF adds additional error checking to the restore function	30	i5/OS Portable Application Solutions Environment (Option 33)	41
DB2 UDB for iSeries changes	30	Chapter 4. Licensed programs	43
BLOB comparisons change	30	Backup, Recovery and Media Services (5722-BR1)	43
Format change of additional parameters that are passed to a procedure defined with PARAMETER STYLE SQL	30	Client Encryption (5722-CE3)	43
Conversion of database file (*FILE) objects for optimal alignment	31	IBM HTTP Server for i5/OS (5722-DG1).	43
DB2_ROW_COUNT_SECONDARY field in the SQL Diagnostics Area changed	32	IBM WebSphere Application Server - Express Version 5.0 for iSeries (5722-IWE)	45
Triggers not allowed on IDDU physical and logical files	32	IBM Developer Kit for Java (5722-JV1)	45
Automatic conversion of *PGM, *SRVPGM, *MODULE, and *SQLPKG objects that include SQL statements	32	IBM i5/OS Integration for Linux on xSeries (5722-LSV).	46
Current optimization degree changes	33	Network Authentication Enablement (5722-NAE).	46
Database Monitor table format changed	33	Performance Tools (5722-PT1)	46
Changes to SQL	33	IBM DB2 Query Manager and SQL Development Kit for iSeries (5722-ST1)	47
		IBM TCP/IP Connectivity Utilities for i5/OS (5722-TC1).	47
		IBM iSeries Integration for Windows Server (5722-WSV)	48
		iSeries Access for Windows (5722-XE1)	48
		iSeries Access for Web (5722-XH2).	51
		IBM iSeries Access for Windows (5722-XE1) and Microsoft Windows Vista	52
		Notices	53

Programming Interface Information	54	Terms and conditions	55
Trademarks	54		

About iSeries Memorandum to Users

This information describes the changes in version 5 release 4 modification 0 (V5R4M0 or V5R4) that could affect your programs or system operations. Use the information in this memorandum to prepare for changes on your current release and to use the new release.

Who should read this memorandum

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

- **Read this first** provides information that should be considered before you install V5R4. This chapter is intended for system and application programmers and for the person responsible for system management.
- **Operating system** contains new release changes to basic operating system functions. This chapter includes changes to systems management functions, such as configuration and tailoring the system, as well as changes that could affect the way things operate or appear in the new release. This chapter is intended for all users of the iSeries™ servers.
- **Options** provides information about new release changes that affect specific program options of the operating system. This chapter is intended for all users of the iSeries servers.
- **Licensed programs** contains new release changes that might affect existing applications. These changes might also affect applications saved on a V5R4 server to be restored on a previous release server. This chapter is intended for application programmers and system programmers who use the iSeries server and its licensed programs, as well as for businesses with complex networks or application development businesses that have systems at different releases.

Additional incompatibility information

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

www.ibm.com/eserver/iseries/infocenter

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

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

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

Click **Technical Databases**.

What's new

The following revisions or additions have been made to the Memorandum to Users since the first V5R4 publication:

- **May 2008 update:**
 - “Host print transform changes” on page 39 was added.
- **February 2008 update:**
 - “File-level backup through QNTC file systems” on page 39 was added.
- **July 2007 update:**

- | – “Capacity donation and processor utilization” on page 38 was added.
- | – “Server firmware PTFs on systems running Licensed Internal Code V5R4M5 or later releases” on page 39 was added.
- | – “IBM iSeries Access for Windows (5722-XE1) and Microsoft Windows Vista” on page 52 was added.
- | • **January 2007 update:**
- | – “Changes to the SQL LOCATE scalar function” on page 36 was added.
- | • **June 2006 update:**
- | – “PTF adds additional error checking to the restore function” on page 30 was added.
- | • **May 2006 update:**
- | – “i5/OS Portable Application Solutions Environment (Option 33)” on page 41 was added.
- | • **April 2006 update:**
- | – “Changes to the SQL LOCATE scalar function” on page 36 was modified.
- | • **February 2006 update:**
- | – “V5R4 installation requirements” on page 2 was modified to include information about server firmware requirements.
- | – “High availability clustering considerations” on page 22 was added.
- | – “Changes to the SQL LOCATE scalar function” on page 36 was added.
- | – “IBM TCP/IP Connectivity Utilities for i5/OS (5722-TC1)” on page 47 was added.

Installing V5R4 over V5R2

Be sure to read the *iSeries Memorandum to Users* for V5R3. This document contains incompatibility-related information about the new functions and enhancements incorporated into V5R3. You can order this document by typing the following command:

SNDPTFORD SF98086

It is also available in the PSP information at this 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 iSeries announcement material. To get the most current information on discontinued products or features along with suggested replacements, go to the iSeries Planning Web site:

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

Select the link for upgrade planning information.

Installation considerations

Information about the installation and installation-related tasks is included in the book *Install, upgrade, or delete i5/OS and related software*. Information regarding planning and preparing for software installation along with licensed program conceptual and reference information can be found in the iSeries Information Center at:

www.ibm.com/eserver/iseries/infocenter 

Click **i5/OS and related software > Install, upgrade, or delete i5/OS and related software**.

PTF numbers in this memorandum

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

Memorandums for previous releases

In addition to ordering previous releases of the *Memorandum to Users* using the SNDPTFORD command, you can also view these documents at the following 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 iSeries Information Center as your starting point for looking up iSeries technical information.

You can access the Information Center two ways:

- From the following Web site:
<http://www.ibm.com/eserver/iseries/infocenter>
- From CD-ROMs that ship with your operating system order:
iSeries Information Center, SK3T-4091-06.

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

With every new hardware order, you receive the *iSeries Setup and Operations CD-ROM*, SK3T-4098-03. This CD-ROM contains IBM @server iSeries Access for Windows and the EZ-Setup wizard. iSeries Access Family offers a powerful set of client and server capabilities for connecting personal computers to iSeries servers. The EZ-Setup wizard automates many of the iSeries setup tasks.

iSeries Navigator

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

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

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

| **For ISVs and Business Partners that have written iSeries Navigator plug-in applications:**

In the future, iSeries Navigator will move to a Web-based user interface, using Java applications and a new user interface rendering technology called AUIML, which allows you to run your plug-in applications on the web as well as on the PC client.

If you currently have non-Java applications that plug into iSeries Navigator, you should plan to migrate them to Java and AUIML to enable them to run in both the PC client and the web. If you have a Java plug-in application, you should plan to migrate to AUIML. You can learn more about this by visiting the iSeries Navigator Plug-in Web site at the following URL:

<http://www.ibm.com/systems/i/software/navigator/plugin/index.html>

How to send your comments

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

- If you prefer to send comments by mail, use the readers' comment form with the address that is printed on the back. If you are mailing a readers' comment form from a country other than the United States, you can give the form to the local IBM branch office or IBM representative for postage-paid mailing.
- If you prefer to send comments by FAX, use either of the following numbers:
 - United States, Canada, and Puerto Rico: 1-800-937-3430
 - Other countries: 1-507-253-5192
- If you prefer to send comments electronically, use one of these e-mail addresses:
 - Comments on books:
RCHCLERK@us.ibm.com
 - Comments on the iSeries Information Center:
RCHINFOC@us.ibm.com

Be sure to include the following:

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

Chapter 1. Read this first

Current customers only — before you install V5R4

Each of the following publications contains additional information you should read and understand before you install this release. All of the sources referenced in this list are located either on the *iSeries Information Center*, SK3T-4091-06, CD-ROM or on the Internet at the following Web site:

www.ibm.com/eserver/iseries/infocenter

Note: After the release of the Information Center, updates will be available in the English Internet version. To review these updates, click **Updates since V5R4 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 iSeries Information Center Install, upgrade, or delete i5/OS and related software topic contains software preinstallation information and information on installing or upgrading the operating system release, part of the release, or related licensed programs. You can also order a printed version of this PDF (SC41-5120; feature code 8004) with your software upgrade orders or new hardware orders.
- The Preventive Service Planning (PSP) Information provides information about software problems you may encounter as you install the new release. You can download PSP by using electronic customer support from the Support for 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 V5R4 is SF98010. Information within the PSP is grouped by product area. To receive this PSP information by using electronic customer support, type the following command on an iSeries command line:

```
SNDPTFORD SF98010
```

- The PSP identifier for information that pertains to problems discovered since the availability of the current cumulative PTF package is SF98540. Information within this PSP describes all PTFs released since the current cumulative PTF package began shipping. It also contains information on all known high-impact and pervasive problems that are not included in the latest cumulative PTF package. To receive this PSP information by using electronic customer support, type this command on an iSeries command line:

```
SNDPTFORD SF98540
```

- The PSP identifier for information that pertains to installing V5R4 hardware is MF98540. Review this PSP information before you install either new iSeries servers or hardware devices. To receive this PSP information by using electronic customer support, use this command:

```
SNDPTFORD MF98540
```

- The PSP identifier for information that pertains to server upgrades and data migrations is SF98168. Information within this PSP describes upgrade and migration corrections. Review this PSP information before you upgrade your server model or migrate your data between servers. To receive this PSP information by using electronic customer support, type the following command on an iSeries command line:

```
SNDPTFORD SF98168
```

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

Notice

Operations Console support is available on V5R2 and later releases of the operating system. In V5R4, the only type of PC console that iSeries Models 270, 800, 810, 820, 825, 830, 840, 870, and 890 support is Operations Console.

V5R4 installation requirements

Successful installation of V5R4 requires the following for each server or logical partition:

- The V5R4M0 Licensed Internal Code requires a load-source disk unit of 17 GB or larger. Before you upgrade, ensure that the server meets disk storage requirements for upgrades; in the Information Center, click **i5/OS and related software > Install, upgrade, or delete i5/OS and related software > Upgrading or replacing i5/OS and related software > Preparing to upgrade or replace software > Performing initial upgrade or replacement tasks > Ensuring the server meets disk storage requirements for upgrades**.
- In addition, all server models with V5R3M0 or earlier installed require additional reserved storage before V5R4 can be installed. Your upgrade will stop during the installation if you do not allocate the additional space. In the Information Center, click **i5/OS and related software > Install, upgrade, or delete i5/OS and related software > Upgrading or replacing i5/OS and related software > Preparing to upgrade or replace software > Performing initial upgrade or replacement tasks > Required: Allocating additional space for Licensed Internal Code** for instructions.
- A minimum memory size of 128 megabytes. Systems with less than this minimum memory fail the installation process for Licensed Internal Code. Memory requirements for partitions are dependent on partition configuration, I/O resources assigned, and applications used. The primary partition needs a minimum of 256 megabytes of memory. Depending on configuration values used, a primary partition may require more than 256 megabytes. Secondary partitions running V5R1 and V5R2 require a minimum of 128 megabytes of memory. Also depending on configuration values used, a secondary partition may require more than 128 megabytes. For more information on memory requirements for logical partitions, see the Logical partition concept: memory topic in the iSeries Information Center.
- Server firmware requirements before you install V5R4: To determine the minimum firmware requirements for your system, you need to understand both the hardware and software requirements. The IBM eServer Prerequisite Web page provides the prerequisite information for features that you currently have or plan to add to your system. If you have an IBM System i5™, System p5™, or IBM eServer™ i5 or p5 system that is managed by Hardware Management Console (HMC), your server firmware must be at the SF235_160 level at a minimum to be able to support V5R4 on logical partitions that are in full-production mode. A server firmware level of SF230_145 is sufficient for logical partitions that are in test mode and do not contain any hardware that requires a more current level of firmware. For systems that do not have logical partitions, the V5R4 Licensed Internal Code contains the minimum level of firmware that is required.
- Some products have specific requirements of their own. Review all product requirement information that you receive with your order.

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 V5R4 (SF99540) from Corrective Service, you will receive the latest cumulative PTF plus DB (Database) Group PTFs and HIPER (High Impact Pervasive) Group PTFs to install with your software upgrade. The cumulative PTF package that in the past may have been delivered with your software order did not include the HIPER or Database Group PTFs.

If an extended 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:

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

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

Refer to the software installation preventative service planning (PSP) document, SF98010, 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 iSeries Support Web site. Click **Technical Databases > PTF Ordering**. Also, you can download the cumulative PTF package using FTP by specifying that option in Fix Central; this option provides faster turnaround than the 7-10 days required for the physical package to arrive.

Plan for your Operations Console installation or upgrade

If you are upgrading to V5R4 and you want to replace an existing console with Operations Console, upgrade the system before migrating the console. This will prevent any conflict between the existing console and the Operations Console. For instructions on upgrading your operating system, click **i5/OS and related software > Install, upgrade, or delete i5/OS and related software > Upgrade or replace i5/OS and related software**.

Prerequisite information for Operations Console users upgrading to, or installing, V5R4:

You must comply with the following before upgrading or installing your software (operating system, Licensed Internal Code) to V5R4:

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

Table 1. Card location

iSeries model	Operations Console asynchronous card location for the cable
270	C07
800 or 810	C07
820	C06

Table 1. Card location (continued)

iSeries model	Operations Console asynchronous card location for the cable
825	C06
830 or SB2	C02
840 or SB3	C02
870 or 890	C02

- For all upgrades and installations, you need to establish a connection between the server and Operations Console PC using the 11111111 (eight 1's) service tools user ID. The default password for this user ID is 11111111; however, this password might have been changed following a previous installation. This default user ID ensures a successful re-authentication of the client connection to the server. When you receive the operating system release upgrade, the shipped service tools user IDs (except 11111111) are expired. To re-authenticate the client connection to the server, enter the service tools user ID of 11111111 (eight 1's) and either the default password of eight 1's or the password you might have previously created for this user ID. This is especially important for automatic installations.
- It is recommended that you update iSeries Access for Windows® to V5R4 before you upgrade the operating system. For more information, see the Install iSeries Access for Windows topic in the iSeries Information Center.

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

Important: During a manual IPL of the server, and if no console has been specified before, you will receive two additional screens to confirm the setting of the console 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 could happen on your first manual IPL of V5R4; for example, the A-mode IPL following the restore of Licensed Internal Code during the upgrade or install when a console value of zero is found.

Migrating to Operations Console before your server model upgrade

If you are going to be utilizing Operations Console on your new iSeries server (migrating from a different console type), it is important that you configure the new Operations Console PC before the beginning of the server model upgrade. At the point in the upgrade instructions where console functions are required on the new iSeries server, you will be able to perform any required functions without the need for your current console device. The Operations Console features matching the connectivity you plan to use should be specified as part of the order for your new iSeries server.

New models

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

Embedded 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 iSeries Model 825 is the only iSeries 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 iSeries Information Center.

Upgrading to V5R4 using image catalog support

When you use image catalog support to upgrade to V5R4 from V5R2 or V5R3, you must apply the following PTFs before you prepare your image catalog for the upgrade:

- For V5R2, use SI19886
- For V5R3, use SI19888

I/O software prerequisites

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

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

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

Support for I/O expansion units

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

<http://www.ibm.com/servers/eserver/support/iseries/planning/upgrade/v5r3/hardware.html> 

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 V5R4, the following topics are particularly important:

- “Load source disk unit must be at least 17 GB” on page 7
- “New IPL attribute for recovery of spooled files” on page 8
- “Job log pending status more common in V5R4” on page 10
- “Data queues and user queues converted on first use” on page 10
- “Object conversions will have marginal effect on system performance after installation” on page 11
- “QLOCALE system value set to default locale during installation” on page 12
- “Changes to Display Save and Restore media panels” on page 29
- “Work with Active Jobs (WRKACTJOB) display changes” on page 17
- “IBM commands copied to QSYS library will be proxy commands” on page 18
- “Changes to Trace commands” on page 27

- "Conversion of database file (*FILE) objects for optimal alignment" on page 31
- "Automatic conversion of *PGM, *SRVPGM, *MODULE, and *SQLPKG objects that include SQL statements" on page 32
- "BRMS database file conversion" on page 43

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 may cause an application with LVLCHK(*YES) to fail with a level check error. If a level check error does occur, please review the application to determine which system file it uses. New fields have been added to IBM-supplied database files in each release of AS/400® and iSeries.

Security audit record changes

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

Programs that use customized versions of IBM-supplied commands

Some i5/OS functions that use IBM-supplied CL commands that are not library-qualified in this release may be changed in a future release to specify a specific library, *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.

Prepare for Install authority change

The **Allocate additional space for LIC** option of Prepare for Install now requires you to have *IOSYSCFG special authority.

Load source disk unit must be at least 17 GB

Systems or partitions on which you want to install V5R4 are required to have a load source disk unit that is at least 17GB. The load source disk unit is disk unit 1 in auxiliary storage pool (ASP) 1.

Booting i5/OS from load source disk unit attached through 2847 I/O processor

If you are planning to boot i5/OS® from the load source disk unit that is attached through the 2847 I/O processor, ensure that you read the IBM Redbook, iSeries and TotalStorage®: A Guide to Implementing External Disk on eServer i5, SG24-7120, for important planning and implementation considerations.

Removal of System Object Model (SOM) related files, parameter values, and MI instructions

Support for the System Object Model® (SOM®) was withdrawn in V3R7, but various related header files, parameter values, and MI instructions have remained on the system for compatibility purposes. In V5R4, these related files, values, and instructions have been removed. You might need to change your programs to remove references to these objects:

- Header file `setobpfp.h` is not being shipped with release V5R4. Any program that includes it will not find it and will receive a compilation error.
- The `MIOBJTYP` header file is provided in file `QSYSINC/MIH`. It contains definitions for all external object types. The definition `WLI_SOMOBJ` is being removed from this header file.
- Use of unblocked MI instruction Lock Object Location (`LOCKOL`), Unlock Object Location (`UNLOCKOL`) or Set Object Pointer From Pointer (`SETOBPFP`) will cause a function check (2003) exception to be signaled in V5R4.
- The Journal code (`JRNCDE`) parameter value 'O' on the Display Journal (`DSPJRN`), Retrieve Journal Entry (`RTVJRNE`), and Receive Journal Entry (`RCVJRNE`) commands has been removed. If you specify this journal code on these commands, the commands will now fail. Likewise, the support for the value 'O' journal code (key 7) is being removed from the `QjoRetrieveJournalEntries` API.

New IPL attribute for recovery of spooled files

In V5R4, the new Spooled file recovery (`SPLFRCY`) IPL attribute specifies what should be done with all spooled files during IPL when a job table is detected as damaged, or if `*YES` is set for these three IPL attributes: Clear job queues (`CLRJOBQ`), Clear output queues (`CLROUTQ`), and Clear incomplete job logs (`CLRINCJOB`). The initial shipped value is `*DETACH`, which changes the default behavior for these situations. In prior releases, all spooled files were removed for these situations. You can specify that this prior release behavior should continue by specifying the value `*REMOVE` for the `SPLFRCY` IPL attribute.

Program QSPGETF removed

The programs `QSPGETF` and `QSPPUTF` have never been documented or supported as operating system APIs. These programs were created for the purpose of capturing spooled files by IBM service for additional debugging.

With the introduction of operating system save and restore support for spooled files, the `QSPGETF` and `QSPPUTF` programs are no longer needed. Program `QSPGETF` has been removed. Program `QSPPUTF` will remain in support of respooling any spooled files that had been archived to media using program `QSPGETF`. In preparation for this, you should remove any references to these programs from your application programs. You should replace calls to `QSPGETF` and `QSPPUTF` with calls to the following documented and supported APIs:

- Spool APIs:
 - `QUSRSPLA`
 - `QSPCRTSP`
 - `QSPOPNSP`
 - `QSPGETSP`
 - `QSPPUTSP`
 - `QSPCLOSP`
 - `QUSLSPL`
- Print Manager/400 APIs
- Save/Restore APIs:
 - `QSRSAVO`

- QSRRSTO

The program source members TSRSV and TSRRST in file QATTSYSC in library QUSRTOOL demonstrate how you can use the supported APIs to perform the same function as programs QSPGETF and QSPPUTF.

Note: QUSRTOOL is optionally installable as option 7 of the base operating system.

IBM Backup Recovery and Media Services for i5/OS is an example of a product that provides support for saving and restoring spooled files.

QAITMON output file format has changed

The format of the QAITMON file has changed substantially. Many individual CPU reporting fields were removed and replaced by the following new fields:

- CONFTOT (Total CPU Configured)
- MINCPU (Minimum CPU Used)
- MAXCPU (Maximum CPU Used)

These new fields contain CPU information for the system, simplify the task of reporting CPU utilization, and more consistently reflect the data reported by the Work with System Activity (WRKSYSACT) command on the character-based interface.

You will need to change any of your applications that process the QAITMON database file to use the new output file format.

Conversion of the CIM repository

Conversion of the Common Information Model (CIM) repository will occur at first startup of the CIM server following installation of V5R4. The CIM repository is located at /QIBM/UserData/OS400/CIM/repository. The CIM schema in the repository will be upgraded to Distributed Management Task Force (DMTF) CIM Schema V2.9, and any class extensions, qualifiers, instances, or provider registrations added by the user will automatically be migrated during this upgrade.

The amount of time required to complete the conversion and upgrade depends on the size of the repository, processor speed, and system utilization during migration. The CIM server will not be available to process CIM requests until after the repository migration completes. The CIMOM server (QYCMCIMOM job) might have the following messages related to the repository migration unless the repository has previously been migrated:

- Message CPIDF80 will appear in the CIM server (QYCMCIMOM) job log when migration begins:
CPIDF80: Common Information Model Object Manager (CIMOM) server information: The CIM server is migrating 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 a loss of data.
- Message CPIDF81 will appear in the CIM server (QYCMCIMOM) job log when migration completes:
CPIDF81: The Common Information Model (CIM) repository has been updated. User data has been successfully preserved.

When the repository migration has successfully completed, the QSHELL commands cimconfig, cimmoftl, and cimprovider are available for use.

Job log pending status more common in V5R4

Starting in V5R4, it will be much more common to have jobs that have a status of job log pending. In previous releases, this status was usually a result of a Power Down System (PWRDWN SYS) command. The **Job log output** (LOGOUTPUT) job attribute now determines how a job log is written. A completed job that has a job log pending has a job state of OUTQ even though the job might currently have no spooled files associated with it.

In previous releases, job logs were written by the SCPF job. In V5R4, the SCPF job no longer writes job logs; instead, they are written by job log server jobs. See the Start Job Log Server (STRLOGSVR) command information for details.

The shipped default for system value QLOGOUTPUT is *JOBEND, but the recommended value for QLOGOUTPUT is *JOBLOGSVR. Applications that require job logs to be produced when the job completes its activity should specify LOGOUTPUT(*JOBEND) in the job description or on the Submit Job (SBMJOB) CL command. You might need to create additional job descriptions so that different applications can easily get different job attributes.

The procedures you use to manage and remove old job logs might need to be changed to handle the pending job log status. In previous releases, if a job had a status of job log pending, a job log would eventually be written. Starting in V5R4, a job log can remain in pending status indefinitely. If you use Operational Assistant to remove old job logs, pending job logs will now also be removed. See the Change Cleanup (CHGCLNUP) command for information on how to set the number of days that job logs should be retained. If you do not use Operational Assistant, you need to add the cleanup of pending job logs to your normal cleanup procedures. See the Remove Pending Job Log (QWTRMVJL) API and the Add Job Schedule Entry (ADDJOBSCDE) command for more information. You can use the Display Job Tables (DSPJOBTL) command to see if your clean up is working. You can also use the Work with Job Logs (WRKJOBLOG) command to help manage both pending and spooled job logs

Constant data items now exported from read-only storage

Prior to V5R4, programs that exported constant data actually exported a writable copy of the data from static storage. Other programs that imported the constant data could modify the data without taking an exception.

When created for a target release of V5R4M0 or later, these programs now properly export the constant data items from read-only storage. A program that imports one of these true constants and attempts to change it will receive exception MCH6802 (Literal values cannot be changed).

Data queues and user queues converted on first use

To enhance the run time performance of journaled data queue (*DTAQ object) synchronization and to increase the debug capabilities of all queues (*DTAQ and *USRQ objects), the first time a queue is touched after V5R4 is installed, an internal conversion takes place. This conversion might take a noticeable amount of time for large journaled data queues but should be relatively quick for other data queues. To avoid delays in a production environment, you might want to touch all queues immediately following the installation.

One way that you can do this is to run the Display Library (DSPLIB) command, and then select either Option 5 or 8 (to display attributes) for each *DTAQ or *USRQ object.

You could also run the following Submit Job (SBMJOB) commands to run the Display Object Description (DSPOBJD) command in batch processing:

```
SBMJOB CMD(DSPOBJD OBJ(mylibname/*ALL) OBJTYPE(*DTAQ) OUTPUT(*OUTFILE)
        OUTFILE(QTEMP/myoutfile)) JOB(myjobname)
SBMJOB CMD(DSPOBJD OBJ(mylibname/*ALL) OBJTYPE(*USRQ) OUTPUT(*OUTFILE)
        OUTFILE(QTEMP/myoutfile)) JOB(myjobname)
```

Alternatively, to touch and convert all queues in the library list, you can run the following SBJJOB commands to run the Display Object Description (DSPOBJD) command in batch processing:

```
SBMJOB CMD(DSPOBJD OBJ(*LIBL/*ALL) OBJTYPE(*DTAQ) OUTPUT(*OUTFILE)
        OUTFILE(QTEMP/myoutfile)) JOB(myjobname)
SBMJOB CMD(DSPOBJD OBJ(*LIBL/*ALL) OBJTYPE(*USRQ) OUTPUT(*OUTFILE)
        OUTFILE(QTEMP/myoutfile)) JOB(myjobname)
```

Object conversions will have marginal effect on system performance after installation

In V5R4, essentially all objects are undergoing a small conversion and most will have enhanced hardware storage protection attributes applied at the same time. These changes will add several minutes to the first IPL (initial program load) following the install of V5R4, plus a small amount of first-touch or background processing overhead.

The first-touch overhead effects cannot be well predicted, because they vary along with each system's unique object usage pattern. The background conversions, running as a number of low-priority input- and output-bound tasks, will complete within a couple days. After the background conversions have completed, no more first-touch conversions are required. System performance will be only marginally affected by these conversions, for at most a couple of days.

These conversions do not cause the modification time stamp of the object to be changed.

If you want to know when the conversions are complete, you can look at the Licensed Internal Code logs, available through service tools. A log entry with major code 1000 and minor code 2030 is added when conversions are complete for each Auxiliary Storage Pool (ASP). One log entry containing '1' is issued when objects in the system and all user ASPs have been converted. A log entry containing the independent ASP number is issued when objects in each independent ASP have been converted.

In addition, many objects are also being changed to align data in associated spaces or space objects. This alignment preserves previous performance when modifications are made to the data in the space. The one-time effects of these alignments cannot be predicted, and will only marginally affect system performance.

The data alignment for affected objects will be handled at different times, but most often will occur the first time the objects are used in any way. The modification time stamp of these objects might be updated when the data alignment occurs.

File delete requests on Microsoft Windows XP SP2

On Microsoft® Windows XP SP2, file delete requests to iSeries NetServer™ now consist of the following steps:

1. Open the integrated file system file.
2. Unlink the open file, removing it from the parent directory (the file is deleted when no job currently has the file open and the file is not linked in a directory).
3. Close the file.

The "root" (/) file system (and others) support this type of delete. Some earlier file systems (QDLS, QSYS.LIB, QOPT) do not support deleting a file that is in use (the open request in this case). To accommodate this type of deletion on these file systems, iSeries NetServer will save the delete for

processing after the file is closed. After the file is closed and until the file is deleted (a narrow window of time), the file can be opened, saved, or copied by another application or user, since the file is not removed from the integrated file system name space.

First Failure Data Capture changes

In V5R4, a new policy-based Service Monitor is used to determine if problem log entries should be written when a software error is detected by the system. Fewer problem log entries might be created on the system, but the amount of data collected when capturing information related to the failure might increase. The Service Monitor, when activated by setting the system value QSFWERRLOG to *LOG, can start several dozen jobs. These jobs will run in the QUSRWRK subsystem and be named SRVMONxxx where xxx is a numeric value.

New printer file DDS keyword RELPOS for relative field positioning

In prior releases, if an AFPDS printer file that specified FONT(*DEVLD), FNTCHRSET, or CDEFNT also specified a font change on one of the fields in a record, and if that field contained only SBCS data and the POSITION keyword was not used, the next field on the line was positioned by inserting spaces without first returning to the left margin.

This inconsistent behavior has been resolved in V5R4. However, if your existing applications encounter problems because of this change, you can use the Relative Position (RELPOS) keyword in the printer file DDS, and recompile the printer file to override the new behavior.

Changes for Electronic Service Agent

In V5R4, Electronic Service Agent™ has been changed in the following ways:

- During an attended initial program load (IPL), if Service Agent has not been previously activated, a Service Agent program is called to start its activation. This eliminates the need to issue GO SERVICE to activate Service Agent. If all required information is available, Service Agent will activate with no user interaction. If all required information is not available, the following CL commands will be prompted, if necessary:
 - CHGCNTINF (Change Contact Information)
 - CRTSRVCFG (Create Service Configuration)
 - CHGSRVAGTA (Change Service Agent Attributes)
- Service Agent reports software problems generated by the Service Monitor, which provides automatic detection of software problems based on policy information. Policy files are updated using a Service Agent connection to IBM.
- Problem reporting retry lets you specify whether service requests placed automatically by Service Agent are to be retried. You can change the following settings:
 - Whether service requests are retried
 - Number of times to retry
 - Interval of time between retries
 - Whether messages are sent to users
- A new Service Agent menu option displays current and historical Service Agent problem reporting information.
- The Display Service Agent (DSPSRVAGT) command no longer supports TYPE(*INV).

QLOCALE system value set to default locale during installation

In V5R4, i5/OS now sets the QLOCALE system value for new installations to a default locale value based on the primary NLV that is installed. In prior releases, this value was set to *NONE. If you want the QLOCALE value set to *NONE, you must make this change separately.

Locales updated to use Euro

In V5R4, those locale environments that have both a Euro version and a non-Euro version now use the Euro version as the default locale.

Special authority failure procedure changes

The following changes to the procedures for auditing authority failures have been made in V5R4

- The CPF2220, CPF4AAE and CPF2246 messages are no longer sent to QHST. You must look for AF-K audit records to determine if a special authority violation has occurred.
- Some AF-A audit records, which were generated for special authority violations, have been changed to the new AF-K audit record so that all special authority violations are audited in the same manner and are not confused with object authority violations.
- For the Start Disk Reorganize (STRDSKRGZ) and End Disk Reorganize (ENDDSKRGZ) commands, message CPF2239 is no longer sent to QHST. You must look for AF-K audit records to determine if a special authority violation has occurred.
- For the Retrieve Job Information (QUSRJOBI) API, message CPF2239 is no longer sent to QHST. You must look for AF-K audit records to determine if a special authority violation has occurred.
- For the Display Job (DSPJOB) command, message CPF2239 is no longer sent to QHST. You must look for AF-K audit records to determine if a special authority violation has occurred.

Dump Tape (DMPTAP) command authority change

If you use the Dump Tape (DMPTAP) command, you must now have all object (*ALLOBJ) special authority when you specify the TYPE(*HEX) parameter as well as when the tape has the secure volume flag or secured file flag set. In prior releases, *ALLOBJ special authority was required only when the tape had the secure volume flag or secure file flag set.

Password authentication failure changes

In V5R4, the Check Password (CHKPWD) command writes a PW-C audit record when it rejects the entered password. Previously, a PW-P audit record might have been written for some situations but not all.

A new PW-Q audit record is written when a user authentication operation detects that a profile is in *DISABLED status and this is the cause of the authentication failing. Previously, other PW audit violation types may have been written for a *DISABLED profile or no audit record was written.

A new PW-R audit record is written when expired passwords are checked and the authentication is rejected because the password is expired. Previously, no audit record was written for this case.

Restricting access to audit settings

Note: The changes described here were also provided in V5R3 Program Temporary Fixes (PTFs), and therefore might not appear to be a new change if you have applied those PTFs on your V5R3 system.

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

These changes have affected the following interfaces:

- Work with Object Links (WRKLNK) command
- Display Object Links (DSPLNK) command
- Get Attributes (Qp0lGetAttr) API interface
- Display Object Description (DSPOBJD) command
- Retrieve Object Description (RTVOBJD) command
- Open List of Objects (QGYOLOBJ) API
- List Objects (QUSLOBJ) API
- Retrieve Object Description (QUSROBJD) API
- Retrieve Library Description (RTVLIBD) command
- Retrieve Library Description (QLIRLIBD) API

In V5R4, message CPF180F is sent to QHST instead of message CPF1806 when the system values QAUDCTL, QAUDENDACN, QAUDFRCLVL, QAUDLVL, QAUDLVL2, and QCRTOBJAUD are changed. CPF180F does not contain the previous and new values in its replacement data.

In prior releases (including V5R3 when the related V5R3 PTF was not applied), if you used Qp0lGetAttr to retrieve only the audit value for an object, it would have worked if you had the appropriate path name resolution authorities. Now, the path name resolution authorities are no longer required. Also, if you are retrieving only the audit value for a QFileSvr.400 object, the QSECOFR profiles on both systems must be enabled and their passwords must match for the operation to succeed.

The Display DLO Auditing Level (DSPDLOAUD) command now returns an error if you do not have *ALLOBJ or AUDIT authority. Previously, if you had *USE authority to the object, you could see the auditing value.

The Change Library (CHGLIB) command prompt returns a value of *SAME rather than the actual create object auditing value when you do not have either audit (*AUDIT) or all object (*ALLOBJ) special authority.

Command menu removed

In V5R4, the Access Group commands (CMDACCGRP) menu object is no longer shipped as part of i5/OS.

Application programs that attempt to use the Go to Menu (GO) command for this removed menu will fail with message CPF6AC7.

System now ignores signals sent to or within system jobs

Some exit programs can potentially run from within a system job. A system job is a job shown as type 'SYS' on the Work with Active Jobs (WRKACTJOB) display.

Beginning in V5R4, the system ignores signals sent to or within system jobs. Furthermore, the system does not allow the signal action to be changed within system jobs. If you use an exit program that relies on the use of signals, you will need to change your exit program to no longer use signals if it is run in a system job.

Default job description changing for some server jobs

In the subsystem description QSYS/QSERVER, the prestart job entries for programs QSYS/QPWFSERVSO, QSYS/QPWFSERVSS, and QSYS/QPWFSERVS2 now use job description QSYS/QPWFSJPJ. In V5R3, these prestart job entries used job description QGPL/QDFTSVR.

In the subsystem description QSYS/QSERVER, the prestart job entry for program QSYS/QZLSFILE now uses job description QSYS/QZLSPJ. In V5R3, this prestart job entry used job description QGPL/QDFTSVR.

Smaller QJOBMSGQMX values for some server jobs

In V5R4, the QJOBMSGQMX system value can now be set to values from 2 to 64 (megabytes). In previous releases, the allowed values were limited to the range 8 to 64 (megabytes).

Default configurations for many server jobs have changed to specify a smaller value for JOBMSGQMX. When this is used in combination with JOBMSGQFL(*PRTWRAP), more spooled QPJOBLOG files might be produced.

Job description QGPL/QDFTSVR now specifies JOBMSGQMX(4). In V5R3, QGPL/QDFTSVR used JOBMSGQMX(8).

Message handler authority level changes

In V5R4, the user profile in effect for default-handling programs and user-defined break-handling programs is the initial user profile for the job.

Prior to V5R4 (or with the application of PTF SI20929 in V5R3), the user profile in effect for break-handling programs and default-handling programs was the current user profile under which the job was running when a message caused the break-handling program or default-handling program to be run. To continue to have the current user profile used in these cases, you must create a data area using one of the following commands.

To affect the behavior of break-handling programs, create a data area with the following command:

```
CRTDTAARA DTAARA(QSYS/QMHBKRPSEC) TYPE(*CHAR) LEN(12)
VALUE(*NOCHGCURUSR)
```

To affect the behavior of default-handling programs, create a data area with the following command:

```
CRTDTAARA DTAARA(QSYS/QMHDFTPSEC) TYPE(*CHAR) LEN(12)
VALUE(*NOCHGCURUSR)
```

Display call stack function enhanced

The display call stack function has been enhanced to show more information in a more usable format. You can access new types of stack entries, including i5/OS PASE stack entries and Java stack entries. These enhancements have resulted in changes to displays and printed listing formats. You can access these new functions using commands or APIs such as the Work with Job (WRKJOB) command and the Retrieve Call Stack (QWVRCSTK) API.

Among these changes, the interactive display and print forms of Display Call Stack present more information and use fewer alternate views. The interactive display now shows the first 29 characters of procedure names (instead of the last 10 characters), and both forms now show activation group numbers as hexadecimal values. Display Call Stack is common support for the Display Jobs (DSPJOB) command, the Work with Jobs (WRKJOB) command (option 11), and the Work with Active Jobs (WRKACTJOB) command (option 10).

Create object auditing attribute restrictions

Note: The changes described here were also provided in V5R3 Program Temporary Fixes (PTFs), and therefore might not appear to be a new change if you have applied those PTFs on your V5R3 system.

Changing the **Create object auditing** attribute using the Change Attribute (CHGATR) command or the Qp0lSetAttr API for QFileSvr.400 file system objects might fail if the QSECOFR profiles on the source and target system are not consistent or are disabled.

API procedures exported by service program QP2USER can run in any activation group

API procedures exported by service program QP2USER (such as Qp2CallPase) can now be used in any activation group, as long as i5/OS PASE is active in the job. In prior releases, these APIs were only usable in the activation group that called Qp2RunPase (to start i5/OS PASE). The system will still automatically end i5/OS PASE when the activation group that calls Qp2RunPase ends.

Service program QP2USER continues to run in the caller's activation group, but API implementations in QP2USER now call procedures in a new service program that runs in the default activation group. You might need to change i5/OS PASE programs that use runtime functions QMHSNDPM, QMHSNDPM2, QMHRCVPM, QMHRCVPM1, or QMHRCVPM2 to account for extra program call levels (including those with *PGMBDY and *CTLB DY attributes) in the stack.

Support for NLV 2950 removed

The national language version (NLV) of 2950 is no longer offered. It is recommended that you change your NLV to 2924.

If you have some need for uppercase-only support (for example, a screen scraper application), you could load NLV 2938 as a secondary language; this alternative keeps the CCSID values and field types of the system cross-reference file set the same as they are on NLV 2950.

Process access group to be removed in a future release

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

Changes for files created by Performance Explorer

In V5R4, Performance Explorer now creates the QAYPE* database files and members only if the file and member are needed for the collection. In prior releases, Performance Explorer created the entire set of QAYPE* files and members even if they were not needed for the collection. Your applications that use Performance Explorer database files now need to check for the existence of each file.

In addition, field QPRPGN in file QAYPEPROCI has been changed from 10 to 30 characters in length.

Note: The field QRECN has often been used as a sort-by-time field. This has never been the recommended sort field, and starting in V5R4 it will yield incorrect results in all Performance Explorer collections if used as a sort-by-time field. You should use the field QTITIMN (nanoseconds from the start of the collection) in file QAYPETIDX for sorting by time. You should use QRECN only to join the files QAYPETIDX, QAYPETIDXL, or QAYPETID2L with the other Performance Explorer files that contain trace records.

Troubleshooting Point-to-Point Protocol (PPP)

The troubleshooting steps and supporting materials for Point-to-Point Protocol (PPP) have been modified. See the Troubleshoot PPP topic in the iSeries Information Center for details.

IBM commands copied to QSYS library will be proxy commands

The command (*CMD) objects of most IBM licensed program products (LPPs) and all options of i5/OS that ship commands in a product library (other than QSYS) are copied to the QSYS library when the LPP or option is installed.

In prior releases, duplicate commands were created in QSYS.

In V5R4, proxy commands are created in library QSYS. A proxy command is a *CMD object that points to a target command; when the proxy *CMD is prompted or run, the operating system uses the target command.

When a regular command by the same name as the command being copied from the product library already exists in library QSYS, any private authorities to the regular command in QSYS are also granted to the command in the product library. If a proxy command by the same name as the command being copied already exists in library QSYS, any private authorities to the proxy command will not be granted to the command in the product library.

Since attempts to prompt or run the proxy command in QSYS will check the authority of the target command in the product library, you should normally grant private authorities for the command only to the regular command in the product library.

Now that some IBM commands in QSYS are proxy commands, if you duplicate commands from QSYS into another library using the Create Duplicate Object (CRTDUPOBJ) command, your duplicate commands might be proxy commands.

Prior to V5R4, using the Change Command (CHGCMD) or Change Command Default (CHGCMDDFT) command to change a command that you duplicated from the QSYS library only affected your duplicate copy of the command.

In V5R4, when a proxy command is specified on the CHGCMD and CHGCMDDFT commands, the actual command changed will be the command that is the target of the proxy command. For LPP proxy commands in QSYS, if you change the proxy command using CHGCMD or CHGCMDDFT, you will change the target command in the LPP library.

If you want CHGCMDDFT and CHGCMD to affect only the duplicate that you create, then do the following:

- Determine if the QSYS command is a proxy command. You can do this interactively using the Display Command (DSPCMD) command or programmatically using the Retrieve Command Information (QCDRCMDI) API.
 - If it is not a proxy command, use CRTDUPOBJ to duplicate it.
 - If it is a proxy command, use DSPCMD or QCDRCMDI to determine the target of the proxy command, and duplicate the target command instead of duplicating the proxy command in QSYS.

If an exit program had been previously registered for the copy of a licensed program command in the QSYS library, the exit program will not be run for the proxy command that replaced it when the V5R4 release of the license program product was installed. You will need to register your exit program against the original command in the product library. For example, if you had an exit program registered for the QIBM_QCA_CHG_COMMAND exit point for the Create SQL ILE RPG Object (CRTSQLRPGI) command in QSYS, you will need to register your exit program for the original command in the QSQL library.

Journaling-related changes

In V5R4, the special value *SYSDFT on the Edit Recovery for Access Paths (EDTRCYAP), Change Recovery for Access Paths (CHGRCYAP), and Display Recovery for Access Paths (DSPRCYAP) commands changed from 60 minutes to 50 minutes. Therefore, the system default access path recovery time (that is, the value that is used if you took no action to change the system-managed access path protection (SMAPP) times) changed from 60 minutes to 50 minutes. If you have explicitly changed the SMAPP times, then you are not affected by this update.

Two displays have been eliminated this release from the Work with Journals (WRKJRN) command flow. In prior releases, these displays were displayed after choosing option 6, **Recover damaged journal**. The **Display Journal Recovery Status** display was displayed after the damaged journal was deleted. A new journal receiver would be created and a new journal would be created. At this point, the **Start journaling for files** and **Start journaling for access paths** displays were displayed. These two displays listed only the file names and library of the files and access paths for which journaling to the newly created journal was restarted. In V5R4, these displays are no longer displayed.

In prior releases, the Apply Journalized Changes (APYJRNCHG) and Apply Journalized Changes Extend (APYJRNCHGX) commands could use an environment variable called QIBM_JO_APPLY_TIMEOUT to increase the time to allow the apply journal changes commands to finish. Typically, some object level changes could take up to 24 hours to complete, such as the SQL statement ALTER TABLE. The default time out for most object level changes was 5 minutes. Starting in V5R4, the following new environment variables exist for specific types of entries, each with different default values:

- QIBM_JO_APPLY_ALTER_TABLE_TIMEOUT (default value is 24 hours)
- QIBM_JO_APPLY_REFRESH_TABLE_TIMEOUT (default value is 24 hours)
- QIBM_JO_APPLY_DROP_TABLE_TIMEOUT (default value is 1 hour)
- QIBM_JO_APPLY_REORG_TIMEOUT (default value is 24 hours)
- QIBM_JO_APPLY_CM_RB_CN_TIMEOUT (default value is 12 hours)

If you were using the QIBM_JO_APPLY_TIMEOUT in a previous release, you should no longer use it for these types of entries. See the APYJRNCHG and APYJRNCHGX command help to see the new default values. These values can now be increased or decreased.

Image catalog byte stream file CCSID changes

In prior releases, the image catalog image files were set to coded character set identifier (CCSID) 37. In V5R4, the image catalog image files are set to CCSID 65535. All image files created on previous releases will remain CCSID 37. All newly created image files will be set to 65535. The Copy From Stream File (CPYFRMSTMF) command will treat these new files differently. For more information, see the Copy From Stream File (CPYFRMSTMF) command documentation.

TCP/IP server configuration information logged in QTCPSTSVRS job log

In V5R4, information regarding the starting of TCP/IP servers that are configured to start as part of the Start TCP/IP (STRTCP) command processing are now logged in the job log for the QTCPSTSVRS job. In previous releases, this information was logged in the job log for the interactive job or the job log for the batch job that was associated with the processing of the STRTCP command.

The job QTCPSTSVRS is submitted by the long running QTCPIP job (running in the QSYSWRK subsystem) during the STRTCP processing. You should now look in the QTCPSTSVRS job log for information about the starting of the TCPIP servers as part of the STRTCP command processing. The information that is logged is the same information that was available in previous releases.

The QTCPSTSVRS job runs in the QSYSWRK subsystem with user profile QTCP for a brief time period during the STRTCP processing.

In previous releases, the following information was displayed to the interactive job or logged in a job log (that is, job log nnnnnn/QPGMR/QSTRTCP), depending on how the STRTCP was processed.

```
CPC1221  Completion      00  08/20/03  19:17:21.436960
QWTCCSBJ  QSYS          0196  QTGSTART  QTCP          0397
          Message . . . . :   Job
068823/QTCP/QTVELNET submitted to job queue QSYSNOMAX
          in library QSYS.
CPC1221  Completion      00  08/20/03  19:17:21.537264
QWTCCSBJ  QSYS          0196  QTGSTART  QTCP          02EC
          Message . . . . :   Job
068824/QTCP/QTVDVICE submitted to job queue QSYSNOMAX
          in library QSYS.
CPC1221  Completion      00  08/20/03  19:17:21.659896
QWTCCSBJ  QSYS          0196  QTGSTART  QTCP          02EC
          Message . . . . :   Job
068825/QTCP/QTVDVICE submitted to job queue QSYSNOMAX
          in library QSYS.
```

This information is now logged in the job log for the QTCPSTSVRS job.

Changed behavior for /QNTC directory

In V5R4, the following changes were made to the /QNTC file system:

- In releases prior to V5R4, the QZLC_SERVERLIST environment variable was used to improve the performance for displaying the /QNTC directory.
In V5R4 the default behavior is the same as setting QZLC_SERVERLIST to 1 in previous releases. If you prefer the previous release default behavior, you need to set the QZLC_SERVERLIST environment variable to 2. If you had previously set the QZLC_SERVERLIST environment variable to 1, you do not need to take any action.
- Normal server connections are attempted using the IP name, rather than the NetBios name. This change is consistent with changes that were made in Windows 2000. In V5R4, NETSTAT shows the default QNTC connection as a **Remote Port** of cifs or 445, rather than netbios or 139.
- You might experience changed behavior in your programs if you attempt to set system attributes that are not supported by the QNTC file system.

Attempts to set an unsupported attribute using a system API will always result in an ENOTSUP(3440) error. Attempts to set an unsupported attribute using the Change Attributes (CHGATR) command will always result in the error message CPFA0AD (Function not supported by file system).

In previous releases message CPFB414 (Attributes changed for 1 object), or a successful return code, would have been received, even though the system attribute was not changed.

The following system attributes are among the system attributes that QNTC does not support:

- *CCSID
- *ALWCKPWRT
- *USECOUNT
- *DISKSTGOPT
- *MAINSTGOPT
- *RSTDRNMUNL
- *CRTOBJAUD
- *SYSARCHIVE
- *CRTOBJSCAN
- *SCAN
- *ALWSAV

Change to journaling of files created with Copy File (CPYF) command

In V5R3, you could specify that journaling should start automatically when a file was created. A data area named QDFTJRN was created in the library that was to contain the newly created file. However, journaling would not be started when the file was created using the CPYF command with the CRTFILE(*YES) option.

With V5R4, when you use the CPYF command with CRTFILE(*YES) option specified, journaling will start for the file if the QDFTJRN data area exists and indicates that *FILE objects are to be journaled at creation time. To learn more about using the QDFTJRN data area for controlling journaling, see the Journal management topic in the iSeries Information Center.

iSeries NetServer changes affecting thread behavior

A Microsoft Windows file-serving client (such as Microsoft Windows operating systems, Linux[®]/Samba, or the iSeries QNTC file system) that accesses a non-threadsafe file system using the iSeries NetServer might receive *access denied* type errors.

As of V5R4, iSeries NetServer services file shares in a multithreaded job by default. The threaded activity for all sessions in a subsystem runs in the pool of threads in the QZLSFILET job for that subsystem.

A client that is running in a threaded environment will receive access-denied type errors when it attempts to access a non-threadsafe file system (such as QDLS, QNetWare, or QSYS.LIB access of save files); similarly, the client will also receive errors when attempting to map a drive to a non-threadsafe file system. For a listing of file systems that are not threadsafe, see "File system considerations for multithreaded programming" in the Multithreaded applications topic in the iSeries Information Center.

A QZLSFILE job in the correct subsystem is still required to launch a threaded session. Whether a client can run threaded is determined when it first maps a drive to the integrated file system. The first phase of mapping the first drive for a client runs in a QZLSFILE job. If the session can run threaded, the session is transferred into the single QZLSFILET job in the subsystem. If the file system is not threadsafe, or the ADDEXITPGM THDSAFE() option for the QIBM_QPWFS_FILE_SERV exit point is specified as *UNKNOWN or *NO, or the QZLSFILET job is not present in the subsystem, the client runs in a QZLSFILE job for this session.

To prevent "access denied" type errors (when you must access a file system that is not threadsafe), the recommended solution is to not have the QZLSFILET job started in the QSERVER subsystem (or other user subsystems). This might involve configuring user subsystems in iSeries Navigator so that some clients run threaded and others non-threaded. Use the following command to remove the prestart job entry for QZLSFILET from the QSERVER subsystem:

```
RMVPJE SBSD(QSYS/QSERVER) PGM(QSYS/QZLSFILET)
```

If a prestart job entry is to be removed from a different subsystem, then that subsystem would need to be specified instead of QSERVER along with its correct library (the program would remain the same). For more information on iSeries NetServer subsystem configuration, see "Specify subsystems for iSeries NetServer" in the iSeries Information Center.

Programs created with the activation group new option (ACTGRP(*NEW)) will cause multithreaded jobs to end when the program returns. Therefore, a program created with ACTGRP(*NEW) should not be registered for exit point QIBM_QPWFS_FILE_SERV when clients might run in a threaded environment (QZLSFILET job). For more information on activation groups and threads, see "Activation groups and threads" in the iSeries Information Center.

System Request authority level changes

In V5R4, the user profile in effect for system request menu options and Presystem Request Program exit programs is the initial user profile for the job.

Prior to V5R4 (or with the application of PTF SE21256 in V5R3), the user profile in effect for system request and Presystem Request Program exit programs was the current user profile under which the job was running when system request was used. To continue to have the current user profile used by system request, you must create a data area with the following command:

```
CRTDTAARA DTAARA(QSYS/QWTSRQSSEC) TYPE(*CHAR) LEN(12)
        VALUE(*NOCHGCURUSR)
```

SSL default cipher specification list will change in the future

In a future release, the system Secure Sockets Layer (SSL) default cipher specification list will no longer contain SSL Version 2 (SSLv2) ciphers or the SSLv2 protocol in general. In addition, the Transport Layer Security (TLS) Protocol Version 1 and SSLv3 ciphers that are less than 128 bits will also be removed from the default list.

High availability clustering considerations

- | Before you upgrade your system to V5R4, if your system uses clustering and is one of the cluster nodes, you should first check the current cluster version using the Display Cluster Information (DSPCLUINF) command to ensure that it is at a level that is compatible with V5R4. At V5R4, the potential cluster version is 5. For a V5R4 node to be able to join the cluster, the current cluster version level must be at least 4. If your V5R3 cluster version is 3, you must increase the version to 4 before you upgrade.
- | See the topic Cluster version for details about this requirement. To learn more about clusters, click **Systems management > Clusters**.

Command and API changes

Removal of Display Access Group (DSPACCGRP) and Analyze Access Group (ANZACCGRP) commands

In V5R4, the Display Access Group (DSPACCGRP) and Analyze Access Group (ANZACCGRP) commands have been removed. The operating system has not collected process access group (PAG) information since V4R4.

Send PTF Order (SNDPTFORD) command and Universal Connection configuration changes

In V5R4, you can configure Universal Connection to connect with HTTP, HTTPS and HTTP proxies. When you use these new connectivity options for Electronic Customer Support (ECS) PTF ordering, there are a few changes in the way SNDPTFORD works:

- A new delivery format (*IMAGE) allows an optical image to be downloaded containing the PTFs ordered. The images will be downloaded to a directory specified in the new parameter IMGDIR.
- The limits for downloading PTFs electronically have changed. If LAN connectivity has been configured, the default is to download PTF orders of any size. If Dial connectivity is used then the default maximum size is 100 MB. You can change these values with the Set User Policy (QPDETPOL) API. You can retrieve these values with the Retrieve User Policy (QPDETRTV) API.
- You can now download the cumulative PTF package as PTF save files or as optical images. The default delivery method for the cumulative PTF package will now be electronic rather than on physical CD media. When the total size of the PTFs ordered from the cumulative PTF package is lower than the

value specified by the maximum download size limit, they will be downloaded immediately. Previously, ordering SF99vrn would always result in the cumulative PTF package being delivered on physical media.

- You can download the entire cumulative PTF package as optical images by specifying DLVRYFMT(*IMAGE) and REORDER(*YES). This lets you distribute the cumulative PTF package to other systems or create your own physical CD images.
- When you order a cumulative PTF package for the current release, and Delivery Format is set to *IMAGE and Reorder is set to *NO, an image is sent that contains any missing PTFs in the system from the current cumulative PTF package (custom cumulative PTF package).
- Downloading the cumulative PTF package by specifying DLVRYFMT(*SAVF) downloads only the PTFs that are not already on the system that is placing the order.

If Universal Connection is not configured or any problems occur when you attempt to connect with it, the previous connectivity options for Electronic Customer Support (ECS) or the SDLC configuration (if configured) will be used instead; in both cases, the new functions and behavior for Universal Connection will not be available (that is, there will be no image support and cumulative PTF packages will always be delivered as physical media).

Note: You must enable user profile QSECOFR to use the new connectivity options for Electronic Customer Support (ECS) PTF ordering.

New special values have been added to the PTFID parameter to allow for easier ordering of PTF groups:

- *CUMPKG – Order the latest level of the cumulative PTF package (SF99vrn) for the operating system release that is installed on the system. HIPER and DB2 Universal Database™ (UDB) group PTFs are automatically requested with this order. It cannot be ordered with any other PTF identifier or special value.
- *HIPERGRP – Order the latest level of HIPER PTF group for the operating system release that is installed on the system.
- *DB2GRP – Order the latest level of DB2® UDB PTF group for the operating system release that is installed on the system.
- *BRSGRP – Order the latest level of Backup Recovery Solutions PTF group for the operating system release that is installed on the system.
- *HTTPGRP – Order the latest level of IBM HTTP Server PTF group for the operating system release that is installed on the system.
- *JVAGRP – Order the latest level of Java PTF group for the operating system release that is installed on the system.
- *PFRGRP – Order the latest level of Performance Tools PTF group for the operating system release that is installed on the system.

Vary Configuration (VRYCFG) command changes

The following parameters have been removed from the Vary Configuration (VRYCFG) command:

- Reset configuration file (RESETCFGF)
- Start TCP/IP interfaces (STRTCPIFC)

Remove Server Storage Link (RMVNWSSTGL) command changes

The default value for the RENUMBER parameter has changed on the Remove Server Storage Link (RMVNWSSTGL) command. In previous releases, the default value was *YES; in V5R4, the default value is *NO.

Create Network Server Description (CRTNWSD) command changes

The TYPE parameter has changed on the Create Network Server Description (CRTNWSD) command. It now consists of two elements: Server connection and Server operating system.

The default for the Network server type (TYPE) parameter has been changed from the single value *WINDOWSNT to the two-element value (*IXSVR *WIN32). TYPE(*WINDOWSNT) is equivalent to TYPE(*IXSVR *WIN32). The value *WINDOWSNT will continue to be supported for compatibility purposes.

Service Configuration commands changed

The Create Service Configuration (CRTSRVCFG), Change Service Configuration (CHGSRVCFG), Verify Service Configuration (VFYSRVCFG), and Delete Service Configuration (DLTSRVCFG) commands have been changed in V5R4. These changes make previous uses of the command syntax incompatible with V5R4. Among these changes is the replacement of the choices for the Connection type (CNNTYPE) parameter and removal of the PTPTYPE, VPNTYPE, SERVICE, and REFSRVCFG parameters (where applicable).

If you had previously configured a direct Internet connection for the Universal Connection using either the Universal Connection Wizard, GO SERVICE, or the CRTSRVCFG command, you had to ensure that specific ports were open to the IBM VPN Gateways in any intervening firewalls. With the support added in V5R4, HTTP traffic (TCP ports 80 and 443) must be opened to additional IBM servers if this new support is to be used. For details, see the IP Packet Filter Firewall topic to learn how to determine the list of ports and addresses.

All configuration is migrated on first touch. This migration will merge your existing connectivity configuration into one common configuration. If you had previously configured different connectivity for different service applications, an attempt is made to select the highest speed available option. If all options are equal, the Electronic Customer Support (ECS) configuration is used.

Authority changes for the image catalog commands

In prior releases, you needed all object (*ALLOBJ) and security administrator (*SECADM) special authorities in order to use the image catalog CL commands. In V5R4, the image catalog commands continue to have restricted authorization (public authority is *EXCLUDE) but the *ALLOBJ and *SECADM special authorities are not required. You need *USE authority to the commands and additional authority to each image catalog and image catalog directory. Refer to each command for the authority requirements for the command. The following image catalog commands are affected by this change:

- Add Image Catalog Entry (ADDIMGCLGE)
- Change Image Catalog (CHGIMGCLG)
- Change Image Catalog Entry (CHGIMGCLGE)
- Create Image Catalog (CRTIMGCLG)
- Delete Image Catalog (DLTIMGCLG)
- Load or Unload Image Catalog (LODIMGCLG)
- Load/Unload/Mount Image Catalog Entry (LODIMGCLGE)
- Remove Image Catalog Entry (RMVIMGCLGE)
- Retrieve Image Catalog (RTVIMGCLG)
- Verify Image Catalog (VFYIMGCLG)
- Work with Image Catalogs (WRKIMGCLG)
- Work with Catalog Entries (WRKIMGCLGE)

Create Image Catalog (CRTIMGCLG) command changes

In prior releases, when you used the Create Image Catalog (CRTIMGCLG) command with the CRTDIR(*YES) parameter, the command created the image catalog directory with *RW public authority. In V5R4, the new default image catalog directory authority will be *EXCLUDE public authority. Also, the default value for the CRTDIR parameter of the Create Image Catalog (CRTIMGCLG) command has been changed from *NO to *YES.

Changes to Change Contact Information (CHGCNTINF) command

In V5R4, the Mailing address (MAILADDR) parameter was changed to split the City and state element into separate elements for City or locality and State or province. MAILADDR was also changed to split the Street address into three discrete address lines. Any CL programs or batch job streams that use the CHGCNTINF CL command must be updated to accommodate these changed MAILADDR parameter values.

Also, the second element of the Contact telephone numbers (TELNBR) parameter was redefined from the contact person's alternative telephone number to the help desk telephone number or pager telephone number for the contact person. You should check any CL programs or batch job streams that use the CHGCNTINF CL command to determine if you should specify a different telephone number for the second element of the TELNBR parameter.

Add Contact Information (ADDCNTINF) command removed

The Add Contact Information (ADDCNTINF) command was introduced in V5R3, but has been removed in V5R4. You should change any CL programs or batch job streams that use the ADDCNTINF command to use the CHGCNTINF command, which you can use to initially add contact information or change the existing contact information for the system.

Return format for Retrieve Service Attributes (QESRSRVA) API changed

In V5R4, the Retrieve Service Attributes (QESRSRVA) API Service Table Attribute Format layout for the receiver variable was changed. The Length of service attribute field was moved after the Reserved field.

Open List of ASPs (QYASPOL) API with mirrored disks

The Open List of ASPs (QYASPOL) API can be used with record format YASP0300 to return disk information in the Disk capacity and Disk storage available fields.

In V5R4, mirrored disk support changed to not require both halves of the mirrored pair to have identical capacities. The information returned in these two fields is only the disk information for this disk and should not be assumed to match the disk information for the other half of the mirrored pair.

Work with Spooled Files (WRKSPLF) command supports generic user data for SELECT

In V5R4, the Work with Spooled Files (WRKSPLF) command was enhanced to support the selection of spooled files by specifying a generic value for the User data element of the SELECT parameter. Because the user data field for a spooled file is a character field, it can contain one or more asterisk characters.

For example, a spooled file could have the four characters of user data ABC*. In releases before V5R4, specifying ABC* for the user data element of the SELECT parameter on WRKSPLF would list only those spooled files with those four characters of user data. In V5R4, specifying ABC* causes the WRKSPLF command to list all spooled files for which the first three characters of user data are ABC.

You might need to change any of your applications that make use of the asterisk (*) character in a spooled file's user data to handle additional records for spooled files selected using the generic user data support.

Changes to CHGAUD, CHGAUT, CHGOWN, and CHGPGP command processing

When a group of objects is being changed, either because a pattern was specified in the object name or SUBTREE(*ALL) was specified, an informational message is no longer sent for each changed object. Instead, one completion message is sent that indicates the number of objects changed. If any objects cannot be changed, a diagnostic message is sent for each object that was not changed, and a summary escape message is sent at the end of the command processing.

Escape message CPFA093 (Name matching pattern not found) is no longer sent. Instead, escape message CPFBC50 (Path name or path names not found) is sent.

Reclaim Activation Group (RCLACTGRP) command does not end some IBM activation groups

After running RCLACTGRP ACTGRP(*ELIGIBLE), you might notice (using option 18 from the Work with Jobs (WRKJOB) command) that some named activation groups remain on the list. These are IBM activation groups. If you want to reclaim one of these IBM activation groups, run the RCLACTGRP and specify the name of the activation group for the ACTGRP parameter.

***WINDOWS value replaces *WINDOWSNT value on network server commands**

Several parameters on the Change Network Server Attributes (CHGNWSA), Display Network Server User Attributes, Submit Network Server Command (SBMNWSCMD), Work with Network Server User Enrollment (WRKNWSENR), and Work with Network Server Status (WRKNWSSTS) CL commands that allowed a value of *WINDOWSNT now accept the special value *WINDOWS. *WINDOWSNT is still accepted for compatibility with prior releases, but *WINDOWS is the preferred equivalent value. Support for the *WINDOWSNT special value could be removed in a future release of the operating system.

CMDTYPE parameter ignored on Submit Network Server Command (SBMNWSCMD)

Starting with V5R4, the Command type (CMDTYPE) parameter is ignored on the Submit Network Server Command (SBMNWSCMD) command. Support for the CMDTYPE parameter could be removed in a future release of the operating system.

Install Windows Server (INSWNTSVR) command changes

In V5R4, the following changes have been made to the Install Windows Server (INSWNTSVR) command:

- The SPD-attached hardware (6617 and 2850) is no longer supported and is the only hardware that allowed an external host LAN to be installed and configured using the PORT1 and PORT2 parameters. These two parameters have been removed.
- The Internal LAN port (INTLANPORT) parameter was used to install on SPD-attached hardware or Windows version 4.0 servers. Both are no longer supported for new Windows installs. This parameter has been removed. The Virtual PTP Ethernet port (VRTPTPPORT) parameter is the functional replacement that is used for all supported configurations.
- The Domain role (DMNROLE) parameter is no longer needed on Windows 2000 and Windows Server 2003 installs. All servers are installed as *SERVER and later changed to *DMNCTL during post-installation steps that are manually performed by the user by installing Active Directory using the Windows DCPRMO command. This parameter has been removed from the INSWNTSVR command.

- The default value for the Convert to NTFS (CVTNTFS) parameter has changed to *YES in order to promote the additional security provided by the NTFS file system, as well as to enhance the virtual disk performance by creating virtual disk images that align logical file system clusters with host pages.

Changes to Trace commands

In V5R4, the Trace Job (TRCJOB) command was changed to run the Start Trace (STRTRC), End Trace (ENDTRC), and Print Trace (PRTRC) CL commands that were added to the operating system in V5R1. Because there is no equivalent function on the STRTRC command, the **Program to call before trace** (EXITPGM) and **Select procedures to trace** (SLTPRC) parameters are not supported. Any values specified for these parameters are ignored. You might need to change any CL programs that used the TRCJOB command and specified the EXITPGM or SLTPRC parameters.

Also, the output spooled file and the database output file produced by the Print Trace (PRTRC) command have a different format from those produced in prior releases when the TRCJOB command was run specifying SET(*OFF). You will need to change any of your programs that are dependent on the layout of the spooled file produced by the TRCJOB command.

The default value of **Length of time to watch** (WCHTIMO) parameter on the Start Trace (STRTRC), Trace Internal (TRCINT), Trace Connection (TRCCNN), Trace TCP/IP Application (TRCTCPAPP), and Start Communications Trace (STRCMNTRC) commands has changed from one day to *NOMAX. This change keeps the behavior of the trace commands the same, regardless of whether the watch parameters are specified.

In addition, processing of the watch comparison data has changed. Comparisons can now span fields and are done without any conversion of the comparison data (for instance, character to integer). This change was previously made available in V5R3 with PTF SI14367.

New TCP/IP API formats to support IPv6

New API formats were added in support of IPv6. The existing formats can still be used, but they will not be enhanced in future releases. If you use IPv6, you are encouraged to use the new formats. New fields have been added to the NIFC0100, NIFC0200, NRTE0200, and IFCD0300 formats.

Also, a new format, named TCPA1300, was added to the QtocRtvTCPA API. You should use this format instead of the TCPA1200 format, which is not supported starting in V5R4. TCPA1200 continues to return TCP-related information, but returned IP-related data is set to 0.

In addition to the new fields in the returned data, some data that is no longer available returns only a default value.

The following table shows the APIs and formats (along with the fields that are no longer available) that are affected by this change.

Table 2.

API name	Format	Fields that are no longer available
QtocLstNetIfc	NIFC0200	Address state Duplicate address detection transmits Multicast - number of references
QtocLstNetRte	NRTE0200	Is on-link Next hop address family Next hop IPv4 Next hop IPv4 binary Route lifetime at creation Route lifetime remaining

Table 2. (continued)

API name	Format	Fields that are no longer available
QtocLstPhyIfcDta	IFCD0300	Accept redirects Accept router advertisements Local tunnel endpoint IPv4 address Local tunnel endpoint IPv4 address binary Neighbor discovery base reachable time - configured Neighbor discovery base reachable time - current Neighbor discovery reachable time Neighbor solicitation retransmit interval - configured Neighbor solicitation retransmit interval - current Use other stateful configuration Use stateful address configuration
QtocRtvTCPA	TCPA1200	ICMP error message send rate time Router solicitation max delay Router solicitation interval Router solicitation max transmits Neighbor advertisement max transmits Neighbor solicitation delay first probe time Neighbor solicitation max unicast solicits Neighbor solicitation max multicast solicits

List Database File Members (QUSLMBR) API change

The List Database File Members (QUSLMBR) API now returns an offset of zero for the "Offset to member description information" when the MBRL0300 format is requested, and the member's description information cannot be retrieved. In previous releases, the member would not have been returned in the list if its description information could not be retrieved.

Default will change for QUSCRTUS parameter

In the release after V5R4 of i5/OS, the default value for the **Optimum alignment** parameter of the Create User Space (QUSCRTUS) API will change from 0 to 1, so optimum alignment will become the default. This change will reduce the maximum amount of space available for user spaces created with the default value. The maximum size of an optimally aligned space can be programmatically retrieved using option Hex 0003 of the MATMDATA (Materialize Machine Data) MI instruction.

Changes to the Copy From Import File (CPYFRMIMPF) and Copy To Import File (CPYTOIMPF) commands

In previous releases, string escape character processing did not exist.

In V5R4, this functionality is being introduced with the **String escape character** (STRESCCHR) parameter on the Copy From Import File (CPYFRMIMPF) and Copy To Import File (CPYTOIMPF) commands. Exported data will be affected by the character specified as the escape character following each string delimiter that is part of the field data. For imported data, string escape character processing will remove this character when the string delimiter is found inside the field data following this.

When the character specified as the escape character and the string delimiter are different, the escape character inside the field data will be followed by another escape character when exporting data, and will be removed during the import.

Changes to Change Job (CHGJOB) and End Job (ENDJOB) commands

Change Job (CHGJOB) and End Job (ENDJOB) commands now work on any completed job

In V5R4, the Change Job (CHGJOB) and End Job (ENDJOB) commands now work on any completed job. In previous releases, these commands would not work on completed system jobs, completed subsystem monitor jobs, completed spool reader jobs, or completed spool writer jobs.

End Job (ENDJOB) command sends fewer messages

When you use the End Job (ENDJOB) command to end spooled files or change the maximum log entries (LOGLMT) for a completed job, you will receive fewer messages. CPF1362 (Job has completed) is now sent as a completion message from ENDJOB processing rather than being sent as an exception. This allows multiple ENDJOB commands from the Work with Submitted Jobs (WRKSBJOB) panel or from the Work with User Jobs (WRKUSRJOB) panel to affect completed jobs. ENDJOB LOGLMT(0) will now remove a job from the job log pending state. If you have programs that monitor for CPF1362, you might need to modify your programs. For example, if a program uses the CPF1362 escape message to determine the state of the job, the program should be changed to use the Retrieve Job Status (QWCRJBST) API.

Changes to Create Job Description (CRTJOB) and Change Job Description (CHGJOB) commands

For job description QDFTJOB in library QGPL and job description QDFTSVR in library QGPL, an ASP group name cannot be specified for the Initial ASP group (INLASGRP) parameter of the Change Job Description (CHGJOB) and Create Job Description (CRTJOB) commands.

Backup and recovery changes

Changes to Display Save and Restore media panels

In V5R4, data displayed and printed for the DSPTAP DATA(*SAVRST), DSPOPT DATA(*SAVRST), and DSPSAVF commands differs from previous releases.

- These operations now use the print file QPSRODSP. Previously, the print files QPSRODSP or QSYSPRT could have been used. The print file QSYSPRT had been used by the DSPOPT command and by DSPTAP and DSPSAVF when integrated file system data had been saved.
- Each of these commands now produces one spooled file. Previously, the number of spooled files produced depended on whether integrated file system data had been saved. A separate spooled file was produced for each tape file or optical file that contained saved integrated file system data.
- Header information has been rearranged to a single column format. For displayed output, a function key is used to display the header information. Previously, DSPTAP and DSPSAVF showed header information in a double column format, and header information and data were shown on the same display panel.
- When displaying objects in a library, the Option (Opt) field is active for all objects, but valid only for database files and output queues. The text for Option 5 changed to *Display*. Previously, the Option 5 text was *Display saved database file members*, and the Option field was only active for database files for the DSPTAP and DSPSAVF commands.
- Database members are printed in a separate section following the list of objects. A text column was added to the list of objects where the members previously appeared.
- The number of objects and access paths is printed in a summary section. Previously, these values were printed in the header information.

Changes for parallel save operations to tape

V5R4M0 is the last release to use the existence of data area QTEMP/QSRPARFMT to indicate that multiple libraries are to be saved in parallel format. The data area is used with any media definition created with format TAPE0100, and with media definitions created with format TAPE0200 that contain the value -2 for the Save format field. In future releases, to save multiple libraries in parallel format, you must use a media definition created with format TAPE0200 that contains the value 0 for the field Save format.

If you use dynamic device allocation, some potential operational differences might exist:

- A save operation might fail with message CPF387C (Minimum device resources not available) later than it did before, requiring more operator attendance.
- Inquiry messages regarding tape devices and media might surface later than they did before, requiring more operator attendance.
- The *Total files* value shown by the Display Tape (DSPTAP) command, for data saved in parallel format, will be the maximum potential number that the save operation could have created. The actual number could be less if the save operation did not allocate all of the drives.
- The tape density on the first device that is allocated will be used instead of the most common density found. If differing densities are found among the specified devices, this could result in fewer devices being used than before.

| **PTF adds additional error checking to the restore function**

| In V5R4, the program temporary fix (PTF) SI23983 adds additional error checking to the restore function.
| If you try to restore a device file in a job that already has the file open, the file will not be restored, and
| the following messages are sent:

- | • CPF0675 – Device file is in use.
| • CPD3737 – Object not restored.

| This additional error checking has also been added to the restore function in V5R2 via PTF SI24174 and
| in V5R3 via PTF SI22948.

DB2 UDB for iSeries changes

BLOB comparisons change

In releases prior to V5R4, query comparisons of BLOB operands that depend on order (for example, >, >=, <, <=) returned incorrect results if one operand was longer or shorter than the other. This is not always correct. In V5R4, these comparisons use binary string comparison rules. If two binary strings are equal up to the length of the shorter string, the shorter string is considered less than the longer string even when the remaining bytes in the longer string are hexadecimal zeros.

The following SQL queries illustrate this change. In V5R4, both queries correctly return all records in the file; in releases prior to V5R4, however, no records are returned.

```
SELECT * FROM QSQPTABL  
WHERE BLOB(X'0102030400')>BLOB(X'01020304')
```

```
SELECT * FROM QSQPTABL  
WHERE BLOB(X'01020305')>BLOB(X'0102030400')
```

Format change of additional parameters that are passed to a procedure defined with PARAMETER STYLE SQL

C and C++ programs that are invoked using SQL CALL statements and that were defined using a CREATE PROCEDURE statement which specified PARAMETER STYLE SQL will see a change in the format of the additional parameters that are passed to the procedure.

These additional parameters are now correctly passed as C null-terminated strings instead of varying character strings. This change was made to provide a consistent interface for functions and procedures defined with PARAMETER STYLE SQL.

Conversion of database file (*FILE) objects for optimal alignment

In V5R4, existing database files will be converted to optimally align the underlying machine interface (MI) objects. Database files consist of multiple MI objects, so the time to convert them could be significant when an application is started the first time after V5R4 is installed.

You can allow the operating system to perform these conversions as your applications use your database files. However, if you would like to avoid these delays when starting applications after V5R4 is installed, you can force the conversion by running either of the following CL commands:

- Use the Display Library (DSPLIB) command to convert all database files in a particular library or a group of libraries.
- Use the Display Object Description (DSPOBJD) command to convert a particular database file, or to convert all or a generic set of database files in a particular library or a group of libraries. Specify OBJTYPE(*FILE) DETAIL(*BASIC) when using the DSPOBJD command.

These database file conversions will cause the modification time stamp of each affected object to be updated.

Conversion scenarios

The following tables show the timing details of the database file conversion process on a V5R3 database that consisted of 45 517 database files that had 254 379 MI objects that needed to be converted.

Note: The results shown are meant to provide only a rough guideline for the amount of time the conversion process might take on similar configurations. The approximate time for the conversion was estimated based on the total time the DSPOBJD command consumed when the conversion took place minus the total time the DPSOBJD command consumed after all conversions had taken place. If you use the DSPOBJD command to force the conversion, you do not need to run the command a second time.

The following commands were used for the measurements:

```
DSPOBJD OBJ(LIBRARY/ONE_FILE) OBJTYPE(*FILE) DETAIL(*BASIC)
        OUTPUT(*OUTFILE) OUTFILE(QTEMP/DSPOBJDATA)
```

```
CHGPF FILE(QTEMP/DSPOBJDATA) SIZE(*NOMAX)
```

```
RMVM FILE(QTEMP/DSPOBJDATA) MBR(DSPOBJDATA)
```

```
DSPOBJD OBJ(LIBRARY/*ALL) OBJTYPE(*FILE) DETAIL(*BASIC)
        OUTPUT(*OUTFILE) OUTFILE(QTEMP/DSPOBJDATA)
```

On a dedicated model 830 system with six processors in a 6402 M pool, the DSPOBJD command that converted all of the remaining database files in the library took approximately 91 minutes. The same DSPOBJD command run after all of the files had been converted took approximately 42 minutes. The conversion of the 254 379 MI objects therefore took approximately 49 minutes.

Table 3. Conversion process on a dedicated model 830 system with six processors in a 6402 M pool

	Clock seconds	CPU seconds	Sync reads (DB)	Sync reads (non-DB)	Sync writes (DB)	Sync writes (non-DB)	Async reads (DB)	Async reads (non-DB)	Async writes (DB)	Async writes (non-DB)
DSPOBJD command that converted all of the remaining database files	5446.8	188.98	300675	451k	153574	2794k	71057	31010	1120	417636
DSPOBJD command that was run after the file conversion	2526.3	23.158	2	565k	163	174	73744	6350	1201	4

On a dedicated model 825 system with six processors in a 3271 M pool, the DSPOBJD command that converted all of the remaining database files in the library took approximately 68 minutes. The same

DSPOBJD command run after all of the files had been converted took approximately 27 minutes. The conversion of the 254 379 MI objects therefore took approximately 41 minutes.

Table 4. Conversion process on a dedicated model 825 system with six processors in a 3271 M pool

	Clock seconds	CPU seconds	Sync reads (DB)	Sync reads (non-DB)	Sync writes (DB)	Sync writes (non-DB)	Async reads (DB)	Async reads (non-DB)	Async writes (DB)	Async writes (non-DB)
DSPOBJD command that converted all of the remaining database files	4104.8	190.24	302411	460k	153574	2795k	71876	30999	1513	417818
DSPOBJD command that was run after the file conversion	1648.4	19.383	2	558k	163	176	73744	6350	1139	8

On a dedicated model 570 System with one processor in a 512 M pool, the DSPOBJD command that converted all of the remaining Database files in the library took approximately 47 minutes. The same DSPOBJD command run after all of the files had been converted took approximately 21 minutes. The conversion of the 254 379 MI objects therefore took approximately 26 minutes.

Table 5. Conversion process on a dedicated model 570 System with one processor in a 512 M pool

	Clock seconds	CPU seconds	Sync reads (DB)	Sync reads (non-DB)	Sync writes (DB)	Sync writes (non-DB)	Async reads (DB)	Async reads (non-DB)	Async writes (DB)	Async writes (non-DB)
DSPOBJD command that converted all of the remaining database files	2797.1	83.440	304184	533k	153574	2794k	74316	35318	2063	418029
DSPOBJD command that was run after the file conversion	1242.2	11.749	2	573k	163	179	73744	6350	2062	4

DB2_ROW_COUNT_SECONDARY field in the SQL Diagnostics Area changed

The meaning of the DB2_ROW_COUNT_SECONDARY field in the SQL Diagnostics Area changed in V5R4.

In prior releases, the field indicated the number of rows affected by cascaded delete operations that were needed to maintain referential integrity.

In V5R4, the field also includes the number of rows affected by processing INSERT, UPDATE, and DELETE triggers. This includes BEFORE and AFTER triggers, as well as INSTEAD OF triggers.

Triggers not allowed on IDDU physical and logical files

In previous releases, there was no restriction to prevent users from creating triggers on interactive data definition utility (IDDU) files. In V5R4, triggers cannot be created on physical or logical IDDU files.

Automatic conversion of *PGM, *SRVPGM, *MODULE, and *SQLPKG objects that include SQL statements

The internal representation of SQL statements in *PGM, *SRVPGM, *MODULE, and *SQLPKG objects changed in V5R4.

*PGM, *SRVPGM, and *SQLPKG objects that were created prior to V5R4 will be converted automatically to the new format on first use. You can also initiate object conversions using the Start Object Conversion (STROBJCVN) command.

*MODULE objects that were created prior to V5R4 cannot be permanently converted, and a conversion will be performed each time the *MODULE is bound into a program or service program.

The amount of time to convert a single object is small, less than 160 milliseconds of CPU time on a model 890 system for a program approaching the system limit on number and size of SQL statements. On a model 840 system, it is less than 350 milliseconds. The conversion time is proportional to the number of SQL statements in the object and to the cumulative size of the SQL statements. If large numbers of objects need to be converted, the conversion time could be noticeable.

For objects created in V5R2 or later that approach the system limit on number and size of SQL statements, the conversion might fail because of space limitations. These programs will perform poorly on releases prior to V5R4 because of insufficient space to store access plans, which need to be rebuilt in temporary storage each time an SQL statement is executed. These situations are expected to be rare. If you encounter this situation, you should recreate the object after reducing the number or size of SQL statements.

Current optimization degree changes

In prior releases, you set the optimization degree for queries using the DEGREE parameter on the Change Query Attributes (CHGQRYA) command.

In V5R4, you can also set the optimization degree in SQL procedures. If the change to the degree is done in an SQL procedure, the procedure exit code forces the value to what it was when the procedure was started. If the change to the degree is made outside an SQL procedure (for example, by a CHGQRYA command run from a command line), the change is persistent.

Database Monitor table format changed

In V5R4, 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 iSeries Navigator to update the file to the new format. The following iSeries Navigator functions update the file to the new format:

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

Also in V5R4, a column of one of the files used by the memory-based SQL Database Monitor was changed. The User-defined field (QQUDEF) in the Query Information file was changed from a character field of length 18 to a character field of length 100. Any existing versions of this file will be automatically updated the first time the Dump SQL Database Monitor (QQQDSDBM) API is used with that file.

Changes to SQL

Recursive common table expressions

If a common table expression is specified in an SQL statement and the common table expression table-identifier has the same name as an unqualified table referenced in the definition of the common-table-expression, the common table expression is considered to be recursive.

Recursive references were not allowed in V5R3; SQLCODE -346 SQLSTATE '42836' was issued for this in V5R3 unless a default collection-id was specified. If a default collection-id was specified, the table in the default collection is used in the implementation of the query.

In V5R4, a table-name that is referenced in a common table expression and has the same name as the common table expression will be considered recursive. If the recursive reference is not valid, an error will be returned.

Non-deterministic functions

Functions that are non-deterministic cannot be specified in a JOIN clause or a GROUP BY clause. This error was not issued on prior releases. In V5R4, SQLCODE +583 SQLSTATE '01565' is returned when functions that are NOT DETERMINISTIC are specified in a JOIN clause or a GROUP BY clause.

The default for user defined functions is NOT DETERMINISTIC. When user defined functions that are NOT DETERMINISTIC are specified in a JOIN or GROUP BY clause, SQLCODE +583 will be returned. The scalar functions RAND and IDENTITY_VAL_LOCAL are considered non-deterministic and also should not be specified in the JOIN clause and GROUP BY clause.

Partitioning key values

If MINVALUE or MAXVALUE is specified for a partitioning key value, all subsequent column values for that range must be the same value. In V5R3, (minvalue, 1) was allowed for the starting-clause of a partitioning-clause. In V5R4, when minvalue is specified for the first column, it must also be specified for the second column. If MINVALUE or MAXVALUE is specified followed by a literal, SQLCODE -636 SQLSTATE '56016' is returned.

Partitioning key values without overlap

In V5R4, the algorithm for determining the table partition that rows are inserted into has changed from V5R3. This change affects partitioned tables that use range partitioning, where the partitioning key is made up of multiple columns. This enhanced algorithm is less restrictive than the algorithm that was used in V5R3. If none of the key ranges for a given partition overlap with corresponding key ranges of other partitions, then the new algorithm checks only enough of the new row's key values, beginning with the leading value, as is necessary to uniquely identify which partition the row belongs in.

For example, consider a partitioned table created with the following statement for V5R4:

```
CREATE TABLE NOOVERLAP (COL1 INT ,COL2 INT) PARTITION BY (COL1,COL2)
  (STARTING FROM (1,1) INCLUSIVE ENDING AT (10,10) INCLUSIVE,
   STARTING FROM (11,11) INCLUSIVE ENDING AT (20,20) INCLUSIVE)
```

Data allowed into the first partition would be determined by the following condition:

```
COL1 > 1 AND COL1 < 10 OR
COL1 = 1 AND COL2 >= 1 AND COL2 <= 10 OR
COL1 = 10 AND COL2 >= 1 AND COL2 <=10
```

Data allowed into the second partition would be determined by the following condition:

```
COL1 > 11 AND COL1 < 20 OR
COL1 = 11 AND COL2 >= 11 AND COL2 <= 20 OR
COL1 = 20 AND COL2 >= 11 AND COL2 <=20
```

This means that if a partitioned table is restored to V5R4 from V5R3, or the system migrates to V5R4 with a partitioned table created on the previous release, data allowed to be inserted into a partitioned table will not be consistent.

Compare the previous example to the V5R3 behavior:

```
CREATE TABLE NOOVERLAP (COL1 INT ,COL2 INT) PARTITION BY (COL1,COL2)
  (STARTING FROM (1,1) INCLUSIVE ENDING AT (10,10) INCLUSIVE,
   STARTING FROM (11,11) INCLUSIVE ENDING AT (20,20) INCLUSIVE)
```

Data allowed into the first partition would be determined by the following condition:

```
COL1 => 1 AND COL1 =< 10 and COL2 => 1 AND COL2 =< 10
```

Data allowed into the second partition would be determined by the following condition:

```
COL1 => 11 AND COL1 =< 20 and COL2 => 11 AND COL2 =< 20
```

Unnamed result columns

If a result column of a fullselect specified in a derived table or common table expression is unnamed, the column cannot be referenced. SQLCODE -206 SQLSTATE '42703' is returned. The result column is unnamed when the corresponding items of all the subselects are not the same name. In V5R3 SQLCODE -153 SQLSTATE '42503' was returned when the columns were unnamed and column names were specified. If the subselect contained constants or expressions, an error was not issued. An AS clause must be added to the select list to make the result named if the column is referenced.

Procedure result set scoping

In V5R3 and prior releases, cursors that are opened in a procedure, left open when the procedure exits, but whose result set was not returned as a procedure result set might not have been accessible as an open cursor on subsequent invocations of the procedure. This problem was corrected in V5R4. As a result of this change, application programs with SQL that were precompiled with a CLOSQLCSR option of *ENDJOB, *ENDSQL, or *ENDACTGRP, are called as stored procedures, and which do not return result sets for all cursors as stored procedure result sets might see a new error message in some scenarios.

The new error message is returned only in cases where the same procedure is called under multiple client statement handles and attempts to open a cursor that was also left open by a previous invocation of the procedure that was called under a different client statement handle. The new error message, SQL7055, indicates that the cursor is ambiguous within the job.

In the unlikely event that you receive this message, you can correct the problem in one of the following ways:

- Change the CLOSQLCSR option that the program was compiled with to *ENDMOD or *ENDPGM
- Explicitly close the cursor in the procedure before exiting the procedure

Parameter marker attributes

In a DESCRIBE input operation, parameter marker attributes, such as type, length, and name, are sometimes derived based on the context of the expression in which they are specified in the SQL statement.

In V5R3 and prior releases, a DESCRIBE input operation could incorrectly return a column name in addition to the other attributes of the parameter marker when the column name was used in the expression with the parameter marker.

This was corrected in V5R4 to follow the ANS/ISO SQL standard behavior, which is to return an implementation-defined value that indicates that the name is derived. In addition, an indication is returned that the parameter marker is unnamed.

Client interfaces such as ODBC, JDBC, and OLE DB, provide APIs to return these parameter marker attributes back to the calling application. These applications will now see this corrected behavior.

Database host server and DRDA/DDM server jobs set the current directory

In V5R4, the database host server and the DRDA/DDM server sets the current directory to the home directory as specified for the user profile that is connecting to the server job. Previously, the servers would only set the home directory for the job, and did not set the current directory. This could affect any commands run in these server jobs that used a non-qualified integrated file system path.

The server jobs now also set the following attributes:

- Country or region ID
- Character identifier control
- Language ID

- Locale
- Status message handling

Temporary storage

SQL query processing continues to be enhanced with new capabilities. Some of these capabilities involve the creation of temporary objects to either satisfy the request or to improve the query's performance. To further improve overall performance, the database might also retain temporary objects for a duration of time after the query request has finished when there is a high likelihood that the result can be reused in a subsequent query request.

With the additional capabilities of temporary objects and their extended retention duration, the amount of temporary storage used for SQL queries might increase. This, in turn, might increase the amount of temporary storage used on the system compared to prior releases. The database monitors its use of temporary storage to prevent the amount used from becoming critical to the system. However, if minimal temporary storage usage is an issue for your environment, you can minimize the amount of temporary storage used by specifying a QAQQINI file with option `CACHE_RESULTS` set to either `*JOB` or `*NONE`. The value of `*JOB` lessens the retention of temporary objects to only the life of the job using the objects. A value of `*NONE` causes temporary objects to be freed up as soon as the query request has completed. Note that if the `CACHE_RESULTS` option is overridden to `*JOB` or `*NONE` (from the default of `*SYSTEM`), query performance can be negatively impacted. See the Database performance and query optimization topic for more information on using database QAQQINI files.

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

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

Changes to the SQL `LOCATE` scalar function

►► `LOCATE`—(*—search-string—*, *—source-string—* , *—start—*)—

In releases prior to V5R4, when the *start* position was specified for the `LOCATE` scalar function, the resulting answer was with respect to the *start* position within the source-string. The formula was:

POSSTR(SUBSTR(*source-string*,*start*) , *search-string*)

For example:

```
SELECT LOCATE('C','CABC',2)
FROM QSYS2.QSQPTABL
```

Is equivalent to:

```
SELECT POSSTR(SUBSTR('CABC',2))
FROM QSYS2.QSQPTABL
```

Which reduces to:

```
SELECT POSSTR('ABC', 'C')
FROM QSYS2.QSQPTABL
```

The result is 3.

Program Temporary Fix (PTF) SI24667 was created in V5R4 to make `LOCATE` work as documented in the DB2 UDB for iSeries SQL Reference topic, as well as to be compatible with other DB2 implementations. The new formula is:

```
| POSSTR( SUBSTR(source-string,start) , search-string ) + start - 1
```

| For example:

```
| SELECT LOCATE('C','CABC',2)
| FROM QSYS2.QSQPTABL
```

| Is equivalent to:

```
| SELECT POSSTR(SUBSTR('CABC',2,'C')) + 2 - 1
| FROM QSYS2.QSQPTABL
```

| Which reduces to:

```
| SELECT POSSTR('ABC', 'C') + 2 - 1
| FROM QSYS2.QSQPTABL
```

| The result is 4.

Machine interface (MI) instruction changes

Header file for the DESMTX MI instruction changed

The header file for the Destroy Pointer-Based Mutex (DESMTX) MI instruction has changed. The file `mih/desmtx.h` contains an enum symbol called `_Destroy`. The symbol `_Destroy` is now qualified as `_DESMTX_Destroy` to prevent namespace collisions with `std/xmemory.h`. If the enum symbol `_Destroy` has been referenced in any code that includes `desmtx.h`, future compiles will fail. You should change code that includes the `mih/desmtx.h` header file to use `_DESMTX_Destroy` instead of `_Destroy`.

Maximum number of outstanding allocations field removed from the output of the MATHSAT and MATHSAT2 MI instructions

The Maximum number of outstanding allocations (`Max_Outstand`) field has been removed from the output of the Materialize Activation Group-Based Heap Space Attributes (MATHSAT and MATHSAT2) MI instructions.

The MATSOBJ field for returning information on the storage protection of encapsulated objects made obsolete

The MATSOBJ field for returning information on the storage protection of encapsulated objects has been made obsolete in V5R4, because the enhanced protections now used for most objects cannot be represented in this field. The value returned for this field will be binary zero. However, because the field is obsolete, a binary zero value no longer means "reference and modify allowed for user state programs," as it did on previous releases.

In the unlikely event that you have used this value in any of your programs to take some specific action, the program might behave differently.

LOCKOL, UNLOCKOL, and SETOBPFP MI instructions not supported

In V5R4, use of the unblocked MI instructions Lock Object Location (LOCKOL), Unlock Object Location (UNLOCKOL), or Set Object Pointer From Pointer (SETOBPFP) causes a function check (MCH3203) exception to be signaled.

New option added to the MATRMD MI instruction

In V5R4, a new option `X'2D'` was added to the Materialize Resource Management Data (MATRMD) MI instruction to support larger main storage sizes in the future. If you use the MATRMD option `X'09'`, no immediate changes are required to use this new option; however, there are fields in option `X'09'` that will

overflow if the main storage size does not fit in a 4-byte field. For example, the main storage size field will overflow if the main storage size is 8 terabytes in a future release.

MATMIF MI instruction changed

The name of the field `Unalloc_Pool_Proc_Cap` in the `_MMIF_OPT_0002_T` structure of the `matmif.mih` header file changed in V5R4 to `Unalloc_LP_Grp_Proc_Cap`. This field returns the amount of unallocated processing capacity in the partition group.

Header file for the MATMTX MI instruction changed

The header file for the Materialize Mutex (MATMTX) MI instruction has changed. The file `mih/matmtx.h` contains a field called `Mutex_Type` of the type `_Mutex_Type_T`. This field is reserved and always contains binary zero. The `Mutex_Type` field is now removed, and the space formerly used for `Mutex_Type` is now reserved in `matmtx.h`.

If you have referred to the `Mutex_Type` field in any code that includes `matmtx.h`, future compiles will fail. The reserved space formerly referred to by `Mutex_Type` will continue to contain binary zero.

Restriction on LOCKTSL MI instruction

Starting with the V5R4 release of i5/OS, the `Lock_Mapped_Location` field in the `milckcom.mih` header file is ignored. This header file is for the Lock Teraspace Storage Location (LOCKTSL) MI instruction. The ability to use a space location lock on a single-level store object mapped into teraspace is no longer supported.

Capacity donation and processor utilization

Capacity donation by dedicated processor logical partitions (LPARs) does not affect the reported processor utilization of the donating LPAR because donation occurs only when a processor is idle. Over any interval, neither the processor available time nor the processor utilized time is affected by donation. However, for shared processor LPARs, donation may affect the results returned by the following instructions:

- MATRMD option Hex 27 - Shared Processor Pool Information

The processors of LPARs configured as donation capable are counted in the shared processor pool. Consequently, the processor count, processor available time, and processor utilized time of the shared processor pool includes the contribution of the processors of dedicated processor partitions configured for donation. For example, when an active dedicated processor LPAR is changed from non-donating to donating, the shared processor pool's processor count increases by the dedicated processor LPAR's processor count, and the rate of accumulation of the pool's processor available time will increase accordingly. The pool's processor utilized time will be affected in the usual way. This means that it will include all processor time that the pool's processors are used by one of the member LPARs, shared and donating dedicated LPARs alike.

Note that a donating dedicated LPAR is free to limit the amount of idle time it donates to the shared pool, so in the extreme it is possible for a donating dedicated LPAR to be 100% idle (reporting 0% utilization) while all of its processor time is reported as utilized time in the shared pool. The same is true for shared processor LPARs, but by convention, they try to minimize idle processor time, so this phenomenon is rarely noticed.

- MATRMD option Hex 26 - Processor Utilization Data

The processors of LPARs configured as donation capable are counted in the shared processor pool. Consequently, the uncapped available time and processor utilized time may include contributions from the processors of dedicated processor partitions configured for donation.

- MATRMD option Hex 28 - Multiprocessor Utilizations

The processors of LPARs configured as donation capable are counted in the shared processor pool. Consequently, the uncapped available time and processor utilized time may include contributions from the processors of dedicated processor partitions configured for donation.

Server firmware PTFs on systems running Licensed Internal Code 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 releases.

For POWER5 and POWER6™ hardware running Licensed Internal Code V5R4M5 or later releases, server firmware PTFs will no longer be associated with Product ID 5722999 (Licensed Internal Code). The PTFs will be associated with a new Product ID, 57339xx, where xx depends upon the hardware platform. The operational changes to specify a different Product ID, 57339xx, 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 (57339xx).

File-level backup through QNTC file systems

If your PC has Windows 2000 installed and you upgrade to V5R4, 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.

Host print transform changes

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

Chapter 3. Options

GDDM (Option 14)

APAR SE17198 reported that, in V5R2, the Graphic Symbol Set objects (*GSS objects) shipped with the operating system option GDDM® (Option 14) that were named as if they were Simplified Chinese were actually Traditional Chinese.

For V5R4, these objects have been fixed as follows. The current objects are renamed using the naming convention defined in the GDDM documentation for Traditional Chinese objects. New Simplified Chinese objects have been added to V5R4 and named correctly. Therefore, if you are now getting Traditional Chinese symbols as your application expects, you will get Simplified Chinese after the install of V5R4 i5/OS Option 14. You can correct this by changing your Load Graphic Symbol Set (GSLSS) calls to use "ADMXT*" for the object name instead of "ADMXC*".

NetWare Enhanced Integration (Option 25)

NetWare Enhanced Integration (option 25) will be removed from i5/OS in a future release. Option 25 is still available in V5R4 as a separately-ordered no-charge feature that you can select from the configurator. This option has a non-warranted license with no committed support for V5R4.

The NetWare Loadable Module (NLM) provided with Option 25 interoperates with NetWare 6.0, but is not going to be enhanced to run on current versions of NetWare. To continue using NetWare with iSeries integration, consider migrating your NetWare server to Linux on an integrated xSeries® solution (IXS) and take advantage of Novell NetWare services that are available on Linux. Linux on the IXS provides some of the integration features previously available with Option 25, along with other integration features that can help manage your servers. For more information on integrated xSeries solutions, see the following Web site:

<http://www.ibm.com/servers/eserver/iseries/integratedxseries>

i5/OS Portable Application Solutions Environment (Option 33)

- | The program temporary fixes (PTFs) MF39106 and SI23443 provide full support for the following changes to the i5/OS PASE environment.
- | **i5/OS PASE support for execution priority changed**
 - | The mapping between values specified on the i5/OS PASE nice command and those specified on the i5/OS Run Priority (RUNPTY) parameter has changed. A single unit specified on the nice command now always maps to two units on the RUNPTY parameter, instead of the previous variable mapping of 1 to 4 units on the RUNPTY parameter. Thus, in interactive jobs, the default run priority for the i5/OS PASE nice command becomes RUNPTY(40) instead of RUNPTY(60). If you run the nice command from an interactive job, which is running at RUNPTY(20), you must specify a nice value of 15 to have the background work run at RUNPTY(50) to match the RUNPTY default for batch work.
- | 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 the pthread_setschedparam() function or an attributes object passed to the pthread_create() API. i5/OS restricts thread priority, so applications can only set priority values that make a thread equal to or less favored than the RUNPTY value for the job.

| **New option on _ILECALLX and _PGMCALL**

- | An option was added to the i5/OS PASE runtime functions _ILECALLX and _PGMCALL to suppress
- | signals for i5/OS exceptions. Adding X'20' (ILECALL_EXCP_NOSIGNAL or
- | PGMCALL_EXCP_NOSIGNAL, defined in the as400_types.h header file) to the "flags" argument causes
- | these functions to return -1 and set the error number (errno) variable to 3474 without raising a signal if
- | an exception occurs during the call.

Chapter 4. Licensed programs

This chapter contains information about individual licensed programs that have been modified or changed for V5R4.

Backup, Recovery and Media Services (5722-BR1)

The following topics describe changes to the Backup, Recovery and Media Services (5722-BR1) product that you need to be aware of for V5R4. To learn more about changes to this product in V5R4, see the Backup, Recovery and Media Services Web site at the following URL:

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

Start Maintenance Using BRM (STRMNTBRM) command changed

The Audit system media (AUDSYSMED) parameter will no longer update other system databases, but will create a report.

BRMS database file conversion

In V5R4, you must convert all BRMS logical and physical files to the new file formats after installation. A job to complete this conversion is submitted to the batch job queue following a successful installation. You cannot use the BRMS menu and command interfaces until the conversion job completes. If the system is in restricted state during install, the conversion job remains on the job queue.

BRMS printer files

The header records of some BRMS printed reports that did not previously contain a system name were changed to include the system name.

Restore operations using iSeries Navigator

Restore operations that specify *MEDCLS processing using iSeries Navigator on V5R4M0 have been changed to provide more flexibility when you insert the desired cartridges into stand-alone devices. The location of the stand-alone devices is no longer restricted to the location of the volume, and the operation always checks the mounted volumes on stand-alone devices first. However, as this might become a performance issue in some configurations, you can skip the stand-alone device check; if the data area QUSRBRM/Q1ANOMNTCK exists, the stand-alone device check is skipped.

Client Encryption (5722-CE3)

The Client Encryption (5722-CE3) product is no longer available in V5R4. The client encryption function, Secure Sockets Layer (SSL), is now included in the base of the V5R4 iSeries Access for Windows (product 5722-XE1). You can manually remove previous releases of the Client Encryption product manually if they are no longer needed for old clients.

IBM HTTP Server for i5/OS (5722-DG1)

Last release for Apache Tomcat support

IBM plans for the V5R4 release to be the final release to support Apache Tomcat in the HTTP Server. If you believe you will continue to need support for Apache Tomcat beyond V5R4, the following alternative solutions will be available:

- Use an open source version of Tomcat, which can be obtained from the following Web site:
<http://jakarta.apache.org>
- Use WebSphere® Application Server Express, which is shipped with i5/OS V5R3

User modules

For user-written Apache modules, some considerations are common to most, if not all, modules. Modules must be recompiled with locale type *LOCALEUTF. This will create an environment where locale-dependent C runtime functions will assume that string data is encoded in UTF-8. Any hardcoded constants will be encoded in UTF-8. Additionally, input data from the client will no longer be converted to EBCDIC but will be passed as-is. Output data sent from the module will not be converted either, and must be encoded in UTF-8 or 7-bit ASCII.

The modules must be compiled with different options.

For the Create ILE C Module (CRTCMOD) command or the Create C++ Module (CRTCPPMOD) command, use the following options:

```
MODULE(MYLIB/MOD_TEST)
SRCSTMF('/mydir/mymodule/source/mod_test.c')
DEFINE(AS400 AS400_UTF8)
LOCALETYPE(*LOCALEUTF)
TERASPACE(*YES)
INCDIR('/qibm/proddata/httpa/include')
```

In the options, notice the change in the LOCALETYPE parameter. Using LOCALETYPE(*LOCALEUTF) specifies that program objects are created using the locale support provided by *LOCALE objects. This change enables the locale dependent C run-time functions to work on UTF-8 strings. See the ILE C/C++ Programmer's Guide for details.

Handling data

User-written modules should also be able to read and handle each directive and its parameters from the configuration file as UTF-8 data rather than as EBCDIC data, as was the case in prior releases. You need to consider the data that is passed to the module. Any data that is received from the client through the server will be in the original format of the received data; no data conversion is performed. The headers are in the 7-bit ASCII code page. Any received data that is part of the query string is not converted. Any POST user data is also passed to the module as it is received and without any conversion. You need to consider the data the module passes to the server. The headers that are sent to the server need to be in ASCII format. The server does not convert these headers. Any program data your program sends to the browser is not converted.

Binding considerations

Several Apache Portable Runtime and IBM HTTP Server (powered by Apache) header files have changed. Be sure to check any included files for assumptions about the data received, requirements for certain system functions, and flags with changed settings.

In general, the Apache Portable Runtime and HTTP header files assume that data passed as parameters are in UTF-8 encoding rather than EBCDIC. Note that ASCII code points between X'20' and X'7E' represent the same character in UTF-8.

Logging

Log files that are written to file systems other than QSYS.LIB, such as error, custom (access and reference), script, maintenance, and FRCA logs, are now created with CCSID tagging 1208 (UTF-8). Data that is added to the files is encoded in UTF-8. Most analysis tools can handle files tagged as UTF-8.

Plug-ins

Changes made in the HTTP Server are not compatible with current versions of the Websphere plug-in for the HTTP Server. You will need to get the latest version of the Websphere plug-in with a Program Temporary Fix (PTF).

Error message change

The QzhhCgiParse API will no longer send HTTP error messages to the job log. Instead, errors will be sent to the error log or script log, if one exists. Operating system messages will continue to be sent to the job log.

IBM WebSphere Application Server - Express Version 5.0 for iSeries (5722-IWE)

In V5R3, IBM Web Enablement for iSeries was shipped with i5/OS, and included IBM WebSphere Application Server Express V5.0 for iSeries and IBM WebSphere Application Server Express V5.1 for iSeries.

For V5R4, IBM Web Enablement for i5/OS will be shipped with i5/OS, and it will include IBM WebSphere Application Server Express V5.1 for iSeries and IBM WebSphere Application Server Express V6.0 for OS/400®. IBM WebSphere Application Server Express V5.0 for iSeries will not be included or supported with IBM Web Enablement for i5/OS for V5R4.

IBM WebSphere Application Server V5.0 for iSeries products (IBM WebSphere Application Server Express V5.0 for iSeries, IBM WebSphere Application Server V5.0 for iSeries, and IBM WebSphere Application Server Network Deployment V5.0 for iSeries) will be supported only when used with J D Edwards EnterpriseOne ERP Software Suite, IBM WebSphere Portal, IBM Workplace™ Collaboration Services, and IBM Workplace Services Express. All other customers currently using WebSphere Application Server Version 5.0 products and upgrading to V5R4 should migrate to the WebSphere Application Server V5.1 or V6.0 products. You can do this migration before or after your upgrade to V5R4. See the WebSphere Application Server for OS/400 V6 documentation at the following Web site for more information:

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

IBM Developer Kit for Java (5722-JV1)

V5R4 introduces a new 32-bit JVM called IBM Technology for Java that you can use in addition to the existing iSeries JVM. An IBM Redbook will be provided that gives more details about how to make use of this JVM. To begin using IBM Technology for Java, complete the following steps:

1. Install the IBM Developer Kit for Java (5722-JV1) product (i5/OS Option 8)
2. Run the following command:

```
ADDENVVAR ENVVAR(JAVA_HOME) VALUE('/QOpenSys/QIBM/ProdData/JavaVM/jdk50/32bit')
```

The following topics describe some requirements for the IBM Technology for Java that you need to consider.

Java Native Methods

If you want to use IBM Technology for Java and have programs that use native methods, you must compile these programs with teraspace storage enabled. Since this is not the default setting, it is likely that you need to recompile. This is necessary because the Java object is in i5/OS PASE storage, which is mapped on top of teraspace storage, and a teraspace storage pointer is returned. Also, the JNI function `NewDirectByteBuffer` has a pointer parameter which is the address of the storage for the direct byte buffer. This pointer must point to the area in teraspace storage that is part of the i5/OS PASE address space. The easiest way to obtain this storage is to call `Qp2malloc`. If you have not compiled your program

with teraspace storage enabled and attempt to run the native method with IBM Technology for Java, you will receive the escape message MCH4443 (Invalid storage model for target program LOADLIB).

Adopt Authority

IBM Technology for Java does not support adopted authority for Java programs.

IBM i5/OS Integration for Linux on xSeries (5722-LSV)

In V5R2 and V5R3, support for Linux on integrated xSeries hardware was made available with a program temporary fix (PTF) to the IBM iSeries Integration for Windows Server (5722-WSV) product.

In V5R4, support for Linux on integrated xSeries hardware requires the following i5/OS option and licensed program product:

- 5722-SS1 Option 29: Integrated Server Support
- IBM i5/OS Integration for Linux on xSeries (5722-LSV)

For migration information, see the following Web site:

| <http://www.ibm.com/systems/i/advantages/integratedserver/linux/v5r4migration.html>

Network Authentication Enablement (5722-NAE)

In V5R3, the Kerberos server was included as part of the 5722-AC3 product.

In V5R4, the 5722-AC3 product is no longer available. The Kerberos server is now shipped in the Network Authentication Enablement (5722-NAE) product.

If V5R4 is installed over V5R3, and the 5722-AC3 product is currently installed, then the 5722-NAE product is automatically installed to ensure that the Kerberos server that was part of the 5722-AC3 product is installed. If V5R4 is installed over V5R2, and the 5722-AC3 product is currently installed, then the 5722-NAE product is not automatically installed, since the Kerberos server was not part of 5722-AC3 in V5R2.

Performance Tools (5722-PT1)

Disk Space Used calculation on **Disk Utilization Summary** section of **Print System Report (PRTSYSRPT)** command changed

In V5R4, the **Disk Space Used** calculation on the **Disk Utilization Summary** section of Print System Report (PRTSYSRPT) command is displayed in GB (gigabytes); in previous releases, it was shown in MB (megabytes).

Additionally, the following fields in the performance reports show MB (megabytes) in V5R4 instead of KB (kilobytes):

Table 6.

Report	Section	Field
System	Storage Pool Utilization	Size
Component	Storage Pool Activity	Pool size
Pool	Pool Activity	Size

Also, the size of some fields in the reports have increased to show greater values. The change affects the following reports:

Table 7.

Report	Section	Field
System	Storage Pool Utilization	DB Pages, Non-DB Pages
Component	Component Interval Activity	User Pool Faults/Sec
Pool	Pool Activity	CPU Util Synchronous DBR Synchronous DBW Synchronous NDBR Synchronous NDBW Asynchronous DBR Asynchronous DBW Asynchronous NDBR Asynchronous NDBW Most active-to-wait transitions

Changes to Job Trace commands

The following changes to the Performance Tools Job Trace commands were made in V5R4:

- The System model code (MODEL) parameter was removed from the Print Job Trace (PRTJOBTRC) command.
- The default size of the trace buffer for the Maximum Storage (MAXSTG) parameter in the Start Job Trace (STRJOBTRC) command was changed to 10000 kilobytes in order to be consistent with the default size of the trace buffer in the Start Trace (STRTRC) command.
- The trace database file created by the End Job Trace (ENDJOBTRC) command has changed. In V5R4, ENDJOBTRC uses the Print Trace (PRTTRC) command specifying OUTPUT(*OUTFILE) to create the new trace database file, instead of using the Trace Job (TRCJOB) command. Applications that process the trace database file from ENDJOBTRC must be changed to use the new output file format.

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

In V5R4, the ILE RPG precompiler will enforce the rule that the SET OPTION statement must precede any other SQL statement.

This rule is enforced by the other precompilers available with the IBM DB2 Query Manager and SQL Development Kit for iSeries product, and has been previously documented in the DB2 UDB for iSeries SQL Reference topic: "If embedded in an application program, [the SET OPTION statement] is not executable and must precede any other SQL statements."

Prior to V5R4, you could have used the SET OPTION statement anywhere in your application program.

IBM TCP/IP Connectivity Utilities for i5/OS (5722-TC1)

FTP reply code changed

For V5R4, the FTP server was changed to return a reply code of 226 instead of 250 when a successful data transfer has been completed. This made the FTP server compliant with RFC 959, which states that if the data transfer connection is closed after a successful data transfer, the reply code should be 226. If you have FTP scripts that are expecting the 250 reply code, you should change the scripts to expect a 226 reply code for a successful data transfer.

IBM iSeries Integration for Windows Server (5722-WSV)

The product IBM iSeries Integration for Windows Server (5722-WSV) has been repackaged as i5/OS Integrated Server Support (i5/OS Option 29).

When you upgrade from a prior release to i5/OS V5R4, product 5722-WSV is automatically removed and product 5722-SS1 option 29 is installed in its place.

iSeries Access for Windows (5722-XE1)

Data Transfer support for Microsoft Excel 95 and Microsoft Excel 98 will change in the future

iSeries Access for Windows Data Transfer provides a Microsoft Excel Add-in that lets you transfer data to and from an open Excel worksheet. In a future release of iSeries Access for Windows, the Data Transfer Excel Add-in will be rewritten to use a more robust Excel interface. This interface is not available in Microsoft Excel 95 and Microsoft Excel 98. Support for those versions of Excel will therefore be dropped from the Add-in. This change will not affect the Data Transfer from iSeries and Data Transfer to iSeries graphical user interfaces. You will still be able to transfer data to and from Excel 95 and Excel 98 documents using that interface.

Automatic upgrade and automatic activation of the AFP™ Printer Driver no longer required

A change has been made to eliminate automatic upgrade and automatic activation of the Advanced Function Printing™ (AFP) Printer Driver, even if the Windows Driver Signing Options are set to Ignore. An Ignore setting will now work the same way with AFP Printer Driver upgrade or activation as the Warn and Block settings.

This change pertains to new installs, reinstalls, and upgrade installs of iSeries Access for Windows.

The AFP Printer Driver files will be copied to the \CWBAFP subdirectory under the iSeries Access for Windows installation directory. You will need to manually upgrade existing printers by opening the properties of the printer; select the Advanced tab and then click **New Driver**. On the Printer Driver Selection page click **Have Disk** and navigate to the \CWBAFP directory.

Windows Administration renamed to Integrated Server Administration

In V5R4, the iSeries Navigator **Windows Administration** folder was renamed to **Integrated Server Administration**, and it was moved to a new location in the iSeries Navigator folder hierarchy. The Integrated Server Administration folder is now located under the i5/OS system folder. Previously, the Windows Administration folder was located under the Network folder.

Changes to the IPv6 configuration files and methods in V5R4 are not compatible with previous releases of iSeries Navigator

The changes to the IPv6 configuration files and methods in V5R4 are not compatible with the graphical user interface (GUI) provided by iSeries Navigator in previous releases. If you attempt to configure IPv6 on a V5R4 i5 server using a prior version of iSeries Navigator, the configuration fails..

The IPv6 configuration files in V5R2 and V5R3 are /qibm/proddata/os400/TCPIP/IP6/QTOCTCPIP6CONFIG.IP6ML and /qibm/userdata/os400/TCPIP/IP6/QTOCTCPIP6CONFIG.IP6ML. These files are no longer used in V5R4.

You must manually recreate any IPv6 configuration when you migrate from a previous release to V5R4.

PC5250 considerations

On previous releases of iSeries Access for Windows, if a user configured an Arabic or Hebrew PC5250 session, the default host code page for the session was 037. Starting with V5R3, the default code page is 420 for Arabic and 424 for Hebrew. This change has no effect on previously-configured sessions.

Cannot install on PC running Microsoft Windows NT®

In V5R4, you can no longer install iSeries Access for Windows on a personal computer (PC) running the Windows NT operating system. You can install iSeries Access for Windows on PCs that are running Windows 2000, Windows XP, or Windows Server 2003.

If you want to upgrade your PC to one of the supported Windows operating systems, you must perform the following steps:

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

Data Transfer considerations

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

New AFP file viewer configuration changes

The new version of the Advanced Function Printing (AFP) viewer installed with V5R4 iSeries Access for Windows has been enhanced to include an interface that allows you to change the code page and font being used to display SCS and ASCII spooled files on i5/OS. If you migrate from the old AFP viewer to the new one, you will find that the settings you had configured to view SCS spooled files that contained national language specific characters no longer work, and the spooled file does not display properly. You can access the new viewer configuration options by opening a spooled file with the AFP viewer and, once the viewer window is open, select **Options > Code Pages and Fonts**. Select the code page that was used to create the spooled file and the font that contains the characters that will be displayed.

Distributing iSeries Navigator plug-ins and add-ins

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

Also, future releases will not support installation of add-ins.

Remote Command/Remote Program Call messages

The remote command (RMTCMD.EXE) and remote program interfaces (CWBRC.H) have been changed to allow a more complete list of job messages to be returned. For RMTCMD.EXE, this might cause a different exit ERRORLEVEL to be returned depending on the severity of any newly included job messages. To enable the old behavior for a given PC process, set up an environment variable called CWB_RMTCMD_V5R2_MSG and set it to the value *ALL or a process name. For example,

CWB_RMTCMD_V5R2_MSG=rmtcmd.exe enables the old behavior for rmtcmd.exe. *ALL enables the old behavior for all processes within this environment variable scope.

Unicode to mixed-EBCDIC conversions

The conversion method for Unicode to mixed-EBCDIC conversions has been changed to allow more characters to be successfully converted. However, this also means that some existing character conversions will be different.

iSeries Access for Windows Service API changes

The Service API set (prefixed by cwbsv), which are used to log data to and retrieve data from iSeries Access for Windows log and trace files, have been changed slightly to be consistent with documentation and expectations. The changes, and list of which APIs each change applies to, are as follows:

- Some APIs that log information take a length as a parameter. In some cases, this length was being ignored, and the data to log were being assumed to be NULL-terminated. This has been corrected, so that no assumption about NULL-termination is made, and the length passed in is honored. The following APIs are affected:
 - cwbsv_LogMessageTextW
 - cwbsv_LogTraceDataW
 - cwbsv_LogAPIEntryW
 - cwbsv_LogAPIExitW
 - cwbsv_LogSPIEntryW
 - cwbsv_LogSPIExitW
- Some APIs that retrieve information take both an input length parameter and an output length parameter. The output length parameter was documented as being a byte count, but the input length parameter was not documented explicitly as being either a byte count or a character count. The two length parameters were not always consistently treated as either both byte counts, or both character counts. This could have resulted in the truncation of the returned data. The documentation has been updated (in cwbsv.h) to explicitly state that both counts are byte counts, and the implementation has been modified to be consistent with the documentation. The following APIs are affected:
 - cwbsv_GetServiceFileNameW
 - cwbsv_GetProduct
 - cwbsv_GetProductW
 - cwbsv_GetComponent
 - cwbsv_GetComponentW
 - cwbsv_GetDateStamp
 - cwbsv_GetDateStampW
 - cwbsv_GetTimeStamp
 - cwbsv_GetTimeStampW
 - cwbsv_GetMessageTextW
 - cwbsv_GetErrTextW
 - cwbsv_GetErrTextIndexedW
 - cwbsv_GetErrFileNameW
 - cwbsv_GetErrFileNameIndexedW
 - cwbsv_GetErrLibNameW
 - cwbsv_GetErrLibNameIndexedW
- Some APIs that retrieve data were storing the data correctly, but were returning a data length that was potentially greater than that of the buffer passed in, even when the buffer was large enough to hold all

the data. This has been corrected; the data length returned is now always the size of the buffer needed to hold all of the requested data, in bytes. The following APIs are affected:

- cwbSV_GetProduct
- cwbSV_GetProductW
- cwbSV_GetComponent
- cwbSV_GetComponentW
- cwbSV_GetDateStamp
- cwbSV_GetDateStampW
- cwbSV_GetTimeStamp
- cwbSV_GetTimeStampW

iSeries Access for Web (5722-XH2)

Before you install V5R4 iSeries Access for Web

If you currently have a configured, running version of iSeries Access for Web on your iSeries server, you must run the iSeries Access for Web configuration commands again when you install the new version before you can use the product.

When running CFGACCWEB2 for the WebSphere Application Server environments, be aware that the value input to the WASINST and WASPRF parameters is case sensitive and should be entered exactly as shown in the instances.properties file.

For detailed instructions on installing and upgrading the product and running the configuration commands, see the V5R4 iSeries Information Center: click **Connecting to iSeries > iSeries Access > iSeries Access for Web**.

WebSphere Application Server V4.0

V5R4 iSeries Access for Web does not support the WebSphere Application Server V4.0 (Advanced and Advanced Single Server Editions) web application server. V5R4 iSeries Access for Web can migrate user-generated data from a WebSphere Application Server V4.0 configuration to a supported web application server environment when new web application servers are configured. V5R4 iSeries Access for Web will also allow you to remove the iSeries Access for Web configuration from the WebSphere Application Server V4.0 environment using the RMVACCWEB2 command.

Be sure you have loaded and applied the latest V5R4 iSeries Access for Web PTF(s) before attempting to migrate from or remove WebSphere Application Server V4.0 configurations.

For additional information, see the V5R4 iSeries Information Center: click **Connecting to iSeries > iSeries Access > iSeries Access for Web**.

Style sheets

The V5R4 look of iSeries Access for Web has changed. The look of the iSeries Access for Web page content is now controlled by external style sheets. If you use customized templates for display by iSeries Access for Web, you will need to modify the templates to include the new style information. You can provide your own style sheets if you want a look that is different from the default iSeries Access for Web look.

If you prefer the V5R3 look, you can use the customization function within iSeries Access for Web to make this change.

For additional information, see the V5R4 iSeries Information Center: click **Connecting to iSeries > iSeries Access > iSeries Access for Web**.

Configuration commands

In previous releases of iSeries Access for Web, only CL commands were provided to manage iSeries Access for Web configurations. QShell script commands were later added to those previous releases using PTFs to support newer web application servers.

In V5R4 iSeries Access for Web, the CL commands and QShell script commands are provided as part of the base release. This allows you to use either CL commands or QShell script commands to manage your iSeries Access for Web configuration.

For information on how to use the commands, see the V5R4 iSeries Information Center: click **Connecting to iSeries > iSeries Access > iSeries Access for Web**.

IBM iSeries Access for Windows (5722-XE1) and Microsoft Windows Vista

Starting with service level SI25949, iSeries Access for Windows can run on a PC with certain editions of Microsoft Windows Vista installed. However, the design of Windows Vista causes the following changes in the functionality and support of some features of iSeries Access for Windows.

- The display of some iSeries Access for Windows help will fail unless the WinHelp engine for Windows Vista is downloaded and installed on the PC. At this time, Microsoft provides this WinHelp engine as a free download. See <http://support.microsoft.com/kb/917607> for more details.
- The Check Service Level function works, but cannot be run automatically when each user logs on to Windows, as it can in prior versions of Windows. Though iSeries Access for Windows Properties still allows you to configure when this automated check will run, the configured values have no effect since it cannot run automatically at logon time at all.
- Operations Console is not supported when run on Windows Vista.
- The SCS print driver is currently not supported on Windows Vista. However, it will likely be supported in a future V5R4 service pack.
- The AFP Viewer does not work reliably unless the viewer daemon program, cwbviewd.exe, is started manually before the use of the viewer. This program is located in the folder called AFPViewr, which is within the folder where iSeries Access for Windows is installed. By default, then, cwbviewd.exe is located in the following folder: C:\Program Files\IBM\Client Access\AFPViewr. If you use the AFP Viewer often, you might want to add a shortcut to the viewer daemon program to your Startup folder so that it is started automatically when you log on to Windows.

Details about these and other Windows Vista-related issues when running iSeries Access for Windows are documented in Informational APARs II14239 and II14247.

Note:

iSeries Access for Windows is not supported on all editions of Windows Vista. See <http://www-03.ibm.com/systems/i/software/access/windows/supportedos.html> for support information.

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.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Programming Interface Information

This Memorandum 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:

Advanced Function Printing
AFP
AS/400
DB2
DB2 Universal Database
Electronic Service Agent
eServer
GDDM
i5/OS
IBM
iSeries
NetServer
OS/400
POWER5
POWER6
Redbooks

SOM
System i
System Object Model
System p5
TotalStorage
WebSphere
Workplace
xSeries

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

IBM Systems - iSeries
Memorandum to Users
Version 5 Release 4 (May 2008 update)

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

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

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

Comments:

Thank you for your support.

Submit your comments using one of these channels:

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

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

Name

Address

Company or Organization

Phone No.

E-mail address



Cut or Fold
Along Line

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

Cut or Fold
Along Line



Printed in USA