

IBM i 7.1 Technical Overview with Technology Refresh Updates




Boosted efficiency with enriched virtualization including Live Partition Mobility



Enriched database functionality and new application development tools



Reenergized management and performance tooling web interface



Hernando Bedoya
Octavio Bustos
Clive Griffiths
Tom Grigoleit
Axel Lachmann
Yohichi Nakamura
Sandra Nelson
Xiao Hu Qu
Tom Vernailen

Redbooks



International Technical Support Organization

IBM i 7.1 Technical Overview with Technology Refresh Updates

March 2013

Note: Before using this information and the product it supports, read the information in “Notices” on page xvii.

Third Edition (March 2013)

This edition applies to Version 7, Release 1, Modification 0 of IBM i (5770-SS1) and related licensed programs.

© Copyright International Business Machines Corporation 2010, 2013. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	xvii
Trademarks	xviii
Preface	xix
The team who wrote this book	xix
Now you can become a published author, too!	xxii
Comments welcome	xxii
Stay connected to IBM Redbooks	xxiii
Summary of changes	xxv
March 2013, Third Edition	xxv
Chapter 1. IBM i 7.1 introduction and summary	1
1.1 Communicating IBM i enhancements	2
1.2 Clarification of operating system terminology	4
1.3 What is new in IBM i 7.1	4
1.3.1 IBM i 7.1 on POWER based servers	4
1.3.2 Technology Refreshes	4
1.3.3 Integrated DB2 database for IBM i	5
1.3.4 PowerHA SystemMirror for i	5
1.3.5 IBM i virtualization	6
1.3.6 IBM i storage management	7
1.3.7 Application development on IBM i enhancements	8
1.3.8 IBM Rational support for IBM i 7.1	9
1.3.9 Systems management enhancements	10
1.3.10 Printing	11
1.3.11 IBM i iSCSI integration with BladeCenter and System x	11
1.3.12 Security enhancements in IBM i 7.1	11
1.3.13 Web application serving	11
1.3.14 IBM i Access Family 7.1 enhancements	12
1.3.15 Networking enhancements	12
1.3.16 Upgrade for IBM i enhancements	12
1.3.17 IBM i network installation	13
1.3.18 IBM i Processor Entitlement Transfer	13
1.3.19 Backup, Recovery, and Media Services and basic backup and recovery	13
Chapter 2. Technology Refresh	15
2.1 Overview	16
2.2 What a Technology Refresh is	16
2.3 What a Resave is	17
2.4 What is in a Technology Refresh	18
2.4.1 When to install a Technology Refresh	18
2.4.2 How it all fits together	19
2.4.3 Determining the Technology Refresh PTF Group level installed on a system	20
2.4.4 Determining the Resave level installed on a system	21
2.4.5 Determining the Technology Refresh PTF level installed on a system	22
2.4.6 Planning for an IBM i Technology Refresh update	22
2.4.7 Installing an IBM i Technology Refresh	23
2.4.8 Ordering an IBM i Technology Refresh PTF Group	23

2.4.9	Installing an IBM i Technology Refresh PTF Group	23
2.4.10	Ordering and installing an IBM i Resave.	23
2.4.11	Ordering an IBM i Technology Refresh PTF.	24
2.4.12	Installing an IBM i Technology Refresh PTF.	24
2.4.13	How a Technology Refresh PTF or PTF Group affects other PTFs	24
2.4.14	Preventing or reducing the impact of a double IPL.	24
2.4.15	Why IBM moved to the Technology Refresh process.	25
2.4.16	New releases of IBM i.	25
2.4.17	Is an IBM i Technology Refresh the same as an AIX Technology Level	26
2.4.18	Related publications	26
Chapter 3.	Security enhancements	27
3.1	User Profile enhancements.	28
3.1.1	New USREXPDATE and USREXPITV attributes	28
3.1.2	CL commands to change and display the User Profile expiration schedule	28
3.2	Object Audit enhancements	29
3.2.1	CL retrieve exit programs running after command completion	29
3.2.2	Extended auditing details for running a CL command	30
3.2.3	Database object auditing control with the QIBM_DB_OPEN exit program.	31
3.3	Data encryption enhancements	32
3.3.1	ASP encryption enhancements.	32
3.3.2	Field procedures to encode and decode data.	32
3.4	Security enhancements for DB2 for i	32
3.4.1	Query Manager user profile auditing.	32
3.4.2	Database Monitor and Plan Cache variable values masking	34
3.4.3	Security function usage for a database.	36
3.4.4	Simplified DDM and Distributed Database (DRDA) authentication.	38
3.5	Real-time database activity monitoring	41
3.5.1	Enhanced audit capability for STRDBMON	41
3.5.2	SQL statement auditing with less resource consumption	44
3.5.3	InfoSphere Guardium database activity monitoring support for DB2 for i.	48
3.6	Security enhancements for printing.	50
Chapter 4.	Backup and recovery	51
4.1	New and enhanced system save and restore functions	52
4.1.1	Fast restore using tape position	52
4.1.2	New ALWOBJDIF (*COMPATIBLE) restore option	53
4.1.3	Enhanced save and restore limits.	53
4.1.4	Save While Active support for integrated VMware ESX servers.	54
4.1.5	Miscellaneous enhancements.	54
4.2	New and improved BRMS for IBM i functions.	55
4.2.1	Support for Domino 8.5 Domino Attachment and Object Service (DAOS)	55
4.2.2	Link list enhancements	56
4.2.3	Support of more than 999,999 objects	56
4.2.4	Start recovery using BRM (STRRCYBRM) command enhancements	57
4.2.5	Recovery report enhancements	60
4.2.6	Output file support for BRMS restores	61
4.2.7	Ability to select Saved ASP on RSTLIBBRM and RSTOBJBRM commands	66
4.2.8	Distributed backup support.	67
4.2.9	Maintenance enhancements.	71
4.2.10	Planning Media Movement report.	72
4.2.11	Improved clearing of volumes in *ERR status.	73
4.2.12	Protection of volumes that are marked for duplication	74

4.2.13	Improved recovery times with new media position function	76
4.2.14	BRMS support for the special value *COMPATIBLE for ALWOBJDIF	76
4.2.15	Improved control with running non-BRMS saves	77
4.2.16	RDX and BRMS	77
4.2.17	Allowing volumes that are marked for duplication to be moved	77
4.2.18	Support for expiry dates greater than the year 2038.	79
4.2.19	Setting the maximum size for BRMS flight recorders	79
4.2.20	Using compression for BRMS Save Files.	79
4.2.21	Improving IFS save performance - the ASYNCBRING parameter	81
4.2.22	Using STRMNTBRM to initialize media when it is expired	82
4.2.23	PRTRPTBRM control group status	83
4.3	BRMS enhancements to GUI and web interfaces.	84
4.3.1	Added support for the IBM Systems Director web browser environment	85
4.3.2	IBM Systems Director navigation to BRMS functions	85
4.3.3	IBM Systems Director Navigator for i navigation to BRMS functions	90
4.3.4	Enhancements to the BRMS initial window.	92
4.3.5	BRMS advanced functions window.	93
4.3.6	Scheduling support for BRMS.	93
4.3.7	Added option to the BRMS Log to filter messages by control groups.	102
4.3.8	Ability to mark and unmark volumes for duplication	106
4.3.9	Multiple email address support	112
4.3.10	High availability support for independent ASPs in a BRMS network	115
4.3.11	Enhanced maintenance features in backup policy	118
4.3.12	New BRMS Enterprise function.	121
4.4	More resources	122
Chapter 5. High availability		123
5.1	PowerHA SystemMirror for i	124
5.1.1	New PowerHA packaging	124
5.1.2	PowerHA versions	125
5.1.3	PowerHA SystemMirror for i enhancements.	125
5.1.4	PowerHA SystemMirror for i: Graphical interfaces	130
5.1.5	N_Port ID virtualization support	133
5.1.6	Asynchronous geographic mirroring	133
5.1.7	LUN level switching.	136
5.1.8	Space-Efficient FlashCopy	137
5.1.9	Better detection of cluster node outages	138
5.1.10	Improved geographic mirroring full synchronization performance	140
5.1.11	Improvement of geographic mirroring in an IBM i hosted IBM i client partition environment	141
5.1.12	Support virtualization capabilities on an active cluster node.	141
5.1.13	Cluster administrative domain enhancements	142
5.1.14	IPv6 support	143
5.1.15	New CL commands for programming cluster automation.	143
5.2	Journaling and commitment control enhancements	144
5.2.1	Journal management	144
5.2.2	Remote journaling.	146
5.2.3	DISPLAY_JOURNAL (easier searches of a journal).	146
5.2.4	Commitment control and independent ASPs	148
5.2.5	System Managed Access Path Protection	149
5.2.6	Journaling and disk arm usage	150
5.2.7	Range of journal receivers parameter and *CURAVLCHN.	150
5.2.8	Journal management functions: IBM Navigator for i	151

5.3 Additional information	151
Chapter 6. IBM DB2 for i	153
6.1 Introduction: Getting around with data	154
6.2 SQL data description and data manipulation language	154
6.2.1 XML support	155
6.2.2 The MERGE statement	160
6.2.3 Creating and using global variables	162
6.2.4 Support for arrays in procedures	163
6.2.5 Result set support in embedded SQL	164
6.2.6 FIELDPROC support for encoding and encryption	166
6.2.7 WebSphere MQ integration	169
6.2.8 Miscellaneous	172
6.2.9 OVRDBF SEQONLY(YES, buffer length)	181
6.3 Performance and query optimization	182
6.3.1 Methods and tools for performance optimization	182
6.3.2 Query optimization	183
6.3.3 Global Statistics Cache	184
6.3.4 Adaptive query processing	185
6.3.5 Sparse indexes	186
6.3.6 Encoded vector index	187
6.3.7 Preserving EVI indexes on ALTER enhancement	189
6.3.8 Keeping tables or indexes in memory	189
6.3.9 SQE optimization for indexes on SSD	189
6.3.10 SQE support of simple logical files	189
6.3.11 QSYS2.INDEX_ADVICE procedure	190
6.3.12 SKIP LOCKED DATA and NC or UR	191
6.3.13 SQL routine performance integer arithmetic (requires re-create)	191
6.3.14 Automatic cancellation of QSQSRVR jobs when an application ends	191
6.3.15 QAQQINI properties	192
6.3.16 ALTER TABLE performance	193
6.3.17 Avoiding short name collisions for CREATE PROCEDURE, FUNCTION, or TRIGGER	193
6.3.18 CREATE PROCEDURE (SQL) PROGRAM TYPE SUB	194
6.3.19 Referential integrity and trigger performance	195
6.3.20 QSQBIGPSA data area	195
6.3.21 Validating constraints without checking	196
6.3.22 Limiting the amount of processing on an RGZPFM cancel	196
6.3.23 CPYFRMIMPF performance	196
6.3.24 QJOSJRNE API option to force journal entries without sending an entry	196
6.3.25 QDBRTVSN API performance	197
6.3.26 Adding total DB opens job level instrumentation to Collection Services	197
6.3.27 SYSTOOLS.REMOVE_INDEXES procedure	197
6.3.28 Improved SQE statistics maintenance during INSERT, UPDATE, and DELETE statements	198
6.3.29 QSYS2.Reset_Table_Index_Statistics procedure	198
6.3.30 Performance enhancements for large number of row locks	199
6.3.31 Improved DSPJOB / CHKRCDLCK command behavior for large numbers of row locks	199
6.3.32 Chart-based graphical interface SQL Performance monitors	200
6.3.33 Collection services include physical I/O breakdown by SQL activity	200
6.3.34 Performance Data Investigator: New Database perspective	202
6.3.35 Performance Data Investigator: SQL Overview	203

6.3.36	Performance Data Investigator: SQL Attribute Mix	204
6.3.37	Index Advisor: Show Statements - Improved query identification	204
6.3.38	Performance improvements for temporary tables	205
6.4	New functionality for DB2 developers	205
6.4.1	QSYS2.SYSCOLUMNS2	205
6.4.2	QSYS2.SYSPARTITIONDISK and QSYS2.SYSPARTITIONINDEXDISK	206
6.4.3	QSYS2.OBJECT_STATISTICS table function ease of use	207
6.4.4	EARLIEST_POSSIBLE_RELEASE	208
6.4.5	SIGNAL support for native triggers	208
6.4.6	Hierarchical queries through the CONNECT BY clause	209
6.4.7	Additional parameter marker support (LAND, LOR, XOR, and TRANSLATE)	210
6.4.8	Supporting PROGRAM NAME on CREATE TRIGGER	210
6.4.9	TINYINT in CLI	211
6.4.10	CHGPFM and CHGLFM UNIT support	211
6.4.11	SYSTOOLS procedures	211
6.4.12	Adding a client and server IP address and port number to the QSYS2.TCPIP_INFO view	212
6.4.13	QSYS2.GROUP_PTF_INFO view	212
6.4.14	QSYS2.DUMP_SQL_CURSORS procedure	214
6.4.15	QIBM_SQL_NO_RLA_CANCEL environment variable	215
6.4.16	QSYS2.CANCEL_SQL and QSYS2.FIND_AND_CANCEL_QSQSRVR_SQL procedures	215
6.4.17	QSYS2.FIND_QSQSRVR_JOBS() procedure	217
6.4.18	SQL server mixed mode for batch processing	219
6.4.19	QDBRPLAY() API: Disable or Enable Constraints option	219
6.4.20	SQL0901 log education	220
6.4.21	Retrieving a short name for long schema name (QDBRTVSN)	221
6.4.22	SQL0440 warning on SQL routine build	221
6.4.23	XMLTABLE	222
6.4.24	Run SQL (RUNSQL): A new command	224
6.4.25	Native JDBC driver support added for J2SE 7	224
6.4.26	Adding the ORDERBY parameter to the CPYTOIMPF command	225
6.4.27	System naming convention expanded for permit (/) and (.) qualifiers	226
6.4.28	Modification of global variables within triggers and functions	226
6.4.29	New QAQQINI option: SQL_GVAR_BUILD_RULE	227
6.4.30	CPYTOIMPF and CPYFRMIMPF commands include column headings	227
6.4.31	Improved performance of the privileges catalogs and the ODBC/JDBC metadata APIs	227
6.4.32	Show statements: Index creates and statistics advised	229
6.4.33	The RUNSQLSTM command and the OPTION parameter	229
6.4.34	Improved NULLID package management	230
6.4.35	Java stored procedures and functions: System naming option	230
6.4.36	Named arguments and defaults for parameters: Procedures	231
6.4.37	Improved catalog management for procedures and functions	232
6.4.38	SQE enhancement for Encoded Vector Indexes defined with INCLUDE	233
6.4.39	Navigator: System Name column added to show related and all objects	233
6.4.40	Improved performance for joins over partitioned tables	234
6.4.41	Navigator: Table list totals (members, data size, and valid and deleted rows)	234
6.5	DB2 database management and recovery enhancements	235
6.5.1	Preserving the SQL Plan Cache size across IPLs	235
6.5.2	Prechecking the physical file size during restore	236
6.5.3	Preventing index rebuild on cancel during catch up	236
6.5.4	Adding the QSYS2/SYSDISKSTAT catalog	236

6.5.5	STRDBMON: FETCH statement shows failures and warnings.	237
6.5.6	STRDBMON: QOI2 result rows for more statements	237
6.5.7	Adding result set information to QUSRJOBI() and Systems i Navigator.	237
6.5.8	STRDBMON pre-filtering of QUERY/400 command usage	238
6.5.9	UNIT SSD supported on DECLARE GLOBAL TEMPORARY TABLE	239
6.5.10	Adding Maintained Temporary Indexes that are used for a source of statistics within the index advisor and index advisor condenser	239
6.5.11	Adding the QSYS2.REMOVE_PERFORMANCE_MONITOR procedure	240
6.5.12	STRDBMON: QOI1 fast delete reason code.	240
6.5.13	Automatically increasing the SQE Plan Cache size	241
6.5.14	Navigator for i: A new look and no client to manage.	241
6.6	DB2 for Web Query for i (5733-QU2, 5733-QU3, and 5733-QU4)	242
6.6.1	DB2 Web Query for i (5733-QU2).	242
6.6.2	DB2 Web Query Report Broker (5733-QU3).	244
6.6.3	DB2 Web Query Software Developer Kit (5733-QU4)	244
6.6.4	DB2 Web Query for i Standard Edition	244
6.7	OmniFind Text Search Server for DB2 for i (5733-OMF)	244
6.7.1	OmniFind for IBM i: Searching Multiple Member source physical files	245
Chapter 7. Performance tools		247
7.1	Introduction to performance tooling.	248
7.1.1	Gathering data	248
7.1.2	Analyzing data	249
7.2	Overview of data collection in IBM i 7.1	249
7.2.1	Collection profiles for Collection Services.	250
7.2.2	Information repository: Performance tools tables	251
7.2.3	Additional enhancements	261
7.3	IBM Navigator for i Performance interface	262
7.4	IBM iDoctor for IBM i	263
7.4.1	Components	263
7.4.2	Installation.	263
7.4.3	My Connections View	264
7.4.4	Main window	265
7.4.5	Collection options	269
7.4.6	Data Viewer	270
7.4.7	Collection Services Investigator	273
7.4.8	Job Watcher	292
7.4.9	PEX Analyzer	301
7.4.10	Plan Cache Analyzer	304
7.4.11	VIOS Investigator	305
7.4.12	iDoctor FTP GUI	307
7.4.13	More information	309
Chapter 8. Virtualization		311
8.1	PowerVM enhancements	312
8.1.1	Active memory sharing	312
8.1.2	Enhanced support for IBM System Storage	316
8.1.3	PowerVM Virtualization and I/O enhanced with NPIV.	318
8.1.4	Expanded HBA and switch support for NPIV on Power Blades	321
8.1.5	PowerVM N_Port ID Virtualization attachment of DS5000	321
8.1.6	Enhanced mirroring algorithms	321
8.1.7	PCIe2 Riser Card (Gen2) (#5685) direct support	321
8.1.8	Redundant VIOS partitions support	322

8.1.9 Shared storage pools	323
8.2 Additional OS level combinations of server and client logical partitions	327
8.3 Hardware Management Console virtual device information	329
8.4 IBM i hosting IBM i	332
8.4.1 Disk virtualization	333
8.4.2 Optical virtualization	337
8.4.3 Tape virtualization	340
8.4.4 Networking virtualization	340
8.4.5 Support for embedded media changers	340
8.4.6 Performance consideration	341
8.4.7 Dual hosting	341
8.5 Virtual Partition Manager enhancements	342
8.5.1 Ethernet Layer-2 bridging	342
8.5.2 Preparing for Ethernet Layer-2 bridging	343
8.5.3 Preferred practices	343
8.5.4 Configuring Ethernet Layer-2 bridging	343
8.5.5 Common errors	343
8.5.6 Managing Ethernet Layer-2 bridging	344
8.6 Partition suspend and resume	344
8.6.1 Requirements for Suspend / Resume	344
8.6.2 Considerations for Suspend / Resume	345
8.7 HEA Daughter cards	346
8.8 10 Gb FCoE PCIe Dual Port Adapter	346
8.9 Live Partition Mobility	346
8.9.1 Requirements for Live Partition Mobility	348
8.9.2 Considerations for Live Partition Mobility	348
8.9.3 Preparing Live Partition Mobility	349
8.9.4 Validating Live Partition Mobility	351
8.9.5 Performing Live Partition Mobility	354
8.9.6 Considerations about Live Partition Mobility	366
Chapter 9. Storage and solid-state drives	367
9.1 General and direct-attached storage management enhancements	368
9.1.1 Concurrent removal of disk units	368
9.1.2 Hot spare for mirroring	371
9.1.3 Dual storage I/O adapters	371
9.1.4 Encrypted ASP enhancements	375
9.1.5 Disk response time buckets enhancements	378
9.1.6 Central processor complex node level mirroring	379
9.1.7 EXP24S SFF Gen2-bay drawer (#5887)	382
9.1.8 Higher capacity 10 K RPM SFF SAS disk drives	382
9.1.9 Tape performance instrumentation	382
9.1.10 Tape library resource name changes for IOP-less IOA attachment	383
9.1.11 Tape library unreadable barcode changes for IOP-less IOA attachment	383
9.1.12 DVD / Tape SAS External Storage Unit for Power 795 CPC Rack	384
9.1.13 RDX support	384
9.1.14 POWER7+ 770/780 Native I/O support	386
9.2 SAN storage management enhancements	386
9.2.1 Multipathing for virtual I/O	386
9.2.2 DS5000 native attachment	388
9.2.3 Level of protection reporting for multipath disk units	390
9.2.4 Library control paths for IOP-less Fibre Channel IOA tape attachment	391
9.2.5 External disk storage performance instrumentation	392

9.2.6 Thin provisioning for DS8700 and DS8800 storage servers and for VIOS shared storage pools	397
9.3 SSD storage management enhancements	398
9.3.1 DB2 media preference	398
9.3.2 ASP balancer enhancements for SSDs	403
9.3.3 User-defined file system media preference	407
9.3.4 177 GB SFF SSD with eMLC	409
9.3.5 IBM Disk Sanitizer PRPQ extended to include SSD devices	409
Chapter 10. Networking enhancements	413
10.1 TCP/IP enhancements summary	414
10.1.1 IPv6 support enhancements	414
10.1.2 ISC-based DHCP server supports IPv6 and failover	414
10.1.3 DHCPv6 client	416
10.1.4 SNMP	417
10.1.5 SMTP	418
10.1.6 IPv6 support added for PPP connections	419
10.2 HTTP server support for PASE FastCGI PHP processing	423
10.2.1 IBM i 5.4 required components	424
10.2.2 IBM i 6.1 required components	424
10.3 Telnet client support for SSL encryption	425
10.4 Sockets programming: New user exits	426
10.4.1 Exit points that are defined in the user registry	426
10.4.2 Example: User exit program for QIBM_QSO_ACCEPT	426
10.5 IKEv2 support	428
10.5.1 Enhanced Cryptographic Algorithms	429
10.6 IPv6 TCP/IP applications and V6R1 enablement PTFs	429
10.7 IBM AnyNet support on IBM i 7.1	430
10.8 Ethernet link aggregation	430
10.8.1 Overview of Ethernet link aggregation	430
10.8.2 Prerequisites for Ethernet link aggregation	431
10.8.3 Configuring Ethernet link aggregation	431
10.8.4 Example: Configuring four aggregated network links	433
10.9 Ethernet Layer-2 bridging	434
10.9.1 Introduction to Ethernet Layer-2 bridging	434
10.9.2 How Ethernet Layer-2 bridging works on IBM i	434
10.9.3 IBM i prerequisites for Ethernet Layer-2 bridging	435
10.9.4 Configuring a shared network connection through Ethernet Layer-2 bridging	436
10.10 IBM Portable Utilities for i (5733-SC1) supported versions	439
Chapter 11. Printing enhancements	441
11.1 Printing overview	442
11.2 Overview of IBM i 6.1 print enhancements	443
11.3 Enhancements to base print functions	444
11.3.1 CPYSPLF command enhancements for copying to PDF or TIFF stream files	444
11.3.2 New QIBM_QSP_SECURITY exit point and formats	447
11.3.3 Encryption of PDF output (Infoprint Server LP)	449
11.4 Print Services Facility for IBM i enhancements	451
11.4.1 Disabling offset stacking and edge mark printing	451
11.4.2 Specifying public data authority for directories that are created by PSF	451
11.4.3 Improved PSF debugging capabilities	452
11.5 Transform Services enhancements	453
11.5.1 Generating PDFs from existing spooled files	453

11.5.2 PDF transform enhancements	454
11.6 IBM Navigator for i print enhancements	454
11.6.1 Viewing printer output in PDF format	455
11.6.2 Exporting the printer output in PDF format	457
11.6.3 System i Navigator and exporting printer output in PDF format	458
11.7 IBM i Access for web print enhancements	462
11.8 Host Print Transform enhancements	463
11.8.1 The Host Print Transform API and Transform Services	463
11.8.2 Example: Host Print Transform API with Transform Services.	463
11.9 References	464
Chapter 12. Integration with IBM BladeCenter and IBM System x	467
12.1 iSCSI software targets	468
12.1.1 IBM i Integrated server object model with a hardware target	468
12.1.2 IBM i Integrated server object model with a software target	469
12.1.3 Direct connect software targets.	469
12.2 Defining iSCSI software target support	469
12.2.1 CRTDEVNWSH CL command interface	470
12.2.2 IBM Navigator for i changes for iSCSI software target support.	471
12.3 Service Processor Manager function	472
12.4 VMware support changes	472
12.4.1 New NWSD types	473
12.4.2 VMware ESX server management	473
12.4.3 SWA storage spaces for VMware ESX servers	474
12.5 New planning worksheets	474
12.6 IBM Navigator for i	474
12.6.1 Create Server task	475
12.6.2 Clone Integrated Windows Server task.	477
12.6.3 Delete Server task	479
12.6.4 Launch Web Console	480
12.6.5 Simplified Windows File Level Backup (FLB) from IBM i	481
12.7 New IBM i CL commands	481
12.7.1 Install Integrated Server (INSINTSVR)	482
12.7.2 Delete Integrated Server (DLTINTSVR) command.	482
12.8 IBM i changed CL commands.	483
12.8.1 Install Windows Server (INSWNTSVR) CL command	483
12.8.2 Create NWS Configuration (CRTNWSCFG) and Change NWS Configuration (CHGNWSCFG) CL commands	483
12.8.3 Install Linux Server (INSLNXSVR) CL command	484
12.8.4 No new integrated Linux servers.	484
12.9 Fewer IBM i licensed programs are required	484
12.10 Changes to IBM i integration with BladeCenter and System x documentation	484
12.10.1 A new IBM i iSCSI Solution Guide PDF	484
12.10.2 IBM i 7.1 Information Center.	485
12.10.3 New IBM i integration with BladeCenter and System x group on IBM developerWorks	486
12.10.4 New IBM i Technology Updates page on developerWorks.	486
12.10.5 IBM i integration with BladeCenter and System x Marketing website.	486
Chapter 13. IBM Advanced Job Scheduler for i enhancements	487
13.1 Advanced Job Scheduler capabilities	488
13.1.1 Scheduling a job	488
13.1.2 Job groups	488

13.1.3	Notification and report distribution.	488
13.1.4	Remote command jobs.	488
13.1.5	Multiple scheduling environments.	488
13.2	IBM Systems Director Navigator for i AJS support	489
13.2.1	Navigating to the AJS menu	489
13.2.2	AJS menu	496
13.2.3	Start Scheduler action.	496
13.2.4	Stop Scheduler action.	497
13.2.5	AJS properties menu	497
13.2.6	Scheduled Jobs menu	504
13.2.7	Using the Scheduled Jobs action	506
13.2.8	Using the Scheduled Jobs table	508
13.2.9	Adding a new scheduled job	514
13.3	Other AJS enhancements	527
13.3.1	Notification enhancements	527
13.3.2	Multiple Scheduling Environments function	528
13.3.3	Scheduling and running jobs.	529
13.4	References	532
Chapter 14.	Web serving	533
14.1	Triggered cache manager removed	534
14.2	Web server search engine and web crawler support removed	534
14.3	Plug-ins and LoadModule directives	534
14.4	HTTP Server for i with HA IPv6 support	539
14.5	IBM HTTP Server plug-in module	540
Chapter 15.	Web enablers.	541
15.1	IBM Integrated Web Services for i	542
15.1.1	Integrated Web Services Client for ILE.	542
15.2	Changes in IBM WebSphere Application Server	542
15.2.1	Packaging for WebSphere Application Server	542
15.2.2	Java for IBM i 7.1	543
15.2.3	Installation.	544
15.2.4	WebSphere Application Server Installation Management Supports on IBM Web Admin for i.	544
15.2.5	Upgrading to IBM i 7.1	551
15.2.6	IBM Installation Manager for WebSphere Application Server V8	552
15.2.7	Enabling IBM Technology for Java virtual machine	553
15.3	IBM Web Administration for i.	554
15.3.1	Web Log Monitor.	554
15.3.2	Permissions	565
15.3.3	Web Performance Advisor	574
Chapter 16.	Application development.	575
16.1	High-level programming languages	576
16.2	Control Language (CL)	576
16.2.1	New workload capping commands	576
16.2.2	Retrieving CL source support for ILE CL	576
16.2.3	Longer Integer CL variables for ILE CL.	577
16.2.4	Showing DO and SELECT nesting levels in a compiler listing	577
16.2.5	Encrypting the debug listing view for ILE CL.	577
16.2.6	Nested INCLUDE support.	578
16.3	PHP.	578
16.3.1	Zend Server Community Edition for IBM i.	579

16.3.2	Zend Server for IBM i	579
16.3.3	Zend Studio for IBM i	579
16.4	XMLSERVICE for IBM i	580
16.4.1	XMLSERVICE highlights	580
16.4.2	XMLSERVICE start	581
16.5	IBM Lotus products for IBM i	585
16.6	Native archive and unarchive API support	586
16.6.1	QZIPZIP API	586
16.6.2	QZIPUNZIP API	587
16.6.3	QZIPUTIL Service program	588
16.7	IBM Toolbox for Java JDBC	588
16.7.1	JDBC 4.1	589
16.7.2	XML Data type support	589
16.7.3	Database metadata updates	589
16.7.4	Currently committed support	589
16.7.5	Array type support	589
16.7.6	Long schema name support	590
16.8	JTLite/JTOpenLite Java Classes Library	590
16.9	Application Runtime Expert for i	591
16.9.1	Deployment template	591
16.9.2	Customizing plug-ins for templates	592
16.9.3	Application Runtime Expert console	597
16.9.4	Application Runtime Expert requirements	602
16.9.5	More information	602
Chapter 17.	IBM Rational products	603
17.1	IBM Rational Developer for Power Systems Software	604
17.1.1	RPG and COBOL Development Tools	604
17.1.2	Rational Team Concert client integration for IBM i	608
17.1.3	Enhancements in Version 8.0.3	608
17.1.4	Version 8.x Fix Packs	609
17.1.5	Enhancements in Version 8.5	610
17.1.6	Migration to Rational Developer for Power Systems Software V8.5	611
17.2	IBM Rational Team Concert	611
17.2.1	Integration with Rational Developer for Power Systems Software	612
17.2.2	What is new in Rational Team Concert V3.0.1	615
17.2.3	What is new in Rational Team Concert V4.0	616
17.3	IBM Rational Developer for i for SOA Construction	617
17.4	IBM Rational Development Studio for i	618
17.4.1	Source code protection option	618
17.4.2	ILE RPG compiler	619
17.4.3	Sorting and searching data structures	620
17.4.4	ALIAS support	621
17.4.5	Performance improvement when returning large values	622
17.4.6	ILE COBOL compiler	624
17.4.7	ILE C and ILE C++	626
17.5	IBM Rational Open Access: RPG Edition	631
17.5.1	How to use Rational Open Access	632
17.5.2	IBM Rational Open Access: RPG Edition withdrawn	635
17.5.3	Open Access requirements	636
17.6	IBM Rational Application Management Toolset for i	636
17.6.1	Application Management Toolset for i Licensing	638
17.6.2	Application Management Toolset for i Requirements	638

17.6.3 Accessing Application Management Toolset for i	638
Chapter 18. IBM Navigator for i V7.1	639
18.1 New name: IBM Navigator for i	640
18.2 Signing on to IBM Navigator for i	641
18.3 Independent tabs	643
18.4 Dynamic navigation	644
18.5 Modern tables	645
18.6 5250 emulation-related enhancements	646
18.7 Set Target System enhancements	647
18.8 Database enhancements	648
18.8.1 On Demand Performance Center enhancements	649
18.8.2 Health center enhancements	656
18.8.3 Database management enhancements	657
18.8.4 Database Maintenance category enhancements	669
18.8.5 Long SQL schema names enhancements	672
18.8.6 OmniFind text search functions enhancements	672
18.9 New journal management enhancements	673
18.9.1 Show Journalized Objects function	673
18.9.2 Change Receivers function	674
18.9.3 Remote Journals function	675
18.9.4 Add Remote Journal function	676
18.9.5 Properties function	677
18.9.6 Activate function	677
18.9.7 Deactivate function	678
18.9.8 Remove function	678
18.9.9 Properties function	679
18.10 Integrated Server Administration enhancements	679
18.10.1 Create Server function	680
18.10.2 New Based on ... (cloning) a Server function	681
18.10.3 Delete Server function	682
18.10.4 Launch Web Console function	683
18.11 Printer output enhancements	683
18.11.1 View as PDF function	684
18.11.2 Export as PDF function	684
18.12 File system enhancements	685
18.13 Networking enhancements	686
18.14 Disk management enhancements	687
18.14.1 Graphical view enhancements	687
18.14.2 Start / Stop encryption on disk pools enhancements	690
18.14.3 Asynchronous delivery for Geographic Mirroring enhancements	690
18.15 Tape support enhancements	691
18.16 Performance enhancements	692
18.16.1 General health indicators	693
18.16.2 Viewing cross-partition processing	703
18.16.3 New Database package	716
18.16.4 QAPMCONF perspective	723
18.16.5 QAPMJOBMI perspective	725
18.16.6 Create Performance Data Report function	727
18.16.7 Image and file export	728
18.16.8 Sizing the next upgrade	729
18.16.9 Ethernet Link Aggregation Modified perspectives	732
18.16.10 Workload Group perspectives	734

18.16.11	Java perspectives	736
18.16.12	Disk Response Time charts	739
18.16.13	Interactive perspective development.	744
18.16.14	Metric finder	751
18.16.15	Performance Explorer content package	753
18.16.16	New metrics	753
18.16.17	Miscellaneous enhancements.	758
18.17	JS1- Advanced Job Scheduler for i enhancements	768
18.18	Backup Recovery Media Services: Plug-in enhancements.	769
18.19	More information	769
Chapter 19.	IBM i Access	771
19.1	IBM i Access for Windows.	772
19.1.1	Installation enhancements	772
19.1.2	.NET Data Provider enhancements	772
19.1.3	OLE Data Provider enhancements	772
19.1.4	Windows ODBC Driver enhancements.	773
19.1.5	Data Transfer enhancements	773
19.1.6	Personal Communications Emulator enhancements	773
19.1.7	Direct Attach Operations Console withdrawal.	773
19.2	IBM i Access Client Solutions	774
19.2.1	Prerequisites	774
19.3	IBM Navigator for i	775
19.4	System i Navigator	776
19.5	IBM i Access for Web enhancements	776
19.5.1	AFP to PDF Transform	777
19.6	IBM System i Access for Wireless	779
Chapter 20.	Miscellaneous enhancements.	781
20.1	Licensed product structures and sizes	782
20.2	Changed or new CL commands and APIs	783
20.3	Temporary user-defined file systems	783
20.4	Watch for Event function (message enhancements)	785
20.4.1	Advantages of using watches	785
20.4.2	New Watch for Event function for IBM i 7.1	785
20.5	IBM Tivoli Directory Server for IBM i enhancements.	788
20.5.1	Creating suffix entries automatically when they are needed.	788
20.5.2	Administrative roles.	788
20.5.3	User interface enhancements	788
20.5.4	Security enhancements	788
20.5.5	New password encryption options	789
20.5.6	Pass-through authentication	789
20.5.7	Enhanced password policy to use global date and time for initialization	789
20.5.8	Multiple password policies	789
20.5.9	Policy that is enforced for Digest-MD5 binds	789
20.5.10	Persistent search	789
20.5.11	Replication configuration enhancements	790
20.5.12	Filtered replication.	790
20.5.13	Limiting the number of values that are returned by a search	790
20.5.14	Enhanced syntaxes and matching rules	790
20.5.15	IASP enablement for Directory Server on IBM i	790
20.5.16	The ldapcompare utility.	790
20.5.17	Providing a re-entrant LDAP C client library	790

20.6 Automating extra IPLs for PTF installation	791
20.7 IBM i workload groups.	791
20.7.1 Example of how Workload Group works.	792
20.7.2 IBM Passport Advantage Sub-capacity Licensing support	793
20.8 IBM i on a Flex Compute Node	794
Chapter 21. Upgrading IBM i	795
21.1 Installing or upgrading IBM i	796
21.1.1 IBM i 7.1 support for POWER5 and later hardware	796
21.1.2 Planning statements	797
21.1.3 Supported upgrade paths	797
21.1.4 Program temporary fixes for upgrading to IBM i 7.1	797
21.1.5 Pre-upgrade verification tool	798
21.1.6 License program releases and sizes.	798
21.1.7 Server firmware requirements.	798
21.1.8 IBM i 7.1 Resaves.	798
21.2 Upgrading from i5/OS 5.4 to IBM i 7.1	799
21.2.1 Object conversions	799
21.2.2 Spooled file conversions	799
21.2.3 Integrated file system conversions	800
21.2.4 Backup Recovery and Media Services	800
21.3 Media delivery changes	801
21.3.1 IBM i 7.1 multi-language distribution media	801
21.3.2 Media charges	802
21.4 IBM i network upgrade	802
21.4.1 Client system requirements.	803
21.5 More considerations for upgrading to IBM i 7.1	805
21.5.1 IBM i Console changes	805
21.5.2 Full automatic installation and upgrades from optical media	806
21.5.3 IBM i 7.1 5770-SS1 Option 1 installation actions	807
21.5.4 Adjusting the cluster version of a cluster.	807
21.5.5 Java considerations	807
21.5.6 Miscellaneous changes and withdrawals	808
21.6 IBM i network installation	810
21.6.1 Requirements for IBM i network installation	810
Related publications	813
IBM Redbooks	813
Other publications	814
Online resources	814
How to get Redbooks.	818
Help from IBM	818

Notices

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

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

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

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

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.

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.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

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

Active Memory™	IBM Flex System™	PowerVM®
AIX®	iCluster®	Print Services Facility™
AIX 5L™	IMS™	ProtectTIER®
AnyNet®	InfoSphere®	PureFlex™
AS/400®	Integrated Language Environment®	Quickr®
BladeCenter®	iSeries®	Rational®
DataMirror®	Jazz™	Rational Team Concert™
DB2®	Language Environment®	Redbooks®
DB2 Connect™	Lotus®	Redbooks (logo)  ®
developerWorks®	Lotus Enterprise Integrator®	Sametime®
Domino®	OmniFind®	Service Director™
DRDA®	OS/400®	Storwize®
DS6000™	Passport Advantage®	System i®
DS8000®	POWER®	System p®
EnergyScale™	POWER Hypervisor™	System Storage®
eServer™	Power Systems™	System x®
Express Servers™	Power Systems Software™	System z®
FileNet®	POWER6®	SystemMirror®
FlashCopy®	POWER6+™	Tivoli®
Guardium®	POWER7®	WebSphere®
i5/OS™	POWER7+™	xSeries®
IBM®	PowerHA®	z/OS®

The following terms are trademarks of other companies:

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

Netezza, and N logo are trademarks or registered trademarks of IBM International Group B.V., an IBM Company.

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

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

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

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

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

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

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

Preface

This IBM® Redbooks® publication introduces a technical overview of the new features, functions, and enhancements available in IBM i 7.1. It provides a summary and brief explanation of new capabilities and what changed in the operating system. This publication also describes many of the licensed programs and application development tools that are associated with IBM i. This publication includes all the enhancements that were delivered with the Technology Refresh 4 and 5 that were made available in 2012.

Many new and enhanced functions are described:

- ▶ IBM DB2® for i support for XML and column level encryption
- ▶ IBM PowerHA® for i asynchronous Geographic Mirroring and LUN-level switching
- ▶ Virtualization enhancements for IBM i hosted partitions and IBM PowerVM®
- ▶ Storage management enhancements for solid-state drives
- ▶ Systems management enhancements for Systems Director Navigator for i
- ▶ IBM Rational® software development product enhancements
- ▶ Zend PHP Enhancements
- ▶ Web Application Serving enhancements

The information that is provided in this book is useful for clients, IBM Business Partners, and IBM service professionals that are involved with planning, supporting, upgrading, and implementing IBM i 7.1 solutions.

The team who wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Rochester Center.



Hernando Bedoya is a Senior IT Specialist at STG Lab Services and Training, in Rochester, Minnesota. He writes extensively and teaches IBM classes worldwide in all areas of DB2 for i. Before joining STG Lab Services, he worked in the ITSO for nine years writing multiple IBM Redbooks publications. He also worked for IBM Colombia as an IBM AS/400® IT Specialist doing presales support for the Andean countries. He has 28 years of experience in the computing field and has taught database classes in Colombian universities. His areas of expertise are database technology, performance, and data warehousing. He holds a Master's degree in Computer Science from EAFIT, Colombia.



Octavio Bustos is a Senior IT Specialist in Colombia from Redsis SAS, an IBM Premier Business Partner in Colombia. He has more than 20 years of experience with IBM midrange products starting with AS/400 and now IBM i and Power Systems™. He has broad knowledge in areas such as Performance Management, Security, Business Continuity architecture, SAP sizing and solution design, RPG and CL programming, Rational and PHP development tools for IBM i.



Clive Griffiths is a Technology Consultant at Admiral Insurance services, which is based in the UK and has been with Admiral for 16 years. His skills are varied within the IBM i arena and he has a deep experience of everything about IBM i from hardware to programming. Trained as an RPG programmer over 25 years ago, he has coded in everything from Cobol to Synon/2. His passion is DB2 for i and he possesses extensive experience in in-depth performance analysis to PEX level along with the experience to use all the tools at hand to provide swift problem determination from code to hardware issues. He is an architect on IBM i and an evangelist for the IBM i platform. He delivers extensive in-house education to the staff of Admiral, from basic operations to SQL performance and IBM i futures.



Tom Grigoleit is a Power Systems and IBM i Specialist from Rochester, Minnesota. He works for IBM Premier Business Partner Meridian IT as a Senior Solution Architect, performing pre-sales and post-sales work for many customers across the US. Tom's expertise is in the design and implementation of Power Servers and Storage, including tape, disk, high availability, and virtualization. Before his current assignment, he enjoyed many years with IBM in various positions, including Systems Engineer, Manager, Product Planner, and Program Developer. He holds certifications in Information Systems Architecture, Open System Storage®, Power i Servers, and Enterprise Architecture, and has a Master's Degree from Syracuse University.



Axel Lachmann is a freelance systems architect at FoxCom, an IBM Business Partner in Germany. He has 18 years of experience in the IBM OS/400® / IBM i5/OS™ / IBM i field and more than 10 years of IBM Redbooks experience. He holds multiple IBM technical certifications. His areas of expertise include system design, performance analysis, BRMS, IBM WebSphere® on IBM i, and LPAR. Axel also teaches technical and certification courses for IBM Learning Services in Germany, UK and Scandinavia. You can reach Axel at alachmann@foxcom.de.



Yohichi Nakamura is an Advisory IT Specialist. He joined IBM Japan Systems Engineering Co., Ltd., which is a subsidiary of IBM Japan, as IBM i Advanced Technical Support in 2003. He temporarily moved to IBM Japan as a Technical Sales Specialist for Power Systems and IBM i in 2010. He holds both FTSS and ATS roles, so he provides pre-sales technical support, especially cutting-edge solutions for Power Systems and IBM i. His areas of expertise include virtualization, external storage, iSCSI integration solutions, high availability, performance analysis, backup/recovery, and overall of IBM i operations.



Sandra Nelson is an Advisory Education Specialist at the IBM i Global Support Center in Rochester, MN. She started with IBM in 1992 as a customer support professional specializing in Facsimile Support/400, PM/400, IBM Service Director™, and IBM i Print. For the last 12 years, she has been managing, writing, and teaching technical education for IBM i software and hardware support engineers worldwide. She has taught classes on various IBM Power Systems and IBM i topics and has extensive experience with virtual classroom tools.



Xiao Hu Qu is a Advisory IT Specialist at the STG Advanced Technical Skill, Beijing, China. He works as pre-sale technical support engineer. He has 13 years of experience working with AS/400, IBM eServer™ iSeries®, IBM System i®, and Power Systems. Xiao Hu's areas of expertise include IBM i System Management, Virtualization, Cloud, High Availability/Disaster Recovery, Security, Performance Tuning, PowerVM, and VIOS. Before joining China ATS, he worked as a systems engineer for an IBM Business Partner.



Tom Vernailen is a Senior IT Specialist in Belgium. His areas of expertise include IBM i problem determination, network connectivity, and communications, including TCP/IP. He has participated in previous ITSO residencies, preparing technical overview presentation material and writing TCP/IP communications-related IBM Redbooks publications for IBM i. He is an Enhanced Technical Support remote Account Advocate for several IBM Power Systems customers.

Thanks to the following people for their contributions to this project:

Ann Lund, Linda Robinson, Jenifer Servais

International Technical Support Organization, Poughkeepsie Center

Terry D. Ackman, Mark J Anderson, Stacy L Benfield, Robert J Bestgen, Chris Beyers, David R Bhaskaran, Brian K. Bratager, Kent L Bruinsma, Dan Boyum, Lilo Bucknell, Bunny Chaney, David S Charron, Armin Christofferson, Jason Clegg, Tom Crowley, Jenny Dervin, Jessica Erber-Stark, Jerry Evans, Margaret Fenlon, Jim Flanagan, Scott Forstie, Christopher Francois, Dave Johnson, Robert Gagliardi, Mark Goff, Maryna Granquist, Roger Guderian, Kristi Harney, Stacy L. Haugen, Terry Hennessy, Mark J Hessler, Wayne Holm, Steven M Howe, Allan E Johnson, Randy Johnson, Rodney Klingsporn, Tim Klubertanz, Stephen A Knight, Joe Kochan, Joseph Kochan, Tim Kramer, Dave Legler, Jim Lembke, Al Levering, Terry Luebbe, Ron McCargar, Ginny McCright, Jossie McManus, Marilyn Martin, Dawn May, Cindy Mestad, Kent Milligan, Joe Mulholland, David B Murray, Glen Nelson, Robert Padzieski, Ricky M Peterson, Mary Ann Pierce, Lora Powell, Padma Rao, Steve Ransom, Tim Rowe, Darcy Rucker, Jon Rush, Rick Saltness, Curt Schemmel, Mark Schroeder, Bob Seemann, Gerald Simon, Shauna Smith, Jasmeet Singh, Alan Smith, Shauna Smith, Steve Sparrow, Ellen Stacey, Eric Stec, Ellie B. Streifel, Johnnie R Talamantes, Jim Tilbury, Rodney Traff, Kathryn A Tri, Mark R Vanderwiel, Wesley Varela, Mervyn Venter, Guy Vig, Brian Warford, Duane Wenzel, Kris Whitney, Steve Will

IBM Rochester

Li LL Guo, Sheng Li Li, Jian Sang, Gang Shi, Ting Ting Sun, Meng MS Su, Dong Dong DD Su, Nan Wang

IBM China

Alison Butterill, Phil Coulthard, George G Farr, Philip Mawby, Barbara Morris

IBM Toronto

Thanks to the authors of the previous editions of this book.

Authors of the first edition, *IBM i 7.1 Technical Overview*, which was published in October 2010, were:

Hernando Bedoya, Ivan Berrios, Ingo Dimmer, Robert Jacobsen, Prema Sathasivam, Jos Vermaere, Tom Vernailen, Allyn Walsh.

Authors of the second edition, *IBM i 7.1 Technical Overview with Technology Refresh Updates*, which was published in December 2011, were:

Hernando Bedoya, Jiri Sochr, Michael Harper, Robert Jacobsen, Matias Centeno Lozada, Ahmed Mansour, Tom Vernailen.

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author - all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

ibm.com/redbooks

- Send your comments in an email to:

redbooks@us.ibm.com

- Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/pages/IBM-Redbooks/178023492563?ref=ts>
- ▶ Follow us on twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>

Summary of changes

This section describes the technical changes that were made in this edition of the book and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

Summary of Changes
for SG24-7858-02
for IBM i 7.1 Technical Overview with Technology Refresh Updates
as created or updated on March 20, 2013.

March 2013, Third Edition

The first edition of this book included all the new enhancements that were announced at the General Availability of IBM i 7.1 in early 2010. Following the General Availability of IBM i 7.1, IBM i point / modification releases were replaced by a new release delivery mechanism that is called a Technology Refresh. A Technology Refresh is a collection of operating system software that is developed together, packaged together, tested together, and delivered as a PTF Group for a base release.

The second edition of this book included all the enhancements that were included in the first three Technology Refreshes (TR1, TR2, and TR3) that were announced since the product went to General Availability. There are enhancements basically for every topic that is related to IBM i.

This third edition of this book includes all the enhancements that were included in the next two Technology Refreshes (TR4 and TR5) that were announced in 2012. There are enhancements basically in every topic that is related to IBM i. For this reason, every chapter was modified to include these enhancements.



IBM i 7.1 introduction and summary

This chapter summarizes new functions and enhancements available in IBM i 7.1 and new features that were added in IBM i 6.1.1, which was delivered in October 2009 after *IBM i 6.1 Technical Overview*, SG24-7713 was published. This chapter highlights these enhancements and the new capabilities in the operating system, the licensed programs, and application development tools that are associated with IBM i. 7.1.

The following chapters provide more overview information. As a technical overview document, detailed instructions or examples are beyond the scope of this publication. The purpose of this document is to consolidate into a single reference a summary of the new information that is related to IBM i 7.1.

This chapter describes the following topics:

- ▶ Communicating IBM i enhancements
- ▶ Clarification of operating system terminology
- ▶ What is new in IBM i 7.1

IBM i Memo to Users 7.1 is a valuable starting point for readers, and is available at the following website:

<https://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzaq9/rzaq9.pdf>

More detailed information about the IBM i 7.1 enhancements can be found at various other sources:

- ▶ The IBM i Information Center for IBM i 7.1 website:
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp>
- ▶ The Upgrade Planning website:
<http://www-947.ibm.com/systems/support/i/planning/upgrade/v7r1/planstmts.html>
- ▶ Planning - Customer Notices and information, found at the following website:
<http://www-947.ibm.com/systems/support/planning/notices71.html>

1.1 Communicating IBM i enhancements

The way IBM communicates IBM i enhancements has evolved. It was not all that many years ago that information was communicated through hardcopy documents and books. Electronically, though, the development lab has used a couple of distribution methods for many years.

The primary way that information has been communicated to the user community is through the *IBM i Information Center*. The Information Center contains topics that help you with basic and advanced tasks. It remains the reference for the platform. The Information Center is updated periodically, but is not able to react quickly to change. You can find the Information Center for IBM i 7.1 at:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp>

The other communication method the IBM i lab uses is the Request for Announcement (RFA). The RFA is the formal document that contains information about a function that is being delivered to the market. For more information about this document, go to the IBM Offering Information - Announcement letters website at:

<http://www-01.ibm.com/common/ssi/>

Today, the IBM i zone in IBM developerWorks® (DW) and social media are part of the platform strategy for getting information to partners and clients. The information in developerWorks is for everyone, not just developers. The IBM developerWorks website can be found at:

<http://www.ibm.com/developerworks/ibmi/>

With DW, if there is new information to share with IBM customers, it can be shared immediately. There are several different ways that information is delivered through DW:

- ▶ **Article-based Information:** Much of the information about DW is in the form of short articles that define a task or technology. The content provides not only information about a topic, but also tries to answer the question, “Why is this topic important to me?”. These articles are written by many developers.
- ▶ **Technology Updates Wiki:** The IBM i Technology Updates page in developerWorks is the home for detailed information about the functions that are delivered with each new Technology Refresh, as well as the functions that are delivered through other means between releases. The technology updates section is organized for easy navigation, searches, and subscription. At the bottom of these pages is a button that lets you subscribe to the page so that you are notified when updates are made to it. Because this page is updated every time a new PTF Group is updated, you get to track and monitor new enhancements and fixes as they are delivered.
- ▶ **Community and Forums:** The Community and Forums section is a list of links to help you connect to various things. There are blogs from experts here in the development lab, as well as blogs from others in the industry. They are short articles on a wide array of topics, all aimed at communicating with our IBM i users. Blogs include:
 - “You and i”: Announcements, strategy, and customer references
 - “i Can”: Tips, techniques, and examples
 - “Modern-i-zation”: Modern application development
 - “DB2fori”: IBM i integrated DB2 and related topics

The Forums target different topical categories and give you an opportunity to ask IBM i developers questions. These forums are monitored by experts in the development lab.

DW has many links to other topic and technology areas that IBM i users need. DW is organized to help you get the information that you need when you need it. It also is a great place for helping you stay informed about what is happening with IBM i.

Figure 1-1 illustrates the various methods that are used to communicate IBM i enhancements to IBM i users.




Figure 1-1 Methods for communicating IBM i enhancements

IBM uses Twitter extensively. Steve Will, Chief Architect for IBM i, uses Twitter to notify followers whenever his blog has something new, what is going on in IBM i development, or to point to a webcast, article, or blog that might be useful. Follow @Steve_Will_IBMi on Twitter.

IBM continues to enhance communication with its users to ensure timely updates with the details that you need.

1.2 Clarification of operating system terminology

When referring to operating systems releases, it is important to understand the naming transition from IBM OS/400 to i5/OS to IBM i. When IBM introduced POWER5, OS/400 was renamed i5/OS. When the IBM POWER6® platform became available in January 2008, IBM announced a major new release named i5/OS V6R1. Later that same year, the name was changed to IBM i to disassociate any presumed dependency on the POWER5 hardware platform. The notations 5.4 and 6.1 were introduced to indicate operating systems release levels V5R4 and V6R1. User documentation, web page links, and programmed interfaces use IBM i terminology and others still use the i5/OS nomenclature. In this publication, we use IBM i terminology, but occasionally also use i5/OS, typically where it is part of a product name or appears in a window.

1.3 What is new in IBM i 7.1

IBM i 7.1 delivers database, virtualization, storage management, web application serving, and business resiliency enhancements that clients can use to reduce cost, improve service, and manage risk of their IT infrastructure.

1.3.1 IBM i 7.1 on POWER based servers

IBM i 7.1 is supported on the following platforms:

- ▶ Power Systems servers with IBM POWER7+™ processors
- ▶ Power Systems servers and blades with IBM POWER7® processors
- ▶ Power Systems servers and blades with POWER6/6+ processors
- ▶ System i servers with POWER6 processors
- ▶ System i servers with POWER5/5+ processors

1.3.2 Technology Refreshes

IBM i 7.1 includes an enhancement to make it easier to deliver hardware and firmware enhancements through *Technology Refreshes*. The new hardware-related and firmware-related machine code content is contained within PTFs in a single Technology Refresh PTF Group. The content is referred to as IBM i 7.1 Technology Refresh 1, IBM i 7.1 Technology Refresh 2, and so on.

If you already have IBM i 7.1 on your partition, you can install the Technology Refresh PTF Group (SF99707). Alternatively, you can install a Resave after the normal Resave installation procedure, and then install the Technology Refresh PTF Group.

Future deliveries of new hardware-related functions for IBM i 7.1 are planned as updates to the Technology Refresh PTF Group. Only the latest level of the Technology Refresh PTF Group is available. Each new level of the PTF Group contains all the content from the previous levels, so at any point in time you need to install only the latest level.

This book assumes that you have IBM i 7.1 with the latest Technology Refresh and the latest Group PTFs installed to take advantage of the new enhancements.

For more details about Technology Refreshes, see Chapter 2, “Technology Refresh” on page 15.

1.3.3 Integrated DB2 database for IBM i

Integrated DB2 database for IBM i is enhanced with support for XML, enabling you to store and search XML documents in DB2 for i and to create XML documents from existing relational data. SQL statements use IBM OmniFind® Text Search Server, which adds support for searching XML documents. It is available with IBM i at no additional charge. DB2 Web Query also includes support for querying XML documents with IBM i 7.1.

DB2 for IBM i supports the calling of field procedures, which can allow a partner application to provide column level encryption in a database table transparently.

Performance improvements include adaptive query processing, which can modify a query plan when the query is running to improve performance. An advanced SQE query optimizer now supports native logical files. Also included is improved management tools to monitor long running operations, a new SQL_CANCEL procedure to cancel long running queries, and the ability to monitor random or sequential I/O statistics to identify tables that can benefit from solid-state drives (SSDs).

Here are other significant enhancements:

- ▶ Row locking performance in cases where a job acquires many locks.
- ▶ The introduction of **RUNSQL**, a new command to run SQL directly from the command line or within a CL program.
- ▶ A system naming convention for SQL that now includes permit (/) and (.) qualifiers. There is a new database perspective within System i Navigator.

1.3.4 PowerHA SystemMirror for i

Major enhancements that are available with IBM i 7.1 include asynchronous geographic mirroring, which you can use to extend the distance between locations for disaster recovery implementations, and support for IBM System Storage DS8000® Space-Efficient FlashCopy®, which allows a FlashCopy without requiring double storage capacity.

LUN level switching is also available for a local high availability (HA) solution. You can use this feature to switch an IASP on an IBM System Storage DS8000 or DS6000™ between local systems. This feature includes support for automatic failover. Starting with IBM i 6.1.1 on POWER6 and beyond, N_Port_ID (NPIV) support was added for Fibre Channel adapter virtualization through IBM PowerVM Virtual IO Server (VIOS) attached SANs.

Here are more enhancements:

- ▶ New PowerHA packaging
- ▶ PowerHA versions
- ▶ New CL commands for programming cluster automation
- ▶ Cluster administrative domain enhancements
- ▶ Better detection of cluster node outages
- ▶ Improved geographic mirroring full synchronization performance
- ▶ Improvement of geographic mirroring in an IBM i hosted IBM i client environment
- ▶ Support virtualization capabilities on an active cluster node
- ▶ IPv6 support

The new 5799-HAS Program Request Pricing Quotation (PRPQ), named IBM PowerHA SystemMirror® for i Enhancements, delivers the following new functions:

- ▶ Support for managing IBM System Storage SAN Volume Controller and IBM Storwize® V7000 Copy Services
- ▶ IBM i command **CFGDEVASP** with 5770-SS1 PTF SI44141
- ▶ IBM i command **CFGGEOMIR** with 5770-HAS PTF SI44148
- ▶ New PowerHA GUI

In April 2012, IBM PowerHA SystemMirror for i started to support the following new functions. Some of these functions are covered with 5799-HAS Program Request Pricing Quotation (PRPQ), so the requirement of the 5799-HAS PRPQ for using these functions were eliminated.

- ▶ Support for managing IBM System Storage SAN Volume Controller and IBM Storwize V7000 Copy Services
- ▶ IBM i command **CFGDEVASP** with 5770-SS1 PTF SI44141
- ▶ IBM i command **CFGGEOMIR** with 5770-HAS PTF SI44148
- ▶ New PowerHA GUI

Since October 2012, geographic mirroring of IBM PowerHA SystemMirror for i can eliminate the transfer of the part of temporary storage spaces for the IBM i client partition in an IBM i hosted IBM i client partition environment. This enhancement reduces the amount of a network traffic between IBM i host partition on the production node side and on the backup node side.

Also in 2012, Live Partition Mobility started to support IBM i 7.1. An IBM i 7.1 partition, which is a member of an active cluster, can also move to another Power Systems server through the Live Partition Mobility function.

1.3.5 IBM i virtualization

IBM i virtualization support offers increased flexibility for running multiple IBM i releases on a single server. Specifically, on the POWER6 or POWER7 platform, IBM i 6.1 can be the server and virtualize the storage for IBM i 7.1, IBM AIX®, and Linux client partitions, or IBM i 7.1 can be the server and virtualize the storage for IBM i 6.1, AIX, and Linux partitions. For enhanced high availability options, IBM i 6.1.1 or 7.1 can be hosted by dual PowerVM Virtual IO Servers, which can be most effective when you use external SAN storage.

IBM i 6.1.1 and 7.1 client partitions hosted by VIOS on a POWER6 or POWER7 platform support N_Port_ID Virtualization (NPIV) for SAN storage and tape.

IBM i 7.1 can take advantage of the larger selection of tape media libraries that are supported by PowerVM VIOS through NPIV, which includes:

- ▶ 3573 (TS3100/TS3200) with LTO3 and LTO4 tape drives
- ▶ 3576 (TS3310) with LTO3 and LTO4 tape drives
- ▶ 3577 (TS3400) with TS1120 and TS1130 tape drives
- ▶ 3584 (TS3500) with LTO3, LTO4, TS1120, TS1130, and 3592-J1A tape drives

IBM i 7.1 or 6.1.1 client partitions can also participate in an IBM Active Memory™ Sharing environment when fully hosted by VIOS or fully hosted by IBM i.

The Virtual Partition Manager (VPM) now supports the creation and management of IBM i partitions. The VPM function is available on IBM POWER6 and IBM POWER7 Express Servers™ that do not have an external management console. With this enhancement to IBM i 7.1, you can create up to four IBM i partitions with VPM.

PowerVM now includes support for suspending and resuming an IBM i 7.1 partition. Using suspend / resume, you can perform long-term suspension of partitions, freeing server resources that were in use by that partition, and later resume operation of that partition and its applications on the same server.

1.3.6 IBM i storage management

In Version 5.4, storage management in IBM i was enhanced to provide support for solid-state drive (SSD) technology on POWER6 and beyond. IBM i 7.1 takes this support even further by using SSD technology to collect I/O performance data and move data that is accessed most frequently to SSD storage. This support improves application performance. Several enhancements are also available to help improve efficiency of SSD storage management with the DB2 media preference.

IBM i 6.1.1 and IBM i 7.1 provide new capabilities for tape and external disk storage performance instrumentation. A significant enhancement in this area is that IBM System Storage DS8000 and DS6000 performance data can now be natively collected on IBM i and analyzed using the IBM iDoctor for IBM i tool.

IBM i 7.1 increases system up time by allowing the complete removal of internal disk units or external storage LUNs without an IPL or outage, and is supported on any IBM POWER® hardware that runs IBM i 7.1. Before 7.1, an IPL (or a varyoff of the independent ASP) was required to complete the process of removing storage disk units or LUNs. For more information, see 9.1.1, “Concurrent removal of disk units” on page 368.

IBM i 7.1 and 6.1.1 add support for IBM System Storage DS5100 and DS5300 through native Fibre attachment when running on POWER6 or POWER7 processor-based servers.

For enhanced availability, IBM i 7.1 or 6.1.1 partitions can be configured in multipath configurations where one IBM i partition uses redundant VIOS partitions to connect to the same IBM System Storage device.

For IBM i internal disk solutions, enhanced redundancy is provided by the new dual SAS adapter support and hot-spare for mirroring support, which is available with IBM i 6.1.1 or later.

Support for the following features is included with IBM i 7.1:

- ▶ CPC node level mirroring
- ▶ EXP24S SFF Gen2-bay drawer (#5887)
- ▶ Higher capacity 10K RPM SFF SAS disk drives
- ▶ DVD/Tape SAS External Storage Unit for Power 795 CPC Rack
- ▶ Thin provisioning for DS8700 and DS8800 storage servers and for VIOS shared storage pools
- ▶ 177 GB SFF SSD with eMLC
- ▶ IBM Disk Sanitizer PRPQ extended to include SSD devices

1.3.7 Application development on IBM i enhancements

The application development strategy for IBM i focuses on high performance compilers that use new technologies, provide a common developer desktop across operating systems and languages, and embrace open source solutions to create a complete application development portfolio.

IBM i 7.1 makes several contributions to this strategy:

- ▶ An updated PHP portfolio for several products, such as Zend Server Community Edition for IBM i (included in IBM 7.1), Zend Studio for IBM i (available to IBM i customers at no additional cost), and Zend Server for IBM i.
- ▶ The XMLSERVICE open source tool, developed with RPG, provides easy access to IBM i data and resources in any development environment that supports HTML and XML, such as Java, Groovy, PHP, Python, Ruby, and RPG.
- ▶ Enhanced IBM Toolbox for Java JDBC with support for several features of the latest JDBC 4.1 definitions with DB2 for i and the new JTLite Java Classes Library. The toolbox supports Java application development for mobile devices.
- ▶ IBM Application Runtime Expert for i (ARE), which can help you ensure consistent performance and deployment for any workload running on IBM i.
- ▶ Native API support of the compression and extraction of files.
- ▶ IBM Lotus® Support for IBM i 7.1, which includes a broad set of products, such as IBM Lotus Domino® V8.5.3, IBM Lotus Sametime® V8.5.2, IBM Lotus Quickr® V8.5.1, IBM Lotus Enterprise Integrator® V8.5.3, IBM Forms Server V3.5.1 and V4.0, and IBM Lotus Workflow V7.0.

The following sections highlight these enhancements in more detail.

High-level programming languages

Enhancements in high-level programming languages include:

- ▶ For C/C++, the following additions and modifications were made in IBM i 7.1:
 - ILE C/C++ predefined macros
 - `do_not_instantiate` and `namemanglingrule` pragma
 - Control Language Command options **DBGENCKEY** and **DECFLTFND**
- ▶ There were also enhancements in compiler options, ILE COBOL, ILE RPG, and Control Language CL.

PHP on IBM i

Enhancements were done for PHP on IBM i. There are several changes in Zend PHP products for IBM i.

Zend Solutions for IBM i now include:

- ▶ Zend Server Community Edition for IBM i

This feature is an enhanced version of open source PHP. It is a lightweight version of Zend Server and replaces Zend Core. It is preinstalled on IBM i 7.1.
- ▶ Zend Server for IBM i

This feature is a robust PHP production environment that helps ensure that applications written in PHP run smoothly at all time. It replaced Zend platform. It offers all the features that are provided in Zend Server FileNet® Content Engine for IBM i plus more features.

- ▶ Zend Studio for IBM i

This product includes all the development components necessary for the full PHP application lifecycle and simplifies complex projects. Several new enhancements were added to Zend Studio.

Lotus Support for IBM i 7.1

IBM i offers a great platform for consolidating IBM Lotus collaboration solutions. IBM i 7.1 supports IBM Lotus Domino 8.5.1, IBM Lotus Sametime 8.5.1, IBM Lotus Quickr 8.5, IBM Lotus Enterprise Integrator 8.5.2, Lotus Forms Server 3.5.1 FP2, IBM Forms Server 4.0, and Lotus Workflow 7.0.

Native API support of compression and extraction of files

There is new enhancement for native API support of compression and extraction of files. This support includes the following native APIs and a service program to create archive files:

- ▶ QzipZip API
- ▶ QzipUnzip API
- ▶ QZIPUTIL Service program

IBM Toolbox for Java JDBC enhancements

The Java JDBC interface now supports several features of the latest JDBC 4.1 definitions with DB2 i. Several other enhancements are added, such as:

- ▶ New methods and classes for XML Data type support
- ▶ Database metadata updates
- ▶ Currently committed support
- ▶ Array type support
- ▶ Long schema name support

IBM Application Runtime Expert for i

IBM Application Runtime Expert for i (ARE) is a new product that has the potential to revolutionize how you do application service and support. ARE can help you ensure consistent performance and deployment for any workload running on your system.

ARE provides a GUI that you can use to collect and verify a customized set of information, systems settings, and attributes about:

- ▶ Applications
- ▶ IBM i systems
- ▶ Runtime environments

ARE collects this information and places it into a template. This template can then be used in verifying the application, and its environment, on the same system where the template was built, or any other IBM i system.

1.3.8 IBM Rational support for IBM i 7.1

The following sections describe Rational products for IBM i:

- ▶ IBM Rational Developer for Power Systems Software™
- ▶ IBM Rational Team Concert™ for Power
- ▶ IBM Rational Developer for i for SOA Construction
- ▶ IBM Rational Development Studio for i
- ▶ IBM Rational Open Access: RPG Edition

Rational Developer for Power Systems Software

Rational Developer for Power Systems Software replaces IBM Rational Developer for i (RDi), provides Eclipse-based tools for RPG/COBOL development, and includes a new print designer and support for language enhancements.

Rational Developer for Power Systems Software adds a feature that combines IBM i development tools with IBM Rational Application Developer Standard Edition for IBM WebSphere Software V8. The new Power Systems tools allow developers to have all their development tools integrated into one work environment. Power Systems tools also reduce the operation costs by deploying a single development workbench image to all developers.

Rational Team Concert for Power

Rational Team Concert for Power provides a collaborative software development environment with integrated source control and build management.

Rational Team Concert integrates with Rational Developer for Power Systems Software to provide IBM i developers an integrated collaborative application development environment. This integrated solution provides the value of the team collaboration capabilities of Rational Team Concert with the individual development environment of Rational Developer for Power Systems Software.

Rational Developer for i for SOA Construction

Rational Developer for i for SOA Construction (RDi SOA) is a software bundle that combines IBM Rational Developer for Power Systems Software and IBM Rational Business Developer.

It provides a powerful, flexible, and extensible workbench environment for IBM i development with support for RPG, COBOL, CL, DDS, SQL, C++, Java, and EGL.

Rational Development Studio for i

Rational Development Studio for i V7.1 includes enhancements for the ILE RPG and ILE COBOL compilers.

IBM Rational Open Access: RPG Edition

IBM Rational Open Access: RPG Edition allows organizations to reuse existing skills when creating modern RPG-based applications by providing a way for RPG developers to use the simple and well-understood RPG I/O model to access resources and devices that are not natively supported by RPG.

1.3.9 Systems management enhancements

IBM i 7.1 continues the transition by deploying new features and functions that are focused on System Director Navigator for i and the cross platform System Director product suite. New enhancements in Director Navigator for i include the ability to set a target server that can run IBM i 5.4, 6.1, or 7.1. This ability enables the management server to run in one place with a single browser and to be used to manage multiple OS environments. Support for tape and journal management was added, as well as new GUI calendaring features for the Advanced Job Scheduler feature.

Additional new features in System Director Navigator for IBM i are as follows:

- ▶ 5250 emulation support
- ▶ New charts and tables for the Investigate Data task in Performance Data Investigator
- ▶ Integration between the work management tasks and the performance tasks
- ▶ iSCSI integration with IBM BladeCenter® and IBM System x® installation wizard

1.3.10 Printing

The **CPYSPFLF** command is enhanced to copy to PDF or TIFF stream files.

1.3.11 IBM i iSCSI integration with BladeCenter and System x

IBM i integration with BladeCenter and System x through iSCSI technology is enhanced with software target support. iSCSI Software Target eliminates the need for specialized hardware and supports higher speed connections, up to 10 Gb. It uses standard Ethernet adapters and the network to connect IBM i on Power Systems to BladeCenter and System x. In addition, since August 2012, the requirement for a switch between an IBM i iSCSI software target and a BladeCenter blade or System x iSCSI initiator is eliminated. Additional operating system support was added, which includes Windows Server 2008 R2, VMware VSphere 4.0, and VMware ESXi Embedded 4.0, and a new installation wizard.

The new installation wizard simplifies the creation, cloning, and deletion of servers in the IBM i Integrated Server environment.

1.3.12 Security enhancements in IBM i 7.1

ASP encryption can now be turned on and off dynamically, and the data encryption key can be changed for an existing user ASP. Column level encryption can be accomplished by using a new database feature called *field procedures*. These procedures are user-written exit programs that are run every time that a column is changed or new values are inserted.

A new enhancement for managing user profiles adds the ability to set a date or a time interval for a profile to expire.

This version includes a significant leap in security for IBM i users, newer functions to increase data confidentiality, such as ASP encryption, which can now be turned on and off dynamically, and the Database Field Procedures, which can be used to accomplish transparent column level encryption.

Real-time database monitoring is an important topic, because a database is the main target of hackers. IBM i adds a significant enhancements to the database performance monitoring tool, audit journaling, and integration with the IBM InfoSphere® Guardium® DAM (Database Access Monitor) solution.

1.3.13 Web application serving

IBM i web application serving technologies are enhanced with the latest Apache 2.2 release, which is Payment Card Industry (PCI) compliant. To remain in sync with IBM strategic direction, IBM i 7.1 no longer ships with the Classic JVM; it was replaced by IBM adding support for J9 VM.

The integrated Web Application Server, which ships with IBM i, supports Java 5 and 6 applications for 32-bit and 64-bit JVMs. The Integrated Web Services Server now supports programs in iASP, supports Static WSDL, and has shown up to 2x performance improvements.

The IBM Web Enablement for IBM i product now includes WebSphere Application Server - Express Editions V6.1, V7, and V8 when ordered for IBM i 7.1.

Integrated Web Services for i enables integrated IBM Language Environment® (ILE) applications to work in the web services and service-oriented architecture (SOA) arena with little effort, knowledge, and resources. The convergence of web services and IBM i technologies can help enterprises liberate these core business assets by making it easier to enrich, modernize, extend, and reuse them well beyond their original scope of design.

1.3.14 IBM i Access Family 7.1 enhancements

The IBM i Access Family products have enhancements not only in the IBM i Access for Windows product but also in the IBM i Access for Web. It includes the following enhancements:

- ▶ IBM i Access for Windows 7.1 offers enhancements in:
 - NET, ODBC, and OLE DB providers
 - Data transfer
 - Updated PC5250 display and printer emulation that is based on IBM Personal Communications 6.0
 - Enhancements to install, including install-time support for secondary languages
- ▶ IBM i Access for Web 7.1 includes an additional option for viewing spooled files as PDF documents.

IBM i Access Client Solutions provide new options for easier deployments to clients PCs with no product installation at all. This Java client runs in a broad range of operating systems, such as Windows 7 and later, Mac OS, and Linux, including Red Hat, SUSE, and Ubuntu.

1.3.15 Networking enhancements

Networking enhancements include the new IBM i DHCP server, which is based on the Internet Systems Consortium (ISC) DHCP server, which enables support for IPv6 and DHCP failover. IPv6 is also supported for the DHCP client. IPv6 and Remote Authentication Dial In User Service (RADIUS). Support was added for point-to-point protocol (PPP). Additional enhancements include support for the Internet Key Exchange Version 2 (IKEv2) tunneling protocol for the IBM i VPN support. The telnet client on IBM i is now SSL-enabled for accessing remote hosts that require SSL/TLS connections.

IBM i V7R1 has new support for Ethernet layer-2 bridging and Ethernet link aggregation. The first one allows the sharing of physical Ethernet connections across partitions. The second one allows you to link up to eight Ethernet links together in a single-line description for improved throughput and reliability.

1.3.16 Upgrade for IBM i enhancements

IBM i 6.1 introduced the ability to access virtual media images on a network-attached remote NFS server. Initially, this ability was limited to PTF images or program installation packages. This NFS capability was enhanced to support an OS upgrade using NFS to access installation media on a remote NFS server. IBM i 6.1 or 6.1.1 running on a POWER6 or POWER7 processor-based server can now upgrade to IBM i 7.1 over a network connection.

1.3.17 IBM i network installation

IBM i can now support the installation of new partitions from images in the network. An IBM POWER7 processor-based server can be installed without using physical media, and instead uses IBM i 7.1 optical images that are stored on a network file server.

1.3.18 IBM i Processor Entitlement Transfer

IBM i Processor Entitlement Transfer was announced in April 2010. This function enables the transfer of IBM i processor entitlements from one machine to another in the same user's enterprise under certain qualifying conditions. These selected transfers help provide investment protection for the IBM i operating system entitlements and include eligible IBM i and IBM i Application Server entitlements. IBM i user licenses are not transferable. This new processor entitlement transfer capability is implemented worldwide by the eConfigurator.

Program/Product numbers include:

- ▶ 5722-SS1 IBM i 5.4
- ▶ 5761-SS1 IBM i 6.1
- ▶ 5770-SS1 IBM i 7.1

For more information, see the program summary at the following website:

<http://www-947.ibm.com/systems/support/planning/notices71.html>

1.3.19 Backup, Recovery, and Media Services and basic backup and recovery

Numerous enhancements are available in IBM i 7.1 in both Backup, Recovery, and Media Services (BRMS) and the basic backup and restore functions (such as fast restore using tape positioning). Save operations now track the physical media position of each saved object. This new function enables Object Restore to move to the location of a stored object on tape, saving time by minimizing tape searching for the object.

Functions previously available and new IBM i 7.1 functions accessible through IBM Systems Director Navigator for i are now also available through IBM Systems Director. This feature provides additional management functions for an improved user interface. A full suite of BRMS functions is added, such as a complete set of scheduling functions, including graphical calendar views. The following list details a few of the many enhancements in BRMS:

- ▶ Support for Domino 8.5 Domino Attachment and Object Service (DAOS)
- ▶ Link list enhancements
- ▶ Recovery report enhancements
- ▶ Output file support for BRMS restores

Save While Active Support for Integrated VMware ESX Servers is included with IBM i 7.1.



Technology Refresh

This chapter summarizes the new Technology Refresh functionality within IBM i. 7.1.

This chapter describes the following topics:

- ▶ Overview
- ▶ What is a Technology Refresh
- ▶ What is a Resave
- ▶ What is in a Technology Refresh

2.1 Overview

Following the General Availability (GA) of IBM i 7.1, IBM i point / modification releases were replaced by a new release delivery mechanism that is called a *Technology Refresh*.

The term Technology Refresh refers to the set of PTFs required to support new hardware functionality. The term *Technology Update* refers to multiple PTFs or PTF Groups that provide additional functionality in IBM i and related products. The term Technology Refresh is widely used to refer to both terms together.

For more information about Technology Update, go to the IBM developerWorks site at <http://www.ibm.com/developerworks/ibmi/>. Look in the left navigation pane for Technology Updates.

2.2 What a Technology Refresh is

A Technology Refresh, or Technology Update, is a collection of operating system software that is developed together, packaged together, tested together, and delivered together as a PTF Group for a base release.

There used to be version and release levels, such as i5/OS V5R3 or V5R4. Now, with IBM i V7R1, there is Technology Refresh 1 (TR 1) or 2 (TR 2). When a release was upgraded, it upgraded from, for example, V5R3 to V5R4. Now, an upgrade is done by installing a new Technology Refresh. You can think of a Technology Refresh as a Technology Refresh PTF Group plus the current Cumulative Package (CUM). In that respect, it is similar to a point / modification release plus a CUM.

The benefits of a Technology Refresh versus a point / modification release are:

- ▶ It avoids a business disruption because of major release upgrades.
- ▶ It gives you faster access to new technology.
- ▶ It is simpler to install.

The Technology Refresh PTF Group contains a Technology Refresh PTF plus related PTFs, including the most recent fixes before the GA of the Technical Refresh. The Technology Refresh PTF is Licensed Internal Code (57xx-999) and is part of the code stream of the base release IBM i 7.1, rather than being a separate code stream, as is the case with a point release.

A Resave associated with the Technology Refresh is also available, as shown in Figure 2-1.




Figure 2-1 Technology Refresh overview

The following items are described later in this chapter:

- ▶ Technology Refresh PTF (TR PTF)
- ▶ Technology Refresh PTF Group (TR PTF Group)
- ▶ Technology Refresh Requisite (TR Reqs)
- ▶ General Availability Resave (GA Resave)
- ▶ Cumulative Package (CUM)
- ▶ Managed Add Function PTFs
- ▶ Other PTFs
- ▶ Other functions that were made GA at the same time

2.3 What a Resave is

A Resave, as distinguished from a Technology Refresh, is an updated release of the Machine Code (57xx-999). When IBM introduces new hardware, typically a Resave is needed to support it. New hardware might include new server models, new I/O adapters, new disk drives, and so on. The Resave code is found on the I_BASE_01 DVD that is shipped with a new release of software. The label on this DVD shows the Resave level, such as RS-710-E. This chapter explains how to order and install Resaves.

2.4 What is in a Technology Refresh

A Technology Refresh might provide the following items:

- ▶ Support for new hardware and firmware (for example, new I/O devices or newly announced models)
- ▶ Support for new virtualization functions
- ▶ Performance improvements
- ▶ New function enablement

The Technology Refresh PTF Group, Resave, or PTF may be downloaded, requested on media, and installed using standard procedures and tools. Either install the Technology Refresh PTF Group plus the most recent Cumulative PTF package, or the Technology Refresh Resave plus the most recent Cumulative PTF package plus the Technology Refresh PTF Group. When a Technology Refresh is installed, the system release level does not change.

2.4.1 When to install a Technology Refresh

A Technology Refresh should be installed in the following cases:

- ▶ To provide the support that is needed for new hardware or virtualization functions
- ▶ To take advantage of performance improvements or new function enablement
- ▶ When required by subsequent PTFs (for example, fixes to code in the Technology Refresh PTF)

It is a preferred practice to keep a system up to date with the latest Technology Refresh installed to avoid unnecessary downtime. To determine the Technology Refresh level that is needed for a specific hardware feature code, see the IBM Prerequisite tool, found at:

https://www-912.ibm.com/e_dir/eServerPrereq.nsf

The PTF Cover Letter or Preventive Service Planning (PSP) information identifies whether a Technology Refresh is a requisite for that PTF. If a PTF has a Technology Refresh requisite, the Technology Refresh PTF must be permanently applied before the PTF that requires it can be loaded.

For more information, see Preventive Service Planning website at:

http://www-912.ibm.com/s_dir/sline003.nsf/554c38c4848b77f2862567bd0046e003/85fc5e8ebc0523048625779900742abc?OpenDocument

2.4.2 How it all fits together

Table 2-1 shows the mapping between Technology Refresh PTF group level, Resave level, Marker PTF, and Technology PTF number.

Table 2-1 IBM i 7.1 Technology Refresh history (with 7.1.0 Machine Code)

Technology Refresh release date	Description	Technology Refresh PTF Group Level	Corresponding 5770-999 Resave Level and Marker PTF	5770-999 Technology Refresh PTF
10/12/2012	IBM i 7.1 Technology Refresh 5 (TR 5)	SF99707 Level 5	RS-710-F RE12249	MF99005
05/18/2012	IBM i 7.1 Technology Refresh 4 (TR 4)	SF99707 Level 4	RS-710-E RE12066	MF99004
10/14/2011	IBM i 7.1 Technology Refresh 3 (TR3)	SF99707 Level 3	RS-710-D RE11221	MF99003
05/13/2011	IBM i 7.1 Technology Refresh 2 (TR2)	SF99707 Level 2	RS-710-C RE11067	MF99002
09/10/2010	IBM i 7.1 Technology Refresh 1 (TR1)	SF99707 Level 1	RS 710-B RE10187	MF99001

The following list describes the columns in Table 2-1 in more detail:

- Technology Refresh release date

This column shows the date when the Technology Refresh was made available.

- Description

This column shows the description of the Technology Refresh.

- Technology Refresh PTF Group Level

This column identifies the age of the Technology Refresh PTF Group. The Technology Refresh PTF Group is a PTF Group that contains a Technology Refresh PTF for a particular GA, plus related PTFs, including the most recent fixes before GA of the Technical Refresh.

In general, the Technology Refresh PTF Group should be installed rather than the individual Technology Refresh PTF. The exception to the rule is if the Technology Refresh PTF is required by another PTF or PTF group (for example, a HIPER PTF Group) before the GA of the Technology Refresh.

A higher PTF group level number indicates a more recent group. Only the most recent version of a PTF group is available to be downloaded or ordered.

- Corresponding 5770-999 Resave Level

This column identifies the 5770-999 Resave level that corresponds to this Technology Refresh. The Resave contains the technology refresh plus related code, including fixes. The name of 5770-999 Resave contains the VRM plus a letter to indicate the Resave level (for example, RS710-A and RS710-B).

► Marker PTF

The marker PTF is a PTF that can appear on a system only if it was installed with an IBM supplied Resave. These markers are cumulative, meaning systems may have more than one of the Marker PTFs available. The highest Marker PTF number that is found can be used to determine what Resave level of the product is installed.

► 5770-999 Technology Refresh PTF

This column identifies the age of the Technology Refresh. A Technology Refresh is a PTF for a specific release / VRM. IBM identifies the Technology Refresh PTF level for 5770-999 by a set of reserved PTF numbers, MF99xxx, with high digits indicating more recent Technology Refreshes. Later Technology Refresh PTFs for a particular release are supersets of previous Technology Refresh PTFs for that release. Systems may have more than one Technology Refresh PTF installed. Later versions of the Technology Refresh PTF for a release supersede previous versions. The active version of a Technology Refresh on a system is the PTF with the highest MF99xxx PTF ID in applied status.

2.4.3 Determining the Technology Refresh PTF Group level installed on a system

To determine the Technology Refresh PTF Group level that is installed on a system, you can use Work with PTF Groups (**WRKPTFGRP**). For example, you could run **WRKPTFGRP** to find the PTF group named SF99707, as shown in Figure 2-2. There might be multiple different levels of the group that are installed on the system. The latest level (the one with the highest level number) with the status of Installed is the level of the fix group that is active.

Work with PTF Groups				System:	TOMVER
Type options, press Enter.					
1=Order 4=Delete 5=Display 6=Print 8=Display special handling PTFs					
9=Display related PTF groups					
Opt	PTF Group	Level	Status		
	SF99710	10229	Installed		
	SF99710	11116	Installed		
	SF99709	32	Installed		
	SF99709	40	Installed		
	SF99708	3	Installed		
	SF99708	8	Installed		
	SF99707	1	Installed		
	SF99707	2	Installed		
	SF99701	9	Installed		
	SF99701	10	Installed		
	SF99637	1	Installed		
	SF99637	3	Installed		
	SF99627	5	Installed		

Figure 2-2 Technology Refresh PTF Group that is installed

2.4.4 Determining the Resave level installed on a system

To determine the Resave level of 5770-999, use Display PTF Status (**DSPPTF**) on product 5770-999, as shown in Figure 2-3. Marker PTFs for this product are in the format of REnnnnn. The highest number Marker PTF on your system, matched with Table 2-1 on page 19, indicates the Resave level for this product.

Display PTF Status			System: TOMVER
Product ID	 : 5770999		
IPL source	 : ##MACH#B		
Release of base option	 : V7R1M0 L00		
Type options, press Enter.			
5=Display PTF details 6=Print cover letter 8=Display cover letter			
Opt	PTF ID	Status	IPL Action
	TL11116	Temporarily applied	None
	TL10229	Superseded	None
	TL10096	Superseded	None
	TL10033	Superseded	None
	RE10187	Permanently applied	None
	RE10084	Permanently applied	None
	RE10026	Permanently applied	None
	QLL2924	Permanently applied	None
			More...
F3=Exit F11=Display alternate view F17=Position to F12=Cancel			

Figure 2-3 IBM i Resave level installed

2.4.5 Determining the Technology Refresh PTF level installed on a system

To determine the Technology Refresh PTF level of 5770-999, use Display PTF Status (**DSPPTF**) for product 5770-999, as shown in Figure 2-4. Technology Refresh PTFs for this product are in the format of MF99nnn. The highest number Technology Refresh PTF on your system, matched with Table 2-1 on page 19, indicates the Technology Refresh level for this product.

```

                                Display PTF Status
                                System:  TOMVER

Product ID . . . . . : 5770999
IPL source . . . . . : ##MACH#B
Release of base option . . . . . : V7R1M0 L00

Type options, press Enter.
  5=Display PTF details  6=Print cover letter  8=Display cover letter

   PTF
Opt ID      Status
RE10187    Permanently applied
RE10084    Permanently applied
RE10026    Permanently applied
QLL2924    Permanently applied
MF99003    Permanently applied
MF99002    Permanently applied
MF99001    Permanently applied
MF54045    Temporarily applied
                                IPL
                                Action
                                None
                                None
                                None
                                None
                                None
                                None
                                None
                                None
                                None
                                More...

F3=Exit  F11=Display alternate view  F17=Position to  F12=Cancel
```

Figure 2-4 Technology Refresh PTF level installed

2.4.6 Planning for an IBM i Technology Refresh update

A Technology Refresh should be used in the following situations:

- ▶ Support for new hardware
- ▶ Support for new virtualization functions, performance improvements, or new function enablement
- ▶ General LIC or IBM i updates for defect fixes

To determine the Technology Refresh level that is needed for hardware, see the IBM Prerequisite tool at:

https://www-912.ibm.com/e_dir/eServerPrereq.nsf

It is important to keep systems up to date with the latest Technology Refresh PTF available. Subsequent PTFs may depend on it, and those PTFs cannot be loaded until the requisite Technology Refresh PTF is permanently applied, which requires an IPL. Therefore, it is a preferred practice to keep systems current with the latest Technology Refresh PTFs, whether through the Technology PTF Group, a Resave, or the Technology Refresh PTF itself. Subsequent Technology Refreshes for a release are supersets of previous ones, so you need apply only the latest Technology Refresh to keep the system current.

2.4.7 Installing an IBM i Technology Refresh

The Technology Refresh PTF itself enables support for the new function of the Technology Refresh; it alone is not sufficient to obtain the complete support of the new functions, and might not contain the latest fixes available. Therefore, when you update a system to a new Technology Refresh level, use one of the following methods:

- ▶ Install the Technology Refresh PTF Group plus the latest Cumulative PTF package.
- ▶ Install the 5770-999 Resave that corresponds with the Technology Refresh level, plus the Technology Refresh PTF Group, plus the latest cumulative PTF package.

2.4.8 Ordering an IBM i Technology Refresh PTF Group

A Technology Refresh PTF Group is a PTF Group that can be ordered like any other PTF Group, using voice support, Fix Central, or the Send PTF Order (**SNDPTFORD**) command. Instructions for ordering PTFs can be found in the “Maintaining and managing IBM i and related software” topic in the IBM i 7.1 Information Center at:

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp?topic=%2Frzam8%2Frzam8fixobtain1.htm>

Before you order a Technology Refresh PTF Group, check and verify that the level of the PTF Group you need is not already on your system.

2.4.9 Installing an IBM i Technology Refresh PTF Group

A Technology Refresh PTF Group is a set of PTFs that is installed like any other IBM i PTF Group, using the Install Program Temporary Fix (**INSPTF**) command or by using Option 8 from the GO PTF menu. Instructions for installing PTF Groups can be found in the “Maintaining and managing IBM i and related software topic” in the IBM i 7.1 Information Center.

The Technology Refresh PTF must be permanently applied before subsequent PTFs can be loaded, which requires an IPL.

For more information about how to avoid or reduce the impact of a *double IPL* during PTF installation, see 2.4.14, “Preventing or reducing the impact of a double IPL” on page 24.

2.4.10 Ordering and installing an IBM i Resave

Resaves are ordered by requesting a software upgrade from your IBM marketing representative or your IBM Business Partner. They run the IBM eConfiguration Tool, which generates an order for a software upgrade. The latest Resaves are included with each software order.

You install a Resave by following the instructions in the IBM Software Installation Manual. For instructions and other information that is related to ordering and installing an IBM i Resave, go to:

<http://www-947.ibm.com/systems/support/i/planning/resave/index.html>

2.4.11 Ordering an IBM i Technology Refresh PTF

A Technology Refresh is a PTF that can be ordered like any other PTF, using voice support, Fix Central, or the Send PTF Order (**SENDPTFORD**) command. Instructions for ordering PTFs can be found in the “Maintaining and managing IBM i and related software” topic in the IBM i 7.1 Information Center.

Before you order a Technology Refresh PTF, check and verify that the PTF is not already on your system as a requisite of another PTF.

2.4.12 Installing an IBM i Technology Refresh PTF

A Technology Refresh PTF is a PTF that can be installed just like any other PTF. Instructions for installing PTFs can be found in the “Maintaining and managing IBM i and related software topic” in the IBM i 7.1 Information Center.

A Technology Refresh PTF must be permanently applied before subsequent PTFs that require it can be applied, so it is a preferred practice to apply the Technology Refresh PTF permanently when it is first applied.

For more information about how to avoid or reduce the impact of a double IPL during PTF installation, see 2.4.14, “Preventing or reducing the impact of a double IPL” on page 24.

2.4.13 How a Technology Refresh PTF or PTF Group affects other PTFs

If a later PTF changes a part or module that is contained in the Technology Refresh PTF, the Technology Refresh PTF becomes a special prerequisite for that PTF, called a *TRREQ. The Technology Refresh PTF must be permanently applied on the system before the PTF that requires it can be loaded. It is a preferred practice to keep a system up to date on Technology Refresh PTFs to avoid the extra time it would take to apply the Technology Refresh PTF. PTFs that do not involve parts or modules that are contained in a Technology Refresh PTF do not require the Technology Refresh PTF to be applied before they can be loaded.

2.4.14 Preventing or reducing the impact of a double IPL

Customers can avoid a double IPL by making sure that the Technology Refresh PTF is permanently applied. Then, if a PTF requires it, the Technology Refresh PTF is already on the system. This setup does not reduce the total number of IPLs required, but it allows the scheduling of IPLs when they are most convenient for the operation of the system.

Ordering and installing the Technology Refresh Resave also ensures that the Technology Refresh PTF is permanently applied and that the double IPL is avoided.

The new function PTF SI43585 is available to automate, but not eliminate, any extra IPLs required during PTF installation. When you are installing PTFs, there are two conditions where you must perform an IPL to apply some of the PTFs, which requires a restart of the PTF installation after the first IPL, and then perform another IPL to apply the delayed PTFs:

- ▶ When installing a cumulative PTF package that contains special handling pre-apply PTFs
- ▶ When installing a technology refresh PTF at the same time as a technology refresh requisite PTF

If an extra IPL is required, the PTF installation parameters are saved and used during the next IPL. Instead of seeing the “Confirm IPL for Technology Refresh or Special Handling PTFs” panel, you see a new message CPF362E: “IPL required to complete PTF install processing”. Actually, if you select Automatic IPL=Y on the “Install Options for PTFs” panel, you do not see any other messages or panels, as a power down then occurs. On the next normal IPL, your second “GO PTF” completes during the “PTF Processing” IPL step in the SCPF job, and then a second IPL of the partition is done automatically. So when the system performs the second IPL to sign on, your PTFs are all activated and ready to go.

If an IPL is required for a technology refresh PTF, SI43585 supports installing only from a virtual optical device or *SERVICE (PTFs downloaded electronically to save files). If you are installing from a physical optical device, you must perform the extra IPL and second GO PTF manually. If you received your PTFs on physical DVDs, create an image catalog from the DVDs and use the new support.

2.4.15 Why IBM moved to the Technology Refresh process

The Technology Refresh process allows IBM to deliver new functions and support quicker and more frequently than was possible with point / modification releases.

Moving up to a Technology Refresh is simpler and cheaper than qualifying a point release, so you can take advantage of new functions and support sooner than in the past.

Technology Refreshes provide a simple way to keep systems up to date with PTFs, avoiding rediscovery of problems that are already fixed, unnecessary downtime, and calls to IBM Support.

Technology Refreshes also provide an easy way to keep a group of systems in sync.

Backing out a point / modification release requires a scratch installation of the system; with a Technology Refresh, it is possible to return to an earlier level of IBM i by simply slip installing the Licensed Internal Code only.

2.4.16 New releases of IBM i

There will still be new releases of IBM i. With a Technology Refresh; only the changed parts that are required for the new hardware / firmware support or function are recompiled and included.

Some large and complex projects are better suited for an actual release, where the entire body of code in IBM i is rebuilt. Developers are working on the next release of IBM i and architects are looking at possible content for the next release.

Now, Technology Refreshes include only Licensed Internal Code (LIC). Enhancements for other levels of IBM i require a release. This situation is similar to point / modification releases, which contained only LIC.

2.4.17 Is an IBM i Technology Refresh the same as an AIX Technology Level

These two are not the same. An AIX Technology Level is a separate code stream with its own set of fixes, similar to an IBM i point / modification release. When you install an AIX Technology Level, the release level changes, and a different library of fixes, for that particular release level must be used.

A Technology Refresh is an update of an existing release, through a PTF Group, containing PTFs in that release's code stream. When an IBM i Technology Refresh is installed, the release level of the system does not change, and the system continues to use PTFs for that release.

2.4.18 Related publications

For more information related to Technology Refreshes, see the following resources:

- ▶ IBM i Support - IBM i Technology Refresh - IBM i 7.1 Information:
<http://www-947.ibm.com/systems/support/i/planning/techrefresh/i71.html>
- ▶ Technology Refreshes and additional Enhancements to IBM i 7.1 on IBM developerWorks:
<https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/IBM%20i%20Technology%20Updates/page/Technology%20Refreshes%20and%20additional%20Enhancements%20to%20IBM%20i%207.1>
- ▶ Using software fixes:
<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzam8/rzam8fix1.htm>
- ▶ Planning your fix management strategy:
<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzam8/rzam8fixstrategy.htm>
- ▶ PTFs - Maintenance Strategy:
http://www-912.ibm.com/s_dir/slkbase.NSF/DocNumber/14854405
- ▶ Guide to fixes:
<http://www-947.ibm.com/systems/support/i/fixes/guide/>



Security enhancements

IT security management is under much pressure today. Security management must deal with an explosion in data growth, a need to access data anytime and from anywhere, more sophisticated hackers, cloud initiatives, and increasing demands from law enforcement and other entities to meet compliance requirements.

IBM i is one of the most secure operating systems in the industry. From the beginning of its development, security has been an important part of its design.

IBM i 7.1 provides an important set of enhancements that, with leading-edge security solutions provided by IBM and Business Partners, reduce risk but also simplify security management and facilitate compliance requirements.

This chapter describes the following security enhancements for IBM i 7.1:

- ▶ User Profile enhancements
- ▶ Object Audit enhancements
- ▶ A data encryption enhancement, including disk level encryption and database field procedures support
- ▶ Security enhancements for DB2 for i
- ▶ New capabilities for Monitor Database Activity, including IBM InfoSphere Guardium V9.0 with support for DB2 for i
- ▶ Security enhancements for printing

For more information about IBM and Business Partners Security solutions for IBM i, go to:

http://www-03.ibm.com/systems/power/software/i/security/partner_showcase.html

3.1 User Profile enhancements

IBM i 7.1 includes additional attributes in the User Profile object that set an expiration date or a days interval to disable a user profile.

The User Profile controls can be used with the Display Expiration Schedule (**DSPEXPSCD**) command. This combination simplifies the task of administering temporary user profiles.

3.1.1 New USREXPDATE and USREXPITV attributes

There are two new parameters for the Create User Profile (**CRTUSRPRF**) and Change User Profile (**CHGUSRPRF**) commands:

► User expiration date (**USREXPDATE**)

This parameter is the date when the user profile expires and is automatically disabled. Possible values are shown in Table 3-1.

Table 3-1 Values for the USREXPDATE

Value	Description
*NONE	No expiration date.
*USREXPITV	Calculated based on the value that is entered in the user expiration interval parameter.
Date	Specifies a date when the user profile expires. The date must be in the job date format.

► User expiration interval (**USREXPITV**)

This parameter specifies the number of days (1 - 366) before the user profile is automatically disabled. Possible values are shown in Table 3-2.

Table 3-2 Value for the USREXPITV

Value	Description
1 - 366	Specifies the number of days before the user profile is disabled.

Important: A value must be specified if the **USREXPDATE** parameter has a value of ***USREXPITV**. If the **USREXPDATE** parameter has a value other than ***USREXPITV**, no value is allowed for this parameter.

Remember: The parameters can be seen only when you use the 5250 user interface.

3.1.2 CL commands to change and display the User Profile expiration schedule

The expiration schedule for a user profile can be set with either the Create User Profile (**CRTUSRPRF**) or Change User Profile (**CHGUSRPRF**) commands or the Change Expiration Schedule Entry (**CHGEXPSCDE**) CL command.

Display Expiration Schedule (DSPEXPSCD) shows a list of user profiles and their expiration date (Figure 3-1). If there are no user profiles that are set to automatically expire, an empty panel is generated.

Deleting a profile: The option to delete a profile on an expiration date is only available through CHGEXPSCDE. Be careful when you use the *DELETE option.

User Profile Expiration Schedule				
User Profile	Expiration Date	Action	Owned Object Option	New Owner
CHUA	12/23/10	*DELETE	*CHGOWN	PREMA
MARIE	12/23/10	*DISABLE		

Bottom

F3=Exit F11=Primary group info F12=Cancel F17=Top F18=Bottom

(C) COPYRIGHT IBM CORP. 1980, 2009.

Figure 3-1 Output from DSPEXPSCD

3.2 Object Audit enhancements

Several changes are included to improve the auditing capabilities of CL commands and the usage of database objects. The audit journal provides more specific details about the execution environment for CL commands. The QIBM_DB_OPEN exit program is enhanced to provide filtering capabilities to reduce processing impact.

3.2.1 CL retrieve exit programs running after command completion

The IBM i operating system supports two exit points for control language (CL) commands: QIBM_QCA_CHG_COMMAND and QIBM_QCA_RTV_COMMAND.

For each regular CL command, and proxy CL commands in the QSYS library, *one* exit program can be registered for the *Change Exit Point* and up to *10* exit programs can be registered for the *Retrieve Exit Point*. The differences between these two are:

- Change Exit Point** The exit program is called by the command analyzer before it passes control to the prompter.
- Retrieve Exit Point** The exit program is called by the command analyzer before or after execution of the command.

Before this enhancement, QIBM_QCA_RTV_COMMAND exit programs were always called just before control was transferred to the command processing program (CPP) associated with the command that is being run.

You can use the enhancement to register an exit program for the QIBM_QCA_RTV_COMMAND exit point and indicate that you want the exit program to be called after control returns from the CPP.

3.2.2 Extended auditing details for running a CL command

The IBM i operating system can be used to track the CL commands that are run by a user. After the command-level auditing is turned on for a user profile, a CD (Command String) audit record is generated for each CL command that is run by the specified user profile.

To set up the command-level auditing, complete the following steps:

1. Auditing must be active before command-level auditing can be done. Run the Change Security Auditing (**CHGSECAUD**) command to change the settings for the system values that control what is being audited on the system.
2. Run the Change User Audit (**CHGUSRAUD**) command and specify the user profile name for the **USRPRF** parameter and ***CMD** for the **AUDLVL** parameter.

These commands generate CD (Command String) audit records for each CL command that is run by the specified user profile. The model file QASYCDJ5 describes the fields in the CD audit records. One of these fields, CDCLP, is redefined to convey more information about how the audited CL command was run.

Before this enhancement, CDCLP had only two possible values, as shown in Table 3-3.

Table 3-3 Original CDCLP field values for a QASYCDJ5 model file

Value	Description
Y	If the command was run from a compiled OPM CL program, a compiled ILE CL module that is part of an ILE program or service program, or an interpreted REXX procedure.
N	In all other cases.

Now, the CDCLP field has the values that are shown in Table 3-4.

Table 3-4 New CDCLP field values for a QASYCDJ5 model file

Value	Description
Y	If the CL command is run from a compiled CL object, which could be an OPM CL program or an ILE CL module that is bound into an ILE program or service program.
R	Indicates that the CL command is being run from an interpreted REXX procedure.

Value	Description
E	Indicates that the command was submitted by passing the command string as a parameter to one of the Command Analyzer APIs: QCMDEXC , QCAPCMD , or QCAEXEC .
B	When the command is not being run from compiled CL or interpreted REXX or through a Command Analyzer API, and is in a batch job. The typical cases for a B value would be if the command is in a batch job stream that is run by using the Start Database Reader (STRDBRDR) or Submit Database Job (SBMDBJOB) command, or is specified for the CMD (Command to run) parameter on a Submit Job (SBMJOB) command.
N	Indicates that the command was run interactively from a command line or by choosing a menu option that runs a CL command.

The new values for the CDCLP field map to the values for the **ALLOW** (Where allowed to run) parameter on the Create Command (**CRTCMD**) command as follows:

- ▶ 'Y' maps to *IPGM, *BPGM, *IMOD, and *BMOD
- ▶ 'R' maps to *IREXX and *BREXX
- ▶ 'E' maps to *EXEC
- ▶ 'B' maps to *BATCH

3.2.3 Database object auditing control with the QIBM_DB_OPEN exit program

Various techniques are used by users and third-party software products to add security controls over and above the security functions built into IBM i.

One primary mechanism that is used to provide this added security is to obtain control through various exit points and to parse the SQL statements. SQL parsing, however, is not only complicated, but continually changing as new SQL functionality is included. In some cases, it is impossible for a SQL parsing solution to handle concepts such as aliases, user-defined functions, and user-defined table functions. For those reasons, this approach is not preferred.

Several releases ago, a single open database exit was implemented in IBM i to relieve user and third-party software from having to parse SQL statements by providing a list of the files that are referenced in the SQL statement. Although this exit solved the parsing problem, the exit was started for every full open. Depending on the amount of resources available and the number of full opens per second in a system, performance could be a concern.

IBM i 7.1 added the capability to have the exit program QIBM_DB_OPEN called only when a full open occurs where at least one of the tables that are referenced by the query has object auditing enabled. This capability was also extended to previous versions V6.1 and V5R4. Example 3-1 shows how to add the exit program.

Example 3-1 QIBM_DB_OPEN exit enhancement

```
ADDEXITPGM EXITPNT(QIBM_QDB_OPEN) FORMAT(DBOP0100) PGMNBR(7) PGM(MJATST/OPENEXIT2)
THDSAFE(*YES) TEXT('MJA') REPLACE(*NO) PGMDTA(*JOB *CALC '*OBJAUD')
```

3.3 Data encryption enhancements

This section describes the enhancements to data encryption provided in IBM i 7.1.

3.3.1 ASP encryption enhancements

You can use ASP encryption to protect data while it is on disk or when it is transferred to and from the disk drive. This encryption is valuable in cases such as a disk replacement or to increase protection in SAN or Geographic Mirror data transmissions. More information about this topic can be found in *Security Guide for IBM i V6.1*, SG24-7680.

With IBM i 7.1 enhancements, ASP encryption can now be turned on and off and the data encryption key can be changed for an existing user ASP. These changes take a significant amount of time because all the data in the disk pool must be processed. This task is done in the background at low priority with a minimal impact on performance.

For more information about ASP encryption enhancements, see 9.1.4, “Encrypted ASP enhancements” on page 375.

3.3.2 Field procedures to encode and decode data

Field procedures are user-written exit programs that can modify the contents and size of the data on insert, update, and read operations. This capability can be used in many ways, but the most popular way is to provide transparent encryption and masking of data based on business logic. The term *transparent encryption* is important because it implies the ability to define a data encryption policy without any change in the actual application code.

For more information about field procedures, see 6.2.6, “FIELDPROC support for encoding and encryption” on page 166.

3.4 Security enhancements for DB2 for i

DB2 for i is positioned as self-protecting database because of the object-based structure of IBM i, which DB2 for i uses.

In addition to the self-protecting security features, the operating system and DB2 for i include built-in encryption capabilities that enable customers to add an extra-layer of security around their data.

IBM i 7.1 includes significant enhancements in database security, including simplified authority management for DDM and IBM DRDA® environments, as well as significant improvements in database activity monitoring and auditing.

3.4.1 Query Manager user profile auditing

DB2 Query Manager (QM) is a function that provides a common method for accessing data and reporting the results from a relational database across the different DB2 platforms.

You can use Query Manager to design and format printed reports from processed queries. Those queries can be included in programs that are written in several High Level Programming Languages (HLLs). QM also provides data edit capabilities that you can use to update and add data to local and remote DB2 for i tables without additional application programs.

For this reason, more user profile support is also provided to allow administrators to tailor QM defaults, limits, and privileges for each user. QM is the only interface on IBM i with the option to grant and revoke permissions to a specific SQL statement (per user); this is a different capability and concept because it is not object based.

Note: The IBM licensed program 5770-ST1 “IBM DB2 Query Manager and SQL Development Kit for IBM i” is required only for Query Manager and SQL application development. After the applications are created, they can be run on other servers that are running IBM i that do not have this product installed by using DB2 for i database manager support.

Figure 3-2 shows an example of the Work with Query Manager profiles option from the STRQM menu. To get there, choose a user profile, select Y in the Select allowed SQL statements parameter, and press Enter.

Select Allowed SQL Statements

User profile : OBTEST

Description :

Type option to allow the use of the statement, press Enter.

Press F21 to select all.

1=Select

Opt	SQL	Opt	SQL
	Statement		Statement
	MERGE		SET ENCRYPTION PASSWORD
	REFRESH TABLE		SET PATH
	RELEASE SAVEPOINT		SET SCHEMA
	RENAME		SET TRANSACTION
	REVOKE		UPDATE
	ROLLBACK		
	SAVEPOINT		
1	SELECT		
	SET CURRENT DEGREE		

F3=Exit F12=Cancel F21=Select all F22=QM Statement

Bottom

Figure 3-2 Select Allowed SQL Statements panel

With IBM i 7.1, you can now audit changes that are made to a Query Manager profile if auditing is enabled for **AUDLVL(*SECURITY)**. A new journal entry type of X2 contains the old and new Query Manager profile values.

An outfile is not provided for this journal entry; instead, the QSYS2.SQLQMProfilesAudit view can be queried, as shown in Example 3-2.

Example 3-2 Creating a permanent table that contains the new journal entries

```
CREATE TABLE mytablename AS  
(SELECT * FROM QSYS2.SQLQMProfilesAudit) WITH DATA
```

There are two main parts to this solution:

- ▶ A new journal entry (X2) is sent to the QAUDJRN any time a Query Manager (QM) profile changes. The journal entry contains the before and after profile information.
- ▶ Normally, the way audit journal entries are externalized is that each has an associated output file. A particular audit entry type is output to that output file, which is tailored to the information in that journal entry. However, output files are expensive because of language translation and maintenance effort. Also, the Query Manager profile changes with each release as IBM adds more SQL statements that you can authorize.

The audit journal entry is externalized using a DB2 for i supplied view in QSYS2, similar to how we provided the current values of profiles using the QSYS2.SQLQMprofiles view.

The view entry returns a set of data that is available for all journal entries that identifies when the change was made and by whom:

- ▶ Journal entry time stamp
- ▶ Current user
- ▶ Job name, job user, and job number
- ▶ Thread

Most of the values that are stored in the QM profile have only two possible values. For example, the values for the authority to use the **INSERT** statement are Y or N.

The following QM profile values have more than two possible values:

- ▶ Default Library
- ▶ Default object creation authority
- ▶ Relational database connection
- ▶ Sample size of Query
- ▶ Maximum records that are allowed on an interactive run
- ▶ Default collection for QM tables
- ▶ Query Data Output Type
- ▶ Table and library for output
- ▶ Job description and library for batch run
- ▶ Commitment control lock level
- ▶ Default printer name

3.4.2 Database Monitor and Plan Cache variable values masking

IBM i 7.1 includes a DB2 system procedure to protect sensitive data. After you set the secure attribute for a column in a specific table, DB2 for i creates the masking for SQL Performance Monitors (Database Monitors) and SQL Plan Cache.

When an SQL statement refers to any column that has the **SECURE** attribute that is set to YES, all host variable values appear as “*SECURE” when examined from the database monitor and plan cache, unless the security officer has started the database monitor or the security officer is accessing the plan cache. Example 3-3 illustrates how to secure a column of a particular table.

Access: Only the QSECOFR user can see the variable values. Users that have the *SECOFR user special authority but are *not* the QSECOFR user see “*SECURE” for the variable values.

Example 3-3 Procedure to set a secure column

```
CALL SYSPROC.SET_COLUMN_ATTRIBUTE(Table_Schema, Table_Name, Column_Name,
Attribute);
```

Here is a description of the parameters that are shown in Example 3-3:

- ▶ Table_Schema - VARCHAR(10): Contains the system name of a table's schema (library).
- ▶ Table_Name - VARCHAR(10): Contains the system name of a table (file).
- ▶ Column_Name - VARCHAR(10): Contains the system column name that is secured.
- ▶ Attribute - VARCHAR(10): Contains the attribute that is set for the column. The valid values are:
 - ‘SECURE NO’: This column does not contain data that must be secured in a database monitor or plan cache.
 - ‘SECURE YES’: This column contains data that must be secured in a database monitor or plan cache.

SECURE column added to QSYS2.SYSCOLUMNS2

QSYS2.SYSCOLUMNS2 is a DB2 catalog that returns more information that is not available in QSYS2.SYSCOLUMNS (such as the allocated length of a varying length column).

The SECURE column externalizes the security setting, if SYSPROC.SET_COLUMN_ATTRIBUTE() is used to configure the security of the column for database performance analysis tasks.

The SECURE column values are:

NULL	This column is not the target of a call to SYSPROC.SET_COLUMN_ATTRIBUTE.
'0'	This column does not contain data that must be secured in a database monitor or plan cache. This column was a target of a call to SYSPROC.SET_COLUMN_ATTRIBUTE with the ‘ SECURE NO ’ parameter.
'1'	This column contains data that must be secured in a database monitor or plan cache. This column was the target of a call to SYSPROC.SET_COLUMN_ATTRIBUTE with the ‘ SECURE YES ’ parameter, as shown in Example 3-4.

Example 3-4 Values with ‘SECURE YES’

```
-- sample table creation
call qsys.create_sql_sample('PRODLIB');

-- Change attribute on SALARY column
CALL SYSPROC.SET_COLUMN_ATTRIBUTE('PRODLIB', 'EMPLOYEE', 'SALARY', 'SECURE YES');
```

```
SELECT Column_name, SECURE FROM qsys2.syscolumns2
WHERE system_table_schema = 'PRODLIB' and system_table_name = 'EMPLOYEE' ;
```

Figure 3-3 shows the results of the previous **SELECT** command that shows SECURE = '1' for 'SALARY'.

COLUMN_NAME	SECURE
EMPNO	0
FIRSTNAME	0
MIDINIT	0
LASTNAME	0
WORKDEPT	0
PHONENO	0
HIREDATE	0
JOB	0
EDLEVEL	0
SEX	0
BIRTHDATE	0
SALARY	1
BONUS	0
COMM	0

Figure 3-3 Results of the previous **SELECT** command that shows the SECURE columns

The database performance analyst sees the output that is shown in Figure 3-4 with the performance analysis results for this query:

```
select * from prodlib.employee where salary > 20000;
```

Operation	Statement Text	Variable Values
OPEN	SELECT * FROM prodlib.employee WHERE salary > ?	*SECURE
CLOSE	CLOSE CRSR0004	SELECT * FROM prodlib.employee WHERE salary > ?
CLOSE (Hard)	HARD CLOSE	1 CURSORS

Figure 3-4 Performance analysis results with SECURE columns

3.4.3 Security function usage for a database

Function usage is a powerful approach to implement a granular security model. Granularity is better for users because they can use IBM i capabilities that are consistent with their job or role in the organization.

Function usage support on IBM i has been in place for many releases, providing an alternative security control for Job Watcher, Cluster Management, IBM Tivoli® Directory Server administration, Backup Recovery and Media Services for i (BRMS), and other components.

DB2 for i includes several functions now for database administration, monitoring, and access control.

Function usage can be configured by the Work with Function Usage (**WRKFCNUSG**) and Change Function Usage (**CHGFCNUSG**) CL commands, IBM Navigator for i, or through the System i Navigator Application Administration feature, as shown in Figure 3-5.




Figure 3-5 System i Navigator window to set up function usage authority

New authority options for SQL analysis and tuning

Before you can use this enhancement, the IBM i system administrator must grant the JOBCTL user special authority to DB2 administrators and DB2 performance analysts.

Because *JOBCTL authority allows a user to change many system critical settings that are unrelated to database activity, it is not an easy decision for security officers to grant this authority. In many cases, the request for *JOBCTL is not granted to database analysts, thus prohibiting the usage of the full set of database tools.

IBM i 7.1 provides a new function usage group called QIBM_DB. The function IDs included for SQL analysis and tuning are:

QIBM_DB_SQLADM Function usage for IBM i Database Administrator tasks

QIBM_DB_SQLMON Function usage for IBM i Database Information tasks

These new functions allow settings where users with the *ALLOBJ special authority cannot do the DB2 administration or DB2 monitoring tasks. Also, a group profile can be specified in this function's setting. If a user profile is associated with several supplemental group profiles, the access is granted if at least one of these group profiles is set to *ALLOWED in a particular function. The adopted authority from program owner profiles has no effect on allowing access to DB2 administration and DB2 monitoring. Access is always granted according to the user profile under which a program with adopted authority runs.

New authority options for database server inbound connections

These function usage IDs block database server inbound connections and are not based on the communication protocol. The function usage IDs ship with the default authority = *ALLOWED. The security officer can easily deny access to specific users or groups.

QIBM_DB_ZDA Restrict ODBC and JDBC Toolbox from the server side, including RUNSQL Scripts, System i Navigator, and DB2 specific portions of IBM Navigator for i.

QIBM_DB_DDMDRDA Lock down DRDA and DDM application server access.

This is an alternative to a User Exit Program approach. Using a function usage ID has several advantages:

- ▶ No coding is required.
- ▶ It is easy to change and it is auditable.

3.4.4 Simplified DDM and Distributed Database (DRDA) authentication

In some environments, it is cumbersome to maintain server authentications in DRDA and Relational Database DDM (RDB DDM) file environments. As systems are added to the network topology, per user profile server authentications for every server must be administered. Customers with many users and many servers face the decision of whether to restrict/control access on each server using the Change DDM TCP/IP Attributes (**CHGDMDTCPA**) command or on each client using the Add Server Authentication Entry (**ADDSVRAUTE**) command.

The dilemma that is faced when you choose to enforce password validation on the servers is that every user who needs to connect to the server needs administrative work that is done on their behalf. This work uses Add Server Authentication Entry (**ADDSVRAUTE**) for every target server, or requires every user to supply a user ID and password on each **CONNECT** statement.

You can use the special value **QDDMDRDASERVER**, which is added to the Add Server Authentication Entry (**ADDSVRAUTE**) command **SERVER** parameter, to simplify this process. This special value allows an administrator to configure a user to work with all possible DDM or DRDA connections to any system in the Internet Protocol network through a common user ID and password. After it is configured for a specific user, no additional changes need to be made for that user, as systems are added to the Relational Database Directory.

As before, this setup does not allow a customer to connect over DRDA/DDM unless they specify a valid user ID and password on the server authentication entry or **CONNECT** statement.

Server authentication entries

A server authentication entry is a way to define a user ID and password to send through a connect request over TCP/IP. A server authentication list is associated with every user profile on the system. By default, the list is empty; however, you can add entries by using the Add Server Authentication Entry (**ADDSVRAUTE**) command.

When you attempt a DRDA connection over TCP/IP without specifying a user ID and password, the DB2 for i client (AR) checks the server authentication list for the user profile under which the client job is running. If it finds a match between the RDB name on the **CONNECT** statement and the SERVER name in an authentication entry (which must be in uppercase), the associated **USRID** parameter in the entry is used for the connection user ID. If a **PASSWORD** parameter is stored in the entry, that password is also sent on the connection request.

A server authentication entry can also be used to send a password over TCP/IP for a DDM file I/O operation. When you attempt a DDM connection over TCP/IP, DB2 for i checks the server authentication list for the user profile under which the client job is running. If it finds a match between either the RDB name (if RDB directory entries are used) or QDDMSERVER and the SERVER name in an authentication entry, the associated **USRID** parameter in the entry is used for the connection user ID. If a **PASSWORD** parameter is stored in the entry, that password is also sent on the connection request.

DDM and DRDA authentication entry management using group profiles

DDM and DRDA also can take advantage of a common user ID and password that are defined in a server authentication entry under a group profile name or supplemental group profile name. The group profile name or supplemental group profile name is specified with the **USRPRF** parameter of the **ADDSVRAUTE** command.

System administration of server authentication entries becomes much simpler when a group profile is used. A common user ID and password profile is used when you establish a group profile-based DDM or DRDA connection. Remote connection capability is managed by controlling the users that belong to the group with their authority to the group profile.

QDDMDRDASERVER usage details

The current special value **QDDMSERVER** allows non-RDB DDM file users to make DDM connections to servers with a common user ID and password. The new special value **QDDMDRDASERVER** is a superset of **QDDMSERVER**, expanding this support to RDB DDM files and DRDA connections.

For DRDA connection requests, if a server authentication entry that specifies the system name exists, and a user ID and password are *not passed* in a **CONNECT** statement, the user ID and password that are associated with the server authentication entry takes precedence over the server authentication entry for **QDDMDRDASERVER**.

For DRDA connection requests, if a server authentication entry specifies the system name exists, and a user ID and password are *passed* in a **CONNECT** statement, the user ID and password that are associated with the **CONNECT** statement takes precedence over any server authentication entry.

For RDB DDM file connection requests, the server authentication entry that specifies the system name takes precedence over the server authentication entry for **QDDMDRDASERVER**. For non-RDB DDM file connection requests, the server authentication entry **QDDMSERVER** takes precedence over the server authentication entry for **QDDMDRDASERVER**.

For example, suppose that you have an environment with three systems (SYSA, SYSB, and SYSC), where:

- ▶ SYSA is the application requestor (AR).
- ▶ SYSB and SYSC are the application servers (AS).

You run the following commands on SYSA:

```
ADDSVRAUTE USRPRF(YOURPRF) SERVER(QDDMDRDASERVER) USRID(youruid) PASSWORD(yourpwd)
STRSQL
CONNECT TO SYSB
CONNECT TO SYSC
```

You now have two connections with shared 'youruid' and password 'yourpwd'. You run the following commands on SYSA:

```
ADDSVRAUTE USRPRF(YOURPRF) SERVER(QDDMDRDASERVER) USRID(youruid) PASSWORD(yourpwd)
ADDSVRAUTE USRPRF(YOURPRF) SERVER(SYSB) USRID(yourotheruid) PASSWORD(yourotherpwd)
ADDSVRAUTE USRPRF(YOURPRF) SERVER(SYSC) USRID(yourotheruid) PASSWORD(yourotherpwd)
STRSQL
CONNECT TO SYSB
CONNECT TO SYSC
```

You now have two connections with the user ID *yourotheruid* and password *yourotherpwd*. This situation occurs because server authentication entries that specify the real system name take precedence over server authentication entries that specify **QDDMDRDASERVER**. You run the following commands on SYSA:

```
ADDSVRAUTE USRPRF(YOURPRF) SERVER(QDDMDRDASERVER) USRID(youruid) PASSWORD(yourpwd)
ADDSVRAUTE USRPRF(YOURPRF) SERVER(SYSB) USRID(yourotheruid) PASSWORD(yourotherpwd)
STRSQL
CONNECT TO SYSB user testuserid using 'testpassword'
CONNECT TO SYSC
```

You now have two connections. The connection to SYSB is made with the user ID *testuserid* and password *testpassword*. This situation occurs because specifying the user ID and password on a **CONNECT** statement takes precedence over server authentication entries. The connection to SYSC is made with user ID 'youruid' and password 'yourpwd', because it uses the **QDDMDRDASERVER** authentication entry when no other server authentication entry exists specifying the system name. You run the following commands on SYSA:

```
ADDSVRAUTE USRPRF(YOURPRF) SERVER(QDDMDRDASERVER) USRID(youruid) PASSWORD(yourpwd)
ADDSVRAUTE USRPRF(YOURPRF) SERVER(QDDMSERVER) USRID(youruid2) PASSWORD(yourpwd2)
ADDSVRAUTE USRPRF(YOURPRF) SERVER(SYSC) USRID(yourotheruid) PASSWORD(yourotherpwd)
CRTDDMF FILE(QTEMP/DDMF) RMTFILE(FILE) RMTLOCNAME(SYSB *IP)
CRTDDMF FILE(QTEMP/DDMF2) RMTFILE(FILE) RMTLOCNAME(*RDB) RDB(SYSB)
CRTDDMF FILE(QTEMP/DDMF3) RMTFILE(FILE) RMTLOCNAME(*RDB) RDB(SYSC)
```

Assume that you use the following statements:

- ▶ SBMRMTCMD CMD('DSPLIB YOURLIB') DDMFILE(QTEMP/DDMF)

The connection to SYSB would be made with the user ID *youruid2* and password *yourpwd2*. This situation occurs because non-RDB DDM files use **QDDMSERVER** for the user ID AND password at connection time. If **QDDMSERVER** does not exist, **QDDMDRDASERVER** is used.

- ▶ SBMRMTCMD CMD('DSPLIB YOURLIB') DDMFILE(QTEMP/DDMF2)

The connection to SYSB is made with the user ID *youruid* and password *yourpwd*. This situation occurs because the RDB DDM files use the **QDDMDRDASERVER** for the user ID and password at connection time if a server authentication entry does not exist that specifies the system name.

- SBMRMTCMD CMD('DSPLIB YOURLIB') DDMFILE(QTEMP/DDMF3)

The connection to SYSC is made with the user ID *yourotheruid* and password *yourotherpwd*. This situation occurs because the RDB DDM files do not use **QDDMDRDASERVER** for the user ID and password at connection time if a server authentication entry exists that specifies the system name.

3.5 Real-time database activity monitoring

Databases are the *primary* target for external hackers and insider attacks because databases contain an organization's most valuable information, including customer records, payment card data, and financial results.

Existing security solutions, such as intrusion detection systems (IDSs), lack knowledge about the database protocols and structures, which is required to detect inappropriate activities. For more information about IDS support on IBM i, go to:

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp>

Today, an increasing number of industries ask for compliance mandates that generally require organizations to detect, record, and remediate unauthorized access or changes to sensitive data, including access or changes by privileged users, while providing a secure audit trail to validate compliance.

Information security and database managers struggle to implement these types of controls, especially regarding monitoring privileged users. Heightened focus on business reputation risk and sensitive data protection is also driving closer internal scrutiny of controls. The result is clear: Providing effective database security and compliance is not easy.

DB2 for i includes several enhancements to help administrators in this effort.

3.5.1 Enhanced audit capability for STRDBMON

The database monitor is a valuable tool for database performance analysis, but you can use its capabilities for real-time logging for database auditing and compliance reinforcement. The database monitor can generate significant processor and disk storage impact, so a targeted approach to using the Start Database Monitor (**STRDBMON**) command is always preferred.

Here are the additional logging and prefiltering capabilities that are supplied for Database Monitor.

Terminology: The term *prefilter* is related to filtering capabilities for a database monitor in the recording process and the term *filter* is applied to the selection criteria in the collected data.

Prefilters to logging RUNSQL, RUNSQLSTM, and STRSQL

The database monitor, started by Start Database Monitor (**STRDBMON**), can track Run SQL (**RUNSQL**), Run SQL Statement (**RUNSQLSTM**), or Start SQL Interactive Session (**STRSQL**) activity that is connected to a specific user profile, application, program ID, accounting code, or relational database.

Table 3-5 shows the prefilters available in the Start Database Monitor (**STRDBMON**) command and the corresponding client registers. Database monitor entries are created for SQL statements that are run where the client special register matches the corresponding prefilter value.

Table 3-5 *STRDBMON* prefilter options for SQL-related CL commands

Client register	STRDBMON prefilter parameter	Values
Account string (ACCTNG)	FTRCLTACG	The user's accounting code (the ACGCDE parameter on the User Profile object)
Application Name (APPLNAME)	FTRCLTAPP	START SQL INTERACTIVE SESSION or RUN SQL STATEMENTS
Program name (PROGRAMID)	FTRCLTPGM	STRSQL , RUNSQLSTM , or RUNSQL
Client User (USERID)	FTRCLTUSR	The user's name
Workstation (WRKSTNNAME)	FTRCLTWS	The DB2 for i database name

There is support for a client-specific filter using the **COMMENT** parameter in Start Database Monitor (**STRDBMON**), but this support is limited to only one parameter, which should be a value up to 50 characters long.

Figure 3-6 shows the new filter parameters for **STRDBMON**.

```

                                Start Database Monitor (STRDBMON)

Type choices, press Enter.

Filter by query governor . . . . FTRQRYGOVR      *NONE
Filter by accounting string . . FTRCLTACG        *NONE

Filter by application name . . . FTRCLTAPP        *NONE

Filter by program name . . . . . FTRCLTPGM        *NONE

Filter by client user ID . . . . FTRCLTUSR        *NONE

Filter by work station . . . . . FTRCLTWS         *NONE

More...
F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys
Messages pending on other displays.
```

Figure 3-6 *STRDBMON* - New filter parameters

Prefilter for SQLCODEs

The Start Database Monitor (**STRDBMON**) command is enhanced with a prefilter parameter, **FTRSQC**ODE, for the results of SQL statements.

There are several prefilters by **SQLCODE** (**FTRSQC**ODE) parameter values:

*NONE	Filtering by SQLCODE is not specified.
*NONZERO	Any non-zero SQLCODE .
*ERROR	SQLCODE < 0. The SQL statement has failed.
*WARN	SQLCODE > 0. A warning condition is indicated.
<user specified SQLCODE>	This is your specified SQLCODE .

Example 3-5 shows how to use the **FTRSQC**ODE parameter.

Example 3-5 Collect QQRID=1000 DBMON records for all instances of SQL failures due to lock contention

```
STRDBMON OUTFILE(DBMONLIB/LOCKMON) JOB(*ALL/*ALL/QZDASOINIT) TYPE(*DETAIL)
FTRSQCODE(-913)
```

For **SQLSTATE**s and **SQLCODE**s for DB2 for IBM i 7.1, see the Information Center at:

<http://www.publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp?topic=%2Frza1a%2Frza1acc1.htm>

Extending the IP address prefilter to work with non-database clients

Before this enhancement was introduced, support for the Start Database Monitor (**STRDBMON**) command Filter by IP address (**FTRINTNETA**) parameter was restricted for work against database initiated connections.

This enhancement allows a Filter by IP address to work as a database monitor prefilter against many non-database initiated connections.

TELNET connections are not recognized by this parameter.

Extending support to prefilter against group profiles

The Start Database Monitor (**STRDBMON**) command has been enhanced to recognize group profile names when you use the **FTRUSER(name)** command parameter.

When you determine whether the current user's SQL should be captured in the SQL Performance Monitor (database monitor) output, the command now determines whether the user is a member of the group.

Wildcard group profile names are allowed.

For example, if you specify **FTRUSER(ADMIN*)** and both **ADMINGRP** and **ADMINGRP2** are group profiles, any SQL run by users in either group is captured.

SQL Performance Monitor interface for "User" can be used to specify the group profile.

Identifying Query/400 users

This enhancement delivers a method to easily identify Query/400 users and the queries run. You can use this tool to audit sensitive data consumption and determine queries usage and reporting modernization strategies.

For more information about this topic, see 6.5.8, “STRDBMON pre-filtering of QUERY/400 command usage” on page 238.

3.5.2 SQL statement auditing with less resource consumption

Implementing Start Database Monitoring (STRDBMON) through a viewing mechanism enables faster SQL statement auditing.

There are two view mechanisms that can be used to start a database performance monitor on a view, saving both performance and storage.

Use input only columns to capture only a subset of the monitor data in an underlying table.

In DB2 for i, the database performance monitor table has 276 columns. Assume that an auditing application is interested in collecting only the SQL statement, the variable values, and the information that identifies the user and job information. This information is contained in only 20 out of the 276 columns (the columns QQRID, QQJFLD, and QQI5 must also be added to process the resulting view).

To use the first view mechanism, complete the following steps:

1. Create a table with the 23 columns that you want to view (QAQQDBMN is the model monitor file that is shipped with DB2 for i), as shown in Example .

Example 3-6 Create a table with 23 columns

```
CREATE TABLE mjatst.looptable4 AS (SELECT
  QQSTIM, QQETIM, QQC81, QQ1000L, QQDBCLOB1,
  QVC5001, QVC3001, QVC3002, QVC3003, QVC3005,
  QVC3006, QQJOB, QQUSER, QQJNUM, QVC102,
  QQI9, QQC104, QQC103, QQC183, QCSMINT2,
  QQRID, QQI5, QQJFLD
FROM QAQQDBMN) WITH NO DATA;
```

2. Create a view that has 276 columns that match the database monitor table columns, as shown in Example . Only the 23 wanted columns are input / output columns; the others are input only columns (those columns that are just CAST as NULL). The columns must have the same attributes and be in the same order as in the base database monitor table.

Example 3-7 Create a table with 276 columns

```
CREATE VIEW mjatst.loopview4 AS SELECT

QQRID AS "Record ID", CAST(NULL AS TIMESTAMP) AS QQTIME,
QQJFLD AS "Join Field", CAST(NULL AS CHAR(18)) AS QQRDBN,
CAST(NULL AS CHAR(8)) AS QCSYS, QQJOB AS "Job Name",
QQUSER AS "Job User", QQJNUM AS "Job Number",
CAST(NULL AS DECIMAL(15,0)) AS QQUCNT , CAST(NULL AS VARCHAR(100)) AS QQUCDEF ,
CAST(NULL AS DECIMAL(15,0)) AS QQSTN , CAST(NULL AS DECIMAL(15,0)) AS QQQDTN ,
CAST(NULL AS DECIMAL(15,0)) AS QQQDTL , CAST(NULL AS DECIMAL(15,0)) AS QQMATN ,
CAST(NULL AS DECIMAL(15,0)) AS QQMATL , CAST(NULL AS CHAR(10)) AS QQTLN ,
CAST(NULL AS CHAR(10)) AS QQTFN , CAST(NULL AS CHAR(10)) AS QQTMN ,
CAST(NULL AS CHAR(10)) AS QQPTLN , CAST(NULL AS CHAR(10)) AS QQPTFN ,
CAST(NULL AS CHAR(10)) AS QQPTMN , CAST(NULL AS CHAR(10)) AS QQILNM ,
CAST(NULL AS CHAR(10)) AS QQIFNM , CAST(NULL AS CHAR(10)) AS QQIMNM ,
CAST(NULL AS CHAR(10)) AS QQNTNM , CAST(NULL AS CHAR(10)) AS QQNLNM ,
```

```

QQSTIM AS "Start Time", QQETIM AS "End Time",
CAST(NULL AS CHAR(1)) AS QQKP , CAST(NULL AS CHAR(1)) AS QQKS ,
CAST(NULL AS DECIMAL(15,0)) AS QQTOTR , CAST(NULL AS DECIMAL(15,0)) AS QQTMPR ,
CAST(NULL AS DECIMAL(15,0)) AS QQJNP , CAST(NULL AS DECIMAL(15,0)) AS QQEPT ,
CAST(NULL AS CHAR(1)) AS QQDSS , CAST(NULL AS CHAR(1)) AS QQIDXA ,
CAST(NULL AS CHAR(1)) AS QQORDG , CAST(NULL AS CHAR(1)) AS QQGRPG ,
CAST(NULL AS CHAR(1)) AS QQJNG , CAST(NULL AS CHAR(1)) AS QQUNIN ,
CAST(NULL AS CHAR(1)) AS QQSUBQ , CAST(NULL AS CHAR(1)) AS QQHSTV ,
CAST(NULL AS CHAR(1)) AS QQRCDS , CAST(NULL AS CHAR(2)) AS QQRCOD ,
CAST(NULL AS DECIMAL(15,0)) AS QQRSS , CAST(NULL AS DECIMAL(15,0)) AS QQREST ,
CAST(NULL AS DECIMAL(15,0)) AS QQRIDX , CAST(NULL AS DECIMAL(15,0)) AS QQFKEY ,
CAST(NULL AS DECIMAL(15,0)) AS QQKSEL , CAST(NULL AS DECIMAL(15,0)) AS QQAJN ,
CAST(NULL AS VARCHAR(1000)) AS QQIDXD , CAST(NULL AS CHAR(1)) AS QQC11 ,
CAST(NULL AS CHAR(1)) AS QQC12 , CAST(NULL AS CHAR(1)) AS QQC13 ,
CAST(NULL AS CHAR(1)) AS QQC14 , CAST(NULL AS CHAR(1)) AS QQC15 ,
CAST(NULL AS CHAR(1)) AS QQC16 , CAST(NULL AS CHAR(1)) AS QQC18 ,
CAST(NULL AS CHAR(2)) AS QQC21 , CAST(NULL AS CHAR(2)) AS QQC22 ,
CAST(NULL AS CHAR(2)) AS QQC23 , CAST(NULL AS DECIMAL(15,0)) AS QQI1 ,
CAST(NULL AS DECIMAL(15,0)) AS QQI2 , CAST(NULL AS DECIMAL(15,0)) AS QQI3 ,
CAST(NULL AS DECIMAL(15,0)) AS QQI4 , QQI5 AS "Refresh Count",
CAST(NULL AS DECIMAL(15,0)) AS QQI6 , CAST(NULL AS DECIMAL(15,0)) AS QQI7 ,
CAST(NULL AS DECIMAL(15,0)) AS QQI8 , QQI9 AS "Thread ID",
CAST(NULL AS DECIMAL(15,0)) AS QQIA , CAST(NULL AS DECIMAL(15,0)) AS QQF1 ,
CAST(NULL AS DECIMAL(15,0)) AS QQF2 , CAST(NULL AS DECIMAL(15,0)) AS QQF3 ,
CAST(NULL AS CHAR(6)) AS QQC61 , QQC81 AS SQLSTATE_ ,
CAST(NULL AS CHAR(8)) AS QQC82 , CAST(NULL AS CHAR(8)) AS QQC83 ,
CAST(NULL AS CHAR(8)) AS QQC84 , CAST(NULL AS CHAR(10)) AS QQC101 ,
CAST(NULL AS CHAR(10)) AS QQC102 , QQC103 AS "Program",
QQC104 AS "Program Schema", CAST(NULL AS CHAR(10)) AS QQC105 ,
CAST(NULL AS CHAR(10)) AS QQC106 , CAST(NULL AS VARCHAR(128)) AS QQC181 ,
CAST(NULL AS VARCHAR(128)) AS QQC182 , QQC183 AS "IP Address",
CAST(NULL AS VARCHAR(30)) AS QQC301 , CAST(NULL AS VARCHAR(30)) AS QQC302 ,
CAST(NULL AS VARCHAR(30)) AS QQC303 , CAST(NULL AS VARCHAR(1000)) AS QQ1000 ,
CAST(NULL AS TIMESTAMP) AS QQTIM1 , CAST(NULL AS TIMESTAMP) AS QQTIM2 ,
CAST(NULL AS VARCHAR(128)) AS QVQTBL , CAST(NULL AS VARCHAR(128)) AS QVQLIB ,
CAST(NULL AS VARCHAR(128)) AS QVPTBL , CAST(NULL AS VARCHAR(128)) AS QVPLIB ,
CAST(NULL AS VARCHAR(128)) AS QVINAM , CAST(NULL AS VARCHAR(128)) AS QVILIB ,
CAST(NULL AS CHAR(1)) AS QVQTBLI , CAST(NULL AS CHAR(1)) AS QVPTBLI ,
CAST(NULL AS CHAR(1)) AS QVINAMI , CAST(NULL AS CHAR(1)) AS QVBNDY ,
CAST(NULL AS CHAR(1)) AS QVJFANO , CAST(NULL AS CHAR(1)) AS QVPARPF ,
CAST(NULL AS CHAR(1)) AS QVPARPL , CAST(NULL AS CHAR(1)) AS QVC11 ,
CAST(NULL AS CHAR(1)) AS QVC12 , CAST(NULL AS CHAR(1)) AS QVC13 ,
CAST(NULL AS CHAR(1)) AS QVC14 , CAST(NULL AS CHAR(1)) AS QVC15 ,
CAST(NULL AS CHAR(1)) AS QVC16 , CAST(NULL AS CHAR(1)) AS QVC17 ,
CAST(NULL AS CHAR(1)) AS QVC18 , CAST(NULL AS CHAR(1)) AS QVC19 ,
CAST(NULL AS CHAR(1)) AS QVC1A , CAST(NULL AS CHAR(1)) AS QVC1B ,
CAST(NULL AS CHAR(1)) AS QVC1C , CAST(NULL AS CHAR(1)) AS QVC1D ,
CAST(NULL AS CHAR(1)) AS QVC1E , CAST(NULL AS CHAR(1)) AS QVC1F ,
CAST(NULL AS CHAR(1)) AS QWC11 , CAST(NULL AS CHAR(1)) AS QWC12 ,
CAST(NULL AS CHAR(1)) AS QWC13 , CAST(NULL AS CHAR(1)) AS QWC14 ,
CAST(NULL AS CHAR(1)) AS QWC15 , CAST(NULL AS CHAR(1)) AS QWC16 ,
CAST(NULL AS CHAR(1)) AS QWC17 , CAST(NULL AS CHAR(1)) AS QWC18 ,
CAST(NULL AS CHAR(1)) AS QWC19 , CAST(NULL AS CHAR(1)) AS QWC1A ,
CAST(NULL AS CHAR(1)) AS QWC1B , CAST(NULL AS CHAR(1)) AS QWC1C ,
CAST(NULL AS CHAR(1)) AS QWC1D , CAST(NULL AS CHAR(1)) AS QWC1E ,

```

CAST(NULL AS CHAR(1)) AS QWC1F , CAST(NULL AS CHAR(2)) AS QVC21 ,
 CAST(NULL AS CHAR(2)) AS QVC22 , CAST(NULL AS CHAR(2)) AS QVC23 ,
 CAST(NULL AS CHAR(2)) AS QVC24 , CAST(NULL AS DECIMAL(15,0)) AS QVCTIM ,
 CAST(NULL AS DECIMAL(15,0)) AS QVPARD , CAST(NULL AS DECIMAL(15,0)) AS QVPARU ,
 CAST(NULL AS DECIMAL(15,0)) AS QVPARRC , CAST(NULL AS DECIMAL(15,0)) AS QVRCNT ,
 CAST(NULL AS DECIMAL(15,0)) AS QVFILES , CAST(NULL AS DECIMAL(15,0)) AS QVP151 ,
 CAST(NULL AS DECIMAL(15,0)) AS QVP152 , CAST(NULL AS DECIMAL(15,0)) AS QVP153 ,
 CAST(NULL AS DECIMAL(15,0)) AS QVP154 , CAST(NULL AS DECIMAL(15,0)) AS QVP155 ,
 CAST(NULL AS DECIMAL(15,0)) AS QVP156 , CAST(NULL AS DECIMAL(15,0)) AS QVP157 ,
 CAST(NULL AS DECIMAL(15,0)) AS QVP158 , CAST(NULL AS DECIMAL(15,0)) AS QVP159 ,
 CAST(NULL AS DECIMAL(15,0)) AS QVP15A , CAST(NULL AS DECIMAL(15,0)) AS QVP15B ,
 CAST(NULL AS DECIMAL(15,0)) AS QVP15C , CAST(NULL AS DECIMAL(15,0)) AS QVP15D ,
 CAST(NULL AS DECIMAL(15,0)) AS QVP15E , CAST(NULL AS DECIMAL(15,0)) AS QVP15F ,
 CAST(NULL AS CHAR(4)) AS QVC41 , CAST(NULL AS CHAR(4)) AS QVC42 ,
 CAST(NULL AS CHAR(4)) AS QVC43 , CAST(NULL AS CHAR(4)) AS QVC44 ,
 CAST(NULL AS CHAR(8)) AS QVC81 , CAST(NULL AS CHAR(8)) AS QVC82 ,
 CAST(NULL AS CHAR(8)) AS QVC83 , CAST(NULL AS CHAR(8)) AS QVC84 ,
 CAST(NULL AS CHAR(8)) AS QVC85 , CAST(NULL AS CHAR(8)) AS QVC86 ,
 CAST(NULL AS CHAR(8)) AS QVC87 , CAST(NULL AS CHAR(8)) AS QVC88 ,
 CAST(NULL AS CHAR(10)) AS QVC101 , QVC102 AS "User Name",
 CAST(NULL AS CHAR(10)) AS QVC103 , CAST(NULL AS CHAR(10)) AS QVC104 ,
 CAST(NULL AS CHAR(10)) AS QVC105 , CAST(NULL AS CHAR(10)) AS QVC106 ,
 CAST(NULL AS CHAR(10)) AS QVC107 , CAST(NULL AS CHAR(10)) AS QVC108 ,
 CAST(NULL AS VARCHAR(128)) AS QVC1281 , CAST(NULL AS VARCHAR(128)) AS QVC1282 ,
 CAST(NULL AS VARCHAR(128)) AS QVC1283 , CAST(NULL AS VARCHAR(128)) AS QVC1284 ,
 QVC3001 AS "Client Application Name", QVC3002 AS "Client User ID",
 QVC3003 AS "Client Workstation Name", CAST(NULL AS VARCHAR(300)) AS QVC3004 ,
 QVC3005 AS "Client Accounting", QVC3006 AS "Client Program Name",
 CAST(NULL AS VARCHAR(300)) AS QVC3007 , CAST(NULL AS VARCHAR(300)) AS QVC3008 ,
 QVC5001 AS "Interface", CAST(NULL AS VARCHAR(500)) AS QVC5002 ,
 CAST(NULL AS VARCHAR(1000)) AS QVC1000 , CAST(NULL AS VARCHAR(1000)) AS QWC1000 ,
 CAST(NULL AS INTEGER) AS QQINT01 , CAST(NULL AS INTEGER) AS QQINT02 ,
 CAST(NULL AS INTEGER) AS QQINT03 , CAST(NULL AS INTEGER) AS QQINT04 ,
 CAST(NULL AS SMALLINT) AS QQSMINT1 , QQSMINT2 AS "IP Port Number",
 CAST(NULL AS SMALLINT) AS QQSMINT3 , CAST(NULL AS SMALLINT) AS QQSMINT4 ,
 CAST(NULL AS SMALLINT) AS QQSMINT5 , CAST(NULL AS SMALLINT) AS QQSMINT6 ,
 QQ1000L AS "Statement Text", CAST(NULL AS CHAR(1)) AS QFC11 ,
 CAST(NULL AS CHAR(1)) AS QFC12 , CAST(NULL AS CHAR(1)) AS QFC13 ,
 CAST(NULL AS CLOB(2G)) AS QQCLOB2 , CAST(NULL AS CHAR(1)) AS QFC14 ,
 CAST(NULL AS CHAR(1)) AS QFC15 , CAST(NULL AS CHAR(1)) AS QFC16 ,
 CAST(NULL AS CLOB(2G)) AS QQCLOB3 , CAST(NULL AS CHAR(1)) AS QFC17 ,
 CAST(NULL AS CHAR(1)) AS QFC18 , CAST(NULL AS CHAR(1)) AS QFC19 ,
 QQDBCLOB1 AS "Variable Values", CAST(NULL AS CHAR(1)) AS QFC1A ,
 CAST(NULL AS CHAR(1)) AS QFC1B , CAST(NULL AS CHAR(1)) AS QFC1C ,
 CAST(NULL AS NCLOB(1G)) AS QQDBCLOB2 , CAST(NULL AS CHAR(1)) AS QFC1D ,
 CAST(NULL AS CHAR(1)) AS QFC1E , CAST(NULL AS CHAR(1)) AS QFC1F ,
 CAST(NULL AS BLOB(2G)) AS QQBLOB1 , CAST(NULL AS CHAR(1)) AS QXC11 ,
 CAST(NULL AS CHAR(1)) AS QXC12 , CAST(NULL AS CHAR(1)) AS QXC13 ,
 CAST(NULL AS CHAR(1)) AS QXC14 , CAST(NULL AS CHAR(1)) AS QXC15 ,
 CAST(NULL AS CHAR(1)) AS QXC16 , CAST(NULL AS CHAR(1)) AS QXC17 ,
 CAST(NULL AS CHAR(1)) AS QXC18 , CAST(NULL AS CHAR(1)) AS QXC19 ,
 CAST(NULL AS CHAR(1)) AS QXC1A , CAST(NULL AS CHAR(1)) AS QXC1B ,
 CAST(NULL AS CHAR(1)) AS QXC1C , CAST(NULL AS CHAR(1)) AS QXC1D ,
 CAST(NULL AS CHAR(1)) AS QXC1E , CAST(NULL AS CHAR(2)) AS QXC21 ,
 CAST(NULL AS CHAR(2)) AS QXC22 , CAST(NULL AS CHAR(2)) AS QXC23 ,

```

CAST(NULL AS CHAR(2)) AS QXC24 , CAST(NULL AS CHAR(2)) AS QXC25 ,
CAST(NULL AS CHAR(2)) AS QXC26 , CAST(NULL AS CHAR(2)) AS QXC27 ,
CAST(NULL AS CHAR(2)) AS QXC28 , CAST(NULL AS CHAR(2)) AS QXC29 ,
CAST(NULL AS CHAR(4)) AS QXC41 , CAST(NULL AS CHAR(4)) AS QXC42 ,
CAST(NULL AS CHAR (4) FOR BIT DATA) AS QXC43 , CAST(NULL AS CHAR(4)) AS QXC44 ,
CAST(NULL AS INTEGER) AS QQINT05 , CAST(NULL AS INTEGER) AS QQINT06 ,
CAST(NULL AS INTEGER) AS QQINT07 , CAST(NULL AS INTEGER) AS QQINT08 ,
CAST(NULL AS INTEGER) AS QQINT09 , CAST(NULL AS INTEGER) AS QQINT0A ,
CAST(NULL AS INTEGER) AS QQINT0B , CAST(NULL AS INTEGER) AS QQINT0C ,
CAST(NULL AS INTEGER) AS QQINT0D , CAST(NULL AS INTEGER) AS QQINT0E ,
CAST(NULL AS INTEGER) AS QQINT0F , CAST(NULL AS SMALLINT) AS QQSMINT7 ,
CAST(NULL AS SMALLINT) AS QQSMINT8 , CAST(NULL AS SMALLINT) AS QQSMINT9 ,
CAST(NULL AS SMALLINT) AS QQSMINTA , CAST(NULL AS SMALLINT) AS QQSMINTB ,
CAST(NULL AS SMALLINT) AS QQSMINTC , CAST(NULL AS SMALLINT) AS QQSMINTD ,
CAST(NULL AS SMALLINT) AS QQSMINTE , CAST(NULL AS SMALLINT) AS QQSMINTF

FROM MJATST.looptable4
RCDFMT QQDBMN;

```

3. Start the database monitor view by running the following command:

```
STRDBMON mjatst.loopview4;
```

By enhancing the database product to allow this view, any data that is written to the database performance monitor view results in only 23 columns in the underlying base table (mjatst.looptable4). The storage that is used with this technique is a small fraction of a traditional monitor, so the performance is better. The resulting smaller table contains only the information necessary for auditing.

For the second view mechanism, use an INSTEAD OF trigger on the view to immediately process a row of monitor data without storing *any* data.

Complete the following steps:

1. Create a view using the same technique that is described for the first view mechanism. The view can be simplified because in this example, no data is inserted into a table. Run the following commands:

```

CREATE TABLE mjatst.looptable5 AS (SELECT * FROM qaqqdbmn) WITH NO DATA;
CREATE VIEW mjatst.loopview5 AS SELECT * FROM mjatst.looptable5;

```

2. Create an INSTEAD OF trigger:

```

CREATE TRIGGER instead5 INSTEAD OF INSERT ON loopview5
FOR EACH ROW MODE DB2ROW
BEGIN ... this is the body of the trigger ... END;

```

3. Start the database monitor using the view by running the following command:

```
CL: STRDBMON mjatst/loopview5;
```

By enhancing the database product to allow this view, any rows that are written to the database monitor file are passed directly to the INSTEAD OF trigger so no monitor storage is used.

Comparison of the view mechanisms and the traditional database monitor

Figure 3-7 provides a comparison of the elapsed time and the storage that is used by the traditional database monitor and the two view techniques.




Figure 3-7 Comparison of the two view methods

In the case of the view with an INSTEAD OF trigger, the elapsed time and the amount of storage that is used is under the control of the INSTEAD OF trigger. For example, in Figure 3-7, the INSTEAD OF trigger sends the data to another system. This action takes some processing time, but no persistent storage is used on the system that is monitored.

3.5.3 InfoSphere Guardium database activity monitoring support for DB2 for i

InfoSphere Guardium is an enterprise information database audit and protection solution that helps enterprises protect and audit information across a diverse set of relational and non-relational data sources, such as Oracle, Teradata, IBM IMS™, VSAM, Microsoft Sharepoint, IBM Netezza®, DB2 for z/OS®, and DB2 for Linux, UNIX, and Windows.

With InfoSphere Guardium V9.0, DB2 for i can now be included as a data source. You can use this configuration, plus the database security enhancements included in DB2 for i, to monitor database accesses from native interfaces and through SQL in real time without any changes to databases or applications or impacting performance.

Introduction to InfoSphere Guardium

The IBM InfoSphere Guardium solution evolved to address the particular needs of organizations that must implement more automated and auditable data security practices.

InfoSphere Guardium continuously monitors database transactions through lightweight software probes that are installed on the database server (or file share or Microsoft Sharepoint). These probes (called *software taps* (S-TAPs)) monitor all database transactions, including those transactions of privileged users. The S-TAPs also do not require any changes to the database or its applications.

The probes forward transactions to a hardened collector in the network, where they are compared to previously defined policies to detect violations. The system can respond with various policy-based actions, including generating an alert, and for some databases, can block the transaction in real time. (Blocking is not currently available for DB2 for i database activity monitoring.)

InfoSphere Guardium supports a various deployments of large and geographically distributed infrastructures.

Figure 3-8 shows a typical deployment of InfoSphere Guardium Database Activity Monitoring.




Figure 3-8 InfoSphere Guardium Database Activity Monitoring

New InfoSphere Guardium support for IBM i

InfoSphere Guardium support for IBM i monitoring was previously available through three main methods:

- Importation of audit journal entries (QSYS/QAUDJRN) and subsequent analysis and reporting

Although the audit journal support in IBM i provides good support of auditable events, the amount of detail in the audit entries is minimal compared to other Guardium database product support. For example, SQL statements and variable values are not audited in QAUDJRN. Also, because the support requires an export and import, the support is not optimal as a real-time solution.

- Importation of database monitor entries and subsequent analysis and reporting

A database monitor (**STRDBMON**) can be used to capture SQL statements and write them to a database table. Then, the table could be imported into the Guardium collector. While this method could capture SQL statements, variables, and more, the database monitor support is primarily designed for performance analysis. The result is that a significant amount of data that is only interesting in a performance analysis context is captured resulting in the consumption of significant storage and processing resources. Also, because the support required an import, the support is not optimal as a real-time solution. This method does not provide any support for native access to database objects.

- ▶ Network monitoring to capture SQL access

Although this method can capture SQL statements in real time that flowed over a network, any SQL statements run in programs, procedures, functions, and so on, on the IBM i server cannot be monitored. This method does not provide any support for native access to database objects.

InfoSphere Guardium V9.0

The new method that is introduced in Version 9.0 of InfoSphere Guardium provides an integrated solution that overcomes the limitations of the previous methods:

- ▶ Any SQL access, whether it is initiated on a client or the IBM i server, can be captured and audited.
- ▶ Any native access that is captured in the audit journal can also be captured and sent to the InfoSphere Guardium collector.
- ▶ Both SQL access and native access are sent to the InfoSphere Guardium collector in real time.
- ▶ Much more detail is captured than is available in the audit journal, including SQL statements, variable values, client special registers, interface information, users, jobs, TCP/IP addresses, and ports. However, unlike the traditional database monitoring, only data that is interesting in a security context is captured and sent to the InfoSphere Guardium collector. This situation reduces the storage and resource consumption that is necessary.
- ▶ Filtering can be specified on the IBM i server to capture only the information that is required by auditors. For example, it is simple to set up auditing of any SQL or native access that is performed by privileged users.
- ▶ The data that is collected for InfoSphere Guardium is never written to disk on the IBM i server, providing a level of secure logging.

The new method is primarily for auditing database access. If you require auditing on a greater variety of non-database object access, the existing IBM i auditing support of exporting and importing the audit journal can still be used.

More information about the capabilities of InfoSphere Guardium can be found at:

<http://www-01.ibm.com/software/data/guardium/>

3.6 Security enhancements for printing

There are several security enhancements that are introduced in IBM i 7.1 for printing. For information about security enhancements for printing, see 11.3.2, “New QIBM_QSP_SECURITY exit point and formats” on page 447.



Backup and recovery

This chapter describes enhancements to the IBM i backup and recovery functions. This chapter describes the base operating system save and restore functions that are included in the IBM i operating system in 4.1, “New and enhanced system save and restore functions” on page 52.

After it describes the basics, this chapter describes the more advanced Backup Recovery and Media Services (BRMS) product in 4.2, “New and improved BRMS for IBM i functions” on page 55.

This chapter then addresses the new BRMS functions and capabilities added to the IBM Systems Director Navigator for i and the System Director products in 4.3, “BRMS enhancements to GUI and web interfaces” on page 84.

A list of references to more information about these topics is included at the end of the chapter.

This chapter describes the following topics:

- ▶ New and enhanced system save and restore functions
- ▶ New and improved BRMS for IBM i functions
- ▶ BRMS enhancements to GUI and web interfaces
- ▶ More resources

4.1 New and enhanced system save and restore functions

This section describes new and enhanced functions for the system save and restore functions in the IBM i 7.1 operating system. Topics include:

- ▶ Fast restore using tape position
- ▶ New ALWOBJDIF (*COMPATIBLE) restore option
- ▶ Enhanced save and restore limits
- ▶ Save While Active support for integrated VMware ESX servers
- ▶ Miscellaneous enhancements

4.1.1 Fast restore using tape position

Fast restore using tape position is a new function that enables object restore to move to the location of a stored object on tape, thus saving time by minimizing searching for the object. You can use this function to restore a single object from a backup much more quickly. There are other scenarios where restoring multiple objects might benefit from this function.

Although restore time savings vary depending on the device, media format, and position of the object on tape, tests restoring the last object from a tape that contains 1.1 million IFS objects reduced object restore time from 22 minutes to less than 3 minutes.

Save operations now track the physical media position of each saved object. This media position is a 32 hex character field in the various save commands output files.

Restores commands have a new **POSITION** parameter, which is used to specify the hexadecimal position value that appeared in the output files that were previously mentioned.

The following restore interfaces support the **POSITION** parameter:

- ▶ Restore Library (**RSTLIB**), Restore Object (**RSTOBJ**) and Restore IFS Object (**RST**) commands.
- ▶ QsrRestore and QSRRSTO application programming interfaces.
- ▶ QsrCreateMediaDefinition application programming interface to create a media definition for use by parallel restores.
- ▶ BRMS supports the **POSITION** parameter.

The default value for the **POSITION** parameter is special value named *FIRST, which restores using the current search from the beginning mode. When you use the **POSITION** (object location) parameter and value, you must also specify the **SEQNBR** parameter with the correct sequence number of the saved object.

In Example 4-1, the Restore Object (**RSTOBJ**) command restores the SYSTEMS file to the HARDWARE library. The saved object is sequence number 547 on the tape, the position of the file on tape is 0000000000190490000000AB430009CA, and the tape device name is TAP01.

Example 4-1 RSTOBJ command that specifies the POSITION parameter

```
RSTOBJ OBJ(SYSTEMS)
       SAVLIB(HARDWARE)
       DEV(TAP01)
       OBJTYPE(*FILE)
       SEQNBR(547)
       POSITION(0000000000190490000000AB430009CA)
```

4.1.2 New ALWOBJDIF (*COMPATIBLE) restore option

A new value, *COMPATIBLE, is added to the Allow Object Differences (**ALWOBJDIF**) parameter to make restores less confusing and less error-prone for database files.

Using **ALWOBJDIF (*ALL)** for database files is undesirable for the following reasons:

- ▶ When a file-level difference occurs, the original file is renamed and the saved file is restored.
- ▶ When a member level difference occurs, the existing member is renamed and the saved member is restored.

Because of the duplicated files and members, system resources are wasted and applications might produce unpredictable results. This situation leaves you with a perplexing choice between the renamed data or the restored data and leaves clean-up activities to perform.

For database objects, **ALWOBJDIF (*COMPATIBLE)** is equivalent to specifying **ALWOBJDIF (*AUTL *OWNER *PGP *FILELVL)**, which allows the following differences:

- ▶ All authorization list differences.
- ▶ All ownership differences.
- ▶ All primary group differences.
- ▶ File level differences where file level and member levels are restored *only* when the format level identifiers of the file on media match format level identifiers of the file on the system. In brief, the file formats must match.

For non-database objects, **ALWOBJDIF (*COMPATIBLE)** performs like **ALWOBJDIF (*ALL)**, which allows all object differences to be restored.

The *COMPATIBLE value for the **ALWOBJDIF** parameter is supported as follows:

- ▶ Restore Library (**RSTLIB**) and Restore Object (**RSTOBJ**) commands.
- ▶ Restore Object (**QSRRSTO**) API.
- ▶ Restore menu options that use **RSTLIB** and **RSTOBJ** commands.
- ▶ Restore Library using BRM (**RSTLIBBRM**), Restore Object using BRM (**RSTOBJBRM**), and Start Recovery using BRM (**STRRCYBRM**) commands.
- ▶ Save Restore Library (**SAVRSTLIB**), Save Restore Object (**SAVRSTOBJ**), and Save Restore Changed Objects (**SAVRSTCHG**) commands.

The following restore menu options now default to **ALWOBJDIF (*COMPATIBLE)** when you restore to another system:

- ▶ 21: Restore entire system
- ▶ 22: Restore system data only
- ▶ 23: Restore all user data

The Restore Licensed Program (**RSTLICPGM**) command now internally uses **ALWOBJDIF (*COMPATIBLE)**, but does not include it on the command interface.

4.1.3 Enhanced save and restore limits

Before IBM i 7.1, database files that contained more than 16 MB of descriptive information could not be saved. This limitation was removed, and is the last known limitation for database file save and restore.

4.1.4 Save While Active support for integrated VMware ESX servers

IBM i Integrated Server Support is enhanced to provide Save While Active (SWA) support for integrated VMware ESX servers. Storage spaces for VMware ESX servers can now be saved from IBM i while the ESX server is active. This feature allows a concurrent save of ESX data without requiring the ESX server to be shut down or applications ended.

This function is available in SF99369 - IBM i integration with BladeCenter and System x Group PTF Level 6. For more information, see the IBM i integration with BladeCenter and System x website at:

http://www-03.ibm.com/systems/i/advantages/integratedserver/iscsi/solution_guide.html

4.1.5 Miscellaneous enhancements

There are numerous miscellaneous enhancements:

- ▶ Supported Save Target (TGTRLS) releases are V7R1M0, V6R1M0, and V5R4M0.
- ▶ SAVLIB and SAVCHGOBJ of the QUSRSYS library now perform **RTVSYSINF** and save the results.

If the correct procedures are followed, this information is already saved. Often, this data is inadvertently not saved and cannot be recovered. This change ensures that the system information is saved for later recovery if needed.

This enhancement is available through PTF SI34094 for V5R4 i5/OS and PTF SI34095 for IBM i 6.1.

- ▶ The system name is now included in the headers of the various output files and spool files.
- ▶ Save and restore menu options that place the system in a restricted state are enhanced to gracefully end TCP/IP servers, host servers, and TCP/IP before ending in a restricted state.

Save menu options 21 (Save entire system), 22 (Save system data only), 23 (Save all user data), and 40 (Save all libraries other than the system library), and Restore menu options 21 (Restore entire system), 22 (Restore system data only), 23 (Restore all user data) and 40 (Restore all libraries other than the system library) now run the following commands before the **ENDSBS SBS(*ALL) OPTION(*IMMED)** command is issued:

```
- ENDTCPSVR
- ENDHOSTSVR
- DLYJOB JOB(300)
- ENDTCP
- DLYJOB JOB(300)
```

This enhancement is available through PTF SI35204 for IBM i 6.1.

- ▶ Improved serviceability

Collection Services data is now collected for specific save/restore events.

Save/restore flight recorder enhancements include a larger default size of 5 MB, more entries to be logged, and the ability to adjust the size of the data.

To change the flight recorder size to 10 MB, run the following command:

```
CRTDTAARA DTAARA(QGPL/QSRFRSIZE) TYPE(*CHAR) LEN(4) VALUE('10')
```

In this example, the size is being changed to 10 MB. The QSRFRSIZE data area can be created in either the QTEMP or QGPL library.

4.2 New and improved BRMS for IBM i functions

This section reviews new and improved functions that are part of the BRMS product:

- ▶ Support for Domino 8.5 Domino Attachment and Object Service (DAOS)
- ▶ Link list enhancements
- ▶ Support of more than 999,999 objects
- ▶ Start recovery using BRM (STRRCYBRM) command enhancements
- ▶ Recovery report enhancements
- ▶ Output file support for BRMS restores
- ▶ Ability to select Saved ASP on RSTLIBBRM and RSTOBJBRM commands
- ▶ Distributed backup support
- ▶ Maintenance enhancements
- ▶ Planning Media Movement report
- ▶ Improved clearing of volumes in *ERR status
- ▶ Protection of volumes that are marked for duplication
- ▶ Improved recovery times with new media position function
- ▶ BRMS support for the special value *COMPATIBLE for ALWOBJDIF
- ▶ Improved control with running non-BRMS saves

4.2.1 Support for Domino 8.5 Domino Attachment and Object Service (DAOS)

Domino Release 8.5 contains a new enhancement that is called the Domino Attachment and Object Service (DAOS). DAOS enables more efficient use of storage by reducing the number of redundant attachments by using a separate repository for a Domino server attachment.

Before DAOS, attachments were part of each Domino database (.nsf) file. If a large attachment is sent to 40 mail users, there are 40 occurrences, one in each mail file.

With DAOS, attachments that exceed a configured size are pulled out of the .nsf files and are placed as objects. In the example that is described in the previous paragraph, rather than one occurrence of the attachment being stored in each mail file, there is one NLO stored per Domino server, thus saving storage space.

BRMS is enhanced to handle the NLO objects as follows:

- ▶ During Online Domino backups, the presence of NLO objects is determined and are automatically saved at the end of the SAVDOMBRM backups. No BRMS changes are necessary.
- ▶ A full save includes the NLO objects.
- ▶ An incremental save includes the new and changed NLO objects since the last full save.

BRMS DAOS support was made available through PTFs for V6R1 (SI34918) and V5R4 (SI31916).

When you configure DAOS on Domino servers, be careful with the configuration of attachment sizes that are externalized into NLOs. If you select a small size, many NLO objects can be created, each of which is an IFS object that can significantly lengthen the IFS backup time. The default is 4096, but consider using 1,000,000 or larger.

DAOS references

The following references provide more information about DAOS:

- ▶ DAOS Quick Start Guide:
<http://www.lotus.com/1dd/dominowiki.nsf/dx/daos-quick-start-guide>
- ▶ DAOS Best Practices:
<http://www.lotus.com/1dd/dominowiki.nsf/dx/daos-best-practices>
- ▶ DAOS Estimator:
<http://www.ibm.com/support/docview.wss?rs=463&uid=swg24021920>
- ▶ BRMS Online Lotus Server Backup Reference:
<http://www-03.ibm.com/systems/i/support/brms/domdaos.html>

4.2.2 Link list enhancements

BRMS has enhanced link lists for ease of use and to ensure complete backups.

The QIBMLINK link list for IBM IFS directories is now automatically added to the supplied system backup control group *SYSGRP for new installations only. In V5R4 i5/OS and IBM i 6.1, QIBMLINK existed, but was not automatically added to *SYSGRP. You should add QIBMLINK manually to *SYSGRP in existing installations. QIBMLINK is used to save system IFS files and directories.

List QIBMLINK includes the following directories:

- ▶ /QIBM/ProdData
- ▶ /QOpenSys/QIBM/ProdData

The QALLUSRLNK link list was added in IBM i 7.1. QALLUSRLNK is used to save user IFS directories and files. QALLUSRLNK is used with the QIBMLINK link list. QALLUSRLNK omits the following directories:

- ▶ /QSYS.LIB
- ▶ /QDLS
- ▶ /TMP/BRMS
- ▶ /QIBM/ProdData
- ▶ /QOpenSys/QIBM/ProdData

Usage of QIBMLINK followed by QALLUSRLNK enables more granularity than the *LINK control group entry and ensures that IBM directories are restored before user directories if a system restore is necessary. The usage of the QALLUSRLNK link list with the QIBMLINK link list also avoids the duplication of saved data that occurs with the combination of using QIBMLINK and *LINK.

4.2.3 Support of more than 999,999 objects

To support more than 999,999 objects, BRMS has an enhanced WRKMEDIBRM panel and QP1ARCY recovery report.

The **WRKMEDIBRM** command previously could not show more than 999,999 objects in the saved objects field. In IBM i 7.1, if more than 999,999 objects or files are saved in a single library or save command, BRMS lists the actual number rather than 999,999 objects on the WRKMEDIBRM Object Detail panel.

Figure 4-1 shows a WRKMEDIBRM Object Detail panel. The circled field shows a saved item with more than 999,999 objects.

work with Media Information									
Position to Date									
Type options, press Enter.									
2=Change 4=Remove 5=Display 6=work with media 7=Restore									
9=work with saved objects ...									
Opt	Saved Item	Save Date	Save Time	Item Type	-----objects----- Saved	Not Saved	Control Group	Obj Dt1	
—	*LINK	3/20/10	2:24:39	*LNK	1222795		5 BKUP6	*NO	
—	QUSRSYS	3/20/10	3:15:50	*LIB	2966		1 BKUP6	*YES	
—	QUSRBRM	3/20/10	3:17:16	*LIB	241		0 BKUP6	*YES	
—	Q1WWT	3/20/10	3:17:18	*LIB	7543		0 BKUP6	*YES	
—	QIJS	3/20/10	3:20:49	*LIB	566		0 BKUP6	*YES	
—	QUSRIJS	3/20/10	3:21:03	*LIB	106		0 BKUP6	*YES	
—	QPFRDATA	3/20/10	3:21:05	*LIB	74		43 BKUP6	*YES	
—	RSM773LIB	3/20/10	3:42:30	*LIB	6		0 BKUP6	*YES	
—	Q1WWTLCL	3/20/10	3:42:31	*LIB	80		0 BKUP6	*YES	
—	QUSRBRM	3/20/10	4:01:17	*OBJ	19		0 *NONE	*YES	
									Bottom
F3=Exit F5=Refresh F11=Duplication status F12=Cancel									
F23=More options									

Figure 4-1 WRKMEDIBRM Object Detail panel with more than 999,999 objects

The BRMS recovery report QP1ARCY previously could not show more than 999,999 in the saved objects column. In IBM i 7.1, if more than 999,999 objects or files are saved in a single library or save command, BRMS lists the actual number, rather than 999,999 objects on the BRMS recovery report QP1ARCY.

Figure 4-2 shows an excerpt of the BRMS recovery report. The circled data shows greater than 999,999 objects were saved for the *LINK item.

Display Spooled File										Page/Line	20/49
File	:	:	:	:	:	:	:	:	:	Columns	1 - 130
Control	:	:	:	:	:	:	:	:	:		
Find	:	:	:	:	:	:	:	:	:		
*****1*****2*****3*****4*****5*****6*****7*****8*****9*****0*****1*****2*****3*****4*****5*****6											

Figure 4-2 BRMS recovery report showing more than 999,999 objects saved

4.2.4 Start recovery using BRM (STRRCYBRM) command enhancements

The STRRCYBRM command was enhanced to override specific recovery elements so that they use another time period, which requires that the override recovery element (**VERRIDE**) parameter be set to *YES. This action affects the **ACTION** parameter values of *REPORT and *RESTORE.

The following recovery elements can be selected for override:

► ***ALLSYS**

This element specifies that you want to recover all the system data, which includes *SAVSYS, *SECDTA, and *SAVCFG.

► ***SAVSYS**

This element specifies that you want to recover the operating system that is based on the BRMS media content information.

► ***SECDTA**

This element specifies that you want to recover the security data.

► ***SAVCFG**

This element specifies that you want to recover the configuration data.

The **STRRCYBRM** command keywords that enable overriding recovery elements are as follows:

► **OVERRIDE**

This keyword specifies whether you want to use another time period for a specific recovery element.

– ***NO**

This keyword indicates that you do not want to specify another date and time range for a specific recovery element. Recovery elements and overrides are ignored if specified.

– ***YES**

This keyword indicates that you want to specify another date and time range for a specific recovery element.

► **RCYELEM**

This keyword specifies a recovery element and its override time period. You can specify as many as three sets of them. Each set has the following parameters:

– Recovery element

• ***ALLSYS**

This element specifies that you want to recover the entire system using an override. If this element is selected, you cannot specify other recovery elements.

• ***SAVSYS**

This element specifies that you want to recover the operating system using an override.

• ***SECDTA**

This element specifies that you want to recover the security data using an override.

• ***SAVCFG**

This element specifies that you want to recover the configuration data using an override.

– Beginning time

This parameter specifies the beginning time at which or after recovery items are included. Any items that are created before the specified time and date are not included in the items that are selected for recovery.

This parameter uses the BRMS standard methods for specifying times.

- Beginning date

This parameter specifies the beginning date on or after which the recovery items must be saved. Any entries that are saved before the specified date are not included in the recovery.

This parameter uses the BRMS standard methods for specifying dates.

- Ending time

This parameter specifies the ending time before which recovery items are included. Any items that are created after the specified time and date are not included in the recovery items that are selected for recovery.

This parameter uses the BRMS standard methods for specifying times.

- Ending date

This parameter specifies the ending date on which or before which the recovery items must be saved. Any recovery items that are created after the specified date are not included in the recovery operation.

This parameter uses the BRMS standard methods for specifying dates.

In Example 4-2, the **STRRCYBRM** command selects all restore items that are found regardless of time, except for the operating system restore items, which selects nothing newer than 6 p.m. on 03/01/2010 because of the *SAVSYS override.

*Example 4-2 STRRCYBRM command using recovery element override of *SAVSYS*

```
STRRCYBRM PERIOD((*AVAIL *BEGIN) (*AVAIL *END))  
            OVERRIDE(*YES)  
            RCYELEM((*SAVSYS ((*AVAIL *BEGIN) ('18:00:00' '03/01/10'))))
```

In Example 4-3, the **STRRCYBRM** command selects all restore items found up to and including 03/01/2010, except for security data and configuration data, which is restored through the current date.

*Example 4-3 STRRCYBRM command using recovery element override of *SECDTA and *SYSCFG*

```
STRRCYBRM PERIOD((*AVAIL *BEGIN) (*AVAIL '03/01/10'))  
            RCYELEM((*SECDTA ((*AVAIL *BEGIN) (*AVAIL *END)))  
                    (*SYSCFG ((*AVAIL *BEGIN) (*AVAIL *END))))
```

When overrides are specified, the recovery report QP1ARCY has an attention block that notes the override, which is an indicator to you that the latest saved data is not being used, as shown in Figure 4-3.




Figure 4-3 BRMS recovery report with override attention block

4.2.5 Recovery report enhancements

The “Verify System Information” step in the recovery report is updated to include new instructions to include the Update System Information (**UPDSYSINF**) command now that the system information is now saved with the QUSRBRM library. The new instructions are in the red circled area of the recovery report excerpt that is shown in Figure 4-4.




Figure 4-4 BRMS recovery report with new UPDSYSINF instructions

The “Restoring User Profiles” step is updated to include the new Allow Object Differences *COMPATIBLE value, replacing the *ALL value of previous releases, as shown in the red circle in Figure 4-5.

Display Spooled File

```

File . . . . . : QP1ARCY
Control . . . . : █
Find . . . . . :
*...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8...+...9...+
STEP 009 : Recover User Profiles
  Start date/time _____ Stop date/time _____ Duration _____
  You should restore a current version of your user profiles.
  To do so, type the following command and press "Enter".
    STRRCYBRM OPTION(*SYSTEM) ACTION(*RESTORE)
  ----- Attention -----
  Press F9 on the Select Recovery Items display to go to the Restore
  Command Defaults display.
  ___ Ensure the tape device name or media library device name is
    correct for the Device prompt.
  ___ Ensure *SAVLIB is specified for the Restore to library prompt.
  ___ Ensure *SAVASP is specified for the Auxiliary storage pool prompt.
  If you are recovering to a different system or a different logical
  partition, you must specify the following:
  ___ *ALL for the Data base member option prompt.
  ___ *COMPATIBLE for the Allow object differences prompt.
  ___ *NONE for the System resource management prompt.
  ___ Select *NEW for Restore spooled file data to recover saved
    spooled files concurrently with restored output queues.

F3=Exit  F12=Cancel  F19=Left  F20=Right  F24=More keys
  
```

Figure 4-5 BRMS recovery report updates for support of ALWOBJDIF(*COMPATIBLE) special value

4.2.6 Output file support for BRMS restores

IBM i 7.1 BRMS supports the option to specify an output file when restoring data is saved through BRMS. The *OUTFILE value can be specified with the **OUTPUT** parameter in these functions:

- BRMS recovery policy

Figure 4-6 shows the third panel of the Change Recovery Policy function.

Change Recovery Policy

RCHASA04

Type choices, press Enter.

Key store file	*SAV _____	Name, *SAV, *NONE
Key store library	*SAV _____	Name, *SAV, *NONE
Restore private authorities.	*NO _____	*NO, *YES

Output options for recoveries:

Output.	*NONE _____	*NONE, *OUTFILE
File to receive output.	_____	Name
Library.	*LIBL _____	Name, *LIBL, *CURLIB

Output member options:

Member to receive output	*FIRST _____	Name, *FIRST
Replace or add records	*REPLACE _____	*REPLACE, *ADD

Output options for directory object recoveries:

Output.	*NONE _____	

Bottom

F3=Exit F4=Prompt F5=Refresh F9=System policy
F12=Cancel

Figure 4-6 OUTPUT support parameters in the BRMS Change Recovery Policy function

Note the following information in the red box in Figure 4-6 on page 61:

The new “Output options for recoveries” section:

- The new “Output” parameter
- The new “File to receive output” and “Library” fields
- The new “Output member options” fields

The new “Output options for directory object recoveries” section:

- The new “Output” parameter

This parameter uses IFS directory syntax for the output field.

► BRMS Restore commands

- Restore Library using the BRM (**RSTLIBBRM**) command
- Restore Object using the BRM (**RSTOBJBRM**) command
- Restore (IFS) Object using the BRM (**RSTBRM**) command

Figure 4-7 shows the new output keywords of the **RSTLIBBRM** command. The **RSTOBJBRM** command is not shown because the keywords are identical. Note the following items in the red box:

- The new **OUTPUT** keyword
- The new **OUTFILE** keyword
- The new **OUTMBR** keyword

The example in Figure 4-7 places the restore output in member RS03012010 of the LIBS file in the RESTORE library.

```
Restore Library using BRM (RSTLIBBRM)

Type choices, press Enter.

Private authorities . . . . . PVTAUT          *NO
Restore to library . . . . . RSTLIB          *SAVLIB
Auxiliary storage pool . . . . . RSTASP       *SAVASP
From system . . . . . FROMSYS             *LCL

Additional Parameters

Output . . . . . OUTPUT          > *OUTFILE
File to receive output . . . . . OUTFILE    > LIBS
Library . . . . .                > RESTORE
Output member options:      OUTMBR
  Member to receive output . . . . .        > RS03012010
  Replace or add records . . . . .          > *ADD

F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys
```

Figure 4-7 RSTLIBBRM command panel with new output keywords

- In the recovery defaults of the Work with Media Information (**WRKMEDIBRM**) command.
The **WRKMEDIBRM** recovery defaults apply only to the current session and are not permanent.

The following sequence of **WRKMEDIBRM** panels opens the Recovery Defaults panel, which contains the new output parameters. From the Work with Media Information (**WRKMEDIBRM**) panel, select option 7 (Restore) for the library to be restored, as shown in Figure 4-8.

Work with Media Information						RCHASA04		
Position to Date _____								
Type options, press Enter.								
2=Change 4=Remove 5=Display 6=Work with media 7=Restore								
9=Work with saved objects ...								
Opt	Saved Item	Save Date	Save Time	Save Type	Parallel Devices	Volume Serial	File Sequence	Expire Date
—	QIJS	4/03/10	2:59:05	*FULL		LT0501	339	7/12/10
—	QUSRIJS	4/03/10	2:59:19	*FULL		LT0501	340	7/12/10
—	QPFRDATA	4/03/10	2:59:22	*FULL		LT0501	341	7/12/10
—	Q1WWTLC	4/03/10	3:00:13	*FULL		LT0501	343	7/12/10
—	RSM773LIB	4/03/10	3:00:13	*FULL		LT0501	342	7/12/10
—	*ALLDLO	4/03/10	3:39:40	*FULL		LT0501	344	7/12/10
—	*LINK	4/03/10	3:40:23	*FULL		LT0501	345	7/12/10
—	QUSRBRM	4/03/10	4:26:02	*QBRM		LT0501	346	7/12/10
7	MSRJJ	4/05/10	9:20:19	*FULL		LT0766	1	4/07/10
—	QUSRSYS	4/05/10	10:58:03	*FULL		LT0766	2	4/07/10
Bottom								
F3=Exit F5=Refresh F11=Object detail F12=Cancel								
F23=More options								

Figure 4-8 **WRKMEDIBRM** navigation to Recovery Defaults panel - step 1

This option causes the Select Recovery Items panel to be displayed, as shown in Figure 4-9. On the Select Recovery Items panel, F9 (circled) displays the Recovery Defaults panel.

Select Recovery Items								RCHASA04	
Type options, press Enter.								Select action : <u>*ALL</u>	
								Select volume : _____	
1=Select 4=Remove 5=Display 7=Specify object									
Opt	Item	Saved Date	Save Time	Save Type	Parallel Devices	Volume Serial	File Sequence	Expire Date	
1	MSRJJ	4/05/10	9:20:19	*FULL		LT0766	1	4/07/10	
Bottom F3=Exit F5=Refresh F9=Recovery defaults F11=Object View F12=Cancel F14=Submit to batch F16=Select									

Figure 4-9 WRKMEDIBRM navigation to Recovery Defaults panel - Step 2

Figure 4-10 is the Recovery Defaults panel. The various output selection fields are in the box. The output fields are nearly identical to the output keywords of the **RSTLIBBRM** command.

Restore Command Defaults	
Type information, press Enter.	
Restore private authorities	<u>*NO</u> *NO, *YES
Output options for recoveries:	
Output	<u>*NONE</u> *NONE, *OUTFILE
File to receive output	_____ Name
Library	<u>*LIBL</u> Name, *LIBL, *CURLIB
Output member options:	
Member to receive output.	<u>*FIRST</u> Name, *FIRST
Replace or add records.	<u>*REPLACE</u> *REPLACE, *ADD
Output options for directory object recoveries:	
Output	<u>*NONE</u> _____
Bottom F12=Cancel	

Figure 4-10 WRKMEDIBRM Recovery Defaults panel

- Start Recovery using the BRM (**STRRCYBRM**) command when the ***RESTORE** action is also specified.

The **WRKMEDIBRM** recovery defaults apply only to the current session and are not permanent.

A sequence of **STRRCYBRM** display panels shows how to access the output parameters and the output parameters themselves.

Figure 4-11 shows the Select Recovery Items panel, which is generated by the **STRRCYBRM** command. Pressing F9 on this panel (circled) displays the Restore Command Defaults panel.

RCHASA04

Select Recovery Items

Select action : *ALL
 Select volume : _____

Type options, press Enter.

1=Select 4=Remove 5=Display 7=Specify object -

Opt	Item	Saved Date	Save Time	Save Type	Parallel Devices	Volume Serial	File Sequence	Expire Date
-	#LIBRARY	3/24/10	1:21:23	*FULL		LT0894	6	4/13/10
-	ADTSLAB	3/24/10	1:21:23	*FULL		LT0894	7	4/13/10
-	AGBTEMP	3/24/10	1:21:23	*FULL		LT0894	8	4/13/10
-	ALRMSETUP	3/24/10	1:21:23	*FULL		LT0894	9	4/13/10
-	APSS.COMMT	3/24/10	1:21:23	*FULL		LT0894	10	4/13/10
-	APSS.TIVD	3/24/10	1:21:23	*FULL		LT0894	11	4/13/10
-	APSS.UPST	3/24/10	1:21:23	*FULL		LT0894	12	4/13/10
-	APSSF010	3/24/10	1:21:23	*FULL		LT0894	13	4/13/10
-	APSSQ010	3/24/10	1:21:23	*FULL		LT0894	14	4/13/10
-	APSSP010.O	3/24/10	1:21:23	*FULL		LT0894	15	4/13/10
-	APSSQSYS	3/24/10	1:21:23	*FULL		LT0894	16	4/13/10

More...

F9=Recovery defaults

F3=Exit F5=Refresh F9=Recovery defaults F11=Object View
 F12=Cancel F14=Submit to batch F16=Select

Figure 4-11 **STRRCYBRM** Select Recovery Items panel

On the Restore Command Defaults panel, scroll to the third window, which is shown in Figure 4-12. This panel shows the new output parameters. The output parameters on this panel look and act like the fields in the Change Recovery Policy panel that is shown in Figure 4-6 on page 61, except that they apply only to the current session.

```

Restore Command Defaults

Type information, press Enter.

Restore private authorities . . . . *NO          *NO, *YES

Output options for recoveries:
Output . . . . . *OUTFILE          *NONE, *OUTFILE
File to receive output . . . . . LIB          Name
Library . . . . . RESTORES        Name, *LIBL, *CURLIB
Output member options:
Member to receive output. . . . . AL03202010  Name, *FIRST
Replace or add records. . . . . *REPLACE      *REPLACE, *ADD

Output options for directory object recoveries:
Output . . . . . '/RESTORES/AL03202010'

F12=Cancel
Bottom

```

Figure 4-12 STRRCYBRM Recovery Defaults panel - Restore options

4.2.7 Ability to select Saved ASP on RSTLIBBRM and RSTOBJBRM commands

You can use this new function to specify the auxiliary storage pool (ASP) from which a library is saved when you perform a restore using the Restore Library using BRM (**RSTLIBBRM**) command and the Restore Object using BRM (**RSTOBJBRM**) command.

Suppose that a system has a system ASP and three independent ASPs. Each of the ASPs has library TOOLS and the entire system, including the IASPs, was saved. There are three saves of library TOOLS. You can use this function to select which of the saves to restore.

The new keyword is **SAVASP**. Values for the parameters for the **RSTLIBBRM** command are as follows:

- ▶ ***ANY**
The library and objects that are saved is restored from any ASPs save. This value is the default value, which works as it did before IBM i 7.1.
- ▶ ***SYSTEM**
The saved library and objects are restored from the system ASP save.
- ▶ **ASP number 1 - 32**
The library and objects are restored from the specified user ASP, or the system ASP if 1 is specified.
- ▶ **ASP name**
The library and objects are restored from the specified ASP save.

The keyword and values for the **RSTOBJBRM** command are identical. The function is identical except that only objects are restored.

There are limitations about which objects can be restored to non-system ASPs. These objects are not allowed to be in user or independent ASPs.

The red circle in Figure 4-13 shows the **SAVASP** keyword for the **RSTLIBBRM** command.

Restore Library using BRM (RSTLIBBRM)

Type choices, press Enter.

Library	> <u>TOOLS</u>	Name, *MEDINF
	+ for more values	
Device	> <u>MLD50</u>	Name, *MEDCLS
	+ for more values	
Parallel device resources:		
Minimum resources	<u>*SAV</u>	1-32, *SAV, *NONE, *AVAIL
Maximum resources	<u>*MIN</u>	1-32, *MIN, *AVAIL
Save level	<u>*CURRENT</u>	1-99, *CURRENT, *SAVDATE
Save level time reference:		
Save date		Date
Save time	<u>*LATEST</u>	Time, *LATEST
Saved auxiliary storage pool . .	> <u>TESTASP</u>	Name, 2-32, *ANY, *SYSTEM
End of media option	<u>*REWIND</u>	*REWIND, *LEAVE, *UNLOAD
Option	<u>*ALL</u>	*ALL, *NEW, *OLD, *FREE
Database member option	<u>*MATCH</u>	*MATCH, *ALL, *NEW, *OLD
Spooled file data	<u>*NEW</u>	*NEW, *NONE
More...		
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display		
F24=More keys		

Figure 4-13 Save ASP (SAVASP) keyword of the BRMS RSTLIBBRM command

4.2.8 Distributed backup support

New options were created for BRMS networking support to control distribution of backup history. In IBM i 6.1, all systems in a BRMS network received all history from other systems in the network.

In IBM i 7.1, there are now options to configure which systems receive information about backups and which do not. Distributed backup support reduces the data on systems that have no need to know about the saved history from other systems. The distributed backup function is available through the Change Network Group menu option of the BRMS System Policy (**BRMSSYSPCY**) menu.

For IBM i 6.1, the Change Network Group panel (see Figure 4-14) could not be used to receive media information. If the system was in the network group, it received the media information.

Change Network Group

RCHASA24 USIBMZQ

Network group : *MEDINV

Position to : _____

Text : Centralized media network systems

Receive media info . . : *NONE

*NONE, *LIB

Enable for TCP/IP . . : *YES

*NO, *YES

Notify period : 30

30-99999 seconds

Type options, press Enter.

1=Add 4=Remove 8=Set time

Opt	Location Name	Remote Network ID	Receive Media Info	System Status	Network Status
-	RCHASA25	USIBMZQ	*NONE	Online	Active
-	RCHASA75	USIBMZQ	*NONE	Online	Active

Bottom

F3=Exit

F5=Refresh

F12=Cancel

Figure 4-14 Change Network Group panel from IBM i 6.1

In IBM i 7.1, the Change Network Group panel now has a function key (F11) that displays a new BRMS Media Information panel.

The Change Network Group panel in Figure 4-15 no longer shows the Receive Media Information column.

Change Network Group

RCHASA04 USIBMZQ

Network group : *MEDINV

FlashCopy state : *ENDPRC

Communication method . . : *AVAIL

Notify period : 30

Type options, press Enter.

1=Add 4=Remove 8=Set time

Position to : _____

*AVAIL, *IP, *SNA, *NONE

30-99999 seconds

	Remote	Local	Remote	System	Network
Opt	Location	Name	Network ID	Status	Status
-					
-	RCHASA10		USIBMZQ	Online	Active
-	RCHASA11		USIBMZQ	Online	Active

Bottom

F3=Exit

F5=Refresh

F11=BRMS Media Information

F12=Cancel

Figure 4-15 Change Network Group panel from IBM i 7.1

The new BRMS Media Information panel replaced the single column Receive Media Info value with a three column set of fields, as shown in Figure 4-16.

Change Network Group

RCHASA04 USIBMZQ

Network group : *MEDINV
 Position to : _____

FlashCopy state . . . : *ENDPRC

Communication method : *AVAIL
 *AVAIL, *IP, *SNA, *NONE

Notify period : 30
 30-99999 seconds

Type options, press Enter.

1=Add 4=Remove 8=Set time

-----BRMS Media Information-----

Opt	Remote Location	Local Name	Remote Network ID	Local Receives	Remote Receives	Remote Overrides
-	RCHASA10		USIBMZQ	*LIB	*NONE	*NO
-	RCHASA11		USIBMZQ	*NONE	*NONE	*NO

Bottom

F3=Exit F5=Refresh F11=Main Network Panel F12=Cancel

Figure 4-16 Change Network Group BRMS Information panel

Some of the fields in this panel are:

- Local Receives field

This field specifies whether media content information, and at what level, is received by the current system from the specified remote system. Media content information represents files that are on a media volume. You can specify whether media content information is shared with the current system or specify that the current system is not to receive any media content information from the specified remote system in the network group. Media inventory information is always shared between systems that are members of the network group. Possible values are as follows:

 - *NONE

Media content information is not received from the specified remote system on the current system.
 - *LIB

Media content information is received from the specified remote system on the current system. The information includes library level information only.
- Remote Receives field:

This value displays whether the remote system is to receive media content information, and at what level, from the current system displayed.

To change this value, you must go to that remote system and change the value on that system.

Possible values are as follows:

- *NONE

Media content information is not sent to the specified remote system.

- *LIB

Media content information is sent to the specified remote system. The information includes library level information only.

► Remote overrides field

This field represents media information overrides that are in place for the remote system for IASP high availability support that is described in 4.3.10, “High availability support for independent ASPs in a BRMS network” on page 115.

This field cannot be updated from the character-based interface, and must be set through IBM Systems Director, IBM Systems Director Navigator for i, or System i Navigator.

4.2.9 Maintenance enhancements

BRMS added enhancements to enable more concurrent maintenance than was allowed in previous releases. When you perform media moves through the Start Maintenance for BRM (**STRMNTBRM**) command, you now receive a BRM6717 Volume (volume-ID) was not moved to location warning message when volumes are in use and cannot be moved. The volume move report also lists an alert that the volume was in use.

Suppose that maintenance is running and a second job issues a BRMS command that attempts to use files in the QUSRBRM library that is used by the maintenance job. In this case, a BRM6714 Job (*job-name*) is being held by job (*maintenance-job-name*) message is issued to that second job’s message queue and is displayed (Figure 4-17).

```
Command Entry                                     RCHASA04
                                                    Request level: 1

Previous commands and messages:

(No previous commands or messages)

Type command, press Enter.
===> WRKMEDBRM

F3=Exit   F4=Prompt   F9=Retrieve   F10=Include detailed messages
F11=Display full   F12=Cancel   F13=Information Assistant   F24=More keys
Job QPADEV000B/MSRJJ/374662 is being held by job QPADEV000F/MSSSJ/274660.
```

Figure 4-17 BRMS6714 message that is displayed

When maintenance uses the BRMS files, the held job is sent a BRM6716 BRM restricted procedure ended message. This message is logged in the held job's message queue but is not displayed. The BRM6714 message is no longer displayed, indicating that the job is continuing.

The job that is running the **STRMNTBRM** command during the period where maintenance requires exclusive use of the BRMS files lists, but does not display, message BRM6715 BRM restricted procedure started and message BRM6716 BRM restricted procedure ended, as shown in Figure 4-18.

```

                                Display All Messages
Job . . : JACOBSENB1      User . . : MSRJJ      System: RCHASA04
                                           Number . . . : 374589

-      - RETURN          /* RETURN due to end of CL program */
23800 - RTVSYSVAL SYSVAL(QATNPGM) RTNVAR(&SYSPROG)
24600 - SETATNPGM PGM(QSYS/QEZMAIN)
25600 - CALL PGM(D510A1RLE)          /* The CALL command contains
parameters */
> strmntbrm
Library QBRM added to library list.
Journal receiver QJR1AC1553 created in library QUSRBRM.
Journal receivers QJR1AC1552 and *N detached.
Sequence number not reset. First sequence number is 3124.
BRM restricted procedure started.
BRM restricted procedure ended.
Object Q1AMNTHALT in QUSRBRM type *DTHHRA deleted.
BRM restricted procedure started.

More...

Press Enter to continue.

F3=Exit  F5=Refresh  F12=Cancel  F17=Top  F18=Bottom

```

Figure 4-18 BRM restricted procedure messages in the STRMNTBRM job message queue

In a typical maintenance run, you might see several pairs of these messages.

4.2.10 Planning Media Movement report

You can now print a report that can help you plan future moves of BRMS media.

The Print Media Movement (**PRTMOVBRM**) command has a new ***NEXT** value on its **TYPE** parameter. **TYPE(*NEXT)**, combined with a future date specified in the select date (**SLTDATE**) parameter, generates a report of future media moves.

In Figure 4-19, the Select dates parameters are set to generate the media movement report with moves starting on the current date for seven days. The **TYPE** parameter is set to ***NEXT**. This command invocation generates a report that lists all media moves to the next location that are to occur between today and the next seven days.

```

Print Media Movement (PRTMOVBRM)

Type choices, press Enter.

Select dates:
From date . . . . . > *CURRENT      Date, *BEGIN, *CURRENT, nnnnn
To date   . . . . . > 7                Date, *END, *CURRENT, nnnnn
Type      . . . . . > *NEXT           *ALL, *NEXT, *NOTVFY, *VFY
From location .....
:                                     Type (TYPE) - Help      :
:                                     :                       :
:      *NEXT                                     :
:      Produces the Media Movement report for the next      :
:      location that media will move to rather than completed :
:      moves. The *NEXT option will include all media      :
:      movement to the next location. Ending dates should be :
:      future dates to assure that report results are      :
:      meaningful for the date range selected.              :
:                                                           :
:                                                           :
:                                                           :
:                                                           :
:                                                           :
:                                                           :
:                                                           :
:                                                           :
:      F2=Extended help   F10=Move to top   F11=Search Index   :
F3=Exit  F4=: F12=Cancel   F20=Enlarge     F24=More keys      :
F24=More keys :
:.....

```

Figure 4-19 Print Media Movement panel using the **TYPE** parameter value of ***NEXT**

The **PRTMOVBRM SLTDATE(06010 06150) TYPE(*NEXT)** command prints media moves to occur between June 1, 2010 and June 15, 2010. Note the special date format that is used by BRMS.

4.2.11 Improved clearing of volumes in *ERR status

Improvements were made to simplify and speed the clearing of tape volumes in error status (***ERR**). Prior releases required you to remove the media volume from inventory, add the media volume back in inventory, and then initialize the media.

In IBM i 7.1, instead of the removing and adding the media volume, you can clear the error status using new option 9 (Remove volume error status) of the Work with Media using the BRM (**WRKMEDBRM**) command. This action changes the media volume status to Requires Initialization (*INZ). See Figure 4-20.

Work With Media

System: RCHASA04

Position to Starting characters

Type options, press Enter.

1=Add

2=Change

4=Remove

5=Display

6=Work with serial set

7=Expire

8=Move

9=Remove volume error status

10=Reinitialize ...

Opt	Volume Serial	Status	Creation Date	Expiration Date	Location	Move Date	Media Class	Dup Sts
—	JA3860	*EXP	03/09/10	*NONE	AVAIL	*NONE	AVAIL	
—	JA3861	*EXP	03/09/10	*NONE	AVAIL	*NONE	AVAIL	
—	JA3862	*EXP	03/09/10	*NONE	AVAIL	*NONE	AVAIL	
—	JA3863	*EXP	03/09/10	*NONE	AVAIL	*NONE	AVAIL	
—	JA3864	*EXP	03/09/10	*NONE	AVAIL	*NONE	AVAIL	
—	JA3865	*EXP	03/09/10	*NONE	AVAIL	*NONE	AVAIL	
—	JA3866	*EXP	03/09/10	*NONE	AVAIL	*NONE	AVAIL	

More...

Parameters or command:

==>

F3=Exit

F4=Prompt

F5=Refresh

F11=Volume system

F12=Cancel

F17=Top

F18=Bottom

F23=More options

(C) COPYRIGHT IBM CORP. 1998, 2007. ALL RIGHTS RESERVED.

Figure 4-20 WRKMEDBRM Work with Media Panel with Remove volume error status option

When in the *INZ status, the media volume can be reinitialized by running one of the following commands

- ▶ Work with Media using BRM (**WRKMEDBRM**) command option 10
- ▶ Work with Media Library Media (**WRKMLMBRM**) command option 5
- ▶ Initialize Media using BRM (**INZMEDBRM**) command

You must verify that the media volume is still usable.

This new function is also available through the IBM System Director Navigator for i web interface and IBM i Access graphical user access (GUI) client.

4.2.12 Protection of volumes that are marked for duplication

BRMS no longer expires media volumes marked for duplication.

When media is marked for duplication, BRMS no longer expires the media when the Start Maintenance for BRM (**STRMNTBRM**) command, the Start Expiration for BRM (**STREXPBRM**) command, or the Work with Media using BRM (**WRKMEDBRM**) command option 7 (expire) is run.

The BRM expiration (QP1AEP) report lists warning messages for each volume that did not expire, as shown in Figure 4-21.

```
File .....: QP1AEP                               Page/Line 1/2
Control .....                               Columns 1 - 130
Find .....

*...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8...+...9...+...0...+...1...+...2...+...3

5770BR1 IBM i 7.1 100416      Media Expiration Report      BRMSSYS 3/06/10 3:48:40      Page
Volume              Creation              Media Files
Serial  System      Date      Location  User      Use Count  ACT  EXP
CVT021 BRMSSYS      2/23/10  CVT2      BRMSOPR      3      13      0 Pending volume duplication, cannot expire.
LN4807 BRMSSYS      2/24/10  NET3590    BRMSOPR     18      0     394 Pending volume duplication, cannot expire.
(No volumes expired)

Volumes warned .....:      0
Volumes expired .....:      0
Previously expired .....:    213
Total expired count .....:    213

***** END OF LISTING *****

Bottom

F3=Exit F12=Cancel F19=Left F20=Right F24=More keys
```

Figure 4-21 Warning messages in the BRMS expiration report

When you try to expire a volume marked for duplication through **WRKMEDBRM** option 7, a BRM0010 message (Volume *vol-ID* cannot be expired.) with reason 2 (The volume has been marked for duplication.) is displayed, as shown in Figure 4-22.


System: RCHASA04

Position to _____ Starting characters

Type options, press Enter.

1=Add 2=Change 4=Remove 5=Display 6=Work with serial set 7=Expire
 8=Move 9=Remove volume error status 10=Reinitialize ...

Opt	Volume Serial	Status	Creation Date	Expiration Date	Location	Move Date	Media Class	Dup Sts
—	LT0763	*EXP	01/20/10	02/09/10	MLD50	*NONE	ULTRIUM1	
7	LT0764	*ACT	12/15/09	03/25/10	MLD50	03/04/02	ULTRIUM1	1
—	LT0765	*EXP	12/04/09	03/14/10	MLD50	*NONE	ULTRIUM1	
—	LT0766 +	*ACT	01/31/08	05/09/09	MLD50	*NONE	ULTRIUM1	

Bottom 

Parameters or command:

==> _____

F3=Exit F4=Prompt F5=Refresh F11=Volume system F12=Cancel F17=Top
 F18=Bottom F23=More options

Volume LT0764 cannot be expired. 

Figure 4-22 Error message when you attempt to expire a volume marked for duplication

4.2.13 Improved recovery times with new media position function

BRMS supports the new media position function that is described in 4.1.1, “Fast restore using tape position” on page 52.

The media position function is automatically and invisibly used by BRMS, but requires object level detail (*YES, *OBJ, *MBR) specified for the saved items in the control group, or on the Save Library using BRM (**SAVLIBRM**) command. BRMS saves retain the media positions in the BRMS database files and BRMS restores retrieve the media positions from the BRMS database files.

The media position function is not supported on the following options:

- ▶ **WRKMEDIBRM** option 7 (restore) or 7 (specify object)
- ▶ **WRKMEDIBRM** option 9 (Work with Saved Objects) or option 9 (Work with Saved Members), 7 (Restore), or 7 (Specify object)

4.2.14 BRMS support for the special value *COMPATIBLE for ALWOBJDIF

BRMS restore functions support the *COMPATIBLE special value for the **ALWOBJDIF** parameter described in 4.1.2, “New ALWOBJDIF (*COMPATIBLE) restore option” on page 53.

4.2.15 Improved control with running non-BRMS saves

When you run non-BRMS saves using non-BRMS volumes on a system with BRMS, the save job no longer ends if a BRMS volume is loaded. BRMS posts a BRM1730 BRMS enrolled volume rejected message, giving you the option to load another non-BRMS volume.

4.2.16 RDX and BRMS

RDX (a disk-based technology) drives are positioned for the lower end of performance and capacity requirements (faster than a DAT tape, but slower than LTO). The available capacities range from 160 GB to 1 TB per removable cartridge. RDX is available natively attached either on USB or SATA drives. The SATA support was implemented back to 61.1. through a PTF. The RDX dock is available in either a 5.25-inch internal (SATA or USB) or external USB version. The dock supports all RDX cartridges, which are reliable and rugged. They can be shipped through courier transport. You do not have to clean the drive.

For further details about the RDX standard, visit the RDX consortium website at:

<http://www.rdxstorage.com/rdx-technology>

RDX devices show up as RMSxx devices, but are classified as optical (random access spinning media) devices and can be used by optical, not tape, commands.

BRMS supports these devices as optical devices as well. The following limitations apply:

- ▶ The Media Policy option 'Mark history for duplication' is restricted to *NO. BRMS does not allow the duplication of optical history items.
- ▶ BRMS does not track opposite side volume identifiers or double-sided volumes. Each piece of media is viewed as a single volume.
- ▶ There is no optical support for the following BRMS functions: Dump BRMS (**DMPBRM**), add media information to BRMS (**ADDMEDIBRM**), extract media information (**EXTMEDIBRM**), and print media exceptions (**PRTMEDBRM**). There is no reclaim support for optical / RDX devices.
- ▶ You can specify only one device parameter per BRMS operation; optical does not support cascading or parallel devices.
- ▶ Optical media also cannot be shared on a BRMS network,. Remote duplication is not supported.
- ▶ **DUPMEDBRM** duplicates only entire optical volumes. The output volume must have the exact same physical characteristics as the input volume.
- ▶ Optical devices do not support software encryption.

4.2.17 Allowing volumes that are marked for duplication to be moved

In releases IBM i 7.1 and earlier, media movement is not allowed for volumes that are marked for duplication. Use the following interfaces to override this behavior and allow movement for volumes that are marked for duplication. This method is useful for virtual tape libraries, which often require virtual volumes to be moved during the duplication process.

You can perform the following actions:

- ▶ To override a move policy to allow movement when a volume is marked for duplication, run the following command:

```
CALL QBRM/Q1AOLD PARM('MOVMRKDUP ' '*SET' 'move policy' 'Y')
```

- To remove the override for a move policy that allows movement when a volume is marked for duplication, run the following command:

```
CALL QBRM/Q1AOLD PARM('MOVMRKDUP ' '*SET' 'move policy' 'N')
```

- To display all overrides for move policies that allow movement when a volume is marked for duplication, run the following command, as shown in Figure 4-23:

```
CALL QBRM/Q1AOLD PARM('MOVMRKDUP ' '*DISPLAY')
```

```

Command Entry                                ITSOP1                                Request level:
4
Previous commands and messages:
  > CALL QBRM/Q1AOLD PARM('MOVMRKDUP ' '*DISPLAY')
    MOVE1
    MOVE2

Bottom
Type command, press Enter.
===>

F3=Exit  F4=Prompt  F9=Retrieve  F10=Include detailed messages
F11=Display full  F12=Cancel  F13=Information Assistant  F24=More
keys

```

Figure 4-23 Move policy exceptions display

- To remove all overrides for move policies that allow movement when a volume is marked for duplication, run the following command:

```
CALL QBRM/Q1AOLD PARM('MOVMRKDUP ' '*CLEAR')
```

Consideration: In releases IBM i 7.1 and earlier, there is no synchronization of this behavior to other systems in a BRMS network. For each system for which you want to use this new function, you must run the commands individually.

Statement of direction: In releases that follow IBM i 7.1, this limitation will be removed. The function will be integrated into any relevant CL commands and the GUI.

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

4.2.18 Support for expiry dates greater than the year 2038

For versioning and media expiration, BRMS currently does not support dates beyond the year 2038. If you need an expiry date beyond 2038, you must set the date to *PERM.

Statement of direction: In releases that following IBM i 7.1, this limitation will be removed.

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

4.2.19 Setting the maximum size for BRMS flight recorders

The BRMS flight recorder size for storage can now be changed. To change the maximum flight recorder size, run the following command:

```
CALL PGM(QBRM/Q1AOLD) PARM('FRSIZE ' '*SET ' 'XXX')
```

XXX is equal to 001 - 999 MB.

Leading zeros: XXX must always be three digits, so leading zeros must be added to the front of the numbers.

Statement of direction: In releases that following IBM i 7.1, interfaces will be provided on the WRKPCYBRM TYPE(*SYS) work panel.

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

The following command can be used to display the maximum flight recorder size:

```
CALL PGM(QBRM/Q1AOLD) PARM('FRSIZE ' '*DISPLAY').
```

4.2.20 Using compression for BRMS Save Files

BRMS in V7R1 now supports override for the **DTACPR** parameter for saves to save file saves within BRMS. *HIGH, *MEDIUM, *LOW, *DEV, *NO, and *YES are all valid values. The override is then active for all savefile operations in BRMS.

You can perform the following actions:

- ▶ To set the new save file **DTACPR** parameter to *HIGH, run the following command:
CALL PGM(QBRM/Q1AOLD) PARM('PARMOVR' '*ADD' 'SAVFDTACPR' '*HIGH')
- ▶ To set the new save file **DTACPR** parameter to *MEDIUM, run the following command:
CALL PGM(QBRM/Q1AOLD) PARM('PARMOVR' '*ADD' 'SAVFDTACPR' '*MEDIUM ')
- ▶ To set the new save file **DTACPR** parameter to *LOW, run the following command:
CALL PGM(QBRM/Q1AOLD) PARM('PARMOVR' '*ADD' 'SAVFDTACPR' '*LOW')
- ▶ To set the new save file **DTACPR** parameter to *YES, run the following command:
CALL PGM(QBRM/Q1AOLD) PARM('PARMOVR' '*ADD' 'SAVFDTACPR' '*YES')

- To set the new save file **DTACPR** parameter to *N0, run the following command:
CALL PGM(QBRM/Q1AOLD) PARM('PARMOVR' '*ADD' 'SAVFDTACPR' '*N0')
- To set the new save file **DTACPR** parameter to *DEV, run the following command:
CALL PGM(QBRM/Q1AOLD) PARM('PARMOVR' '*ADD' 'SAVFDTACPR' '*DEV')
- To remove the override, run the following command:
CALL PGM(QBRM/Q1AOLD) PARM('PARMOVR' '*REMOVE' 'SAVFDTACPR')

Our initial tests showed that a significant compression could be reached with the *HIGH setting, albeit at the cost of a longer save time.

To show the current override that is being used, run the following command:

```
CALL PGM(QBRM/Q1AOLD) PARM('PARMOVR' '*DISPLAY' 'SAVFDTACPR')
```

The command output is shown in Figure 4-24.

```

Command Entry                               ITSOP2                                Request level:  4

Previous commands and messages:
> CALL PGM(QBRM/Q1AOLD) PARM('PARMOVR' '*DISPLAY' 'SAVFDTACPR')
SAVFDTACPR: *HIGH.

Type command, press Enter.
===>

F3=Exit   F4=Prompt   F9=Retrieve   F10=Include detailed messages
F11=Display full   F12=Cancel   F13=Information Assistant   F24=More keys
Bottom

```

Figure 4-24 SAVFDTACPR

Statement of direction: In releases that follow IBM i 7.1, the function will be integrated into any relevant CL commands and the GUI.

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

4.2.21 Improving IFS save performance - the ASYNCBRING parameter

In releases V6R1M0 and later, the **ASYNCBRING** parameter can be overridden to help improve IFS save performance.

To use the **ASYNCBRING** parameter, run the following commands:

- ▶ To override the new **ASYNCBRING** parameter, run the following commands:
`CALL PGM(QBRM/Q1AOLD) PARM('PARMOVR' '*ADD' 'ASYNCBRING' '*YES')`
- ▶ To disable the **ASYNCBRING** parameter, run the following commands:
`CALL PGM(QBRM/Q1AOLD) PARM('PARMOVR' '*ADD' 'ASYNCBRING' '*NO ')`
- ▶ To remove the override for the **ASYNCBRING** parameter, run the following command:
`CALL PGM(QBRM/Q1AOLD) PARM('PARMOVR' '*REMOVE ' 'ASYNCBRING')`
- ▶ To show the current override for the **ASYNCBRING** parameter, as shown in Figure 4-25, run the following command:
`CALL PGM(QBRM/Q1AOLD) PARM('PARMOVR' '*DISPLAY' 'ASYNCBRING')`

```
Command Entry                               ITSOP2                                Request level:  4
Previous commands and messages:
> CALL PGM(QBRM/Q1AOLD) PARM('PARMOVR' '*DISPLAY' 'ASYNCBRING')
ASYNCBRING: *YES.

Type command, press Enter.
====>

F3=Exit   F4=Prompt   F9=Retrieve   F10=Include detailed messages
F11=Display full   F12=Cancel   F13=Information Assistant   F24=More keys

Bottom
```

Figure 4-25 Displaying the ASYNCBRING parameter

Statement of direction: In releases that follow IBM i 7.1, interfaces will be provided on BRMS commands and on the control groups to use the new Asynchronous Bring parameters and these calls will not work. If *YES was set in the call statement, then the backup policy reflects that action on a new field after the upgrade finishes.

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

4.2.22 Using STRMNTBRM to initialize media when it is expired

A new option was added to the Start Maintenance for BRM (**STRMNTBRM**) command to initialize the BRMS media when it is expired. When the option is enabled, it requires media that is expired to be mounted on a device. This command helps you free disk storage sooner for virtual media technologies.

You can perform the following actions:

- To enable this option, run the following command:

```
CALL QBRM/Q1AOLD PARM('INZONEXP ' '*SET ' 'xxxxxxxxxx' 'Y')
```

xxxxxxxxxx is a media class name.


- To disable this option, run the following command:

```
CALL QBRM/Q1AOLD PARM('INZONEXP ' '*SET ' 'xxxxxxxxxx' 'N')
```

xxxxxxxxxx is a media class name.

- To display all media classes that have this option turned on, as shown in Figure 4-26, run the following command:

```
CALL QBRM/Q1AOLD PARM('INZONEXP ' '*DISPLAY ')
```



```
Command Entry                               ITSOP2                                Request level:  4

Previous commands and messages:
> CALL QBRM/Q1AOLD PARM('INZONEXP ' '*DISPLAY ')
MEDCLS1
MEDCLS2
MEDCLS3

Type command, press Enter.
====>

F3=Exit  F4=Prompt  F9=Retrieve  F10=Include detailed messages
F11=Display full  F12=Cancel  F13=Information Assistant  F24=More keys

Bottom
```

Figure 4-26 Displaying media classes with initialize on expiry

- To disable this option for all media classes that have this option enabled, run the following command:

```
CALL QBRM/Q1AOLD PARM('INZONEXP ' '*CLEAR ')
```

Consideration: In releases IBM i 7.1 and earlier, there is no synchronization of this behavior to other systems in the BRMS network. If you want each system to use this new function, you must run the commands that are shown above.

Initializing volumes: For IBM ProtecTIER® libraries, if volumes are initialized when they expire, performance for operations using the library may be temporarily degraded as ProtecTIER frees the storage that was being used by the volume.

For native IBM i virtual tape, if volumes are initialized when they expire, system performance might be temporarily degraded as the system frees the storage that was being used by the volume.

For native IBM i virtual tape, if multiple virtual tape devices are registered with BRMS, the **STRMNTBRM** job log may contain CPF41B0 ('Incorrect image catalog name specified.') messages. These messages can be safely ignored.

4.2.23 PRTRPTBRM control group status

The Print Report using BRM (**PRTRPTBRM**) command is enhanced to support the ***CTLGRPSTAT** value on the Report type (**TYPE**) parameter. The ***CTLGRPSTAT** report gives you information about backup control groups, which are specified on the Control group (**CTLGRP**) parameter, that are processed by BRMS, and meet the other selection criteria that are specified with the **PRTRPTBRM** command.

The following values can be specified for the **CTLGRP** parameter:

- ▶ ***ALL:** All backups that are produced by running **STRBKUBRM** are included in this report. The output that is generated by **CTLGRP(*NONE)** is not included in this report.
- ▶ ***SYSTEM:** Backups that are produced by running **STRBKUBRM** for the ***SYSTEM** control group are included in the report.
- ▶ ***BKUGRP:** Backups that are produced by running **STRBKUBRM** for the ***BKUGRP** control group are included in the report.
- ▶ ***SYSGRP:** Backups that are produced by running **STRBKUBRM** for the ***SYSGRP** control group are included in the report.
- ▶ ***NONE:** Backups that are produced by running **SAVxxxBRM** commands are summarized in the report.
- ▶ **Generic name:** Specify the generic name of the control group to be included in the report. A generic name is a character string that contains one or more characters followed by an asterisk (*). If an * is not specified with the name, the name is assumed to be a complete control group name.
- ▶ **Name:** Specify the name of a control group. All backups that are produced by running **STRBKUBRM** for the named control group are included in the report.

Considerations: To run **PRTRPTBRM TYPE(*CTLGRPSTAT)** in release IBM i 7.1, the following PTF or its superseding PTFs are required:

SI46340

PRTRPTBRM TYPE(*CTLGRPSTAT) CTLGRP(<name>) summarizes only information about control groups that are run *after* the PTFs are applied.

PRTRPTBRM TYPE(*CTLGRPSTAT) CTLGRP(*NONE) summarizes information about backups that were run before or after the PTFs are applied.

The **CTLGRP** parameter is ignored if the **TYPE** parameter is not ***CTLGRPSTAT**.

4.3 BRMS enhancements to GUI and web interfaces

As an alternative to the BRMS character-based interface, BRMS provides a full function graphical user interface (GUI).

System i Navigator is a client-based GUI system management console that uses a BRMS plug-in to offer full-featured BRMS capabilities. These capabilities were extended.

IBM Systems Director Navigator for i and IBM Systems Director are web-based interfaces that had limited BRMS function in IBM i 6.1. The capabilities of these interfaces were greatly expanded into a full-featured BRMS interface, bringing these web interfaces into parity with the client-based System i Navigator product.

This section describes the new capabilities and changes in the System Director web interfaces and points out which ones are also new to the System i Navigator product. This section describes the following enhanced functions:

- ▶ Added support for the IBM Systems Director web browser environment
 - IBM Systems Director navigation to BRMS functions
 - IBM Systems Director Navigator for i navigation to BRMS functions
- ▶ Enhancements to the BRMS initial window
- ▶ BRMS advanced functions window
- ▶ Scheduling support for BRMS
- ▶ Added option to the BRMS Log to filter messages by control groups
- ▶ Ability to mark and unmark volumes for duplication
- ▶ Multiple email address support
- ▶ High availability support for independent ASPs in a BRMS network
- ▶ Enhanced maintenance features in backup policy

Required features: To use the IBM i 7.1 enhancements, you must install the BRMS 7.1 plug-ins.

There are instructions at the following link for installing the plug-ins for IBM i Navigator, for installing the plug-ins for Systems Director Navigator for i, and for installing the plug-ins for Systems Director Navigator:

<http://www-03.ibm.com/systems/i/support/brms/pluginfaq.html>

Also in IBM i 7.1, management of tape devices and libraries were added to IBM Systems Director and IBM Systems Director Navigator for i. For more information, see 18.9, “New journal management enhancements” on page 673.

4.3.1 Added support for the IBM Systems Director web browser environment

Functions previously available and new IBM i 7.1 functions accessible through IBM Systems Director Navigator for i are now also available through IBM Systems Director.

Both products' BRMS functions are almost functionally and visually identical. The major differences are the navigation steps to get to the BRMS functions and the main BRMS window.

IBM Systems Director is intended for multiple systems and multiple system platforms. IBM Systems Director Navigator for i is intended for managing a single IBM i.

4.3.2 IBM Systems Director navigation to BRMS functions

IBM Systems Director, being a multiplatform and multiple system environment, requires navigation steps to select an IBM i system and to navigate to the BRMS functions. This section describes how this task is done. This section also describes the main BRMS window functions.

To access the BRMS functions for a particular IBM eServer iSeries or IBM i, complete the following steps:

1. Log on to IBM Systems Director.
2. Select a system resource that is an IBM i system or partition.
3. Access the IBM i resource.
4. Navigate to the BRMS function.

After logging on to the IBM Systems Director, you see the Systems Director Welcome window. In the view pane, at the upper left corner, is the “Navigate Resources” link, as shown in Figure 4-27.




Figure 4-27 IBM Systems Director Welcome window and Navigate Resources link

When you click **Navigate Resources**, a window opens and lists the various resource groups (Figure 4-28). Although you can select several of the groups to select an IBM i system, the “All Operating Systems” resource group is used in this example. The red arrow points to the link for this group.




Figure 4-28 IBM Systems Director Navigate Resources group list window

Click **All Operating Systems** and the window that is shown in Figure 4-29 opens, showing a list of member systems in the group. Note the OS Type column at the right of the display. Identify an IBM i system, and click its link.

The screenshot shows the IBM Systems Director interface. The left sidebar contains a navigation tree with categories like Automation, Availability, Inventory, Release Management, Security, System Configuration, and System Status and Health. The main area displays the 'Navigate Resources' window, which shows a list of operating systems under the group 'All Operating Systems (View Members)'. The table has columns for Select, Name, Access, Problems, Compliance, IP Addresses, and OS Type. A red arrow points to the system 'iszl1p13.rchland.ibm.com', which is an IBM i system. The OS Type column for this system is highlighted with a red box.

Select	Name	Access	Problems	Compliance	IP Addresses	OS Type
☐	elm8a105.beaverton.ibm.com	Offline	OK	OK	9.47.80.105	Unknown
☐	elm8a155.beaverton.ibm.com	Offline	OK	OK	9.47.80.155, 19;	Windows
☐	elm8a167.beaverton.ibm.com	Offline	OK	OK	9.47.80.167	Unknown
☐	elm8a172.beaverton.ibm.com	Offline	OK	OK	9.47.80.172	Unknown
☐	elm8a20.beaverton.ibm.com	Offline	OK	OK	9.47.80.20	Unknown
☐	elm8a22.beaverton.ibm.com	Offline	OK	OK	9.47.80.22	Hypervis
☐	iszl1p13.rchland.ibm.com	OK	OK	OK	9.5.215.103	IBM i
☐	rchasyum.rchland.ibm.com	Offline	Information	OK	9.5.50.82	IBM i
☐	x0716p3.rchland.ibm.com	OK	OK	OK	9.5.53.254, 127;	IBM i
☐	x0716p7.rchland.ibm.com	OK	Minor	OK	9.5.52.190, 200;	Linux
☐	z1014p26.rchland.ibm.com	OK	OK	OK	9.5.103.39	IBM i
☐	z1014p28.rchland.ibm.com	No access	Information	OK	9.5.105.200	IBM i

Figure 4-29 IBM Systems Director Navigate Resources Operating System Group list

Select the system and the resource window that is shown in Figure 4-30 opens. When you select the **Actions** drop-down menu, the first drop-down menu is displayed. When you move your cursor to IBM i Management, another menu of management functions is displayed. Backup Recovery and Media Services is at the top of the list of management functions, as shown in Figure 4-30.




Figure 4-30 Navigation from selected system to the BRMS function

Select the **Backup Recovery and Media Services** link and the BRMS initial window opens, as shown in Figure 4-31.




Figure 4-31 IBM Systems Director BRMS initial window

The initial window is described in 4.3.4, “Enhancements to the BRMS initial window” on page 92.

4.3.3 IBM Systems Director Navigator for i navigation to BRMS functions

IBM Systems Director Navigator for i, being a single IBM i environment, has much simpler navigation for accessing the BRMS function than the IBM Systems Director product. Although this function is not new function, it is included here for reference. This section reviews how this navigation is done and describes changes to the main window of the BRMS functions.

To access the IBM Systems Director Navigator for i welcome window (Figure 4-32 on page 91), you must meet the following conditions:

- ▶ Have the *ADMIN HTTP server started on the IBM i system.
- ▶ Have a web browser open to `http://system-name:2001` (where *system-name* is the IP name of the IBM i system).
- ▶ Be logged on to an IBM i user profile with sufficient privileges.

Upon logging on, the Welcome window opens, as shown in Figure 4-32. At the upper left corner of the navigation area is a minimized item that is labeled IBM i Management. Select this item to expand it.




Figure 4-32 IBM Systems Director Navigator for i welcome window

After you expand the IBM i Management item, you see a list of links to IBM i management functions (Figure 4-33). At the bottom of the list is the Backup, Recovery and Media Services option. Click this link and a tab opens that shows the BRMS initial window.




Figure 4-33 IBM System Director Navigator for i - BRMS initial window

4.3.4 Enhancements to the BRMS initial window

This section describes changes to the initial window of the BRMS functions when it is accessed through IBM Systems Director and IBM Systems Director Navigator for i. The BRMS tab with the initial window is nearly identical for both.

Here is what is new on the window, as shown in Figure 4-34:

1. The date, time, and status of the most recent save of all system and user data. The date, time, and status reported here are for the last run of the *SYSTEM control group. If your total system saves are run under your own backup group, they are not shown here.
2. Buttons to view the BRMS log and to view BRMS history. If the system has never run the *SYSTEM backup control group, the preceding items are not shown.
3. The check box to enable skipping of the initial window and to go directly to the advanced functions window.




Figure 4-34 Enhancements to BRMS web initial window

4.3.5 BRMS advanced functions window

To access the BRMS advanced functions from the initial window, click the **Advanced** button. The BRMS advanced functions window opens, as shown in Figure 4-35.




Figure 4-35 BRMS advanced functions window

The format of the window is a page with a tab. Selecting the tab (top arrow) brings that page to the foreground and places the others in the background.

The arrow at the left points to an object that, when clicked, can hide the left navigation pane. The remaining figures in this chapter do not show the navigation pane.

The small circled icon, when selected, opens a menu of actions. In Figure 4-35, the menu is shown for the BRMS Backup Control Groups field.

4.3.6 Scheduling support for BRMS

Scheduling support was added to IBM Systems Director and IBM Systems Director Navigator for i. Actions that can be scheduled are as follows:

- ▶ Back up and archive control groups
- ▶ Maintenance activities, including:
 - Media maintenance, such as expiration and moves
 - Reporting, such as expired media, media audit, media information, system information, and recovery reports
 - BRMS file management, such as journal receiver changes, data cleanup, and BRMS file reorganization

To schedule a backup control group to run, open the menu for Backup Control Groups from the BRMS advanced menu, as shown in Figure 4-36.




Figure 4-36 Open action of Backup Control Groups menu

Open Backup Control Groups

To open backup control groups, complete the following steps:

1. Select the **Open** action. A table of backup control groups (Figure 4-37) opens.




Figure 4-37 Selecting schedule on a backup control group

2. Open the menu for the control group that is to be scheduled.

3. Select the **Schedule** action. The Schedule Information window opens, as shown in Figure 4-38.

The screenshot shows the 'Schedule Information' window for a task named 'QBRMBKUP'. The window is part of the IBM Systems Director Navigator interface, with a top bar showing 'Welcome msrjj' and navigation tabs for 'Work Manage...', 'System i Na...', 'System Stat...', 'System', and 'Backup, Rec...'. The 'Schedule Information' window has a title bar 'Schedule Information - Rchasa04' and a 'Select Action' dropdown.

Task Name: QBRMBKUP
Task Description: Run Control Group

When to run:

- ☒ Once
- ☐ Weekly
- ☐ Monthly:
 - ☒ Absolute Day
 - ☐ Relative Day
 - ☐ Last Week ☐ Third Week
 - ☐ First Week ☐ Fourth Week
 - ☐ Second Week ☐ Fifth Week

Schedule day(s) or specific date to run

- ☐ Days
 - ☐ Monday ☐ Friday
 - ☐ Tuesday ☐ Saturday
 - ☐ Wednesday ☐ Sunday
 - ☐ Thursday
- ☒ Date to Start: 3/27/2010 (with a calendar icon)
- ☐ Last Day of the Month

Time to start: 21:00:00 PM Example: 12:30:00 PM

Omit dates

Omit Date: 3/27/2010 (with a calendar icon) **Add**

Omit Dates: [Empty] **Remove**

OK **Cancel**

Figure 4-38 Scheduling Information window

4. Schedule the backup control group to run once on 27 March 2010 at 21:00.

Scheduling BRMS maintenance

The following procedure describes one way of scheduling BRMS maintenance:

1. From the BRMS advanced menu, click **Select Action** to display a drop-down menu, as shown in Figure 4-39. From the drop-down menu, click **Run Maintenance** and click **Go**.




Figure 4-39 Selecting Run Maintenance from the Select Action drop-down menu

The Run Maintenance window (Figure 4-40) opens.




Figure 4-40 Run Maintenance window

2. Clicks **Options**. The Run Maintenance Options window (Figure 4-41) opens.




Figure 4-41 Run Maintenance Options - Media options window

The Run Maintenance Options window has three options to choose from:

- Media (shown in Figure 4-41)
- Print (shown in Figure 4-42 on page 99)
- File Management (shown in Figure 4-43 on page 99)




Figure 4-42 Run Maintenance Options - Print options window




Figure 4-43 Run Maintenance Options - File Management options window

3. Select maintenance options as required.
4. Click **OK** on any of the Run Maintenance options windows. The Run Maintenance window (Figure 4-40 on page 97) opens.

5. Click **Schedule**. The Schedule Information window (Figure 4-44) opens.

IBM® Systems Director Navigator for i Welcome msrjj Help | Logout IBM

Work Manage... X System i Na... X System Stat... X System X Backup, Rec... X --- Select Action ---

Schedule Information - Rchasa04

Task Name: QBRMMNT

Task Description: Run Maintenance Task

When to run:

☐ Once

☐ Weekly

☒ Monthly:

☐ Absolute Day

☒ Relative Day ☐ Last Week ☒ Third Week

☒ First Week ☒ Fourth Week

☒ Second Week ☒ Fifth Week

Schedule day(s) or specific date to run

☒ Days ☒ Monday ☐ Friday

☐ Tuesday ☐ Saturday

☐ Wednesday ☐ Sunday

☐ Thursday

☐ Date to Start 3/27/2010 [Calendar Icon]

☐ Last Day of the Month

Time to start: 4:45:25 PM Example: 12:30:00 PM

Omit dates

Omit Date: 3/27/2010 [Calendar Icon] Add

Omit Dates: [Empty] Remove

OK Cancel

Figure 4-44 Schedule Information window

This Schedule Information window looks identical and operates identically to the Schedule Information window that is used when you schedule a backup control group. The differences are with the Task name and Task Description field contents at the upper left of the window.

If the task is already scheduled, the current values are shown. If you make no changes and click **OK**, the maintenance is not scheduled because it is already scheduled. An error window opens.

If you change the schedule, then click **OK**, a Run Maintenance confirmation window (Figure 4-45) opens, giving you a chance to change any options, cancel the request, or accept the Run Maintenance as configured.




Figure 4-45 Run Maintenance Confirmation window

6. Click **OK** to complete scheduling maintenance.

Viewing scheduled tasks

To view scheduled tasks, complete the following steps:

1. Expand the menu for Tasks from the list column of the BRMS advanced menu.
2. From the menu, select **Open**, as shown in Figure 4-46.




Figure 4-46 Task List menu with Open selected

Figure 4-47 shows the window that opens, which is a list of scheduled BRMS maintenance tasks.




Figure 4-47 List of BRMS maintenance tasks

In IBM i 6.1, only active or completed tasks can be viewed or monitored.

In IBM i 7.1, scheduled tasks can be viewed, including those tasks that are scheduled by System i Navigator.

BRMS 5250 support lists System i Navigator, IBM Systems Director, and IBM Systems Director Navigator for i BRMS tasks and jobs.

4.3.7 Added option to the BRMS Log to filter messages by control groups

In IBM i 7.1, the BRMS Log can now be filtered by control group. You can filter by one control group at a time. Similar functionality is now available in the System i Navigator client.

To use this option, you can select **BRMS Log** from the BRMS initial menu that is shown in Figure 4-33 on page 91. You can also navigate to the Task list menu shown in Figure 4-46 on page 101, except that instead of clicking **Open**, click **BRMS Log**.

Another way that you can accomplish the same objective is by choosing **BRMS Log** from the Select Action drop-down menu of the BRMS advanced function menu page, as shown in Figure 4-48.




Figure 4-48 Selecting BRMS Log from the Select Action drop-down menu

Regardless of the method that is used to access the BRMS log, the BRMS Log - Include window opens, as shown in Figure 4-49.

IBM® Systems Director Navigator for i

Welcome msrjj

Help | Logout

Work Manage... x System i Na... x System Stat... x System x Backup, Rec... x

--- Select Action ---

BRMS Log - Include

Entry types: All

Minimum severity (0-99): 0 0,1,2...99

Message ID: All

Dates:

☒ All dates

☐ Specific dates

From: 3/26/2010

To: 3/27/2010

Times:

☒ All times

☐ Specific times

From: 5:45:30 PM 12:00:00 AM-11:59:59 PM

To: 5:45:30 PM 12:00:00 AM-11:59:59 PM

Program: All

User: All

Job name: All

Job number: All

Control group: Use entry from below Browse...

Qaudjrn

☒ Always show this first

OK Cancel

Figure 4-49 New Control group selection of BRMS Log - Include window

The new control group selection parameters are shown in the circle. The **Browse** button displays a list of controls groups from which you can make selections.

A partial view of the resulting control group for BRMS log listing is shown in Figure 4-50.

IBM® Systems Director Navigator for i Welcome msrjj [Help](#) | [Logout](#)

Work Manage... X System i Na... X System Stat... X System X Backup, Rec... X

Backup, Recovery and Media Services

BRMS Log - Rchasa04

[Refresh](#) Control Group: Qaudjrn

List items:

Select	Message ID	Severity	Entry Type	Message	Date Sent	Time Sent	User	Job Name	Job
☐	BRM1096	0	Backup	Job queue BRMMMAINT is released.	Feb 28, 2010	11:49:01 AM	Ms773	Secbkppaud	33
☐	BRM1049	0	Backup	Control group QAUDJRN type *BKU processing is complete.	Feb 28, 2010	11:49:01 AM	Ms773	Secbkppaud	33
☐	BRM1587	10	Backup	Save of BRM media information at level *OBJ complete.	Feb 28, 2010	11:49:00 AM	Ms773	Secbkppaud	33
☐	CPC3701	20	Backup	19 objects saved from library QUSRBRM.	Feb 28, 2010	11:48:59 AM	Ms773	Secbkppaud	33
☐	CPI6705	10	Backup	7485 blocks processed for sequence 333, volume LT0026, on device MLD50.	Feb 28, 2010	11:48:54 AM	Ms773	Secbkppaud	33
☐	BRM1647	10	Backup	Starting save of media information at level *OBJ to device MLD50.	Feb 28, 2010	11:48:12 AM	Ms773	Secbkppaud	33
☐	BRM4201	10	Backup	Change to media controls successful.	Feb 28, 2010	11:48:12 AM	Ms773	Secbkppaud	33
☐	BRM4201	10	Backup	Change to media controls successful.	Feb 28, 2010	11:48:12 AM	Ms773	Secbkppaud	33

Page 1 of 11 [1](#) [Go](#) Rows [8](#) Total: 81

[Close](#)

Figure 4-50 BRMS Log list - Filtered by control group QAUDJRN

4.3.8 Ability to mark and unmark volumes for duplication

IBM Systems Director Navigator for i has the same mark and unmark for duplication functions as **WRKMEDBRM** options 18 and 19. In this section, the navigation is described.

In Figure 4-51, the Volume list menu is displayed with the Open option specified.




Figure 4-51 Opening the volume list

When you click **Open**, the volumes table is displayed. From the volume table, open the menu for a specific volume, as shown in Figure 4-52.

The screenshot shows the IBM Systems Director Navigator for i interface. The main window displays the 'Volumes - Rchasa04' table. A context menu is open over the volume 'Lt0029', showing options: Open, Expire..., Remove..., Initialize..., Volume Set, View Report..., Properties, Mark volume(s) for duplication, and Save as a Report... The table lists various volumes with columns for Select, Volume, Type, Image Catalog, Status, Expiration, and Media Pool. The status of all listed volumes is 'Active'. The media pool for all is 'Ultrium1'. The expiration dates range from Mar 30, 2010 to Jun 14, 2010. The table is on page 1 of 160, showing 15 rows, with a total of 2395 volumes.

Select	Volume	Type	Image Catalog	Status	Expiration	Media Pool
☐	Lt0020	L1		Active	May 17, 2010	Ultrium1
☐	Lt0021	L1		Active	Apr 11, 2010	Ultrium1
☐	Lt0029	L1		Active	Jul 5, 2010	Ultrium1
☐	Lt0502	L1		Active	Jun 27, 2010	Ultrium1
☐	Lt0503	L1		Active	Jul 4, 2010	Ultrium1
☐	Lt0504	Volume Set		Active	Apr 5, 2010	Ultrium1
☐	Lt0890	L1		Active	Apr 5, 2010	Ultrium1
☐	Lt0892	L1		Active	Apr 6, 2010	Ultrium1
☐	Lt0893	L1		Active	Apr 9, 2010	Ultrium1
☐	Lt0894	L1		Active	Apr 13, 2010	Ultrium1
☐	Lt0895	L1		Active	Apr 13, 2010	Ultrium1
☐	Lt0896	L1		Active	Apr 13, 2010	Ultrium1
☐	Lt0897	L1		Active	Apr 13, 2010	Ultrium1
☐	Lt0898	L1		Active	Jun 14, 2010	Ultrium1
☐	Lt0891	L1		Active	Mar 30, 2010	Ultrium1

Figure 4-52 Volumes table with the menu displayed for volume LT0029

You can tell that the volume is not marked for duplication because the “Mark volume(s) for duplication” option exists in the menu. Select the mark option, and the Confirm Volume Action window (Figure 4-53) opens.




Figure 4-53 Confirm Volume Action window

Click the **OK** button and the volume is marked duplication. The Volumes window opens again.

You cannot tell from the Volumes table window whether the volume is marked. One way to verify that the volume is marked is to select the volume, open its menu, and look at its contents, as shown in Figure 4-54.




Figure 4-54 Volume menu that shows Unmark volume(s) for duplication

Because the Unmark volume(s) for duplication option is shown, you know that the volume is marked for duplication.

If you want to unmark the volume, click the **Unmark** option.

Suppose that you do much marking and unmarking for duplication. You can modify the volume table and add the "Marked for Duplication" column to the table view. In the following scenario, you remove the "Image Catalog" column to keep the table narrow enough to fit on the window.

To accomplish this task, open the Select Action drop-down menu, as shown in Figure 4-55, and select the **Columns and Sorts** action.




Figure 4-55 Select Action drop-down menu - Columns and Sorts action

When Columns and Sorts is selected, the Set Columns to Display and Sort window opens. Figure 4-56 shows a “Show Column” column in which the columns are selected for display.

Set Columns to Display and Sort

Move Up Move Down Reset to Default Settings

Columns:

--- Select Action --- Go Filter

Select	Show Column	Column Name	Sort Order	Sort Ascending
☐	☒	Volume		☒
☐	☒	Type		☒
☐	☒	Image Catalog		☒
☐	☒	Status	1	☒
☐	☒	Expiration		☒
☐	☒	Media Pool		☒
☐	☐	Container		☒
☐	☐	Date Added or Reused		☒
☐	☐	Description		☒
☐	☐	Duplicated Volume		☒
☐	☐	Expired		☒
☐	☐	Last Reused		☒
☐	☐	Location		☒
☐	☐	Marked for Duplication		☒
☐	☐	Moved Date		☒

Page 1 of 2 1 Go Rows 15 Total: 24

OK Cancel

Figure 4-56 Set Columns to Display and Sort window

You do not want the Image Catalog column to display, so you clear that column. You want the Marked for Duplication column to display, so you select that column for display. When you click the **OK** button, the volume that is displayed in Figure 4-57 opens.

Select	Volume	Type	Status	Expiration	Media Pool	Marked for Duplication
☐	Lt0020	L1	Active	May 17, 2010	Ultrium1	No
☐	Lt0021	L1	Active	Apr 11, 2010	Ultrium1	No
☐	Lt0029	L1	Active	Jul 5, 2010	Ultrium1	No
☐	Lt0502	L1	Active	Jun 27, 2010	Ultrium1	No
☐	Lt0503	L1	Active	Jul 4, 2010	Ultrium1	No
☐	Lt0504	L1	Active	Apr 5, 2010	Ultrium1	No
☐	Lt0890	L1	Active	Apr 5, 2010	Ultrium1	No
☐	Lt0891	L1	Active	Apr 20, 2010	Ultrium1	No
☐	Lt0892	L1	Active	Apr 6, 2010	Ultrium1	No
☐	Lt0893	L1	Active	Apr 9, 2010	Ultrium1	No
☐	Lt0894	L1	Active	Apr 13, 2010	Ultrium1	No
☐	Lt0895	L1	Active	Apr 13, 2010	Ultrium1	No
☐	Lt0896	L1	Active	Apr 13, 2010	Ultrium1	No
☐	Lt0897	L1	Active	Apr 13, 2010	Ultrium1	No
☐	Lt0898	L1	Active	Oct 7, 2010	Ultrium1	No

Page 1 of 158 1 Go Rows 15 Total: 2370

Figure 4-57 Volume table with column changes

The Image Catalog column is removed and the Marked for Duplication column is shown. You can now see the Marked for Duplication status of each volume without selecting each one.

4.3.9 Multiple email address support

Multiple email addresses can now be specified for sending BRMS alerts. This function is not available under 5250 emulation, but is supported by System i Navigator, the GUI, and IBM Systems Director and IBM Systems Director Navigator for i interfaces.

To configure this support, access the Global Policy Properties menu from the BRMS advanced menu, as shown in Figure 4-58.




Figure 4-58 Selecting Global Policy Properties

The Global Policy Properties window opens, as shown in Figure 4-59.

Global Policy Properties - Rchasa04

General

BRMS system name: Use entry from below
Rchasa04

For saves of Licensed Internal Code and IBM i

Maximum time for stopping all subsystems: 1200 seconds

Maximum time to allow for save: No maximum minutes

☒ Only use interactive console monitor

Note: If the interactive console monitor is not used, the control group run operation must be fully automated and require no user intervention.

☒ Allow BRMS to monitor media

Time to wait for servers to end: 0 0,1,2...9,999 seconds

☒ Allow restores from alternate media

Prefix for volume serials:

BRMS submitted jobs:

Job description: User profile

Job library: Use library list

Job queue: Job description

Job queue library: Use library list

OK Cancel

Figure 4-59 Global Policy Properties window

Click **Network Properties**. In the email address field, multiple email addresses can now be entered, as shown in Figure 4-60.




Figure 4-60 Email address field with multiple entries on the Network policy properties window

4.3.10 High availability support for independent ASPs in a BRMS network

This function runs an incremental save from any system in the BRMS network for IASPs using the IBM i Clustering technology. Suppose System A has run a full backup of an IASP and the save information that is stored in BRMS databases is synchronized to System B in the BRMS network. Because System B has System A's save information, System B can do an incremental save using reference dates from that save information. This feature is only available through the BRMS graphical user interfaces, the IBM Systems Director web environment, or System i Navigator.

To configure this function, complete the following steps:

1. Click **Backup, Recovery, and Media Services** → **Global Policy Properties**. The window that is shown in Figure 4-58 on page 113 opens.

Global Policy Properties - X0716p3.rchland.ibm.com

General
System update interval (30-9999): 60
Offline notification period (30-99,999): 30

Notification
☒ Use TCP/IP for networking
☐ Use SNA for networking
☐ Receive history information

Signoff Exceptions

Power Down

Maintenance

Retrieval

Network
Logging

Distribution:
Primary system: None [Browse...]
Secondary system: None [Browse...]
E-mail address:
Maximum message length: 4999 0,1,2...5,000

Network restricted state interfaces to start:
[Icons] --- Select Action --- [Go]
Interface Address Type
9.10.1.1 Virtual
Page 1 of 1 Total: 1 Filtered: 1 Displayed: 1
[Manage Interfaces to Start...]

BRMS network systems:
[Icons] --- Select Action --- [Go]

System	System Status	Network Status	TCP/IP Name	Relational Database	Local Receives	Remote Receives	Re
Kitkat	Online	Active	Kitkat	Kitkat	Yes	Yes	Ye
Mounds	Online	Inactive	Mounds	Mounds	Yes	Yes	No
Rchasyum	Online	Active	Rchasyum	Rchasyum	Yes	No	No
Reeses	Online	Active	Reeses	Reeses	Yes	No	No
Twix	Online	Active	Twix	Twix	Yes	Yes	No
X0716p2	Online	Active	X0716p2	X0716p2	Yes	No	No

Page 1 of 2 1 [Go] Total: 11 Filtered: 11 Displayed: 6
[Manage Systems...]
[Manage Disk Pool History...]

Figure 4-61 Global Policy Properties - Network Properties

2. On the Global Policy Properties window, select the **Network** tab.
3. Click the **Manage Disk Pool History** button, as shown in Figure 4-61. The BRMS graphical interface refers to IASPs as disk pools.

- On the Manage Disk Pool History to Send window, click **List actions** from the menu bar and select **New**, as shown in Figure 4-62.




Figure 4-62 Send New Disk Pool History

- Enter the disk pool name, the system name, and network identifier of the system to which you want the history information to be sent.

Decide whether to use the BRMS media information remote receives value that is defined by “Send to system”. You can override that systems value by selecting the **Always send** option.





Figure 4-63 Send Disk Pool History Properties

To determine what the Remote Receives value is for the remote system, view it by going back to the Global Policy Properties window, clicking **Network properties**, and clicking **Manage Systems**. The value is listed under the Remote Receives column for that remote system, as shown in Figure 4-64.



Select	System	System Status	Network Status	TCP/IP Name	Relational Database	Local Receives	Remote Receives	Remote Receives O
☐	Kitkat	Online	Active	Kitkat	Kitkat	Yes	Yes	Yes
☐	Mounds	Online	Inactive	Mounds	Mounds	Yes	Yes	No
☐	Rchasyum	Online	Active	Rchasyum	Rchasyum	Yes	No	No
☐	Reeses	Online	Active	Reeses	Reeses	Yes	No	No
☐	Twix	Online	Active	Twix	Twix	Yes	Yes	No
☐	X0716p2	Online	Active	X0716p2	X0716p2	Yes	No	No
☐	Z1014p10	Online	Active	Z1014p10	Z1014p10	No	Yes	No
☐	Z1014p26	Online	Active	Z1014p26	Z1014p26	Yes	Yes	No
☐	Z1014p28	Online	Active	Z1014p28	Z1014p28	Yes	Yes	No
☐	Z1014p6	Online	Active	Z1014p6	Z1014p6	Yes	Yes	No
☐	Z1014p7	Online	Active	Z1014p7	Z1014p7	Yes	No	No

Page 1 of 1 Total: 11 Filtered: 11 Displayed: 11

Close

Figure 4-64 Checking a System's Remote Receives value

- Return to the Send Disk Pool History addition of a new disk pool window and click **OK** to complete the addition.

4.3.11 Enhanced maintenance features in backup policy

You can now use the BRMS maintenance features in the backup policy to run movement, expire partial volume sets, and reorganize BRMS database files.

Figure 4-65 shows the Global Policy Properties - Backup Maintenance Options window. The new Run move policies and Expire partial volume sets are circled.




Figure 4-65 Backup policy properties - Run move policy and Expire partial volume sets options

Click **File Management** and the new Reorganize BRMS database option becomes available, as shown in Figure 4-66.




Figure 4-66 Backup control group Reorganize BRMS database option

4.3.12 New BRMS Enterprise function

BRMS now includes the new BRMS enterprise function that you can use to centralize your daily BRMS chores that you have on multiple systems on one central system. You can do this task with either individual BRMS systems, one BRMS network, multiple BRMS networks, or any combination thereof, as shown in Figure 4-67.




Figure 4-67 BRMS Enterprise network

The central "Enterprise System" (HUB) pulls important information from systems (NODES) defined in its "Enterprise" network. From this information, specific notifications, verifications, and various other functions can be made and help an administrator manage the health of their BRMS backups/recoveries from one central server. This situation is beneficial for customers with multiple BRMS systems or BRMS networks.

This feature is available in the System Navigator only. To use it, you need a valid license for 5770 BR1 features *BASE (BRMS), 1 (BRMS-Network Feature) and 3 (BRMS-Advanced Functions Feature). Your http *Admin Server also must be up and running. To access the System Navigator, point your browser to `http://<systemname>:2001`. Sign on with your IBM i user profile and password, then click **IBM i Management** → **Backup, Recovery and Media Services** → **Advanced** → **Enterprise Services**.

You can see the initial Enterprise Network window in Figure 4-68.




Figure 4-68 BRMS Enterprise initial window

For more information about what you can do with the BRMS Enterprise network, see *BRMS Enterprise Enhancements*, REDP-4926.

4.4 More resources

For more information about the topics in this chapter, see the following resources:

- ▶ IBM i 7.1 Information Center:
<http://publib.boulder.ibm.com/iserics/v7r1m0/index.jsp>
- ▶ *IBM i Backup, Recovery, and Media Services for IBM i 7.1*, SC41-5345-07
- ▶ BRMS for IBM i website:
<http://www-03.ibm.com/system/i/support/brms>
- ▶ BRMS GUI documentation:
<http://www-03.ibm.com/servers/eserver/iserics/service/brms/pluginfaq.html>
- ▶ DAOS Quick Start Guide:
<http://www.lotus.com/ldd/dominowiki.nsf/dx/daos-quick-start-guide>
- ▶ DAOS Best Practices:
<http://www.lotus.com/ldd/dominowiki.nsf/dx/daos-best-practices>
- ▶ DAOS Estimator:
<http://www.ibm.com/support/docview.wss?rs=463&uid=swg24021920>
- ▶ BRMS Online Lotus Server Backup Reference:
<http://www-03.ibm.com/systems/i/support/brms/domdaos.html>



High availability

This chapter covers the following high availability (HA)-related enhancements of IBM i:

- ▶ PowerHA SystemMirror for i
- ▶ Journaling and commitment control enhancements

5.1 PowerHA SystemMirror for i

This section describes the following enhancements included with PowerHA SystemMirror for i, which is the strategic IBM high availability product for IBM i:

- ▶ New PowerHA packaging
- ▶ PowerHA versions
- ▶ PowerHA SystemMirror for i enhancements
- ▶ PowerHA SystemMirror for i: Graphical interfaces
- ▶ Asynchronous geographic mirroring
- ▶ LUN level switching
- ▶ Space-Efficient FlashCopy
- ▶ Better detection of cluster node outages
- ▶ Improved geographic mirroring full synchronization performance
- ▶ Improvement of geographic mirroring in an IBM i hosted IBM i client partition environment
- ▶ Support virtualization capabilities on an active cluster node
- ▶ Cluster administrative domain enhancements
- ▶ IPv6 support
- ▶ New CL commands for programming cluster automation

5.1.1 New PowerHA packaging

IBM PowerHA for i was renamed to IBM PowerHA SystemMirror for i to align with the corresponding Power Systems PowerHA family product PowerHA SystemMirror for AIX.

IBM PowerHA SystemMirror for i is offered in two editions for IBM i 7.1:

- ▶ IBM PowerHA SystemMirror for i Standard Edition (5770-HAS *BASE) for local data center replication only
- ▶ IBM PowerHA SystemMirror for i Enterprise Edition (5770-HAS option 1) for local or multi-site replication

Customers already using PowerHA for i with IBM i 6.1 are entitled to an upgrade to PowerHA SystemMirror for i Enterprise Edition with IBM i 7.1.

The functional differences between the IBM PowerHA SystemMirror for i Standard and Enterprise Edition are summarized in Figure 5-1.

PowerHA SystemMirror for i	Standard Edition	EnterpriseEdition
Centralized cluster management	✓	✓
Cluster resource management	✓	✓
Centralized cluster configuration	✓	✓
Automated cluster validation	✓	✓
Cluster admin domain	✓	✓
Cluster device domain	✓	✓
Integrated heartbeat	✓	✓
Application monitoring	✓	✓
IBM i event/error management	✓	✓
Automated planned switch over	✓	✓
Managed unplanned fail over	✓	✓
Centralized FlashCopy	✓	✓
LUN level switching	✓	✓
GeoMirror sync delivery	✓	✓
GeoMirror async delivery		✓
Multi-Site HA/DR management		✓
DS8000/DS6000 Metro Mirror		✓
DS8000/DS6000 Global Mirror		✓
SVC/V7000 Metro Mirror		✓
SVC/V7000 Global Mirror		✓

Figure 5-1 PowerHA SystemMirror for i editions

5.1.2 PowerHA versions

To use any of the new PowerHA SystemMirror for i enhancements, all nodes in the cluster must be upgraded to IBM i 7.1. Before this upgrade, both the cluster version and the PowerHA version must be updated to the current cluster Version 7 and PowerHA Version 2 by running the following CL command:

```
CHGCLUVER CLUSTER(cluster_name) CLUVER(*UP1VER) HAVER(*UP1VER)
```

As PowerHA SystemMirror for i now has N-2 support for clustering, it is possible to skip one level of IBM i just by running the earlier command twice. As such, a V5R4M0 system within a clustered environment can be upgraded towards IBM i 7.1 by skipping IBM i 6.1.

5.1.3 PowerHA SystemMirror for i enhancements

This section describes the enhancements in IBM i 7.1 regarding Power SystemMirror for i.

Enhancements that are delivered with 5799-HAS PRPQ

In April 2012, IBM PowerHA SystemMirror for i started to support the following new functions. Some of these functions were covered with 5799-HAS Program Request Pricing Quotation (PRPQ). So, the requirement of the 5799-HAS PRPQ for using these functions was eliminated:

- Support for managing IBM System Storage SAN Volume Controller and IBM Storwize V7000 Copy Services with IBM PowerHA SystemMirror for i 5770-HAS PTF SI45741

- ▶ IBM i command-line commands for configuring an independent auxiliary storage pool (**CFGDEVASP**) with 5770-SS1 PTF SI44141
- ▶ IBM i command-line command for configuring geographic mirroring (**CFGGEOMIR**) with PTF SI44148
- ▶ PowerHA GUI in IBM Navigator for i

The following subsections provide a brief overview of these enhancements. For more information, see *PowerHA SystemMirror for IBM i Cookbook*, SG24-7994.

Support for SAN Volume Controller / V7000 Copy Services

PowerHA SystemMirror for i now also supports Metro Mirror, Global Mirror, and FlashCopy for the IBM System Storage SAN Volume Controller and IBM Storwize V7000.

The available commands are similar to the ones that you use for DS8000 Copy Services, but some parameters are different:

- ▶ Add SVC ASP Copy Description (**ADDSVCCPYD**): This command is used to describe a single physical copy of an auxiliary storage pool (ASP) that exists within an SAN Volume Controller and to assign a name to the description.
- ▶ Change SVC Copy Description (**CHGSVCCPYD**): This command changes an existing auxiliary storage pool (ASP) copy description.
- ▶ Remove SVC Copy Description (**RMVSVCCPYD**): This command removes an existing ASP copy description. It does not remove the disk configuration.
- ▶ Display SVC Copy Description (**DSPSVCCPYD**): This command displays an ASP copy description.
- ▶ Work with ASP Copy Description (**WRKASPCPYD**) shows both DS8000 and SAN Volume Controller / V7000 copy descriptions.
- ▶ Start SVC Session (**STRSVCSSN**): This command assigns a name to the Metro Mirror, Global Mirror, or FlashCopy session that links the two ASP copy descriptions for the source and target IASP volumes and starts an ASP session for them.
- ▶ Change SVC Session (**CHGSVCSSN**): This command is used to change an existing Metro Mirror, Global Mirror, or FlashCopy session.
- ▶ End SVC ASP Session (**ENDSVCSSN**): This command ends an existing ASP session.
- ▶ Display SVC Session (**DSPSVCSSN**): This command displays an ASP session.

Important: The PowerHA GUI does not support the SAN Volume Controller / V7000 Copy Services.

The Configure Device ASP (CFGDEVASP) command

The new Configure Device ASP (**CFGDEVASP**) command is part of the IBM i 7.1 base operating system and is available with PTF SI44141. It is used to create or delete an independent auxiliary storage pool (IASP).

If you use this command with the ***CREATE** action, it does the following actions:

- ▶ Creates the IASP using the specified non-configured disk units.
- ▶ Creates an ASP device description with the same name if one does not exist yet.

If you use this command with the ***DELETE** action, it does the following actions:

- ▶ Deletes the IASP.
- ▶ Deletes the ASP device description if it was created by this command.

For more information about creating / deleting an IASP through the command-line interface (CLI), see Figure 5-2.

Configure Device ASP (CFGDEVASP)		
Type choices, press Enter.		
ASP device	ASPDEV	> IASP1
Action	ACTION	> *CREATE
ASP type	TYPE	*PRIMARY
Primary ASP device	PRIASPDEV	
Protection	PROTECT	*NO
Encryption	ENCRYPT	*NO
Disk units	UNITS	*SELECT
+ for more values		
Additional Parameters		
Confirm	CONFIRM	*YES

Figure 5-2 Configure Device ASP (CFGDEVASP) command

CFGGEOMIR command

The Configure Geographic Mirror (CFGGEOMIR) command that is shown in Figure 5-3 can be used to create a geographic mirror copy of an existing IASP in a device cluster resource group (CRG). The command can also create ASP copy descriptions if they do not exist yet and can start an ASP session. It performs all the necessary configuration steps to take an existing stand-alone IASP and create a geographic mirror copy. To obtain this command, the 5770-HAS PTF SI44148 must be on the system that is running IBM i 7.1.

Configure Geographic Mirror (CFGGEOMIR)		
Type choices, press Enter.		
ASP device	ASPDEV	
Action	ACTION	
Source site	SRCSITE	*
Target site	TGTSITE	*
Session	SSN	
Source ASP copy description .		
Target ASP copy description .		
Transmission delivery	DELIVERY	*SYNC
Disk units	UNITS	*SELECT
+ for more values		
Additional Parameters		
Confirm	CONFIRM	*YES
Cluster	CLUSTER	*
Cluster resource group	CRG	*
More...		

Figure 5-3 Configure Geographic Mirror (CFGGEOMIR) command

PowerHA GUI in IBM Navigator for i

PowerHA GUI is available in IBM Navigator for i. It supports switched disk and geographic mirroring environments, and DS8000 FlashCopy, Metro Mirror, and Global Mirror technologies. The PowerHA GUI provides an easy to use interface to configure and manage a PowerHA high availability environment. It can configure and manage cluster nodes, administrative domain, and independent auxiliary storage pools (IASPs).

Section 5.1.4, “PowerHA SystemMirror for i: Graphical interfaces” on page 130 introduces GUIs that are associated with a high availability function, including two existing interfaces and the PowerHA GUI.

Duplicate library error handling

Duplicate libraries and spool files are not allowed between SYSBAS and an IASP. When the IASP is available, the operating system prevents a duplicate library or spool file from being created. However, if a duplicate library or spool file is created in SYSBAS while the IASP is unavailable, the next varyon of the IASP fails.

With this enhancement, the message ID CPDB8EB and CPF9898, shown in Figure 5-4, is displayed in the QSYSOPR message queue at the IASP varyon time when a duplicate library is found in SYSBAS and the IASP. The varyon of the IASP can be continued or canceled after the duplicate library issue is resolved.

Additional Message Information			
Message ID	CPDB8EB	Severity	30
Message type	Information		
Date sent	10/10/12	Time sent	09:30:30
Message : Library TESTLIB01 exists in *SYSBAS and ASP device IASP01. Cause : Auxiliary storage pool (ASP) device IASP01 cannot be varied on to an available status because the ASP contains a library named TESTLIB01 and a library by the same name already exists in the system ASP or a basic user ASP (*SYSBAS). Recovery : Use the Rename Object (RNMOBJ) or Delete Library (DLTLIB) command to rename or delete library TESTLIB01 from ASP IASP01 or *SYSBAS to remove the duplicate library condition. Vary ASP device IASP01 off and retry the vary on.			
			Bottom
Press Enter to continue.			
Additional Message Information			
Message ID	CPF9898	Severity	40
Message type	Inquiry		
Date sent	10/10/12	Time sent	09:30:30
Message : See previous CPDB8EB messages in MSGQ QSYSOPR. Respond (C G). Cause : This message is used by application programs as a general escape message.			
			Bottom
Type reply below, then press Enter.			
Reply			
F3=Exit F6=Print F9=Display message details F12=Cancel F21=Select assistance level			

Figure 5-4 CPDB8EB and CPF9898 messages for duplicate library

This enhancement is available in both 6.1 with 5761-SS1 PTF SI44564 and 7.1 with 5770-SS1 PTF SI44326.

N-2 support for clustering

Before the release of 5770-SS1 PTF SI42297, PowerHA and clustering supported only a difference of one release, which meant that the nodes within the cluster needed to be within one release, and a node within the cluster could be upgraded only one release at a time.

After this PTF is installed, a 7.1 node can be added to a 5.4 cluster. A node can also be upgraded from a 5.4 cluster node directly to a 7.1 cluster node if this PTF is installed during the upgrade.

The main intent of this enhancement is to ensure that nodes can be upgraded directly from 5.4 to 7.1. PowerHA replication of the IASP still does not allow replication to an earlier release, so for a complete high availability solution, other than during an upgrade of the HA environment, all nodes should remain at the same release level.

5.1.4 PowerHA SystemMirror for i: Graphical interfaces

With IBM i 7.1, there are three different GUIs within IBM Navigator for i:

- ▶ Cluster Resource Services GUI
- ▶ High Availability Solutions Manager GUI
- ▶ PowerHA GUI

Cluster Resource Services GUI

You can access this GUI by completing the following steps, as shown in Figure 5-5:

1. Expand **IBM i Management**.
2. Select **Cluster Resource Services**.




Figure 5-5 Cluster Resource Services GUI

The Cluster Resource Services GUI has the following characteristics:

- ▶ Supports the existing environment.
- ▶ Limited Independent ASP (IASP) function.
- ▶ Cannot manage from one node.
- ▶ Difficult to determine the status.

Removed feature: The clustering GUI plug-in for System i Navigator from High Availability Switchable Resources licensed program (IBM i option 41) was removed in IBM i 7.1.

High Availability Solutions Manager GUI

You can access this GUI by completing the following steps, as shown in Figure 5-6:

1. Expand **IBM i Management**.
2. Select **High Availability Solutions Manager**.




Figure 5-6 High Availability Solutions Manager GUI

The High Availability Solutions Manager GUI has the following characteristics:

- ▶ “Dashboard” interface.
- ▶ No support for existing environments.
- ▶ Cannot choose names.
- ▶ Limited to four configurations.

New PowerHA SystemMirror for i GUI

In IBM PowerHA SystemMirror for i, there is a new interface available for PowerHA within IBM Navigator for i. This new PowerHA GUI was already supported by the October 2011 PRPQ announcement of 5799-HAS, but it supported only an English GUI at that level. With the April 2012 enhancement of IBM PowerHA SystemMirror for i, it supports multiple languages in the GUI.

You can access the new GUI by completing the following steps, as shown in Figure 5-7:

1. Expand **IBM i Management**.
2. Select **PowerHA**.




Figure 5-7 PowerHA SystemMirror for i GUI

The PowerHA GUI handles the high availability solution from one single window. It supports the following items:

- ▶ Geographic mirroring
- ▶ Switched disk (IOA)
- ▶ DS6000/DS8000 Metro Mirror
- ▶ DS6000/DS8000 Global Mirror
- ▶ DS6000/DS8000 FlashCopy

For more detailed information about the PowerHA GUI, see Chapter 9, “PowerHA User Interfaces”, in *PowerHA SystemMirror for IBM i Cookbook*, SG24-7994.

Differences between the three graphical interfaces

Figure 5-8 shows the main differences between the three available graphical interfaces.

	Cluster Resource Services GUI	High Availability Solutions Manager GUI	PowerHA GUI
Single Node Management			✓
Quick Problem Determination		✓	✓
Flexible Configuration	✓		✓
IASP Configuration and Management		Limited	✓
Add/Remove Multiple Monitored Resources			✓
Guided Wizards	✓	Limited	✓

Figure 5-8 Main differences between the graphical interfaces

As the PowerHA GUI is a combination of the two other ones, those GUIs will be withdrawn in a later release.

5.1.5 N_Port ID virtualization support

The new N_Port ID Virtualization (NPIV) support made available with IBM i 6.1.1 or later is fully supported by IBM PowerHA SystemMirror for i for IBM System Storage DS8000 series storage-based replication.

Using NPIV with PowerHA SystemMirror for i does not require dedicated Fibre Channel IOAs for each SYSBAS and IASP because the (virtual) IOP reset that occurs when you switch the IASP affects the virtual Fibre Channel client adapter only, instead of all ports of the physical Fibre Channel IOA, which are reset in native-attached storage environment.

For an overview of the new NPIV support by IBM i, see Chapter 8, “Virtualization” on page 311.

For further information about NPIV implementation in an IBM i environment, see *DS8000 Copy Services for IBM i with VIOS*, REDP-4584.

5.1.6 Asynchronous geographic mirroring

Asynchronous geographic mirroring is a new function that is supported by PowerHA SystemMirror for i Enterprise Edition with IBM i 7.1 that extends the previously available synchronous geographic mirroring option, which, for performance reasons, is limited to metro area distances up to 30 km.

The asynchronous delivery of geographic mirroring (not to be confused with the asynchronous mirroring mode of synchronous geographic mirroring) allows IP-based hardware replication beyond synchronous geographic mirroring limits.

Asynchronous delivery, which also requires the asynchronous mirroring mode, works by duplicating any changed IASP disk pages in the *BASE memory pool on the source system and sending them asynchronously while you preserve the write-order to the target system. Therefore, at any time, the data on the target system (though not up-to-date) still represents a so-called crash-consistent copy of the source system.

With the source system available, you can check the currency of the target system and memory impact on the source system because of asynchronous geographic mirroring. Use the Display ASP Session (**DSPASPSSN**) command to show the total data in transit, as shown in Figure 5-9.

Display ASP Session					
				04/09/10	15:53:50
Session	:	GEO			
Type	:	*GEOMIR			
Transmission Delivery	:	*ASYNC			
Mirroring Mode	:	*ASYNC			
Total data in transit	:	0.02 MB			
Suspend timeout	:	240			
Synchronization priority	:	*MEDIUM			
Tracking space allocated	:	100%			
Copy Descriptions					
ASP	ASP			Data	
Device	Copy	Role	State	State	Node
GE0001	GE0001S2	PRODUCTION	AVAILABLE	USABLE	RCHASHAM
GE0001	GE0001S1	MIRROR	ACTIVE	UNUSABLE	RCHASEGS

Figure 5-9 DSPASPSSN command data in transit information

For ASP sessions of type *GEOMIR, the changing of geographic mirroring options requires that the IASP be varied off. The option for asynchronous delivery can be enabled through the Change ASP Session (**CHGASPSSN**) command's new **DELIVERY(*ASYNC)** parameter, as shown in Figure 5-10.

Change ASP Session (CHGASPSSN)		
Type choices, press Enter.		
Session	SSN	
Option	OPTION	
ASP copy:	ASPCPY	
Preferred source		*SAME
Preferred target		*SAME
	+ for more values	
Suspend timeout	SSPTIMO	*SAME
Transmission delivery	DELIVERY	*ASYNC
Mirroring mode	MODE	*SAME
Synchronization priority	PRIORITY	*SAME
Tracking space	TRACKSPACE	*SAME
FlashCopy type	FLASHTYPE	*SAME
Persistent relationship	PERSISTENT	*SAME
ASP device	ASPDEV	*ALL
	+ for more values	
Track	TRACK	*YES
		More...

Figure 5-10 CHGASPSSN command - *ASYNC Transmission delivery parameter

This setting can also be seen and changed through the IBM Navigator for i GUI by clicking **PowerHA → Independent IASPs → Details...** (in the menu that is displayed by clicking next to the IASP name). Click the drop-down menu under the large arrow in the middle of the topic pane and select **Mirroring Properties**, as shown in Figure 5-11 and Figure 5-12 on page 136. You must stop geographic mirroring session by running the **ENDASPSSN** command before you change this setting.




Figure 5-11 Independent ASP details in PowerHA GUI




Figure 5-12 Independent ASP Mirroring Properties in PowerHA GUI

5.1.7 LUN level switching

LUN level switching is a new function that is provided by PowerHA SystemMirror for i in IBM i 7.1. It provides a local high availability solution with IBM System Storage DS8000 or DS6000 series similar to what used to be available as switched disks for IBM i internal storage.

With LUN level switching single-copy, that is, non-replicated, IASPs that are managed by a cluster resource group device domain and in IBM System Storage DS8000 or DS6000 series can be switched between IBM i systems in a cluster.

A typical implementation scenario for LUN level switching is where multi-site replication through Metro Mirror or Global Mirror is used for disaster recovery and protection against storage subsystem outages. When this scenario happens, additional LUN level switching at the production site is used for local high availability protection, eliminating the requirement for a site-switch if there are IBM i server outages.

To implement LUN level switching, create an ASP copy description for each switchable IASP using Add ASP Copy Description (ADDASPCPYD), which is enhanced with recovery domain information for LUN level switching, as shown in Figure 5-13.

Add ASP Copy Description (ADDASPCPYD)

Type choices, press Enter.

Logical unit name: LUN

TotalStorage device *NONE

Logical unit range

+ for more values

Consistency group range

+ for more values

Recovery domain: RCYDMN

Cluster node *NONE

Host identifier

+ for more values

Volume group

+ for more values

+ for more values

Bottom

F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display

F24=More keys

Figure 5-13 IBM i ADDASPCPYD enhancement for LUN level switching

An ASP session is not required for LUN level switching, as there is no replication for the IASP involved.

Important: Setting up an ASP copy description for LUN level switching is supported only from the text-based interface.

For LUN level switching, the backup node host connection on the DS8000 or DS6000 storage system must not have a volume group (VG) assigned. PowerHA automatically unassigns the VG from the production node and assigns it to the backup node at site-switches or failovers.

5.1.8 Space-Efficient FlashCopy

PowerHA for SystemMirror for i with IBM i 7.1 now supports Space-Efficient FlashCopy of the IBM System Storage DS8000 series.

You can use the IBM System Storage DS8000 series FlashCopy SE licensed feature to create Space-Efficient FlashCopy target volumes that can help you reduce the required physical storage space for the FlashCopy target volumes. These volumes are typically needed only for a limited time (such as during a backup to tape).

A Space-Efficient FlashCopy target volume has a virtual storage capacity that is reported to the host that matches the physical capacity of the fully provisioned FlashCopy source volume, but no physical storage space is ever allocated. Physical storage space for Space-Efficient FlashCopy target volumes is allocated in 64-KB track granularity. This allocation is done on demand for host write operations from a configured repository volume that is shared by all Space-Efficient FlashCopy target volumes within the same DS8000 extent pool, as shown in Figure 5-14.




Figure 5-14 DS8000 Space-Efficient FlashCopy

From a user perspective, the PowerHA setup (not the DS8000 FlashCopy setup) for Space-Efficient FlashCopy is identical to the setup for traditional FlashCopy with the **nocopy** option. The reason for this situation is PowerHA SystemMirror for i internally interrogates the DS8000 to determine the type of FlashCopy relationship and makes sure that it uses the corresponding correct DS CLI command syntax. The syntax check is done for either traditional FlashCopy or FlashCopy SE when you run the **mkflash** and **rmflash** commands.

For more information about using IBM System Storage DS8000 FlashCopy SE with IBM i, see *IBM System Storage Copy Services and IBM i: A Guide to Planning and Implementation*, SG24-7103.

5.1.9 Better detection of cluster node outages

There are situations where a sudden cluster node outage, such as a main storage memory dump, an HMC immediate partition power-off, or a system hardware failure, results in a partitioned cluster. If so, you are alerted by the failed cluster communication message CPFBB22 sent to QHST and the automatic failover not started message CPFBB4F sent to the QSYSOPR message queue on the first backup node of the CRG.

You must determine the reason for the cluster partition condition. This condition could be caused either by a network problem or a sudden cluster node outage. You must either solve the network communication problem or declare the cluster node as failed, which can be done by running Change Cluster Node Entry (**CHGCLUNODE**) in preparation of a cluster failover.

With IBM i 7.1, PowerHA SystemMirror for i now allows advanced node failure detection by cluster nodes. This task can be accomplished by registering with an HMC or Virtual I/O Server (VIOS) management partition on IVM managed systems. The clusters are notified when severe partition or system failures trigger a cluster failover event instead of causing a cluster partition condition.

For LPAR failure conditions, it is the IBM POWER Hypervisor™ (PHYP) that notifies the HMC that an LPAR failed. For system failure conditions other than a sudden system power loss, it is the flexible service processor (FSP) that notifies the HMC of the failure. The CIM server on the HMC or VIOS can then generate a power state change CIM event for any registered CIM clients.

Whenever a cluster node is started, for each configured cluster monitor, IBM i CIM client APIs are used to subscribe to the particular power state change CIM event. The HMC CIM server generates such a CIM event and actively sends it to any registered CIM clients (that is, there is no heartbeat polling that is involved with CIM). On the IBM i cluster nodes, the CIM event listener compares the events with available information about the nodes that constitute the cluster to determine whether it is relevant for the cluster to act upon. For relevant power state change CIM events, the cluster heartbeat timer expiration is ignored (that is, IBM i clustering immediately triggers a failover condition in this case).

Using advanced node failure detection requires SSH and CIMOM TCP/IP communication to be set up between the IBM i cluster nodes and the HMC or VIOS. Also, a cluster monitor must be added to the IBM i cluster nodes, for example, through the new Add Cluster Monitor (**ADDCLUMON**) command, as shown in Figure 5-15. This command enables communication with the CIM server on the HMC or VIOS.

Add Cluster Monitor (ADDCLUMON)

Type choices, press Enter.

Cluster

Node identifier

Monitor type

CIM server:

CIM server host name

CIM server user id

CIM server user password . . .

NAME

NAME

*CIMSVR

HMC1

hmcuser

password

Bottom

F3=Exit

F4=Prompt

F5=Refresh

F12=Cancel

F13=How to use this display

F24=More keys

Figure 5-15 IBM i ADDCLUMOD command

For more information about configuring clustering advanced node failure detection, see the IBM i 7.1 Information Center at:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=%2Frzaue%2Frzaueconceptnodefailureddetection.htm>

5.1.10 Improved geographic mirroring full synchronization performance

Performance improvements were implemented in IBM i 7.1 for geographic mirroring full synchronization.

Changes within the System Licensed Internal Code (SLIC) provide more efficient processing of data that is sent to the target system if there is a full resynchronization. Even with source and target side tracking, there are instances that require a full synchronization of the production copy, such as any time that the IASP cannot be normally varied off, because of a sudden cluster node outage.

The achievable performance improvement varies based on the IASP data. IASPs with many small objects see more benefit than those IASPs with a smaller number of large objects.

5.1.11 Improvement of geographic mirroring in an IBM i hosted IBM i client partition environment

In an IBM i hosted IBM i client partition environment, you can replicate the whole storage spaces of the IBM i client partition by configuring *NWSSTG for the client partition in an IASP of the IBM i host partition and using the geographic mirroring between the IBM i host partitions, as shown in Figure 5-16.




Figure 5-16 Geographic mirroring for an IBM i hosted IBM i client partition environment

The advantage of this solution is that no IASP is needed on the production (client) partition, so no application changes are required. Conversely, the following two considerations are part of this solution:

- ▶ It is a cold-standby solution, so an IPL on the backup IBM i client partition is required after you switch to the backup node.
- ▶ Temporary storage spaces on the IBM i client partition are also transferred to the backup node side.

In October 2012, geographic mirroring of PowerHA SystemMirror for i can eliminate the transfer of temporary storage spaces for the IBM i client partition. This enhancement reduces the amount of a network traffic between IBM i host partition on the production node side and on the backup node side.

5.1.12 Support virtualization capabilities on an active cluster node

IBM i 7.1 supports Suspend/Resume and Live Partition Mobility, which are provided by the PowerVM virtualization technology stack. If your IBM i 7.1 partition is a member of an active cluster, you cannot suspend this partition, but can move it to another Power Systems server by using Live Partition Mobility.

For more information about Suspend/Resume and Live Partition Mobility, see Chapter 8, "Virtualization" on page 311.

5.1.13 Cluster administrative domain enhancements

The IBM cluster administrative domain support was enhanced in IBM i 7.1 with the following two new monitored resource entries (MREs):

- ▶ Authorization lists (*AUTL)
- ▶ Printer device descriptions (*PRTDEV) for LAN or virtual printers

PowerHA SystemMirror for i is required to support these two new administration domain monitored resource entries.

For a complete list of attributes that can be monitored and synchronized among cluster nodes by the cluster administrative domain in IBM i 7.1, see the IBM i 7.1 Information Center at:

<http://public.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzaig/rzaigrefattribmon.htm>

Another enhancement for cluster administrative domains is available through PTF SI39263 for IBM i 7.1 and PTF SI39264 for IBM i 6.1, which you can use to add MREs even if the object cannot be created. An inconsistent state of the MRE reminds you that the object still must be manually created.

5.1.14 IPv6 support

PowerHA SystemMirror for i on IBM i 7.1 now fully supports IPv6 or a mix of IPv6 and IPv4. All HA-related APIs, commands, and GUIs are extended for field names that hold either a 32-bit IPv4 or a 128-bit IPv6 address, as shown for the Change Cluster Node Entry (**CHGCLUNODE**) command in Figure 5-17. An IPv6 IP address is specified in the form x:x:x:x:x:x:x, with x being a hexadecimal number 0 - FFFF, and “::” can be used once in the IPv6 address to indicate one or more groups of 16-bit zeros.

Change Cluster Node Entry (CHGCLUNODE)

Type choices, press Enter.

Cluster > *HASM_CLU*

Name

Node identifier > *CTCV71*

Name

Option > **CHGIFC*

*ADDIFC, *RMVIFC, *CHGIFC...

Old IP address

New IP address

: New IP address (NEWINTNETA) - Help :

: :

: Specifies the cluster interface address which is being :

: added to the node information or replacing an old cluster :

: interface address. The interface address may be an IPv4 :

: address (for any cluster version) or **IPv6 address** (if :

: current cluster version is 7 or greater). :

: :

: More... :

: F2=Extended help F10=Move to top F12=Cancel :

F3=Exit F4=: F13=Information Assistant F20=Enlarge F24=More keys :

F24=More keys :

:.....:

Figure 5-17 IBM i change cluster node entry

5.1.15 New CL commands for programming cluster automation

With PowerHA SystemMirror for i, the following new CL commands are introduced in IBM i 7.1 to better support CL programming for cluster automation management:

- ▶ Retrieve Cluster (**RTVCLU**) command
- ▶ Retrieve Cluster Resource Group (**RTVCRG**) command
- ▶ Retrieve ASP Copy Description (**RTVASPCPYD**) command
- ▶ Retrieve ASP Session (**RTVASPSSN**) command
- ▶ Print Cluster Administrative Domain Managed Resource Entry (**PRTCADMRE**) command

For more information about the new PowerHA CL commands, see the IBM i 7.1 Information Center at:

<http://public.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp?topic=/rzaue/rzauewhatnew.htm>

5.2 Journaling and commitment control enhancements

Journaling and commitment control are the base building blocks for any HA solution, as they ensure database consistency and recoverability.

Several enhancements were made in the area of the integrity preservation and journaling. The main objectives of these enhancements are to provide easier interfaces for the setup and monitoring of the database's persistence, including HA setups.

5.2.1 Journal management

Journals (more familiarly known as *logs* on other platforms) are used to track changes to various objects. Although the OS has built-in functions to protect the integrity of certain objects, use journaling to protect the changes to objects, to reduce the recovery time of a system after an abnormal end, to provide powerful recovery and audit functions, and to enable the replication of journal entries on a remote system.

The Start Journal Library (**STRJRNLIB**) command was introduced in IBM i 6.1. This command defines one or more rules at a library or schema level. These rules are used, or *inherited*, for journaling objects.

In the IBM i 7.1 release, **STRJRNLIB** (see Figure 5-18) now provides two new rules:

- If these objects are eligible for remote journal filtering by object (*OBJDFT, *NO, or *YES).
- A name filter to associate with the inherit rule. This filter can be specified with a specific or generic name. The default is to apply the rule to all objects that match the other criteria that are specified in the inherit rule regardless of the object name. You can use this filter to start journaling on new production work files, but not journal temporary work files if they have unique names.

Start Journal Library (STRJRNLIB)		
Type choices, press Enter.		
Library	> LIBA	Name, generic*
+ for more values		
Journal	> QSQJRN	Name
Library	> AJRNLIB	Name, *LIBL, *CURLIB
Inherit rules:		
Object type	> *FILE	*ALL, *FILE, *DTAARA, *DTAQ
Operation	*ALLOPR	*ALLOPR, *CREATE, *MOVE...
Rule action	*INCLUDE	*INCLUDE, *OMIT
Images	*OBJDFT	*OBJDFT, *AFTER, *BOTH
Omit journal entry	*OBJDFT	*OBJDFT, *NONE, *OPNCLO
Remote journal filter	*OBJDFT	*OBJDFT, *NO, *YES
Name filter	*ALL	Name, generic*, *ALL
+ for more values		
Logging level	*ERRORS	*ERRORS, *ALL

Figure 5-18 STRJRNLIB command prompt

If the library is journaled already and you want to define one of the new inherit rules, run the Change Journaled Object (CHGJRNOBJ) command. If the library is not already journaled, the new rules can be specified by running the Start Journal Library (STRJRNLIB) command. To view the inherit rules that are associated with a journaled library, run the Display Library Description (DSPLIBD) command, and then click F10 - Display inherit rules.

There is an equivalent in the IBM Navigator for i to do the same task. Click **Expand File Systems** → **Select Integrated File System** → **Select QSYS.LIB**. Select the library that you want to journal, as shown in Figure 5-19.




Figure 5-19 Select a library for Journaling

Select the Journaling action, as shown in Figure 5-20.




Figure 5-20 Setting a rule

5.2.2 Remote journaling

When a remote journal connection ends with a recoverable error, you can now specify that the operating system try to restart the connection automatically. This action is done by identifying the number of attempts and the time, expressed in seconds, between restart attempts. These settings can be set by running the Change Remote Journal (**CHGRMTJRN**) command or with the Change Journal State (QJOCHANGEJOURNALSTATE) API. For a list of errors for which an automatic restart attempt is made, see the Journal management topic collection in the IBM i Information Center at:

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

You can also run this command to filter the remote journals. Filtering out journal entries that are not needed on the target system can decrease the amount of data that is sent across the communication line.

This remote journal filtering feature is available with option 42 of IBM i, that is, feature 5117 (HA Journal Performance). Ensure that critical data is not filtered when you define remote journal filtering. Three criteria can be used to filter entries sent to the remote system:

- ▶ Before images
- ▶ Individual objects
- ▶ Name of the program that deposited the journal entry on the source system

The filtering criteria are specified when you activate a remote journal. Different remote journals that are associated with the same local journal can have different filtering criteria. Remote journal filtering can be specified only for asynchronous remote journal connections. Because journal entries might be missing, filtered remote journal receivers cannot be used with the Remove Journalized Changes (**RMVJRNCHG**) command. Similarly, journal receivers that filtered journal entries by object or by program cannot be used with the Apply Journalized Change (**APYJRNCHG**) command or the Apply Journalized Change Extend (**APYJRNCHGX**) command.

The Work with Journal Attributes (**WRKJRNA**) command can now monitor, from the target side, how many seconds the target is behind in receiving journal entries from the source system. Also new in IBM i 7.1 is the ability, from the source side, to view the number of retransmissions that occur for a remote journal connection.

5.2.3 DISPLAY_JOURNAL (easier searches of a journal)

Displaying a journal entry from a GUI interface either requires using APIs or writing the journal entries to an outfile. The APIs are labor-intensive and the outfile is restrictive and slower because a copy of the data required.

QSYS2.Display_Journal is a new table function that you can use to view entries in a journal by running a query.

There are many input parameters of the table function that can (and should) be used for best performance to return only those journal entries that are of interest. For more information about the special values, see the “QjoRetrieveJournalEntries API” topic in the Information Center at:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topics/apis/QJ0RJRNE.htm>

Unlike many other UDTFs in QSYS2, this one has no DB2 for i provided view.

Here is a brief summary of the parameters:

- ▶ **Journal_Library and Journal_Name**

The Journal_Library and Journal_Name must identify a valid journal. *LIBL and *CURLIB are *not* allowed as values of the Journal_Library.

- ▶ **Starting_Receiver_Library and Starting_Receiver_Name**

If the specified Starting_Receiver_Name is the null value, an empty string, or a blank string, *CURRENT is used and the Starting_Receiver_Library is ignored. If the specified Starting_Receiver_Name contains the special values *CURRENT, *CURCHAIN, or *CURAVLCHN, the Starting_Receiver_Library is ignored. Otherwise, the Starting_Receiver_Name and Starting_Receiver_Library must identify a valid journal receiver. *LIBL and *CURLIB may be used as a value of the Starting_Receiver_Library. The ending journal receiver cannot be specified and is always *CURRENT.

- ▶ **Starting_Timestamp**

If the specified Starting_Timestamp is the null value, no starting time stamp is used. A value for Starting_Timestamp and Starting_Sequence cannot both be specified at the same time. However, both values can be queried when querying the table function.

- ▶ **Starting_Sequence**

If the specified Starting_Sequence is the null value, no starting sequence number is used. If the specified Starting_Sequence is not found in the receiver range, an error is returned. A value for Starting_Timestamp and Starting_Sequence cannot both be specified at the same time. However, both values can be queried when querying the table function.

- ▶ **Journal_Codes**

If the specified Journal_Codes is the null value, an empty string, or a blank string, *ALL is used. Otherwise, the string may consist of the special value *ALL, the special value *CTL, or a string that contains one or more journal codes. Journal codes may be separated by one or more separators. The separator characters are the blank and comma. For example, a valid string may be 'RJ', 'R J', 'R,J', or 'R, J'.

- ▶ **Journal_Entry_Types**

If the specified Journal_Entry_Types is the null value, an empty string, or a blank string, *ALL is used. Otherwise, the string may consist of the special value *ALL, the special value *RCD, or a string that contains one or more journal entry types. Journal entry types may be separated by one or more separators. The separator characters are the blank and comma. For example, a valid string may be 'RJ', 'R J', 'R,J', or 'R, J'.

- ▶ **Object_Library, Object_Name, Object_ObjType, and Object_Member**

If the specified Object_Name is the null value, an empty string, or a blank string, no object name is used and the Object_Library, Object_ObjType, and Object_Member are ignored.

Otherwise, if the specified Object_Name contains the special value *ALL, the Object_Library must contain a library name and Object_ObjType must contain a valid object type (for example, *FILE).

Otherwise, only one object may be specified and the Object_Library, Object_Name, Object_ObjType, and Object_Member must identify a valid object. *LIBL and *CURLIB may be used as a value of the Object_Library.

The Object_ObjType must be one of *DTAARA, *DTAQ, *FILE, or *LIB (*LIB is 6.1 only). The Object_Member may be *FIRST, *ALL, *NONE, or a valid member name. If the specified object type is not *FILE, the member name is ignored.

- User

If the specified user is the null value, an empty string, or a blank string, *ALL is used. Otherwise, the user must identify a valid user profile name.
- Job

If the specified job is the null value, an empty string, or a blank string, *ALL is used. Otherwise, the job must identify a valid job name, that is, a specific job where the first 10 characters are the job name, the second 10 characters are the user name, and the last six characters are the job number.
- Program

If the specified program is the null value, an empty string, or a blank string, *ALL is used. Otherwise, the program must identify a valid program name.

Example 5-1 gives a possible usage of the DISPLAY_JOURNAL function.

Example 5-1 Possible usage of DISPLAY_JOURNAL function

```

set path system path, jsochr; -- Change jsochr to your library you chose above

-- Select all entries from the *CURRENT receiver of journal mjatst/qsqjrn.

select * from table (
Display_Journal(
  'JSOCHR', 'QSQJRN', -- Journal library and name
  '', '', -- Receiver library and name
  CAST(null as TIMESTAMP), -- Starting timestamp
  CAST(null as DECIMAL(21,0)), -- Starting sequence number
  '', -- Journal codes
  '', -- Journal entries
  '', '', '', -- Object library, Object name, Object type, Object member
  '', -- User
  '', -- Job
  '' -- Program
) ) as x;

```

This function provides a result table with data similar to what you get from using the Display Journal Command (DSPJRN) command.

5.2.4 Commitment control and independent ASPs

You can use commitment control to define the boundaries of a business or logical transaction, identifying when it starts and where it ends, and to ensure that all the database changes are either applied permanently or removed permanently. Furthermore, if any process or even a complete system performs such transactions ends abnormally, commitment control provides recovery of pending transactions by bringing the database contents to a committed status, and identifies the last transactions that were pending and recovered.

With commitment control, you have assurance that when the application starts again, no partial updates are in the database because of incomplete transactions from a prior failure. As such, it is one of the building blocks of any highly available setup and it identifies the recovery point for any business process.

If your application was deployed using independent ASPs (IASPs), you are using a database instance that is in that IASP. This situation has an impact on how commitment control works.

When a process starts commitment control, a commitment definition is created in a schema (QRECOVERY) that is stored in the database to which the process is connected. Assuming that your process is connected to an IASP, commitment control is started in the database that is managed by the IASP. When your process is running commitment control from an IASP (that is, it has its resources registered with commitment control on that disk pool), switching to another disk pool fails and throws message CPDB8EC (The thread has an uncommitted transaction).

However, if you switch from the system disk pool (ASP group *NONE), commitment control is not affected. The commitment definitions stay on the system disk pool. A new feature in IBM i 7.1 is that if you later place independent disk pool resources under commitment control before system disk pool resources, the commitment definition is moved to the independent disk pool. This situation means that if your job is not associated with an independent ASP, the commitment definition is created in *SYSBAS; otherwise, it is created in the independent ASP. If the job is associated with an independent ASP, you can open files under commitment control that are in the current library name space. For example, they can be in the independent ASP or *SYSBAS.

If the first resource that is placed under commitment control is not in the same ASP as the commitment definition, the commitment definition is moved to the resource's ASP. If both *SYSBAS and independent ASP resources are registered in the same commitment definition, the system implicitly uses a two-phase commit protocol to ensure that the resources are committed atomically in the event of a system failure. Therefore, transactions that involve data in both *SYSBAS and an independent ASP have a small performance degradation versus transactions that are isolated to a single ASP group.

When recovery is required for a commitment definition that contains resources that are in both *SYSBAS and an independent ASP, the commitment definition is split into two commitment definitions during the recovery, one in *SYSBAS and one in the independent ASP, as though there were a remote database connection between the two ASP groups. Resynchronization can be initiated by the system during the recovery to ensure that the data in both ASP groups is committed or rolled back atomically.

5.2.5 System Managed Access Path Protection

You can use System Managed Access Path Protection (SMAPP) to reduce the time for the system or independent disk pool to restart after an abnormal end. When the system must rebuild access paths, the next restart takes longer to complete than if the system ended normally. When you use SMAPP, the system protects the access paths implicitly and eliminates the rebuild of the access paths after an abnormal end.

SMAPP affects the overall system performance. The lower the target recovery time that you specify for access paths, the greater this effect can be. Typically, the effect is not noticeable, unless the processor is nearing its capacity.

Another situation that can cause an increase in processor consumption is when local journals are placed in standby state and large access paths that are built over files that are journaled to the local journal are modified. Using the F16=Display details function from the Display Recovery for Access Paths (DSPRCYAP) or Edit Recovery for Access Paths (EDTRCYAP) panels shows the internal threshold that is used by SMAPP (see Figure 5-21). This panel was added in IBM i 7.1. All access paths with estimated rebuild times greater than the internal threshold are protected by SMAPP. The internal threshold value might change if the number of exposed access paths changes, the estimated rebuild times for exposed access paths changes, or if the target recovery time changes.

Display Details		CTCV71	
		03/15/10	12:46:18
ASP	: *SYSTEM		
Internal threshold	: 00:52:14		
Last retune:			
Date	: 03/09/10		
Time	: 06:54:58		
Last recalibrate:			
Date	: 02/24/10		
Time	: 08:19:44		

Figure 5-21 Display Details from the Edit and Display Recovery for Access Paths panel

5.2.6 Journaling and disk arm usage

Starting in IBM i 7.1 with PTF MF51614, journal receivers are spread across all disk arms in the disk pool. Journaling no longer directs writes to specific arms.

The journal receiver threshold value influences the number of parallel writes that journal allows. The higher the journal receiver threshold value, the more parallel I/O requests are allowed. Allowing more parallel I/O requests may improve performance.

For more information, see the TechDocs “IBM i 7.1 and changes for journaling” at:

<http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD105661>

5.2.7 Range of journal receivers parameter and *CURAVLCHN

In IBM i 6.1 with DB2 PTF Group SF99601 Level 28 and IBM i 7.1 with DB2 PTF Group SF99701 Level18, the Range of Journal Receivers parameter (**RCVRNG**) is extended to allow a ***CURAVLCHN** option with the **RTVJRNE**, **RCVJRNE**, and **DSPJRN** commands.

If you saved any journal receivers using **SAVOBJ** or **SAVLIB** with the **STG(*FREE)** option, the receiver chain is effectively broken and the ***CURCHAIN** option fails to retrieve journal entries.

By specifying the ***CURAVLCHN** option, if journal receivers exist in the receiver chain that are not available because they were saved with the storage freed option, those journal receivers are ignored and the entries are retrieved starting with the first available journal receiver in the chain.

The ***CURAVLCHN** option is already supported on the **QjoRetrieveJournalEntries()** API. With these enhancements, the three journal commands have the same capability.

5.2.8 Journal management functions: IBM Navigator for i

IBM Navigator for i now supports additional journal management functions. With IBM i 7.1, the following functions were added:

- ▶ Change journal receivers and attributes that are associated with a journal.
- ▶ View the properties that are associated with a journal receiver.
- ▶ View the objects that are journaled to a specific journal.
- ▶ Add and remove remote journals.
- ▶ View the list of remote journals that are associated with a specific journal.
- ▶ Activate and deactivate remote journals.
- ▶ View the details of a remote journal connection.

For more information, see Chapter 18, “IBM Navigator for i V7.1” on page 639.

5.3 Additional information

For more information and in-depth details about the latest enhancements in IBM PowerHA SystemMirror for i, see *PowerHA SystemMirror for IBM i Cookbook*, SG24-7994.



IBM DB2 for i

This chapter describes what is new for DB2 for i.

This chapter describes the following topics:

- ▶ SQL data description and data manipulation language
- ▶ Performance and query optimization
- ▶ New functionality for DB2 developers
- ▶ DB2 management and recovery enhancements
- ▶ Enhancements to products related to DB2 for i

6.1 Introduction: Getting around with data

DB2 for i is a member of the IBM leading-edge family of DB2 products. It has always been known and appreciated for its ease of use and simplicity. It supports a broad range of applications and development environments.

Because of its unique and self-managing computing features, the cost of ownership of DB2 for i is a valuable asset. The sophisticated cost-based query optimizer, the unique single level store architecture of the OS, and the database parallelism feature of DB2 for i allow it to scale almost linearly. Rich SQL support makes not only it easier for software vendors to port their applications and tools to IBM i, but it also enables developers to use industry-standard SQL for their data access and programming. The IBM DB2 Family has this focus on SQL standards with DB2 for i, so investment in SQL enables DB2 for i to use the relational database technology leadership position of IBM and maintain close compatibility with the other DB2 Family products.

Reading through this chapter, you find many modifications and improvements as part of the new release. All of these features are available to any of the development and deployment environments that are supported by the IBM Power platforms on which IBM i 7.1 can be installed.

Many DB2 enhancements for IBM i 7.1 are also available for Version 6.1. If you must verify their availability, go to:

<https://www.ibm.com/developerworks/ibmi/techupdates/db2>

This link takes you to the DB2 for i section of the IBM i Technology Updates wiki.

6.2 SQL data description and data manipulation language

There are several changes and additions to the SQL language:

- ▶ XML support
- ▶ Support for the **MERGE** statement
- ▶ Global variables
- ▶ Support for arrays in procedures
- ▶ Result set support in embedded SQL
- ▶ Encryption enhancements (**FIELDPROCS**)
- ▶ Removal of identity column and constraint restrictions on partitioned tables
- ▶ IBM WebSphere MQ integration functions
- ▶ Parameter marker enhancements
- ▶ Expressions in the **CALL** statement
- ▶ Three-part names in statements and aliases
- ▶ Currently committed concurrent access resolution
- ▶ **REPLACE** option on **CREATE** commands
- ▶ BIT scalar functions
- ▶ Encoded vector indexes **INCLUDE** of aggregate functions
- ▶ Inlining of SQL scalar function
- ▶ **INSERT** with remote **SUBSELECT**
- ▶ SQL Procedure and function obfuscation
- ▶ **CREATE TABLE** with remote **SUBSELECT**
- ▶ Qualified name option added to Generate SQL

6.2.1 XML support

Extensible Markup Language (XML) is a simple and flexible text format that is derived from SGML (ISO 8879). Originally designed to meet the challenges of large-scale electronic publishing, XML is also playing an increasingly important role in the exchange of a wide variety of data on the web and elsewhere.

For more information about XML, go to:

<http://www.w3.org/XML/>

Until this release, XML data types were supported only through user-defined types. At the same time, any handling of XML data must be done using user-defined functions. In IBM i 7.1, the DB2 component is complemented with the support for XML data types and publishing functions. It also supports XML document and annotation, document search (OmniFind) without decomposition, and client and language API support for XML (CLI, ODBC, JDBC, and so on).

XML data type

An XML value represents well-formed XML in the form of an XML document, XML content, or an XML sequence. An XML value that is stored in a table as a value of a column that is defined with the XML data type must be a well-formed XML document. XML values are processed in an internal representation that is not comparable to any string value, including another XML value. The only predicate that can be applied to the XML data type is the IS NULL predicate.

An XML value can be transformed into a serialized string value that represents an XML document using the XMLSERIALIZE (see “XML serialization” on page 157) function. Similarly, a string value that represents an XML document can be transformed into an XML value using the XMLPARSE (see “XML publishing functions” on page 156) function. An XML value can be implicitly parsed or serialized when exchanged with application string and binary data types.

The XML data type has no defined maximum length. It does have an effective maximum length of 2 GB when treated as a serialized string value that represents XML, which is the same as the limit for Large Object (LOB) data types. Like LOBs, there are also XML locators and XML file reference variables.

With a few exceptions, you can use XML values in the same contexts in which you can use other data types. XML values are valid in the following circumstances:

- ▶ **CAST** a parameter marker, XML, or **NULL** to XML
- ▶ **XMLCAST** a parameter marker, XML, or **NULL** to XML
- ▶ **IS NULL** predicate
- ▶ **COUNT** and **COUNT_BIG** aggregate functions
- ▶ **COALESCE**, **IFNULL**, **HEX**, **LENGTH**, **CONTAINS**, and **SCORE** scalar functions
- ▶ XML scalar functions
- ▶ A **SELECT** list without **DISTINCT**
- ▶ **INSERT VALUES** clause, **UPDATE SET** clause, and **MERGE**
- ▶ **SET** and **VALUES INTO**
- ▶ Procedure parameters
- ▶ User-defined function arguments and result
- ▶ Trigger correlation variables
- ▶ Parameter marker values for a dynamically prepared statement

XML values cannot be used directly in the following places. Where expressions are allowed, an XML value can be used, for example, as the argument of **XMLSERIALIZE**.

- ▶ A **SELECT** list that contains the **DISTINCT** keyword
- ▶ A **GROUP BY** clause
- ▶ An **ORDER BY** clause
- ▶ A subselect of a fullselect that is not **UNION ALL**
- ▶ A basic, quantified, **BETWEEN**, **DISTINCT**, **IN**, or **LIKE** predicate
- ▶ An aggregate function with the **DISTINCT** keyword
- ▶ A primary, unique, or foreign key
- ▶ A check constraint
- ▶ An index column

No host languages have a built-in data type for the XML data type.

XML data can be defined with any EBCDIC single byte or mixed CCSID or a Unicode CCSID of 1208 (UTF-8), 1200 (UTF-16), or 13488 (Unicode-specific version). 65535 (no conversion) is not allowed as a CCSID value for XML data. The CCSID can be explicitly specified when you define an XML data type. If it is not explicitly specified, the CCSID is assigned using the value of the **SQL_XML_DATA_CCSID QAQQINI** file parameter (see 6.3.15, “QAQQINI properties” on page 192). If this value is not set, the default is 1208. The CCSID is established for XML data types that are used in SQL schema statements when the statement is run.

XML host variables that do not have a **DECLARE VARIABLE** that assigns a CCSID have their CCSID assigned as follows:

- ▶ If it is XML **AS DBCLOB**, the CCSID is 1200.
- ▶ If it is XML **AS CLOB** and the **SQL_XML_DATA_CCSID QAQQINI** value is 1200 or 13488, the CCSID is 1208.
- ▶ Otherwise, the **SQL_XML_DATA_CCSID QAQQINI** value is used as the CCSID.

Because all implicit and explicit **XMLPARSE** functions are performed using UTF-8 (1208), defining data in this CCSID removes the need to convert the data to UTF-8.

XML publishing functions

Table 6-1 describes the functions that are directly used in a SQL query.

Table 6-1 XML publishing functions

Function	Description
xmlagg	Combines a collection of rows, each containing a single XML value to create an XML sequence that contains an item for each non-null value in a set of XML values.
xmlattributes	Returns XML attributes from columns, using the name of each column as the name of the corresponding attribute.
xmlcomment	Returns an XML value with the input argument as the content.
xmlconcat	Returns a sequence that contains the concatenation of a variable number of XML input arguments.
xmldocument	Returns an XML document.
xmlelement	Returns an XML element.
xmlforest	Returns an XML value that is a sequence of XML element nodes.
xmlgroup	Returns a single top-level element to represent a table or the result of a query.

Function	Description
xmlnamespaces	Constructs namespace declarations from the arguments.
xmlparse	Parses the arguments as an XML document and returns an XML value.
xmlpi	Returns an XML value with a single processing instruction.
xmlrow	Returns a sequence of row elements to represent a table or the result of a query.
xmlserialize	Returns a serialized XML value of the specified data type generated from the XML-expression argument.
xmltext	Returns an XML value that has the input argument as the content.
xmlvalidate	Returns a copy of the input XML value that is augmented with information obtained from XML schema validation, including default values and type annotations.
xslttransform	Converts XML data into other forms, accessible for the XSLT processor, including but not limited to XML, HTML, or plain text.

You can use the **SET CURRENT IMPLICIT XMLPARSE OPTION** statement to change the value of the **CURRENT IMPLICIT XMLPARSE OPTION** special register to **STRIP WHITESPACE** or to **PRESERVE WHITESPACE** for your connection. You can either remove or maintain any white space on an implicit **XMLPARSE** function. This statement is not a committable operation.

XML serialization

XML serialization is the process of converting XML data from the format that it has in a DB2 database to the serialized string format that it has in an application.

You can let the DB2 database manager perform serialization implicitly, or you can start the **XMLSERIALIZE** function to request XML serialization explicitly. The most common usage of XML serialization is when XML data is sent from the database server to the client.

Implicit serialization is the preferred method in most cases because it is simpler to code, and sending XML data to the client allows the DB2 client to handle the XML data properly. Explicit serialization requires additional handling, which is automatically handled by the client during implicit serialization.

In general, implicit serialization is preferable because it is more efficient to send data to the client as XML data. However, under certain circumstances (for example, if the client does not support XML data) it might be better to do an explicit **XMLSERIALIZE**.

With implicit serialization for DB2 CLI and embedded SQL applications, the DB2 database server adds an XML declaration with the appropriate encoding specified to the data. For .NET applications, the DB2 database server also adds an XML declaration. For Java applications, depending on the SQLXML object methods that are called to retrieve the data from the SQLXML object, the data with an XML declaration added by the DB2 database server is returned.

After an explicit **XMLSERIALIZE** invocation, the data has a non-XML data type in the database server, and is sent to the client as that data type. You can use the **XMLSERIALIZE** scalar function to specify the SQL data type to which the data is converted when it is serialized (character, graphic, or binary data type) and whether the output data includes the explicit encoding specification (**EXCLUDING XMLDECLARATION** or **INCLUDING XMLDECLARATION**).

The best data type to which to convert XML data is the BLOB data type, because retrieval of binary data results in fewer encoding issues. If you retrieve the serialized data into a non-binary data type, the data is converted to the application encoding, but the encoding specification is not modified. Therefore, the encoding of the data most likely does not agree with the encoding specification. This situation results in XML data that cannot be parsed by application processes that rely on the encoding name.

Although implicit serialization is preferable because it is more efficient, you can send data to the client as XML data. When the client does not support XML data, you can consider doing an explicit **XMLSERIALIZE**. If you use implicit XML serialization for this type of client, the DB2 database server then converts the data to a **CLOB** (Example 6-1) or **DBCLOB** before it sends the data to the client.

Example 6-1 XMLSERIALIZE

```
SELECT e.empno, e.firstnme, e.lastname,  
       XMLSERIALIZE(XMLELEMENT(NAME "xmp:Emp",  
                               XMLNAMESPACES('http://www.xmp.com' as "xmp"),  
                               XMLATTRIBUTES(e.empno as "serial"),  
                               e.firstnme, e.lastname  
                               OPTION NULL ON NULL))  
       AS CLOB(1000) CCSID 1208  
       INCLUDING XMLDECLARATION) AS "Result"  
FROM employees e WHERE e.empno = 'A0001'
```

Managing XML schema repositories (XSR)

The XML schema repository (XSR) is a set of tables that contain information about XML schemas. XML instance documents might contain a reference to a Uniform Resource Identifier (URI) that points to an associated XML schema. This URI is required to process the instance documents. The DB2 database system manages dependencies on externally referenced XML artifacts with the XSR without requiring changes to the URI location reference.

Without this mechanism to store associated XML schemas, an external resource might not be accessible when needed by the database. The XSR also removes the additional impact that is required to locate external documents, along with the possible performance impact.

An XML schema consists of a set of XML schema documents. To add an XML schema to the DB2 XSR, you register XML schema documents to DB2 by calling the DB2 supplied stored procedure **SYSPROC.XSR_REGISTER** to begin registration of an XML schema.

The **SYSPROC.XSR_ADDSCHEMADOC** procedure adds more XML schema documents to an XML schema that you are registering. You can call this procedure only for an existing XML schema that is not yet complete.

Calling the **SYSPROC.XSR_COMPLETE** procedure completes the registration of an XML schema. During XML schema completion, DB2 resolves references inside XML schema documents to other XML schema documents. An XML schema document is not checked for correctness when you register or add documents. Document checks are performed only when you complete the XML schema registration.

To remove an XML schema from the DB2 XML schema repository, you can call the **SYSPROC.XSR_REMOVE** stored procedure or use the **DROP XSROBJECT** SQL statement.

More considerations: Because an independent auxiliary storage pool (IASP) can be switched between multiple systems, there are more considerations for administering XML schemas on an IASP. Use of an XML schema must be contained on the independent ASP where it was registered. You cannot reference an XML schema that is defined in an independent ASP group or in the system ASP when the job is connected to the independent ASP.

Annotated XML schema decomposition

Annotated XML schema decomposition, also referred to as *decomposition* or *shredding*, is the process of storing content from an XML document in columns of relational tables. Annotated XML schema decomposition operates based on annotations that are specified in an XML schema. After an XML document is decomposed, the inserted data has the SQL data type of the column into which it is inserted.

An XML schema consists of one or more XML schema documents. In annotated XML schema decomposition, or schema-based decomposition, you control decomposition by annotating a document's XML schema with decomposition annotations. These annotations specify the following details:

- ▶ The name of the target table and column in which the XML data is to be stored
- ▶ The default SQL schema for when an SQL schema is not identified
- ▶ Any transformation of the content before it is stored

The annotated schema documents must be stored in and registered with the XSR. The schema must then be enabled for decomposition. After the successful registration of the annotated schema, decomposition can be performed by calling the decomposition stored procedure `SYSPROC.XDBDECOMPXML`.

The data from the XML document is always validated during decomposition. If information in an XML document does not comply with its specification in an XML schema, the data is not inserted into the table.

Annotated XML schema decomposition can become complex. To make the task more manageable, take several things into consideration. Annotated XML schema decomposition requires you to map possible multiple XML elements and attributes to multiple columns and tables in the database. This mapping can also involve transforming the XML data before you insert it, or apply conditions for insertion.

Here are items to consider when you annotate your XML schema:

- ▶ Understand what decomposition annotations are available to you.
- ▶ Ensure, during mapping, that the type of the column is compatible with the XML schema type of the element or attribute to which it is being mapped.
- ▶ Ensure complex types that are derived by restriction or extension are properly annotated.
- ▶ Confirm that no decomposition limits and restrictions are violated.
- ▶ Ensure that the tables and columns that are referenced in the annotation exist at the time the schema is registered with the XSR.

XML decomposition enhancements (order of result rows)

In IBM i 7.1, a series of decomposition annotations are provided to define how to decompose an XML document into relational database tables, such as `db2-xdb:defaultSQLSchema` or `db2-xdb:rowSet`, `db2-xdb:column`.

In one XSR, multiple target tables can be specified, so data in an XML document can be shredded to more than one target tables using one XSR. But the order of insertion into tables cannot be specified with existing decomposition annotations, so if the target tables have a reference relationship, the insertion of dependent row fails if its parent row is not inserted before it.

Two new annotations are supported:

- db2-xdb:order

The db2-xdb:order annotation specifies the insertion order of rows among different tables.

- db2-xdb:rowSetOperationOrder

The db2-xdb:rowSetOperationOrder annotation is a parent for one or more db2-xdb:order elements.

Using db2-xdb:order and db2-xdb:rowSetOperationOrder is needed only when referential integrity constraints exist in target tables and you try to decompose to them using one XSR.

6.2.2 The MERGE statement

This statement enables the simplification of matching rows in tables so that you can use a single statement that updates a target (a table or view) using data from a source (result of a table reference). Rows might be inserted, updated, or deleted in the target row, as specified by the matching rules. If you insert, update, or delete rows in a view, without an **INSTEAD OF** trigger, it updates, deletes, or inserts the row into the tables on which the view is based.

More than one modification-operation (**UPDATE**, **DELETE**, or **INSERT**) or signal-statement can be specified in a single **MERGE** statement. However, each row in the target can be operated on only once. A row in the target can be identified only as **MATCHED** with one row in the result table of the table-reference. A nested SQL operation (RI or trigger except **INSTEAD OF** trigger) cannot specify the target table (or a table within the same hierarchy) as a target of an **UPDATE**, **DELETE**, **INSERT**, or **MERGE** statement. This statement is also often referred to as an *upsert*.

Using the **MERGE** statement is potentially good in a Business Intelligence data load scenario, where it can be used to populate the data in both the fact and the dimension tables upon a refresh of the data warehouse. It can also be used for archiving data.

In Example 6-2, the **MERGE** statement updates the list of activities that are organized by Group A in the archive table. It deletes all outdated activities and updates the activities information (description and date) in the archive table if they were changed. It inserts new upcoming activities into the archive, signals an error if the date of the activity is not known, and requires that the date of the activities in the archive table be specified.

Example 6-2 UPDATE or INSERT activities

```

MERGE INTO archive ar
  USING (SELECT activity, description, date, last_modified
        FROM activities_groupA) ac
  ON (ar.activity = ac.activity) AND ar.group = 'A'
  WHEN MATCHED AND ac.date IS NULL THEN
    SIGNAL SQLSTATE '70001'
    SET MESSAGE_TEXT =
      ac.activity CONCAT ' cannot be modified. Reason: Date is not known'
  WHEN MATCHED AND ac.date < CURRENT DATE THEN
    DELETE
  WHEN MATCHED AND ar.last_modified < ac.last_modified THEN

```



```

UPDATE SET
    (description, date, last_modified) = (ac.description, ac.date, DEFAULT)
WHEN NOT MATCHED AND ac.date IS NULL THEN
    SIGNAL SQLSTATE '70002'
    SET MESSAGE_TEXT =
        ac.activity CONCAT ' cannot be inserted. Reason: Date is not known'
WHEN NOT MATCHED AND ac.date >= CURRENT DATE THEN
    INSERT
        (group, activity, description, date)
        VALUES ('A', ac.activity, ac.description, ac.date)
ELSE IGNORE

```

Each group has an activities table. For example, activities_groupA contains all activities Group A organizes, and the archive table contains all upcoming activities that are organized by groups in the company. The archive table has (group, activity) as the primary key, and date is not nullable. All activities tables have activity as the primary key. The last_modified column in the archive is defined with CURRENT_TIMESTAMP as the default value.

There is a difference in how many updates are done depending on whether a **NOT ATOMIC MERGE** or an **ATOMIC MERGE** was specified:

- In an **ATOMIC MERGE**, the source rows are processed as though a set of rows is processed by each **WHEN** clause. Thus, if five rows are updated, any row level update trigger is fired five times for each **WHEN** clause. This situation means that n statement level update triggers are fired, where n is the number of **WHEN** clauses that contain an **UPDATE**, including any **WHEN** clause that contains an **UPDATE** that did not process any of the source rows.
- In a **NOT ATOMIC MERGE** setting, each source row is processed independently as though a separate **MERGE** statement ran for each source row, meaning that, in the previous case, the triggers are fired only five times.

After running a **MERGE** statement, the **ROW_COUNT** statement information item in the SQL Diagnostics Area (or SQLERRD(3) of the SQLCA) is the number of rows that are operated on by the **MERGE** statement, excluding rows that are identified by the **ELSE IGNORE** clause.

The **ROW_COUNT** item and SQLERRD(3) do not include the number of rows that were operated on as a result of triggers. The value in the **DB2_ROW_COUNT_SECONDARY** statement information item (or SQLERRD(5) of the SQLCA) includes the number of these rows.

No attempt is made to update a row in the target that did not exist before the **MERGE** statement ran. There are no updates of rows that were inserted by the **MERGE** statement.

If **COMMIT(*RR)**, **COMMIT(*ALL)**, **COMMIT(*CS)**, or **COMMIT(*CHG)** is specified, one or more exclusive locks are acquired during the execution of a successful insert, update, or delete. Until the locks are released by a commit or rollback operation, an inserted or updated row can be accessed only by either the application process that performed the insert or update or by another application process using **COMMIT(*NONE)** or **COMMIT(*CHG)** through a read-only operation.

If an error occurs during the operation for a row of source data, the row being processed at the time of the error is not inserted, updated, or deleted. Processing of an individual row is an atomic operation. Any other changes that are previously made during the processing of the **MERGE** statement are not rolled back. If **CONTINUE ON EXCEPTION** is specified, execution continues with the next row to be processed.

6.2.3 Creating and using global variables

You can use global variables to assign specific variable values for a session. Use the **CREATE VARIABLE** statement to create a global variable at the server level.

Global variables have a session scope, which means that although they are available to all sessions that are active on the database, their value is private for each session. Modifications to the value of a global variable are not under transaction control. The value of the global variable is preserved when a transaction ends with either a **COMMIT** or a **ROLLBACK** statement.

When a global variable is instantiated for a session, changes to the global variable in another session (such as **DROP** or **GRANT**) might not affect the variable that is instantiated. An attempt to read from or to write to a global variable created by this statement requires that the authorization ID attempting this action holds the appropriate privilege on the global variable. The definer of the variable is implicitly granted all privileges on the variable.

A created global variable is instantiated to its default value when it is first referenced within its given scope. If a global variable is referenced in a statement, it is instantiated independently of the control flow for that statement.

A global variable is created as a *SRVPGM object. If the variable name is a valid system name but a *SRVPGM exists with that name, an error is issued. If the variable name is not a valid system name, a unique name is generated using the rules for generating system table names.

If a global variable is created within a session, it cannot be used by other sessions until the unit of work is committed. However, the new global variable can be used within the session that created the variable before the unit of work commits.

Example 6-3 creates a global variable that defines a user class. This variable has its initial value set based on the result of starting a function that is called **CLASS_FUNC**. This function is assumed to assign a class value, such as administrator or clerk that is based on the **USER** special register value. The **SELECT** clause in this example lists all employees from department A00. Only a session that has a global variable with a **USER_CLASS** value of 1 sees the salaries for these employees.

Example 6-3 Creating and using global variables

```
CREATE VARIABLE USER_CLASS INT DEFAULT (CLASS_FUNC(USER))
GRANT READ ON VARIABLE USER_CLASS TO PUBLIC
SELECT
    EMPNO,
    LASTNAME,
    CASE
        WHEN USER_CLASS = 1
        THEN SALARY
        ELSE NULL
    END
FROM
    EMPLOYEE
WHERE
    WORKDEPT = 'A00'
```

6.2.4 Support for arrays in procedures

An *array* is a structure that contains an ordered collection of data elements in which each element can be referenced by its ordinal position in the collection. If N is the cardinality (number of elements) of an array, the ordinal position that is associated with each element is an integer value greater than or equal to 1 and less than or equal to N . All elements in an array have the same data type.

An *array type* is a data type that is defined as an array of another data type. Every array type has a maximum cardinality, which is specified on the **CREATE TYPE** (Array) statement. If A is an array type with maximum cardinality M , the cardinality of a value of type A can be any value 0 - M inclusive. Unlike the maximum cardinality of arrays in programming languages such as C, the maximum cardinality of SQL arrays is not related to their physical representation. Instead, the maximum cardinality is used by the system at run time to ensure that subscripts are within bounds. The amount of memory that is required to represent an array value is proportional to its cardinality, and not to the maximum cardinality of its type.

SQL procedures support parameters and variables of array types. Arrays are a convenient way of passing transient collections of data between an application and a stored procedure or between two stored procedures.

Within SQL stored procedures, arrays can be manipulated as arrays in conventional programming languages. Furthermore, arrays are integrated within the relational model in such a way that data represented as an array can be easily converted into a table and data in a table column can be aggregated into an array.

In Example 6-4, we use two array data types (`intArray` and `stringArray`), and a `persons` table with two columns (`ID` and `name`). The `processPersons` procedure adds three more persons to the table, and returns an array with the person names that contain the letter 'a', ordered by `ID`. The `IDs` and `names` of the three persons to be added are represented as two arrays (`IDs` and `names`). These arrays are used as arguments to the **UNNEST** function, which turns the arrays into a two-column table, whose elements are then inserted into the `persons` table. Finally, the last set statement in the procedure uses the **ARRAY_AGG** aggregate function to compute the value of the output parameter.

Example 6-4 Support for arrays in procedures

```
CREATE TYPE intArray AS INTEGER ARRAY[100]
CREATE TYPE stringArray AS VARCHAR(10) ARRAY[100]
CREATE TABLE persons (id INTEGER, name VARCHAR(10))
INSERT INTO persons VALUES(2, 'Tom'),
                           (4, 'Gina'),
                           (1, 'Kathy'),
                           (3, 'John')

CREATE PROCEDURE processPersons(OUT witha stringArray)
BEGIN
  DECLARE ids intArray;
  DECLARE names stringArray;
  SET ids = ARRAY[5,6,7];
  SET names = ARRAY['Denise', 'Randy', 'Sue'];
  INSERT INTO persons(id, name)
    (SELECT t.i, t.n FROM UNNEST(ids, names) AS t(i, n));
  SET witha = (SELECT ARRAY_AGG(name ORDER BY id)
    FROM persons
    WHERE name LIKE '%a%');
END
```

If **WITH ORDINALITY** is specified, an extra counter column of type BIGINT is appended to the temporary table. The ordinality column contains the index position of the elements in the arrays. See Example 6-5.

The ARRAY UNNEST temporary table is an internal data structure and can be created only by the database manager.

Example 6-5 UNNEST temporary table WITH ORDINALITY

```
CREATE PROCEDURE processCustomers()
BEGIN
  DECLARE ids INTEGER ARRAY[100];
  DECLARE names VARCHAR(10) ARRAY[100];
  set ids = ARRAY[5,6,7];
  set names = ARRAY['Ann', 'Bob', 'Sue'];
  INSERT INTO customerTable(id, name, order)
    (SELECT Customers.id, Customers.name, Customers.order
     FROM UNNEST(ids, names) WITH ORDINALITY
     AS Customers(id, name, order) );
END
```

6.2.5 Result set support in embedded SQL

You can write a program in a high-level language (C, RPG, COBOL, and so on) to receive results sets from a stored procedure for either a fixed number of result sets, for which you know the contents, or a variable number of result sets, for which you do not know the contents.

Returning a known number of result sets is simpler, but if you write the code to handle a varying number of result sets, you do not need to make major modifications to your program if the stored procedure changes.

The basic steps for receiving result sets are as follows:

1. Declare a locator variable for each result set that is returned. If you do not know how many result sets are returned, declare enough result set locators for the maximum number of result sets that might be returned.
2. Call the stored procedure and check the SQL return code. If the SQLCODE from the CALL statement is +466, the stored procedure returned result sets.
3. Determine how many result sets the stored procedure is returning. If you already know how many result sets the stored procedure returns, you can skip this step.

Use the SQL statement **DESCRIBE PROCEDURE** to determine the number of result sets. The **DESCRIBE PROCEDURE** places information about the result sets in an SQLDA or SQL descriptor.

For an SQL descriptor, when the **DESCRIBE PROCEDURE** statement completes, the following values can be retrieved:

- DB2_RESULT_SETS_COUNT contains the number of result sets returned by the stored procedure.
- One descriptor area item is returned for each result set:
 - DB2_CURSOR_NAME

This item contains the name of the cursor that is used by the stored procedure to return the result set.

- **DB2_RESULT_SET_ROWS**

This item contains the estimated number of rows in the result set. A value of -1 indicates that no estimate of the number of rows in the result set is available.

- **DB2_RESULT_SET_LOCATOR**

This item contains the value of the result set locator that is associated with the result set.

For an SQLDA, make the SQLDA large enough to hold the maximum number of result sets that the stored procedure might return. When the **DESCRIBE PROCEDURE** statement completes, the fields in the SQLDA contain the following values:

- SQLDA contains the number of result sets returned by the stored procedure.
- Each SQLVAR entry gives information about a result set. In an SQLVAR entry, the following information is in effect:
 - The SQLNAME field contains the name of the cursor that is used by the stored procedure to return the result set.
 - The SQLIND field contains the estimated number of rows in the result set. A value of -1 indicates that no estimate of the number of rows in the result set is available.
 - The SQLDATA field contains the value of the result set locator, which is the address of the result set.

4. Link result set locators to result sets.

You can use the SQL statement **ASSOCIATE LOCATORS** to link result set locators to result sets. The **ASSOCIATE LOCATORS** statement assigns values to the result set locator variables. If you specify more locators than the number of result sets returned, the extra locators are ignored.

If you ran the **DESCRIBE PROCEDURE** statement previously, the result set locator values can be retrieved from the **DB2_RESULT_SET_LOCATOR** in the SQL descriptor or from the **SQLDATA** fields of the SQLDA. You can copy the values from these fields to the result set locator variables manually, or you can run the **ASSOCIATE LOCATORS** statement to do it for you.

The stored procedure name that you specify in an **ASSOCIATE LOCATORS** or **DESCRIBE PROCEDURE** statement must be a procedure name that was used in the **CALL** statement that returns the result sets.

5. Allocate cursors for fetching rows from the result sets.

Use the SQL statement **ALLOCATE CURSOR** to link each result set with a cursor. Run one **ALLOCATE CURSOR** statement for each result set. The cursor names can differ from the cursor names in the stored procedure.

6. Determine the contents of the result sets. If you already know the format of the result set, you can skip this step.

Use the SQL statement **DESCRIBE CURSOR** to determine the format of a result set and put this information in an SQL descriptor or an SQLDA. For each result set, you need an SQLDA large enough to hold descriptions of all columns in the result set.

You can use **DESCRIBE CURSOR** only for cursors for which you ran **ALLOCATE CURSOR** previously.

After you run **DESCRIBE CURSOR**, if the cursor for the result set is declared **WITH HOLD**, for an SQL descriptor **DB2_CURSOR_HOLD** can be checked. For an SQLDA, the high-order bit of the eighth byte of field **SQLDAID** in the SQLDA is set to 1.

7. Fetch rows from the result sets into host variables by using the cursors that you allocated with the **ALLOCATE CURSOR** statements. If you ran the **DESCRIBE CURSOR** statement, complete these steps before you fetch the rows:
 - a. Allocate storage for host variables and indicator variables. Use the contents of the SQL descriptor or SQLDA from the **DESCRIBE CURSOR** statement to determine how much storage you need for each host variable.
 - b. Put the address of the storage for each host variable in the appropriate SQLDATA field of the SQLDA.
 - c. Put the address of the storage for each indicator variable in the appropriate SQLIND field of the SQLDA.

Fetching rows from a result set is the same as fetching rows from a table.

8. Close the cursors.

Example 6-6 gives you an idea on how to implement this process in an RPG program:

Example 6-6 Result set support in an RPG program

```

D MYRS1                S                SQLTYPE(RESET_SET_LOCATOR)
D MYRS2                S                SQLTYPE(RESET_SET_LOCATOR)
...
C/EXEC SQL CALL P1(:parm1, :parm2, ...)
C/END-EXEC

...
C/EXEC SQL DESCRIBE PROCEDURE P1 USING DESCRIPTOR :MYRS2
C/END-EXEC

...
C/EXEC SQL ASSOCIATE LOCATORS (:MYRS1,:MYRS2) WITH PROCEDURE P1
C/END-EXEC
C/EXEC SQL ALLOCATE C1 CURSOR FOR RESULT SET :MYRS1
C/END-EXEC
C/EXEC SQL ALLOCATE C2 CURSOR FOR RESULT SET :MYRS2
C/END-EXEC

...
C/EXEC SQL ALLOCATE DESCRIPTOR 'SQLDES1'
C/END-EXEC
C/EXEC SQL DESCRIBE CURSOR C1 INTO SQL DESCRIPTOR 'SQLDES1'
C/END-EXEC

```

6.2.6 FIELDPROC support for encoding and encryption

You can now specify a **FIELDPROC** attribute for a column, designating an external program name as the field procedure exit routine for that column. It must be an ILE program that does not contain SQL. It cannot be a *SRVPGM, OPM *PGMs, or a Java object. Field procedures are assigned to a table by the **FIELDPROC** clause of the **CREATE TABLE** and **ALTER TABLE** statements. A field procedure is a user-written exit routine that transforms values in a single column.

This procedure allows for transparent encryption / decryption or encoding / decoding of data that is accessed through SQL or any other interface. It allows for transparent encryption or encoding of data that is accessed through SQL or natively.

When values in the column are changed, or new values are inserted, the field procedure is started for each value, and can transform that value (encode it) in any way. The encoded value is then stored. When values are retrieved from the column, the field procedure is started for each value, which is encoded, and must decode it back to the original value. Any indexes that are defined on a non-derived column that uses a field procedure are built with encoded values.

The transformation your field procedure performs on a value is called *field-encoding*. The same routine is used to undo the transformation when values are retrieved, which is called *field-decoding*. Values in columns with a field procedure are described to DB2 in two ways:

- ▶ The description of the column as defined in **CREATE TABLE** or **ALTER TABLE** appears in the catalog table QSYS2.SYSCOLUMNS. This description is the description of the field-decoded value, and is called the column description.
- ▶ The description of the encoded value, as it is stored in the database, appears in the catalog table QSYS2.SYSFIELDS. This description is the description of the field-encoded value, and is called the field description.

The field-decoding function must be the exact inverse of the field-encoding function. For example, if a routine encodes ALABAMA to 01, it must decode 01 to ALABAMA. A violation of this rule can lead to unpredictable results.

The field procedure is also started during the processing of the **CREATE TABLE** or **ALTER TABLE** statement. That operation is called a *field-definition*. When so started, the procedure provides DB2 with the column's field description. The field description defines the data characteristics of the encoded values. By contrast, the information that is supplied for the column in the **CREATE TABLE** or **ALTER TABLE** statement defines the data characteristics of the decoded values.

The data type of the encoded value can be any valid SQL data type except ROWID or DATALINK. Also, a field procedure cannot be associated with any column that has values generated by **IDENTITY** or **ROW CHANGE TIMESTAMP**.

If a DDS-created physical file is altered to add a field procedure, the encoded attribute data type cannot be a LOB type or DataLink. If an SQL table is altered to add a field procedure, the encoded attribute precision field must be 0 if the encoded attribute data type is any of the integer types.

A field procedure cannot be added to a column that has a default value of **CURRENT DATE**, **CURRENT TIME**, **CURRENT TIMESTAMP**, or **USER**. A column that is defined with a user-defined data type can have a field procedure if the source type of the user-defined data type is any of the allowed SQL data types. DB2 casts the value of the column to the source type before it passes it to the field procedure.

Masking support in FIELDPROCs

FIELDPROCs were originally designed to transparently encode or decode data. Several third-party products use the support in 7.1 to provide transparent column level encryption. For example, to allow a credit card number or social security number to be transparently encrypted on disk.

The FIELDPROC support is extended to allow masking to occur to that same column data (typically based on what user is accessing the data). For example, only users that need to see the actual credit card number see the value, while other users may see masked data. For example, XXXX XXXX XXXX 1234.

The new support is enabled by allowing the FIELDPROC program to detect masked data on an update or write operation and returning that indication to the database manager. The database manager then ignores the update of that specific column value on an update operation and replaces it with the default value on a write.

A new parameter is also passed to the FIELDPROC program. For field procedures that mask data, the parameter indicates whether the caller is a system function that requires that the data be decoded without masking. For example, in some cases, **RGZPFM** and **ALTER TABLE** might need to copy data. If the field procedure ignores this parameter and masks data when these operations are performed, the column data is lost. Hence, it is critical that a field procedure that masks data properly handles this parameter.

Parameter list for execution of field procedures

The field procedure parameter list communicates general information to a field procedure. It signals what operation is to be done and allows the field procedure to signal errors. DB2 provides storage for all parameters that are passed to the field procedure. Therefore, parameters are passed to the field procedure by address.

When you define and use the parameters in the field procedure, ensure that no more storage is referenced for a parameter than is defined for that parameter. The parameters are all stored in the same space and exceeding a parameter's storage space can overwrite another parameter's value. This action, in turn, can cause the field procedure to see invalid input data or cause the value returned to the database to be invalid. The following list details the parameters you can pass:

- ▶ A 2-byte integer that describes the function to be performed. This parameter is input only.
- ▶ A structure that defines the field procedure parameter value list (FPPVL).
- ▶ The decoded data attribute that is defined by the Column Value Descriptor (CVD). These attributes are the column attributes that were specified at **CREATE TABLE** or **ALTER TABLE** time. This parameter is input only.
- ▶ The decoded data.
 - The exact structure depends on the function code.
 - If the function code is 8, then the **NULL** value is used. This parameter is input only.
 - If the function code is 0, then the data to be encoded is used. This parameter is input only.
 - If the function code is 4, then the location to place the decoded data is used. This parameter is output only.
- ▶ The encoded data attribute that is defined by the Field Value Descriptor (FVD). This parameter is input only.
- ▶ The encoded data that is defined by the FVD. The exact structure is dependent on the function code. This parameter is input only.
- ▶ The **SQLSTATE** (character(5)). This parameter is input/output. This parameter is set by DB2 to 00000 before it calls the field procedure. It can be set by the field procedure. Although the **SQLSTATE** is not normally set by a field procedure, it can be used to signal an error to the database.
- ▶ The message text area (varchar(1000)). This parameter is input / output.

6.2.7 WebSphere MQ integration

IBM WebSphere MQ is a family of network communication software products that allow independent and potentially non-concurrent applications on a distributed system to communicate with each other.

The implementation described here provides a set of scalar functions and table functions to provide the integration with DB2.

Scalar functions

The **MQREAD** function returns a message in a **VARCHAR** variable from a specified WebSphere MQ location, which is specified by *receive-service*, using the policy that is defined in *service-policy*, starting at the beginning of the queue but without removing the message from the queue. If no messages are available to be returned, a null value is returned.

Example 6-7 reads the first message with a correlation ID that matches 1234 from the head of the queue that is specified by the MYSERVICE service using the MYPOLICY policy.

Example 6-7 MQREAD Scalar

```
SELECT MQREAD ('MYSERVICE','MYPOLICY','1234')  
FROM SYSIBM.SYSDUMMY1
```

The **MQREADCLOB** function returns a message in a **CLOB** variable from a specified WebSphere MQ location, which is specified by *receive-service*, using the policy that is defined in *service-policy*, starting at the beginning of the queue but without removing the message from the queue. If no messages are available to be returned, a null value is returned.

Example 6-8 reads the first message with a correlation ID that matches 1234 from the head of the queue that is specified by the MYSERVICE service using the MYPOLICY policy.

Example 6-8 MQREADCLOB Scalar

```
SELECT MQREADCLOB ('MYSERVICE','MYPOLICY','1234')  
FROM SYSIBM.SYSDUMMY1
```

The **MQRECEIVE** function returns a message in a **VARCHAR** variable from a specified WebSphere MQ location, which is specified by *receive-service*, using the policy that is defined in *service-policy*. This operation removes the message from the queue. If a correlation-id is specified, the first message with a matching correlation identifier is returned. If a correlation-id is not specified, the message at the beginning of queue is returned. If no messages are available to be returned, a null value is returned.

Example 6-9 receives the first message with a correlation-id that matches 1234 from the head of the queue that is specified by the MYSERVICE service using the MYPOLICY policy.

Example 6-9 MQRECEIVE Scalar

```
SELECT MQRECEIVE ('MYSERVICE','MYPOLICY','1234')  
FROM SYSIBM.SYSDUMMY1
```

The **MQRECEIVECLOB** function returns a message in a **CLOB** variable from a specified WebSphere MQ location, which is specified by receive-service, using the policy that is defined in service-policy. This operation removes the message from the queue. If a correlation-id is specified, the first message with a matching correlation identifier is returned. If a correlation-id is not specified, the message at the beginning of queue is returned. If no messages are available to be returned, a null value is returned.

Example 6-10 receives the first message with a correlation-id that matches 1234 from the head of the queue that is specified by the MYSERVICE service using the MYPOLICY policy.

Example 6-10 MQRECEIVECLOB Scalar

```
SELECT MQRECEIVECLOB ('MYSERVICE', 'MYPOLICY', '1234')
FROM SYSIBM.SYSDUMMY1
```

If for all of the previously mentioned scalars the receive-service is not specified or the null value is used, the DB2.DEFAULT.SERVICE is used.

The **MQSEND** function sends the data in a **VARCHAR** or **CLOB** variable msg-data to the WebSphere MQ location specified by send-service, using the policy that is defined in service-policy. An optional user-defined message correlation identifier can be specified by correlation-id. The return value is 1 if successful or 0 if not successful. If the send-service is not specified or the null value is used, the DB2.DEFAULT.SERVICE is used.

On all of these functions, you can specify a correlation-id (correl-id) expression. The value of the expression specifies the correlation identifier that is associated with this message. A correlation identifier is often specified in request-and-reply scenarios to associate requests with replies. The first message with a matching correlation identifier is returned.

Table functions

The **MQREADALL** function returns a table that contains the messages and message metadata in **VARCHAR** variables from the WebSphere MQ location that is specified by receive-service, using the policy that is defined in service-policy. This operation does not remove the messages from the queue. If num-rows is specified, a maximum of num-rows messages is returned. If num-rows is not specified, all available messages are returned.

Example 6-11 reads the head of the queue that is specified by the default service (DB2.DEFAULT.SERVICE) using the default policy (DB2.DEFAULT.POLICY). Only messages with a CORRELID of 1234 are returned. All columns are returned.

Example 6-11 MQREADALL table function

```
SELECT *
FROM TABLE (MQREADALL ()) AS T
WHERE T.CORRELID = '1234'
```

The **MQREADALLCLOB** function returns a table that contains the messages and message metadata in **CLOB** variables from the WebSphere MQ location that is specified by receive-service, using the policy that is defined in service-policy. This operation does not remove the messages from the queue. If num-rows is specified, a maximum of num-rows messages is returned. If num-rows is not specified, all available messages are returned.

Example 6-12 receives the first 10 messages from the head of the queue that is specified by the default service (DB2.DEFAULT.SERVICE), using the default policy (DB2.DEFAULT.POLICY). All columns are returned.

Example 6-12 MQREADALLCLOB table function

```
SELECT *  
FROM TABLE (MQREADALLCLOB (10)) AS T
```

The **MQRECEIVEALL** function returns a table that contains the messages and message metadata in **VARCHAR** variables from the WebSphere MQ location that is specified by receive-service, using the policy that is defined in service-policy. This operation removes the messages from the queue. If a correlation-id is specified, only those messages with a matching correlation identifier are returned. If a correlation-id is not specified, all available messages are returned. If num-rows is specified, a maximum of num-rows messages is returned. If num-rows is not specified, all available messages are returned.

Example 6-13 receives all the messages from the head of the queue that is specified by the service MYSERVICE, using the default policy (DB2.DEFAULT.POLICY). Only the MSG and CORRELID columns are returned.

Example 6-13 MQRECEIVEALL table function

```
SELECT T.MSG, T.CORRELID  
FROM TABLE (MQRECEIVEALL ('MYSERVICE')) AS T
```

The **MQRECEIVEALLCLOB** function returns a table that contains the messages and message metadata in **CLOB** variables from the WebSphere MQ location that is specified by receive-service, using the policy that is defined in service-policy. This operation removes the messages from the queue. If a correlation-id is specified, only those messages with a matching correlation identifier are returned. If correlation-id is not specified, all available messages are returned. If num-rows is specified, a maximum of num-rows messages is returned. If num-rows is not specified, all available messages are returned.

Example 6-14 receives all the messages from the queue that is specified by the default service (DB2.DEFAULT.SERVICE), using the default policy (DB2.DEFAULT.POLICY). The messages and all the metadata are returned as a table.

Example 6-14 MQRECEIVEALLCLOB table function

```
SELECT *  
FROM TABLE (MQRECEIVEALLCLOB ()) AS T
```

If for all of the previously mentioned table functions the receive-service is not specified or the null value is used, the DB2.DEFAULT.SERVICE is used.

DB2 WebSphere MQ tables

The DB2 WebSphere MQ tables contain service and policy definitions that are used by the DB2 WebSphere MQ functions. The DB2 WebSphere MQ tables are SYSIBM.MQSERVICE_TABLE and SYSIBM.MQPOLICY_TABLE. These tables are user-managed. The tables are initially created by DB2 and populated with one default service (DB2.DEFAULT.SERVICE) and one default policy (DB2.DEFAULT.POLICY). You can modify the attributes of the default service and policy by updating the rows in the tables. You can add more services and policies by inserting more rows in the tables.

DB2 WebSphere MQ CCSID conversion

When a message is sent, the message can be converted to the job CCSID by DB2. When a message is read or received, it can be converted to a specified CCSID by WebSphere MQ.

The msg-data parameter on the **MQSEND** function is in the job CCSID. If a string is passed for msg-data, it is converted to the job CCSID. For example, if a string is passed for msg-data that has a CCSID 1200, it is converted to the job CCSID before the message data is passed to WebSphere MQ. If the string is defined to be bit data or the CCSID of the string is the CCSID of the job, no conversion occurs.

WebSphere MQ does not perform CCSID conversions of the message data when **MQSEND** is ran. The message data that is passed from DB2 is sent unchanged along with a CCSID that informs the receiver of the message and how to interpret the message data. The CCSID that is sent depends on the value that is specified for the CODEDCHARSETID of the service that is used on the **MQSEND** function. The default for CODEDCHARSETID is -3, which indicates that the CCSID passed is the job default CCSID. If a value other than -3 is used for CODEDCHARSETID, the invoker must ensure that the message data passed to **MQSEND** is not converted to the job CCSID by DB2 and that the string is encoded in that specified CCSID.

When a message is read or received by a WebSphere MQ scalar or table function, the msg-data return parameter (and the MSG result column for the WebSphere MQ table functions) is also defined by in job default CCSID. DB2 does no conversions and relies on WebSphere MQ to perform any necessary conversions. Whether WebSphere MQ converts the message data can be controlled by setting the **RCV_CONVERT** value to *N* in the specified policy.

If the specified service has a value for CODEDCHARSETID of -3, DB2 instructs WebSphere MQ to convert any message that is read or received into the job CCSID. If a value other than -3 is used for CODEDCHARSETID, DB2 instructs WebSphere MQ to convert any message that is read or received into that CCSID. Specifying something other than -3 for CODEDCHARSETID in a service that is used to read or receive messages is not a preferred practice, because the msg-data return parameter and MSG result column are defined by DB2 in job default CCSID.

When you read or receive a message, truncation can occur. If the specified policy has a value for **RCV_ACCEPT_TRUNC_MSG** of *Y*, the message can be truncated without any warning. If the value for **RCV_ACCEPT_TRUNC_MSG** is *N* and the message is too long, the function ends with an error.

6.2.8 Miscellaneous

There are a number of functions that are aggregated under this heading. Most are aimed at upscaling or improving the ease of use for existing functions.

Partitioned table support

A partitioned table is a table where the data is contained in one or more local partitions (members). This release permits you to partition tables that use referential integrity or identity columns.

If you specify a referential constraint where the parent is a partitioned table, the unique index that is used for the unique index that enforces the parent unique constraint must be non-partitioned. Likewise, the identity column cannot be a partitioned key.

Partitioned tables with referential constraints or identity columns cannot be restored to a previous release.

Parameter markers

You can use this function to simplify the definition of variables in a program. Example 6-15 shows how you can write it.

Example 6-15 Parameter markers

```
SELECT stmt1 =  
    'SELECT * FROM t1  
      WHERE c1 = CAST(? AS DECFLOAT(34)) + CAST(? AS DECFLOAT(34));'  
PREPARE prestmt1 FROM STMT1;  
#Replace this with:  
SET STMT1 = 'SELECT * FROM T1 WHERE C1 > ? + ? '  
PREPARE PREPSTMT1 FROM STMT;
```

Expressions in a CALL statement

You can now call a procedure and pass as arguments an expression that does not include an aggregate function or column name. If extended indicator variables are enabled, the extended indicator variable values of **DEFAULT** and **UNASSIGNED** must not be used for that expression. In Example 6-16, **PARAMETER1** is folded and **PARAMETER2** is divided by 100.

Example 6-16 Expressions in a CALL statement

```
CALL PROC1 ( UPPER(PARAMETER1), PARAMETER2/100 )
```

Three-part names support

You can use three-part names to bypass the explicit **CONNECT** or **SET CONNECTION**. Statements that use three-part names and see distributed data result in IBM DRDA access to the remote relational database. When an application program uses three-part name aliases for remote objects and DRDA access, the application program must be bound at each location that is specified in the three-part names. Also, each alias must be defined at the local site. An alias at a remote site can see yet another server if a referenced alias eventually refers to a table or view.

All object references in a single SQL statement must be in a single relational database. When you create an alias for a table on a remote database, the alias name must be the same as the remote name, but can point to another alias on the remote database. See Example 6-17.

Example 6-17 Three-part alias

```
CREATE ALIAS shkspr.ph1 FOR wllm.shkspr.ph1  
SELECT * FROM shkspr.ph1
```

Concurrent access resolution

The concurrent access resolution option can be used to minimize transaction wait time. This option directs the database manager how to handle record lock conflicts under certain isolation levels.

The concurrent access resolution option can have one of the following values:

- Wait for outcome

This value is the default. This value directs the database manager to wait for the commit or rollback when it encounters locked data that is being updated or deleted. Locked rows that are being inserted are not skipped. This option does not apply for read-only queries that are running under **COMMIT(*NONE)** or **COMMIT(*CHG)**.

- Use currently committed

This value allows the database manager to use the currently committed version of the data for read-only queries when it encounters locked data being updated or deleted. Locked rows being inserted can be skipped. This option applies where possible when it is running under **COMMIT(*CS)** and is ignored otherwise. It is what is referred to as “Readers do not block writers and writers do not block readers.”

- Skip locked data

This value directs the database manager to skip rows in the case of record lock conflicts. This option applies only when the query is running under **COMMIT(*CS)** or **COMMIT(*ALL)**.

The concurrent access resolution values of **USE CURRENTLY COMMITTED** and **SKIP LOCKED DATA** can be used to improve concurrency by avoiding lock waits. However, care must be used when you use these options because they might affect application functionality.

You can specify the usage for concurrent access resolution in several ways:

- By using the concurrent-access-resolution clause at the statement level for a select-statement, **SELECT INTO**, searched **UPDATE**, or searched **DELETE**
- By using the **CONACC** keyword on the **CRTSQLxxx** or **RUNSQLSTM** commands
- With the **CONACC** value in the **SET OPTION** statement
- In the attribute-string of a **PREPARE** statement
- Using the **CREATE** or **ALTER** statement for a **FUNCTION**, **PROCEDURE**, or **TRIGGER**

If the concurrent access resolution option is not directly set by the application, it takes on the value of the **SQL_CONCURRENT_ACCESS_RESOLUTION** option in the **QAQQINI** query options file.

CREATE statement

Specifying the **CREATE OR REPLACE** statement makes it easier to create an object without having to drop it when it exists. This statement can be applied to the following objects:

- ALIAS
- FUNCTION
- PROCEDURE
- SEQUENCE
- TRIGGER
- VARIABLE
- VIEW

To replace an object, the user must have both ***OBJEXIST** rights to the object and ***EXECUTE** rights for the schema or library, and privileges to create the object. All existing privileges on the replaced object are preserved.

BIT scalar functions

The bitwise scalar functions **BITAND**, **BITANDNOT**, **BITOR**, **BITXOR**, and **BITNOT** operate on the “two’s complement” representation of the integer value of the input arguments and return the result as a corresponding base 10 integer value in a data type based on the data type of the input arguments. See Table 6-2.

Table 6-2 Bit scalar functions

Function	Description	A bit in the two's complement representation of the result is:
BITAND	Performs a bitwise AND operation.	1 only if the corresponding bits in both arguments are 1
BITANDNOT	Clears any bit in the first argument that is in the second argument.	Zero if the corresponding bit in the second argument is 1; otherwise, the result is copied from the corresponding bit in the first argument
BITOR	Performs a bitwise OR operation.	1 unless the corresponding bits in both arguments are zero
BITXOR	Performs a bitwise exclusive OR operation.	1 unless the corresponding bits in both arguments are the same
BITNOT	Performs a bitwise NOT operation.	Opposite of the corresponding bit in the argument

The arguments must be integer values that are represented by the data types **SMALLINT**, **INTEGER**, **BIGINT**, or **DECFLOAT**. Arguments of type **DECIMAL**, **REAL**, or **DOUBLE** are cast to **DECFLOAT**. The value is truncated to a whole number.

The bit manipulation functions can operate on up to 16 bits for **SMALLINT**, 32 bits for **INTEGER**, 64 bits for **BIGINT**, and 113 bits for **DECFLOAT**. The range of supported **DECFLOAT** values includes integers -2^{112} - 2^{112} - 1, and special values such as *NaN* (Not a Number) or *INFINITY* are not supported (SQLSTATE 42815). If the two arguments have different data types, the argument that is supporting fewer bits is cast to a value with the data type of the argument that is supporting more bits. This cast affects the bits that are set for negative values. For example, -1 as a **SMALLINT** value has 16 bits set to 1, which when cast to an **INTEGER** value has 32 bits set to 1.

The result of the functions with two arguments has the data type of the argument that is highest in the data type precedence list for promotion. If either argument is **DECFLOAT**, the data type of the result is **DECFLOAT(34)**. If either argument can be null, the result can be null. If either argument is null, the result is the null value.

The result of the **BITNOT** function has the same data type as the input argument, except that **DECIMAL**, **REAL**, **DOUBLE**, or **DECFLOAT(16)** returns **DECFLOAT(34)**. If the argument can be null, the result can be null. If the argument is null, the result is the null value.

Because of differences in internal representation between data types and on different hardware platforms, using functions (such as **HEX**) or host language constructs to view or compare internal representations of **BIT** function results and arguments is data type-dependent and not portable. The data type- and platform-independent way to view or compare **BIT** function results and arguments is to use the actual integer values.

Use the **BITXOR** function to toggle bits in a value. Use the **BITANDNOT** function to clear bits. **BITANDNOT(val, pattern)** operates more efficiently than **BITAND(val, BITNOT(pattern))**.

Example 6-18 is an example of the result of these operations.

Example 6-18 BIT scalar functions

```
# Return all items for which the third property bit is set.
SELECT ITEMID FROM ITEM
  WHERE BITAND(PROPERTIES, 4) = 4
# Return all items for which the fourth or the sixth property bit is set.
SELECT ITEMID FROM ITEM
  WHERE BITAND(PROPERTIES, 40) <> 0
# Clear the twelfth property of the item whose ID is 3412.
UPDATE ITEM
  SET PROPERTIES = BITANDNOT(PROPERTIES, 2048)
  WHERE ITEMID = 3412
# Set the fifth property of the item whose ID is 3412.
UPDATE ITEM
  SET PROPERTIES = BITOR(PROPERTIES, 16)
  WHERE ITEMID = 3412
# Toggle the eleventh property of the item whose ID is 3412.
UPDATE ITEM
  SET PROPERTIES = BITXOR(PROPERTIES, 1024)
  WHERE ITEMID = 3412
# Switch all the bits in a 16-bit value that has only the second bit on.
VALUES BITNOT(CAST(2 AS SMALLINT))
#returns -3 (with a data type of SMALLINT)
```

Encoded vector index

When you create an encoded vector index (EVI), you can now use an **INCLUDE** statement in the index option of the **CREATE ENCODED VECTOR INDEX** command, specifying an aggregate function to be included in the index. These aggregates make it possible for the index to be used directly to return aggregate results for a query. The aggregate function name must be one of the built-in functions **AVG**, **COUNT**, **COUNT_BIG**, **SUM**, **STDDEV**, **STDDEV_SAMP**, **VARIANCE**, or **VARIANCE_SAMP**, or a sourced function that is based on one of these built-in functions.

INCLUDE is only allowed for an encoded vector index.

This change has the potential of improving performance on queries that make this type of calculations. Example 6-19 shows the syntax for constructing a simple **INCLUDE** statement when you create such an index.

Example 6-19 Aggregate function support for EVI

```
CREATE ENCODED VECTOR INDEX GLDSTRN.RSNKRNZ_EVI1
  ON GLDSTRN.HMLT (JOB_TYPE, JOB_CATEGORY)
  INCLUDE (AVG(WORK_TIME))
```

Inlining of scalar functions

In cases of simple SQL scalar functions, instead of starting the function as part of a query, the expression in the **RETURN** statement of the function can be copied (inlined) into the query itself. Such a function is called an *inline function*. A function is an inline function if the following criteria are met:

- ▶ The SQL function is deterministic.
- ▶ The SQL-routine-body contains only a **RETURN** statement.
- ▶ The **RETURN** statement does not contain a scalar subselect or fullselect.

INSERT with remote SUBSELECT

INSERT with **SUBSELECT** is enhanced to allow the select to reference a single remote database that is different from the current server connection

An implicit remote connection is established and used by DB2 for i.

This enhancement is the second installment in extending DB2 for i on 7.1 to use implicit or explicit remote three-part names within SQL.


Example 6-20 declares the global temporary table from a remote subselect, which is followed by the insert. Figure 6-1 displays the output that is generated from the example in Example 6-20.

Example 6-20 Usage of INSERT with remote SUBSELECT

```
DECLARE GLOBAL TEMPORARY TABLE SESSION.MY_TEMP_TABLE (SERVER_NAME VARCHAR(40),  
DATA_VALUE CHAR(1)) WITH REPLACE
```

```
INSERT INTO SESSION.MY_TEMP_TABLE (SELECT CURRENT_SERVER CONCAT 'is the server  
name', IBMREQD FROM EUT72P1.SYSIBM.SYSDUMMY1)
```

```
SELECT * FROM SESSION.MY_TEMP_TABLE;
```



SERVER_NAME	DATA_VALUE
EUT72P1 is the Server Name	Y

Figure 6-1 Clearly showing the result was from a remote subselect

SQL procedure and function obfuscation

Obfuscation provides the capability of optionally obscuring proprietary SQL statements and logic within SQL procedures or functions.

ISVs can use this support to prevent their customers from seeing or changing SQL routines that are delivered as part of their solution.

Example 6-21 demonstrates how obfuscation is performed.

Example 6-21 Example of obfuscation

```
CALL SYSIBMADM.CREATE_WRAPPED('CREATE PROCEDURE UPDATE_STAFF (  
  IN P_EmpNo CHAR(6),  
  IN P_NewJob CHAR(5),  
                                IN P_Dept INTEGER)  
  LANGUAGE SQL  
  TR : BEGIN  
    UPDATE STAFF  
      SET JOB = P_NewJob  
      WHERE DEPT = P_Dept and ID = P_EmpNo; END TR ');  
  
SELECT ROUTINE_DEFINITION FROM QSYS2.SYSROUTINE WHERE  
  ROUTINE_NAME = 'UPDATE_STAFF';
```

The output in Figure 6-2 clearly shows the result of obfuscation.




Figure 6-2 Output of routine clearly showing obfuscation

System i Navigator can also be used to obfuscate SQL through the Generate SQL option by checking the **Obfuscate (for SQL function, and procedures)** check box, as shown in Figure 6-3.




Figure 6-3 Obfuscate (for SQL function and procedure objects) check box

Important: The obfuscate option is not available for triggers.

Figure 6-4 shows an example of obfuscated code after the option to obfuscate is selected.



```
-- Generate SQL
-- Version: V7R1M0 100423
-- Generated on: 10/03/12 08:42:11
-- Relational Database: ITSOP1
-- Standards Option: DB2 for i
CREATE SCHEMA GRIFFIC ;
-- SQL150C 10 CRTAUT for schema GRIFFIC ignored.
SET PATH "QSYS","QSYS2","SYSPROC","SYSIBMADM","HBEDOYA" ;

CREATE OR REPLACE PROCEDURE GRIFFIC.ADD_ARG (
    IN P_EMPNO CHAR(6) ,
    IN P_NEWJOB CHAR(6) ,
    IN P_DEPT INTEGER )
    WRAPPED QSQ07010
aaccxW8plW8FjG8FbG8pnG8pH69pj68Ffl_Ft58:x5qpdW8pdW8pd08Vhv1Kbaqebazeb8vaGBUrJYW5ArLnwBNtylcOIfK5
hHLlboHcEy3pw4mthmZ0E0yOXmnRGz5pTEqKuv9lKnWwIibfc1kqWJhIUR299I76t:X8YUh72kI8nTpzFYIW6::mTwDq
ZwHucxSginMeyT6NDHh4bx6by7oTnOJBQnJT0AUqmYFuxWeBdyMgD3zPAX_5B1rkBPjoOAaxi8ZuZWerrLJ648mFrJo
NKa9v5Jls9RROBYTV8O6f1Q1ue3vZackjgYOQvtK9E95wfM4dTIEF9fdOx53nSL4wNyA0bXxaa0y8drgGelNzaSrU9Zz0Du
TKmPUGCep4EcsCkKER9Vm:5cI52DI7BR29H3Z9Fbp5PI080ZNsUl2etz2w_a_n:Y_rspejTRGtjIWcJXI2BBX8fE0EGfSny
o6ZAcx7YUvRxEbTcCmRyewfX0jCgn3OwjtgUJABxN15QNFveScqfj4mGqeZ_8ia6A4eAeig5SwVwT02QZtCergPk82dZ
5TjeycemdDoscufTln11NZSjxV9JieDwMhmUM24pPvsE3lfeMXOGMzdPTmYhtWaa; GRANT ALTER , EXECUTE
ON SPECIFIC PROCEDURE GRIFFIC.ADD_ARG
TO GRIFFIC ;

SET PATH "QSYS","QSYS2","SYSPROC","SYSIBMADM","HBEDOYA" ;

CREATE OR REPLACE PROCEDURE GRIFFIC.ADD_CAT (
    IN P_EMPNO CHAR(6) ,
    IN P_NEWJOB CHAR(6) ,
    IN P_DEPT INTEGER )
    WRAPPED QSQ07010
aaccxW8plW8FjG8FbG8pnG8pH69pj68Ffl8VpZ8Vl1qpdW8pdW8pd08Vhv1KbaqebaGmYCYcbQ2gSeqaxdetsaHrVfdtoH
MC57MqaE9F5o7F8uuz5:Kk_5m2DRmVBRqhf4AnGreo1tted5honoWADBP89BPD85C6NdnO5JMNpXNyNuc270G_H:OT
LOV7qicwZq6FbkJOKdp3:NQly4r2Imf5k55xos9ivEsjTRgs46YQwDnk:UPolFsMlZZC3T8I_caZgn9plMJ8x1109U46xnUmF
VRn58wdAbaXvoQdJk3lkNpXzAG2XowNLFymsRmhmiFmk6OW:MHpolvJS_FuDS7ScOEKU8NfnFOMUA_4g9plzNJtOjGAt
pe3mMEo6YtMdZUaJMCsk2yEdkeOGtGzhCfnZ7Ke0EKM0JkLn5h8b5Y4:Vlc6psqrAIRqgB:SRrU:2FrJTuzYT6o_uJOykGIw
ypt3JX3e8y4s1Y7Tx3ZeXUHQPTDH9N_:qGc_Lkb_:DPygbwS8ebr53UtbpuUdfpzNDwyORwTpLwOvZlqJGbcJDbQ:YFr
TgykNAtdBHC16jHLzmVYyZ7lKwEr0oBd8LcmueqjSlacjExrcDOOmYlh5YE4r5YE4aaa; GRANT ALTER , EXECUTE
ON SPECIFIC PROCEDURE GRIFFIC.ADD_CAT
TO GRIFFIC ;

SET PATH "QSYS","QSYS2","SYSPROC","SYSIBMADM","HBEDOYA" ;

CREATE OR REPLACE PROCEDURE GRIFFIC.ADD_CUSTOMER (
    IN P_EMPNO CHAR(6) ,
    IN P_NEWJOB CHAR(6) ,
    IN P_DEPT INTEGER )
    WRAPPED QSQ07010
```

Figure 6-4 Example of obfuscated code

CREATE TABLE with remote SUBSELECT

One of the enhancements in the **CREATE TABLE** command in IBM i 7.1 is the capability of creating a table based on a remote table.

- ▶ **CREATE TABLE AS** and **DECLARE GLOBAL TEMPORARY TABLE** are enhanced to allow the select to reference a single remote database that is different from the current server connection.
- ▶ An implicit remote connection is established and used by DB2 for i.

- The remote query can reference a single remote homogeneous or heterogeneous table.


This enhancement is the third installment for extending DB2 for i on 7.1 to use implicit or explicit remote three-part names within SQL.

A table is created based on a remote table, as shown in Example 6-22. Running the example SQL produces the output that shows that a remote table was accessed, as shown in Figure 6-5.

Example 6-22 Create a table in the local database that references a remote database with the AS clause

```
CREATE TABLE DATALIB.MY_TEMP_TABLE AS (SELECT CURRENT_SERVER CONCAT ' is the
Server Name', IBMREQD
FROM X1423P2.SYSIBM.SYSDUMMY1) WITH DATA

SELECT * FROM DATALIB.MY_TEMP_TABLE
```



SERVER_NAME	DATA_VALUE
X1423P2 is the Server Name	Y

Figure 6-5 Output from the SQL showing that the remote table was accessed

Qualified name option added to generate SQL

Today, generated SQL includes schema/library qualification of objects.

System i Navigator and the **QSQGNDDL()** are enhanced to include the qualified name option, making it easier to redeploy generated SQL. Any three-part names (object and column) are left unchanged. Any schema qualification within the object that does not match the database object library name are left unchanged.

The qualified name option specifies whether qualified or unqualified names should be generated for the specified database object. The valid values are:

- ‘0’ Qualified object names should be generated. Unqualified names within the body of SQL routines remain unqualified (by default).
- ‘1’ Unqualified object names should be generated when a library is found that matches the database object library name. Any SQL object or column reference that is RDB qualified is generated in its fully qualified form. For example, rdb-name.schema-name.table-name and rdb-name.schema-name.table-name.column-name references retain their full qualification.

This option also appears on the Generate SQL dialog within System i Navigator, as shown in Figure 6-6. The default behavior is to continue to generate SQL with schema qualification.




Figure 6-6 Schema qualified names for objects check box

6.2.9 OVRDBF SEQONLY(YES, buffer length)

OVDBF adds support to allow the user to specify the buffer length rather than the number of records for **OVDBF SEQONLY(*YES N)**. N can be:

- ▶ *BUF32KB
- ▶ *BUF64KB
- ▶ *BUF128KB
- ▶ *BUF256KB

This setting means that the number of records are the number of records that fit into a 32 KB, 64 KB, 128 KB, or 256 KB buffer.

6.3 Performance and query optimization

In the IBM i 7.1 release of DB2 for i, a considerable effort was undertaken to enhance the runtime performance of the database, either by extending existing functions or by introducing new mechanisms.

Runtime performance is affected by many issues, such as the database design (the entity-relationship model, which is a conceptual schema or semantic data model of a relational database), the redundancy between functional environments in composite application environment, the level of normalization, and the size and volumes processed. All of these items influence the run time, throughput, or response time, which is supported by the IT components and is defined by the needs of the business. Performance optimization for database access must address all the components that are used in obtained acceptable and sustainable results, covering the functional aspects and the technical components that support them.

This section describes the query optimization method. It describes what is behind the changes that are implemented in the database management components to relieve the burden that is associated with the tools and processes a database administrator uses or follows to realize the non-functional requirements about performance and scalability, including:

- ▶ Global Statistics Cache (GSC)
- ▶ Adaptive Query Processing
- ▶ Sparse indexes
- ▶ Encoded vector index-only access, symbol table scan, symbol table probe, and **INCLUDE** aggregates
- ▶ Keeping tables or indexes in memory

6.3.1 Methods and tools for performance optimization

Typically, the autonomous functions in IBM i, and the new ones in IBM i 7.1, all strive to obtain the best possible performance and throughput. You can, however, tweak settings to pre-emptively enhance the tooling of IBM i.

In today's business world, the dynamics of a business environment demand quick adaptation to changes. You might face issues by using a too generic approach in using these facilities. Consider that you made the architectural decision for a new application to use a stateless runtime environment and that your detailed component model has the infrastructure for it. If the business processes it supports are changing and require a more stateful design, you might face an issue if you want to preserve information to track the statefulness in your transactions. Then, the database where you store information about these transactions might quickly become the heaviest consumer of I/O operations. If your infrastructure model did not consider this factor, you have a serious issue. Having high volumes with a low latency is good. However, this situation must be balanced against the effort it takes to make it sustainable and manageable throughout all of the IT components you need to support the business.

When you define components for a database support, develop a methodology and use preferred practices to obtain the best results. Any methodology must be consistent, acceptable, measurable, and sustainable. You want to stay away from *ad hoc* measures or simple bypasses.

IBM i provides statistics about I/O operations, provided by the database management function. These statistics show accumulated values, from which you can derive averages, on the I/O operations on tables and indexes. These statistics do not take into account the variability and the dynamic nature of the business functions these objects support. So if you want to use these statistics to define those objects to be placed either in memory or on faster disks, you must consider a larger scope.

For example: Since the introduction of solid-state drives (SSD), which have a low latency, the IBM i storage manager has awareness about this technology and uses it as appropriate. Since release 6.1, you can specify the media preference on the **CREATE TABLE/INDEX** and **ALTER TABLE/INDEX** commands along with the **DECLARE GLOBAL TEMPORARY TABLE** (see 6.3.9, “SQE optimization for indexes on SSD” on page 189). The SYSTABLESTAT and SYSINDEXSTAT catalog tables provide more I/O statistics (SEQUENTIAL_READS and RANDOM_READS) in release 7.1 on these objects. These statistics, generated by the database manager, indicate only possible candidates to be housed on SSD hardware. Further investigation of the run time and the contribution to the performance and capacity or the infrastructure reveals whether they are eligible for those settings.

For more information about SSDs, see Chapter 9, “Storage and solid-state drives” on page 367.

Finally, and as a last resort, there is now a stored procedure available that you can use to cancel long running SQL jobs using the **QSYS2.CANCEL_SQL** procedure.

6.3.2 Query optimization

Whenever a query is submitted, the database engine creates an artifact that allows the query to trigger a set of events and processes that allows it to run the request with the lowest cost. In this context, cost is expressed as the shortest time possible to run the query. This cost calculation is done on a number of both fixed and variable elements. The fixed cost elements are attributes, such as both the hardware components (processor, memory, and disks) and in the instruments or methods that can be used to handle rows and columns in a (set of) database files. These methods are known as using indexes (binary radix index or encoded vector index), index or table scan, hashing, sorting, and so on. The variable elements are typically the volume of data (that is, the number or rows) to be handled and the join functions that are required by the query. Based on these methods, the database query engine builds an access plan that targets reduction of cost.

Even with all the technologies that are used, the access plans might still yield an incorrect (that is, not obeying the rule of capping the cost) result. This situation can, for example, be the result of not having an index to navigate correctly through the data. For that reason, IBM i supports the technology to create temporary indexes autonomically until the system undergoes an IPL. This index can be used by any query that might benefit from its existence. These autonomous indexes can be viewed and carry information that a database administrator can use to decide whether to make it a permanent object by using the definition of the temporary index.

Other elements that can contribute to incorrect access plans are as follows:

- ▶ Inclusion of complex or derivated predicates, which are hard to predict without running the query about the existence of stale statistics on busy systems
- ▶ Hidden correlations in the data, often because of a poor design, data skew, and data volatility
- ▶ Changes in the business or infrastructure environment

In the last case, this situation is more likely to happen with variations in both memory and processor allocations on partitioned systems, which are reconfigured using dynamic partitioning. It can also be caused when the data is changed frequently in bulk.

If you want to read more about the database query engine, see *Preparing for and Tuning the SQL Query Engine on DB2 for i5/OS*, SG24-6598.

6.3.3 Global Statistics Cache

There are several process models to reduce the impact of managing the dynamics of a database structure and its content. Moreover, this database is often deployed on a system that is subject to many changes. These tasks can be a wide array of non-automated interventions, including the setup of a validation process of access plans, manually tuning the query, or having the access plans invalidated and re-created. It can also include a reset of the statistics information or an extensive review of the query functions to achieve a higher degree of expected consumability by the system. These actions are typically post-mortem actions and are labor-intensive.

To reduce this labor-intensive work, the DB2 Statistics Manager was revised. By default, it now collects data about observed statistics in the database and from partially or fully completed queries. This data is stored in the Global Statistics Cache (GSC), which is a system-wide repository, containing those complex statistics. The adaptive query processing (AQP) (see 6.3.4, “Adaptive query processing” on page 185) inspects the results of queries and compares the estimated row counts with the actual row counts. All of the queries that are processed by the SQL Query Engine (SQE) use this information to increase overall efficiency. One of the typical actions the SQE can take is to use the live statistics in the GSC, compare the estimated row count with the actual row count, and reoptimize and restart the query using the new query plan. Furthermore, if another query asks for the same or a similar row count, the Storage Manager (SM) can return the stored actual row count from the GSC. This action allows generating faster query plans by the query optimizer.

Typically, observed statistics are for complex predicates, such as a join. A simple example is a query that joins three files, A, B, and C. There is a discrepancy between the estimate and actual row count of the join of A and B. The SM stores an observed statistic into the GSC. Later, if a join query of A, B, and Z is submitted, SM recalls the observed statistic of the A and B join. The SM considers that observed statistic in its estimate of the A, B, and Z join.

The GSC is an internal DB2 object, and the contents of it are not directly observable. You can harvest the I/O statistics in the database catalog tables SYSTABLESTAT and SYSINDEXSTAT or by looking at the I/O statistics using the Display File Description (**DSPFD**) command. This command provides only a limited number of I/O operations. Both counters (catalog tables and the object description) are reset at IPL time.

The statistics collection is defined by the system value Data Base file statistics collection (**QDBFSTCCOL**). The SM jobs that update the statistics carry the same name.

6.3.4 Adaptive query processing

The SQE uses statistics to build the mechanism to perform an SQL statement. These statistics come from two sources:

- Information that is contained in the indexes on the tables that are used in the statement
- Information that is contained in the statistics tables (the GSC)

When the query compiler optimizes the query plans, its decisions are heavily influenced by statistical information about the size of the database tables, indexes, and statistical views. The optimizer also uses information about the distribution of data in specific columns of tables, indexes, and statistical views if these columns are used to select rows or join tables. The optimizer uses this information to estimate the costs of alternative access plans for each query.

In IBM i 7.1, the SQE query engine uses a technique called adaptive query processing (AQP). AQP analyzes actual query runtime statistics and uses that information to correct previous estimates. These updated estimates can provide better information for subsequent optimizations. It also focuses on optimizing join statements to improve the join orders and minimizing the creation of large joins for sparsely populated join results. This inspection is done during the run of a query request and observes its progress. The AQP handler wakes up after a query runs for at least 2 seconds without returning any rows. Its mission is to analyze the actual statistics from the partial query run, diagnose, and possibly recover from join order problems. These join order problems are because of inaccurate statistical estimates. This process is referred to as the *AQP Handler*.

After a query completes, another task, the AQP Request Support, starts and runs in a system task so that it does not affect the performance of user applications. Estimated record counts are compared to the actual values. If significant discrepancies are noted, the AQP Request Support stores the observed statistic in the GSC. The AQP Request Support might also make specific recommendations for improving the query plan the next time the query runs.

Both tasks collect enough information to reoptimize the query using partially observed statistics or specific join order recommendations or both. If this optimization results in a new plan, the old plan is terminated and the query is restarted with the new plan, provided the query has not returned any results. The restart can be done for long running queries during the run time itself.

AQP looks for an unexpected *starvation join* condition when it analyzes join performance. Starvation join is a condition where a table late in the join order eliminates many records from the result set. In general, the query can perform better if the table that eliminates the large number of rows is first in the join order. When AQP identifies a table that causes an unexpected starvation join condition, the table is noted as the *forced primary table*. The forced primary table is saved for a subsequent optimization of the query. That optimization with the forced primary recommendation can be used in two ways:

- The forced primary table is placed first in the join order, overriding the join order that is implied by the statistical estimates. The rest of the join order is defined using existing techniques.
- The forced primary table can be used for LPG preselection against a large fact table in the join.

The database monitor has a new set of records to identify the action that is undertaken with by the AQP.

Figure 6-7 provides a sample of how a join can be optimized. The estimated return of rows on table C proved to be much smaller during the execution of the query, forcing the SQE to recalculate the number of rows that are returned and dramatically reduced the size of the result set.




Figure 6-7 AQP Optimization on join

6.3.5 Sparse indexes

Starting from the IBM i 6.1 release, you can create a sparse index using a **WHERE** condition. In IBM i 7.1, the query optimizer inspects those indexes and uses them where appropriate.

The reason for creating a sparse index is to provide performance enhancements for your queries. The performance enhancement is done by precomputing and storing results of the **WHERE** selection in the sparse index. The database engine can use these results instead of recomputing them for a user-specified query. The query optimizer looks for any applicable sparse index and can choose to implement the query using a sparse index. The decision is based on whether using a sparse index is a faster implementation choice.

To use a sparse index, the **WHERE** selection in the query must be a subset of the **WHERE** selection in the sparse index; the set of records in the sparse index must contain all the records to be selected by the query. It might contain extra records, but it must contain all the records to be selected by the query. This comparison of the **WHERE** selection is performed by the query optimizer during optimization. It is like the comparison that is performed for Materialized Query Tables (MQT).

Besides the comparison of the **WHERE** selection, the optimization of a sparse index is identical to the optimization that is performed for any Binary Radix index.

In Example 6-23, we create a sparse index over a table in which events are stored. These events can be of four types:

- ▶ On-stage shows (type OSS)
- ▶ Movies (type MOV)
- ▶ Broadcasts (BRO)
- ▶ Forums (FOR)

Example 6-23 Sparse indexes

```
CREATE INDEX EVENTS/OSS_MOV_BRO on EVENTS/OSS_MOV_BRO_FOR (EVTYPE)
  WHERE EVTYPE in ('OSS', 'MOV', BRO');
CREATE INDEX EVENTS/OSS_MOV_BRO_FOR on EVENTS/OSS_MOV_BRO_FOR (EVTYPE)
  WHERE EVTYPE in ('OSS', 'MOV', BRO', 'FOR');
```

In the first index, we select type OSS, MOV, and BRO, and in the second index, all of the types. In the first index, the query selection is a subset of the sparse index selection and an index scan over the sparse index is used. The remaining query selection (EVTYPE=FOR) is run following the index scan. For the second index, the query selection is not a subset of the sparse index selection and the sparse index cannot be used.

6.3.6 Encoded vector index

Section “Encoded vector index” on page 176 described the enhancements for encoded vector indexes (EVIs). The EVI can be used for more than generating a bitmap or row number list to provide an asynchronous I/O map to the wanted table rows. The EVI can also be used by two index-only access methods that can be applied specifically to the symbol table itself. These two index-only access methods are the EVI symbol table scan and the EVI symbol table probe.

These two methods can be used with **GROUP BY** or **DISTINCT** queries that can be satisfied by the symbol table. This symbol table-only access can be further employed in aggregate queries by adding **INCLUDE** values to the encoded vector index.

EVI symbol table scan

An encoded vector index symbol table scan operation is used to retrieve the entries from the symbol table portion of the index. All entries (symbols) in the symbol table are sequentially scanned if a scan is chosen. The symbol table can be used by the optimizer to satisfy **GROUP BY** or **DISTINCT** portions of a query request.

Selection is applied to every entry in the symbol table. The selection must be applied to the symbol table keys unless the EVI was created as a sparse index with a **WHERE** clause. In that case, a portion of the selection is applied as the symbol table is built and maintained. The query request must include matching predicates to use the sparse EVI.

All entries are retrieved directly from the symbol table portion of the index without any access to the vector portion of the index. There is also no access to the records in the associated table over which the EVI is built.

The advantages of this setup are obvious:

- ▶ Pre-summarized results are readily available.
- ▶ There is a need to process only the unique values in the symbol table, thus avoiding processing table records.
- ▶ It extracts all the data from the index unique key values or **INCLUDE** values, thus eliminating the need for a Table Probe or vector scan.
- ▶ With **INCLUDE** providing ready-made numeric aggregates, it eliminates the need to access corresponding table rows to perform the aggregation.

However, for grouping queries where the resulting number of groups is relatively small compared to the number of records in the underlying table, the performance improvement is low. Even more, it can perform poorly when there are many groups that are involved, making the symbol table large. You are likely to experience poor performance if a large portion of the symbol table is put into the overflow area. Alternatively, you experience a significant performance improvement for grouping queries when the aggregate is specified as an **INCLUDE** value of the symbol table.

INCLUDE aggregates

To enhance the ability of the EVI symbol table to provide aggregate answers, the symbol table can be created to contain more **INCLUDE** values. These results are ready-made numeric aggregate results, such as **SUM**, **COUNT**, **AVG**, or **VARIANCE** values that are requested over non-key data. These aggregates are specified using the **INCLUDE** keyword on the **CREATE ENCODED VECTOR INDEX** request.

These included aggregates are maintained in real time as rows are inserted, updated, or deleted from the corresponding table. The symbol table maintains these additional aggregate values in addition to the EVI keys for each symbol table entry. Because these results are numeric results and finite in size, the symbol table is still a desirable compact size.

The included aggregates are over non-key columns in the table where the grouping is over the corresponding EVI symbol table defined keys. The aggregate can be over a single column or a derivation.

Encoded vector index symbol table probe

The encoded vector index symbol table probe operation is used to retrieve entries from the symbol table portion of the index, which avoids scanning the entire symbol table. The symbol table can be used by the optimizer to satisfy **GROUP BY** or **DISTINCT** portions of a query request.

The optimizer attempts to match the columns that are used for the selection against the leading keys of the EVI index. It then rewrites the selection into a series of ranges that can be used to probe directly into the symbol table. Only those symbol table pages from the series of ranges are paged into main memory. The resulting symbol table entries that are generated by the probe operation can then be further processed by any remaining selection against EVI keys. This strategy provides for quick access to only the entries of the symbol table that satisfy the selection.

Similar to an encoded vector symbol table scan, a symbol table probe can return ready-made aggregate results if **INCLUDE** is specified when the EVI is created. All entries are retrieved directly from the symbol table portion of the index without any access to the vector portion of the index. In addition, it is unnecessary to access the records in the associated table over which the EVI is built.

6.3.7 Preserving EVI indexes on ALTER enhancement

Before this enhancement, an **ALTER TABLE** or fast delete under commitment control would require any encoded vector indexes on the table being altered to be rebuilt.

This enhancement allows encoded vector indexes on the table being altered to be preserved if the data type or other attribute of a key column of the index is not changed by the **ALTER**.

6.3.8 Keeping tables or indexes in memory

The **KEEPINMEM** parameter specifies whether the data or an access for a file member is brought into a main storage pool by the SQL Query Engine (SQE) when the data is used in the query to improve the performance. When you specify ***YES** for this parameter, the Query Options File (QAQQINI) parameter **MEMORY_POOL_PREFERENCE** (see 6.3.15, “QAQQINI properties” on page 192) specifies the preferred main storage pool to be used.

This function applies only during the run time of a query, and might therefore be substituted for the Set Object Access (**SETOBJACC**) command that puts the table or index in memory in a static function. After the query completes, the memory might be freed again, contrary to the effects of the Set Object Access (**SETOBJACC**), where you must clear it using the ***PURGE** option on the Storage Pool (**POOL**) parameter of the command.

Similarly, the DB2 database manager reduces the storage that is occupied by a table that does not contain any data, thus reducing the storage space that is needed for unused objects. This situation is also referred to as *deflated table support*.

6.3.9 SQE optimization for indexes on SSD

The query optimizer now recognizes that indexes may potentially be on SSDs and prioritizes usage of those indexes higher than indexes on spinning disk when you order the indexes during optimization.

Indexes must have the SSD attribute specified through the **UNIT(*SSD)** parameter on the Create Logical File (**CRTL**) or Change Logical File (**CHGLF**) CL commands or by using the **UNIT SSD** clause on the **SQL CREATE INDEX** statement. For more information about this topic, see 6.4.10, “CHGPFM and CHGLFM UNIT support” on page 211.

6.3.10 SQE support of simple logical files

SQE supports simple logical files in IBM i 7.1. SQE support of simple logical files has the following restrictions:

- ▶ No SQE support of OmniFind using logical files.
- ▶ No SQE support of multi-data space logical files.
- ▶ No SQE support of logical files over a partition table.
- ▶ SQE supports only read-only queries. There is no SQE support of insert, update, or delete using logical files.

The QAQQINI file option 'IGNORE_DERIVED_INDEX' continues to be supported. If **IGNORE_DERIVED_INDEX(*NO)** is specified, and a select / omit logical file exists based on file of the simple logical file, then SQE does not process the query of the simple logical file.

6.3.11 QSYS2.INDEX_ADVICE procedure

This procedure is useful to anyone who wants to analyze index advice from different machines or from different points in time.

The DB2 for i index advice condenser is externalized through the QSYS2.CONDENSEDINDEXADVICE view. The view and underlying user-defined table function are hardwired to use the raw index advice that is stored within the QSYS2/SYSIXADV file. Some users must use the index advice condenser against a file that was saved and restored from a different machine.

A new database supplied procedure (**QSYS2.INDEX_ADVICE**) was added. The procedure establishes the QTEMP/CONDENSEDINDEXADVICE view over a user supplied library and file name. Once established, the user can query QTEMP/CONDENSEDINDEXADVICE to condense the index advice against the target index advice file.

The **QSYS2.INDEX_ADVICE** procedure also has options to return the index advice as a result set, either in raw advice format or in condensed format. When the job ends or disconnects, the objects in QTEMP are automatically removed. The **QSYS2.INDEX_ADVICE** procedure also has options to return the index advice as a result set, either in raw advice format or in condensed format.

Example 6-24 illustrates the usage of the INDEX_ADVICE procedure.

Example 6-24 Usage of QSYS2.INDEX_ADVICE procedure

#Procedure definition:

```
create procedure QSYS2.INDEX_ADVICE(
in advice_library_name char(10),
in advice_file_name char(10),
in advice_option integer)
#Advice_option values:
#if advice_option=0 then setup for targeted condensed index advice, do not return
a result set
#if advice_option=1 return condensed index advice as a result set
#if advice_option=2 return raw index advice as a result set
#Example usage:
call qsys2.index_advice('ADVICELIB', 'SYSIXADV', 0);

-- Count the rows of raw advice
select count(*) from QTEMP.SYSIXADV where table_schema = 'PRODLIB' ;

-- Count the rows of condensed advice
select count(*) from QTEMP.CONDENSEDINDEXADVICE where table_schema = 'PRODLIB';

-- Review an overview of the most frequently advised, using condensed advice
select table_name, times_advised, key_columns_advised from
QTEMP.CONDENSEDINDEXADVICE where table_schema = 'PRODLIB' order by times_advised
desc;
```

When the procedure is called with `advise_option=0`, the index advice level of the target file is determined. If the advice file originated from an IBM i 5.4 or 6.1 machine, the file is altered to match the 7.1 advice format. This alteration is a one time conversion of the advice file. Once established, the user can query `QTEMP.CONDENSEDINDEXADVICE` to condense the index advice against the target index advice file.

6.3.12 SKIP LOCKED DATA and NC or UR

You can use the **SKIP LOCKED DATA** clause to fetch rows from a table or view without waiting for row locks. When the option is specified, any row that is already locked by another job is skipped. This behavior is typically wanted for tables or views that are used as a queue. **SKIP LOCKED DATA** can be used only when isolation level NC, UR, CS, or RS is in effect. The **SKIP LOCKED DATA** clause is ignored when used when isolation level RR is in effect.

Before this enhancement, **SKIP LOCKED DATA** was allowed only when the isolation level was CS or RS.

For more information about this topic, go to:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/db2/rbafz/skiplocked>

6.3.13 SQL routine performance integer arithmetic (requires re-create)

This procedure improves the performance of the generated code within **LANGUAGE SQL** routines. When you add and subtract a positive integer from an SQL smallint, integer, and bigint variable in a **LANGUAGE SQL** procedure, function, or trigger, the assignment is accomplished directly within the generated ILE C code. This improvement is not observed when you build SQL routines for previous releases (for example, when you use **SET OPTION TGTRLS=V6R1M0** or **V5R4M0**).

To achieve the improved code generation, SQL procedures, functions, and trigger routines must be re-created after you upgrade the operating system to IBM i 7.1.

This improvement applies to the following usage of the SQL **SET** statement:

`SET v1 = v1 + <in lit>` where v1 is a smallint, in, and bigint

`SET v1 = v1 - <in lit>` where v1 is a smallint, in, and bigint

The following statements will generate inline ILE C code:

`SET v1 = v1 + <integer literal>`

`SET v1 = v1 + <bigint literal>`

`SET v1 = v1 + <negative integer literal>`

`SET v1 = <any literal> + v1`

`SET v1 = <any literal> +/- <any literal>`

These statements do not generate inline ILE C code:

`SET v1 = v2 + <integer literal>`

`SET v2 = v1 + <bigint literal>`

6.3.14 Automatic cancellation of QSQRVR jobs when an application ends

When an application using SQL Server Mode is terminated and the SQL Server Mode connections are not ended, the termination of the application job may completely end while the QSQRVR job remains active performing a long running system operation.

Job termination is improved to signal an SQL Cancel request to any QSQRVR jobs that being used by the application. The cancellation interrupts some long running operations, allowing the QSQRVR job to observe that the application is ending.

SQL Server Mode users must apply only the PTF to receive the improved cancel handling support.

6.3.15 QAQQINI properties

Table 6-3 lists the new parameters and their values that can be put in the QAQQINI file, which is used to define settings for processes that perform queries. These changes are in line with the new features in the DB2 for i.

Table 6-3 QAQQINI file

Parameter	Description
ALLOW_ADAPTIVE_QUERY_PROCESSING (For more information, see 6.3.4, "Adaptive query processing" on page 185.)	Specifies whether AQP processing is done for a query.
ALLOW_ARRAY_VALUE_CHANGES	Specifies whether changes to the values of array elements are visible to the query when the query is running.
DETERMINISTIC_UDF_SCOPE	Specifies the scope or lifetime of the deterministic setting for User Defined Functions (UDFs) and User Defined Table Functions (UDTFs).
FIELDPROC_ENCODED_COMPARISON (For more information, see 6.2.6, "FIELDPROC support for encoding and encryption" on page 166.)	Specifies the amount of optimization that the optimizer might use when queried columns have attached field procedures.
MEMORY_POOL_PREFERENCE	Specifies the preferred memory pool that database operations use. This option does not ensure usage of the specified pool, but directs database to perform its paging into this pool when supported by the database operation.
PSEUDO_OPEN_CHECK_HOST_VARS	This parameter can be used to allow SQE to check the selectivity of the host variable values at pseudo-open time. If the new set of host variable values requires a different plan to perform well, SQE reoptimizes the query. The possible values are: ▶ *DEFAULT: The default value is *NO. ▶ *NO: Do not check host variable selectivity at pseudo-open time. This behavior is compatible with the previous behavior. ▶ *OPTIMIZE: The optimizer determines when host variable selectivity should be checked. In general, the SQE engine monitors the query. If, after a certain number runs, the engine determines that there is no advantage to checking host variable values, that is, the selectivity is not changing enough or selectivity changes result in the same plan, the optimizer will stop checking for host variable selectivity changes at pseudo-open time. Full opens do the normal plan validation. ▶ *YES: Always check host variable selectivity at pseudo-open time. If the REOPTIMIZE_ACCESS_PLAN QAQQINI option is set to *ONLY_REQUIRED, the PSEUDO_OPEN_CHECK_HOST_VARS option has no effect.
SQL_CONCURRENT_ACCESS_RESOLUTION (For more information, see "Concurrent access resolution" on page 173.)	Specifies the concurrent access resolution to use for an SQL query.

Parameter	Description
SQL_XML_DATA_CCSID (For more information, see “XML data type” on page 155.)	Specifies the CCSID to be used for XML columns, host variables, parameter markers, and expressions, if not explicitly specified.
TEXT_SEARCH_DEFAULT_TIMEZONE	Specifies the time zone to apply to any date or dateTime value that is specified in an XML text search using the CONTAINS or SCORE function. The time zone is the offset from Coordinated Universal Time (Greenwich Mean Time). It is only applicable when a specific time zone is not given for the value.
SQL_GVAR_BUILD_RULE	Influences whether global variables must exist when you build SQL procedures, functions, triggers, or run SQL precompiles. For more information, see 6.4.29, “New QAQQINI option: SQL_GVAR_BUILD_RULE” on page 227.

6.3.16 ALTER TABLE performance

ALTER TABLE can be a long running operation. The general performance of **ALTER TABLE** was improved (although it can still be long running) by reducing the path length of the operation and by reducing lock contention. Lock contention is reduced when multiple tables are referenced by one or more views and the related tables are altered or dropped concurrently in different jobs.

6.3.17 Avoiding short name collisions for CREATE PROCEDURE, FUNCTION, or TRIGGER

When the SQL routines **CREATE PROCEDURE (SQL)**, **CREATE FUNCTION (SQL)**, and **CREATE TRIGGER** are created using a long name, the database generates the system name of the routine. For long names, the first five characters of the long name are combined with '00001'. If an object with that system name exists, the second half of the name is incremented by one and the create is tried again.

If you have many SQL routines whose names begin with common first five characters, the creation of the routines is slowed down by name conflicts and rebuild attempts that determine whether a system name has not been used.

The QGENOBJNAM data area can be used to control the system name that is generated by DB2 for i for SQL routines. Through use of the data area, the performance of the SQL routine creation can be greatly improved.

To be effective, the data area must be created as CHAR(10) and must be within a library that is in the library list.

The user that creates the routine must have ***USE** authority to the data area.

When the **PROGRAM NAME** clause is used on **CREATE TRIGGER** to specify the system name of the program, the data area has no effect on the operation.

In Example 6-25, MNAME123 is always used for the system name of the trigger program.

Example 6-25 Using the system name of the program in CREATE TRIGGER command

```
create trigger newlib/longname_trig123 after insert on newlib/longname_table123
program name mname123 begin end
```

There are two ways to use the QGENOBJNAM data area:

1. Use question marks and a starting value ('?????xxxx'), where x is a number digit, and the generated name begins with xxxxx instead of 00001. For example, if the value of the data area was '?????50000' and a procedure named ProductionProcedure1 was being created, the first generated system name would be PRODU50000.
2. Use '*GEN00' for the data area value to direct the database to use the first three characters of the long name, and the last four digits from the job number and '000'. For example, if the value of the data area was '*GEN00', the job number was 098435, and a procedure named ProductionProcedure1 was being created, the first generated system name would be PRO8435000. *GEN00 can be used to improve SQL routine creation throughput by spreading the creates across multiple jobs.

You can automatically assign trigger system programs according to the value of QGETNOBJNAM, as shown in Example 6-26.

Example 6-26 Automatically assigned trigger system programs according to the value of QGETNOBJNAM

```
create schema newlib;
c1: CRTDTAARA DTAARA(NEWLIB/QGENOBJNAM) TYPE(*CHAR) LEN(10) ;
c1: CHGDTAARA DTAARA(NEWLIB/QGENOBJNAM *ALL) VALUE('?????50000');

create procedure newlib.longname_proc123 () language sql begin end;
create procedure newlib.longname_proc123a () language sql begin end;
create procedure newlib.longname_proc123b () language sql begin end;

create procedure newlib.longname_proc123_srv () PROGRAM TYPE SUB language sql begin end;
create procedure newlib.longname_proc123_srva () PROGRAM TYPE SUB language sql begin end;
create procedure newlib.longname_proc123_srvb () PROGRAM TYPE SUB language sql begin end;

create function newlib.longname_func123() returns int language sql begin return(10); end;
create function newlib.longname_func123a() returns int language sql begin return(10); end;
create function newlib.longname_func123b() returns int language sql begin return(10); end;

create table newlib.longname_table123 (c1 int);
create trigger newlib.longname_trig123 after insert on newlib.longname_table123 begin end;
create trigger newlib.longname_trig123a after insert on newlib.longname_table123 begin end;
create trigger newlib.longname_trig123b after insert on newlib.longname_table123 begin end;

select routine_name, external_name from qsys2.sysroutines where specific_schema = 'NEWLIB';

select TRIGGER_NAME, TRIGGER_PROGRAM_NAME from qsys2.systriggers where TRIGGER_SCHEMA = 'NEWLIB';
```

This command shows you the short names that are assigned to trigger programs.

6.3.18 CREATE PROCEDURE (SQL) PROGRAM TYPE SUB

A simple action that improves the performance of SQL procedures is using the **PROGRAM TYPE SUB** clause. When omitted or **PROGRAM TYPE MAIN** is used with the **CREATE PROCEDURE (SQL)** statement, an ILE C program (*PGM) is built for the procedure. **PROGRAM TYPE SUB** results in an ILE C service program (*SRVPGM) being built for the procedure. The usage of **PROGRAM TYPE SUB** is most relevant for procedures that are frequently called within a performance critical application.

PROGRAM TYPE SUB procedures perform better because ILE service programs are activated a single time per activation group, while ILE programs are activated on every call. The cost of an ILE activation is related to the procedure size, complexity, number of parameters, number of variables, and the size of the parameters and variables.

The only functional difference to be noted when you use **PROGRAM TYPE SUB** is that the QSYS2.SYSROUTINES catalog entry for the EXTERNAL_NAME column is formatted to show an export name along with the service program name.

6.3.19 Referential integrity and trigger performance

When a database **DELETE**, **UPDATE**, or **INSERT** operation is performed on behalf of a referential constraint or a trigger, the operation runs in a nested transaction. Before this enhancement, if many operations and nested transactions were performed as part of the outer transaction (because of multiple levels of cascading constraints), performance could suffer. With this enhancement, the larger the number of operations and nested transactions, the larger the performance improvement.

6.3.20 QSQBIGPSA data area

For some DB2 for i SQL applications, it is natural to accumulate and reuse *DUMMY cursors.

The default threshold for *DUMMY cursors is 150, but can be configured to be a higher threshold through the *QSQCSRTH* data area.

*DUMMY cursors exist when unique SQL statements are prepared using a statement name that is not unique. The SQL cursor name is changed to 'DUMMY' to allow the possibility of the cursor being reused in the future.

Prepared SQL statements are maintained within a thread scoped internal data structure that is called the Prepared Statement Area (PSA). This structure is managed by the database and can be compressed. The initial threshold of the PSA is small and gradually grows through use. For an application with heavy *DUMMY cursor use, you observe *DUMMY cursors being hard closed at each PSA compression.

This type of application is gaining little value from the PSA compression and must endure the performance penalty of its *DUMMY cursors being hard closed.

A new data area control is being provided for this type of user. QSQBIGPSA indicates that the application wants to start with a large size for the PSA threshold. By using this option, the application skips all the PSA compressions that it takes to reach a large PSA capacity. This control should be used with care, as PSA compression has value for most SQL users.

One way to determine the value of this data area for an application would be to use the Database Monitor and look for occurrences of QQRID=1000 & QQC21='HC' & QQC15 = 'N'. To use this control, the QSQBIGPSA data area must exist within the library list for a job when the first SQL **PREPARE** statement is ran. The data area merely needs to exist; it does not need to be set to any value.

6.3.21 Validating constraints without checking

In IBM i 7.1, a new **CHECK** parameter was added to the Change PF Constraint (**CHGPF CST**) command to allow a user to enable a constraint without checking. By default, when a referential or check constraint that is in a disabled state is enabled, DB2 verifies that the table's data conforms to the constraint definition.

This operation can be a long running one. **CHECK(*NO)** enables the constraint without checking. If the data is not checked when the constraint is enabled, it is the responsibility of the user to ensure that the data in the file is valid for the constraint.

Before Version 7.1, a data area can be created to enable a constraint without checking. When Change PF Constraint (**CHGPF CST**) is performed, DB2 searches for a data area in QTEMP called QDB_CHGPF CST. If the data area is found and its length is exactly nine characters and contains the value 'UNCHECKED', DB2 enables the constraint without validation.

6.3.22 Limiting the amount of processing on an RGZPFM cancel

A Reorganize Physical File Member (**RGZPFM**) command with an **ALWCANCEL(*YES)** parameter can be canceled and then later restarted where it left off.

Before this enhancement, a significant amount of processing was performed during the cancel to allow the Reorganize Physical File Member (**RGZPFM**) to be restarted later and to return as much storage to the system as possible.

With this enhancement, the amount of time processing that is performed at cancel time is minimized, allowing the Reorganize Physical File Member (**RGZPFM**) to be canceled in a reasonable amount of time. The processing that is bypassed is performed later when the Reorganize Physical File Member (**RGZPFM**) is restarted.

6.3.23 CPYFRMIMPF performance

Before this enhancement, when you issue the Copy from Import File (**CPYFRMIMPF**) command from an IFS file into a database file, the data in the IFS file for each character-string column of the database file was converted separately. Typically, all character-string columns of a table or physical file have the same CCSID. With this enhancement, the data for all the character-string columns of such a table or physical file can be converted in one operation rather than separately. This action can drastically reduce the processor that is used and the elapsed time for the Copy from Import File (**CPYFRMIMPF**). The more columns the table or physical file has, the larger the performance benefit.

6.3.24 QJOSJRNE API option to force journal entries without sending an entry

This enhancement provides a new option to force the journal receiver without sending an entry. If key 4 (FORCE) has a value of 2, the journal receiver is forced without sending an entry. If option 2 is specified, then key 4 must be the only key specified and the length of the entry data must be zero.

A *force journal entry* is an entry where the journal receiver is forced to auxiliary storage after the user entry is written to it. Possible values are:

0	The journal receiver is not forced to the auxiliary storage. This value is the default value if the key is not specified.
1	The journal receiver is forced to the auxiliary storage.

2	The journal receiver is forced to the auxiliary storage, but no journal entry is sent. When this value is specified, key 4 can be the only key specified and zero must be specified for the length of entry data. Specifying any other additional keys or a value other than zero for the length of entry data results in an error.

6.3.25 QDBRTVSN API performance

Before this enhancement, finding the short name for a given long name of a table or view was processed by enqueueing a request to the database cross-reference job, which would look up the short name in the cross-reference.

The QDBRTVSN() API now finds the short name in most cases without enqueueing a request to the database cross-reference.

6.3.26 Adding total DB opens job level instrumentation to Collection Services

Collection Services can be used to observe the total number of SQL full opens, SQL pseudo-opens, and the total number of databases full opens (SQL and native I/O). The database has instrumented the number of full opens that occur within a job. This metric is reported by Collection Services in the QAPMJOB05 file.

The new and existing fields contain the total number of times the specific operation occurred within the job during the Collection Services time interval.

Here are the new fields:

Field Name - JBNUS	The number of native database (non-SQL) files and SQL cursors that are fully opened. Subtracting the value within field JBLBO from JBNUS yields the number of non-SQL full opens.

Here are the existing fields (for SQL Cursors):

Field Name - JBLBO	The cumulative number of SQL cursors that are fully opened.
Field Name - JBLBS	The cumulative number of SQL cursors that are pseudo-opened. Pseudo-opens are also known as reused SQL cursors.

6.3.27 SYSTOOLS.REMOVE_INDEXES procedure

You can automate changes to the index strategy by using DB2 for i supplied examples within SYSTOOLS. You can use the DB2 for i procedures directly and pass hardcoded criteria or use the SYSTOOLS source as a fast-start to building your own procedures to manage indexes.

Creating indexes using **ACT_ON_INDEX_ADVICE** is demonstrated in Example 6-27, which creates permanent SQL indexes from instances of index advice where an MTI was used more than 1000 times.

Example 6-27 Example of index creation

```
CALL SYSTOOLS.ACT_ON_INDEX_ADVICE('PRODLIB',NULL,NULL,1000,NULL)
```

The call in Example 6-28 finds indexes that are created by **ACT_ON_INDEX_ADVICE** that are at least 7 days old. For any index that was used less than 500 times by the Query engine, drop the index.

Example 6-28 Example of index removal

```
CALL SYSTOOLS.REMOVE_INDEXES('PRODLIB', 500, ' 7 days ')
```

6.3.28 Improved SQE statistics maintenance during INSERT, UPDATE, and DELETE statements

The SQL Query Engine statistics processing now includes a proactive response to data changes as they happen to a database file, rather than just when the file is queried by SQE. The query engine checks for stale statistics during file inserts, updates, or deletes, including **INSERT**, **UPDATE**, or **DELETE** SQL statements.

When stale statistics are detected, a background statistics refresh is initiated, and statistics are refreshed before subsequent query processing, avoiding performance degradation that might occur because of stale statistics being used during query optimization.

This improvement is most beneficial in batched data change environments, such as a data warehouse, where many data change operations occur at one time and are followed by the execution of performance critical SQL queries.

6.3.29 QSYS2.Reset_Table_Index_Statistics procedure

This procedure zeros the **QUERY_USE_COUNT** and **QUERY_STATISTICS_COUNT** usage statistics for indexes over the specified tables.

These counts are also zeroed by the **CHGOBJD** command, but the command requires an exclusive lock. This procedure does *not* require an exclusive lock.

LAST_QUERY_USE, **LAST_STATISTICS_USE**, **LAST_USE_DATE**, and **NUMBER_DAYS_USED** are not affected.

The same wildcard characters (**_** and **%**) allowed in the SQL **LIKE** predicate are supported.

The procedure writes information that is related to any index processed into an SQL global temporary table.

The following query displays the results of the last call to the procedure:

```
select * from session.SQL_index_reset;
```

In Example 6-29, calls are made to zero the statistics for all indexes over a table, followed by a call to zero the statistics for all indexes over a table, starting with CAT and using the wildcard %.

Example 6-29 Using Reset_Table_Index_Statistics

The following call will zero the statistics for all indexes over table
STATST.CATHDANAA
call QSYS2.Reset_Table_Index_Statistics ('STATST', 'CATHDANAA')

The following call will zero the statistics for all indexes over any table in
schema STATST whose name starts with the letters CAT
call QSYS2.Reset_Table_Index_Statistics ('STATST', 'CAT%')

6.3.30 Performance enhancements for large number of row locks

Performance of row locking is enhanced for cases where a transaction or job acquires many locks.

Before this enhancement, the more row locks that were acquired on a table, the slower each additional row lock was acquired.

A marked improvement is seen in performance, as shown in Figure 6-8.




Figure 6-8 Row lock performance enhancement

6.3.31 Improved DSPJOB / CHKRCDLCK command behavior for large numbers of row locks.

The Check Record Locks (**CHKRCDLCK**) command returns the number of record locks in the current job. The performance of **CHKRCDLCK** is improved when many record locks are held by the job.

The Display Job (**DSPJOB**) command provides the ability to return the locks that are held by a job. If more records are held than can be displayed, a CPF9898 message is sent that indicates the number of record locks that are held by the job.

When a job holds more than 100,000 record locks, both of these commands run for a long time before they fail. The enhancement quickly recognizes the existence of a great number of record locks and returns the record lock count.

6.3.32 Chart-based graphical interface SQL Performance monitors

This enhancement provides the ability to view several high-level charts for specific SQL performance monitors, SQL plan cache snapshots, or SQL plan cache event monitors.

Starting the Investigate Performance Data action from System i Navigator or IBM Navigator for i displays the new graphical interface for SQL Performance monitors, as shown in Figure 6-9.




Figure 6-9 Start the “Investigate Performance Data” action from System i Navigator or IBM Navigator for i

6.3.33 Collection services include physical I/O breakdown by SQL activity

These metrics are included in the *JOBMI collection category.

An updated template file for QAPMJOBMI is shipped in QSYS with the additional fields, as shown in Figure 6-10. This template file is used only when you create a QAPMJOBMI file in libraries where it does not exist.

These metrics enable improved native versus SQL database performance isolation.




Figure 6-10 Native versus SQL

Here are some of the new metrics that are included:

- ▶ SQL clock time (total time in SQL and below) per thread (microseconds)
- ▶ High Level SQL Statement count
- ▶ SQL unscaled processor per thread (microseconds)
- ▶ SQL scaled processor per thread (microseconds)
- ▶ SQL synchronous database reads per thread
- ▶ SQL synchronous non-database reads per thread
- ▶ SQL synchronous database writes per thread
- ▶ SQL synchronous non-database writes per thread
- ▶ SQL asynchronous database reads per thread
- ▶ SQL asynchronous non-database reads per thread
- ▶ SQL asynchronous database writes per thread
- ▶ SQL asynchronous non-database writes per thread

6.3.34 Performance Data Investigator: New Database perspective

The New Database perspective is a robust graphical interface that is enabled by Collection Services Support. The new breakdown makes it easier to view changes, as shown in Figure 6-11.




Figure 6-11 New Database Perspective

6.3.35 Performance Data Investigator: SQL Overview

The SQL Overview contains 11 views. Figure 6-12 illustrates the Query Time Summary view, which is one of the 11 views.




Figure 6-12 SQL Overview - Query Time Summary view

6.3.36 Performance Data Investigator: SQL Attribute Mix

The SQL Attribute Mix contains nine views. Figure 6-13 shows the SQL Statement view.




Figure 6-13 SQL Attribute Mix - Statement summary view

6.3.37 Index Advisor: Show Statements - Improved query identification

The Show Statements view of the Index Advisor now uses precise filtering of the SQL Plan Cache. Before this enhancement, the Show Statements view showed all the queries that referenced the target table where the index advice was generated.

Show Statements finds queries based on how it is started:

- ▶ If you start it from Index Advice (exact match), it shows Keys Advised, Leading Order Independent Keys, NLSS Table and Schema, Index Type, and Page Size.
- ▶ If you use direct Show Statements (table match), you can select the index advice and table name filters.

An advanced form of the SQL Plan cache statements filter is populated by IBM i Navigator, as shown in Figure 6-14.




Figure 6-14 An advanced form of the SQL Plan Cache Statements filter is populated by IBM i Navigator

These attributes cannot be directly modified by the user.

6.3.38 Performance improvements for temporary tables

The **DECLARE GLOBAL TEMPORARY TABLE WITH REPLACE** SQL statement might complete much faster when the same table is created again in the same job. In most cases, if the SQL statement text is identical to the previous table, the table is cleared instead of dropped and re-created.

Also, the SQL Plan Cache is improved to recognize the cases where the temporary table is reused with an identical table format. The plan cache plan and statistics management is improved to retain and reuse plans for temporary tables.

While the overall performance characteristics of temporary tables continue to improve, it is not the best performing choice when you deal with large amounts of data in a performance critical environment.

6.4 New functionality for DB2 developers

This section covers the new functionality for DB2 for i developers.

6.4.1 QSYS2.SYSCOLUMNS2

QSYS2.SYSCOLUMNS2 is a view that is based on a table function that returns more information that is not available in SYSCOLUMNS (such as the allocated length of a varying length column). Because it is based on a table function, it typically returns results faster if a specific table is specified when querying it.

For more information about the view, go to:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=%2Fdb2%2Ffrbafzcatsyscol2.htm>

6.4.2 QSYS2.SYSPARTITIONDISK and QSYS2.SYSPARTITIONINDEXDISK

These two catalog views return allocation information for tables and indexes. The views can be useful in determining how much storage for a partition of index is allocated on an SSD.

For more information, see the *SQL Reference* at:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=%2Fdb2%2Ffrbafzcatsyspdisk.htm>

Example 6-30 shows the return allocation information for DB2 tables and physical files in MJATST.

Example 6-30 Return allocation information for DB2 tables and physical files in MJATST

```
SELECT MAX(table_schema) AS table_schema, MAX(table_name) AS table_name,
MAX(table_partition) AS table_partition,
SUM(CASE WHEN unit_type = 1 THEN unit_space_used ELSE null END) AS ssd_space,
SUM(CASE WHEN unit_type = 0 THEN unit_space_used ELSE null END) AS non_ssd_space
FROM qsys2.syspartitiondisk a
WHERE system_table_schema = 'MJATST'
GROUP BY a.table_schema, a.table_name, table_partition
ORDER BY 1,2,3;
```

Figure 6-15 shows the results of Example 6-30.

TABLE_SCHEMA	TABLE_NAME	TABLE_PARTITION	SSD_SPACE	NON_SSD_SPACE
MJATST	/BIC/FZAWPARTX	PART000001	-	4096
MJATST	/BIC/FZAWPARTX	PART000002	-	4096
MJATST	t_	t_	-	4096
MJATST	A	A	-	4096
MJATST	ABADMON	ABADMON	-	1613824
MJATST	ABCSRC	ABCSRC	-	4096
MJATST	ACHRIS	ACHRIS	-	552960
MJATST	ADATES	ADATES	-	4096
MJATST	ADC_MANO	ADC_MANO	-	7273971712
MJATST	ADC_SHORT	ADC_SHORT	-	7143948288
MJATST	ADC_010606	ADC_010606	-	1091567616
MJATST	AGDBWE02	AGDBWE02	-	57671680
MJATST	AGTDTL	AGTDTL	-	4096

Figure 6-15 Results of Example 6-30

Example 6-31 shows the return allocation information for DB2 indexes (keyed files, constraint, and SQL indexes) in MJATST.

Example 6-31 Return allocation information for DB2 indexes (keyed files, constraint, and SQL indexes) in MJATST

```
SELECT index_schema, index_name, index_member, index_type,
SUM(CASE unit_type WHEN 1 THEN unit_space_used ELSE 0 END)/COUNT(*) AS ssd_space,
SUM(CASE unit_type WHEN 0 THEN unit_space_used ELSE 0 END)/COUNT(*) AS nonssd_space
```

```
FROM qsys2.syspartitionindexdisk b
WHERE system_table_schema = 'MJATST'
GROUP BY index_schema, index_name, index_member, index_type;
```

Figure 6-16 shows the results of Example 6-31 on page 206.

INDEX_SCHEMA	INDEX_NAME	INDEX_MEMBER	SSD_SPACE	NON_SSD_SPACE
MJATST	QZG0000155_QINX4	QZG0010014	0	139264
MJATST	Q_MJATST_MYSPRI2_C1_00001	-	0	139264
MJATST	RPT_INDEX	PART000001	0	139264
MJATST	QZG0001566_QINX6	QZG0010004	0	139264
MJATST	MYSI	MYSI	0	139264
MJATST	QIDCTP31	QIDCTP31	0	24576
MJATST	Q_MJATST_UNT1_C1_00001	-	0	139264
MJATST	I2BRCDFMT	I2BRCDFMT	0	139264
MJATST	END_JI	END_JI	0	139264
MJATST	PK_KEY	-	0	139264
MJATST	KEYSTRRN_FK1	-	0	139264
MJATST	Q_MJATST_DFRCT2_C1_00001	-	0	69632

Figure 6-16 Results of Example 6-31 on page 206

6.4.3 QSYS2.OBJECT_STATISTICS table function ease of use

The **OBJECT_STATISTICS** table function returns information about objects in a schema (library).

The first argument is the long or short library name. The second argument is a list of object types to be filter the objects in the library. The second parameter is enhanced to allow a list of object types. The object types in the list can be separated by blanks or commas or a combination of both.

The specified object types can include or exclude the leading *. For example, either FILE or *FILE can be specified.

Example 6-32 shows the allowed formats of the object type lists.

Example 6-32 Allowed formats of the object type lists

```
-- Find all journals in schema MJATST.
select * from table (qsys2.OBJECT_STATISTICS('MJATST ','JRN') ) as x ;
select * from table (qsys2.OBJECT_STATISTICS('MJATST ','*JRN') ) as x ;

-- Find all journals and journal receivers in schema MJATST.
select * from table (qsys2.OBJECT_STATISTICS('MJATST ','JRN JRNRCV') ) as x ;
select * from table (qsys2.OBJECT_STATISTICS('MJATST ','*JRN *JRNRCV') ) as x ;
```

6.4.4 EARLIEST_POSSIBLE_RELEASE

DB2 for i now provides a way to see the earliest IBM i release that could be used for any SQL statement or program. The SQL statement level detail is available through Database Monitor. The program level detail is available through the QSYS2.SYSPROGRAMSTAT and QSYS2.SYSPACKAGESTAT catalogs. In both cases, you must capture the DBMON or rebuild the program after you apply the latest DB Group PTFs.

Database Monitor and the QSYS2.SYSPROGRAMSTAT and QSYS2.SYSPACKAGESTAT catalogs can be used to evaluate SQL application deployment possibilities per operating system releases. The QQC82 column contains the earliest IBM i OS release level where this SQL statement is supported. This information can be used to assess whether applications can be deployed on earlier IBM i releases or whether they are using SQL functionality unique to IBM i 6.1 or 7.1.

This field applies only if the SQL statement is dynamic (QQC12= 'D').

Possible values for QQC82 are:

'	The statement release level is not determined.
'ANY'	The statement is valid on any supported IBM i OS release,
'V6R1M0'	The statement is valid on IBM i 6.1 or later.
'V7R1M0'	The statement is valid on IBM i 7.1 or later.

The QSYS2.SYSPROGRAMSTAT and QSYS2.SYSPACKAGESTAT column name is EARLIEST_POSSIBLE_RELEASE. The System column name is MINRLS.

6.4.5 SIGNAL support for native triggers

INSERT, **UPDATE**, and **DELETE** SQL statements were changed to recognize when system triggers use the **SIGNAL** SQL statement to communicate failure details to the application.

If the system trigger runs the **SIGNAL** statement and sends an escape message to its caller, the SQL **INSERT**, **UPDATE**, or **DELETE** statement fails with MSGSQL0438 (SQLCODE=-438) instead of MSGSQL0443.

The **SQLSTATE**, **MSG**, and other values within the SQL diagnostics area or SQLCA contain the values that are passed into the **SIGNAL** statement.

For more information, go to:

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp?topic=/dbp/rbaforzahftrm.htm>

The website contains recommendations for native trigger programs. Here is an example:

“Signal an exception if an error occurs or is detected in the trigger program. If an error message is not signaled from the trigger program, the database assumes that the trigger ran successfully. This might cause the user data to end up in an inconsistent state.”

The **SIGNAL** SQL statement provides the SQL linkage between the native trigger and the application that caused the trigger to be fired by using SQL.

The **SIGNAL** SQL statement does not signal an exception, so be sure to use the `QMHSNDPM()` API to send an escape message after you run the **SIGNAL** statement, as shown in Example 6-33.

Example 6-33 JDBC failure

```
*** SQLException caught ***
Statement was insert into mylib.mytable values(1)
SQLState: IWF99
Message: [SQL0438] DOCUMENT NOT FOUND
Vendor: -438
java.sql.SQLException: [SQL0438] DOCUMENT NOT FOUND
at com.ibm.as400.access.JDError.throwSQLException(JDError.java:650)
etc.....
```

6.4.6 Hierarchical queries through the **CONNECT BY** clause

DB2 for i has had recursive query support since V5R4. Another recursive query technique called *hierarchical query* was added. This technique is a more concise method of representing a recursive query.

For more information, see the IBM i 7.1 Information Center at the following websites:

- ▶ <http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/sqlp/rbafyreursivequeries.htm>
- ▶ <http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/db2/rbafzhierquery.htm>

Example 6-34 shows a hierarchal query example.

Example 6-34 Hierarchical query example

```
CALL QSYS.CREATE_SQL_SAMPLE('MYDB');
SET CURRENT SCHEMA MYDB;
SET CURRENT PATH MYDB;

SELECT LEVEL,
CAST(SPACE((LEVEL - 1) * 4) || '/' || DEPTNAME AS VARCHAR(40)) AS DEPTNAME
FROM DEPARTMENT
START WITH DEPTNO = 'A00'
CONNECT BY NOCYCLE PRIOR DEPTNO = ADMRDEPT
```

Figure 6-17 shows the result of Example 6-34 on page 209.

LEVEL	DEPTNAME
1	/SPIFFY COMPUTER SERVICE DIV.
2	/SUPPORT SERVICES
3	/BRANCH OFFICE J2
3	/BRANCH OFFICE I2
3	/BRANCH OFFICE H2
3	/BRANCH OFFICE G2
3	/BRANCH OFFICE F2
3	/SOFTWARE SUPPORT
3	/OPERATIONS
2	/DEVELOPMENT CENTER
3	/ADMINISTRATION SYSTEMS
3	/MANUFACTURING SYSTEMS
2	/INFORMATION CENTER
2	/PLANNING
2	/SPIFFY COMPUTER SERVICE DIV.

Figure 6-17 Result of hierarchical query

6.4.7 Additional parameter marker support (LAND, LOR, XOR, and TRANSLATE)

Before IBM i 7.1, there were many restrictions on where a parameter marker was allowed in an SQL statement. Many of these restrictions were removed in IBM i 7.1.

The **LAND**, **LOR**, **XOR**, and **TRANSLATE** scalar functions were enhanced by removing similar restrictions.

Example 6-35 shows the possible usage of parameter markers.

Example 6-35 Possible usage of parameter markers

```
PREPARE s1 FROM 'SELECT TRANSLATE(c1,?,?,?) FROM t1'
PREPARE s1 FROM 'SELECT LAND(c2,?,?,?), LOR(c2,?,?,?), XOR(c2,?,?,?) FROM t1'
```

6.4.8 Supporting PROGRAM NAME on CREATE TRIGGER

It is now possible to specify a short name for the created trigger program. When this specification is not supplied, the database determines the system name, which could lead to differences in the system name for the trigger program across different machines.

Example 6-36 shows the usage of PROGRAM NAME in the CREATE TRIGGER definition. If a program name is not specified, then the system generates one, such as TR1_U00001 or TR1_U00002.

Example 6-36 Usage of PROGRAM NAME in the CREATE TRIGGER definition

```
CREATE TRIGGER TR1_UPDATE_TRACKER
AFTER UPDATE OF c1 ON TR1
REFERENCING OLD AS o NEW AS n
FOR EACH ROW MODE DB2ROW
PROGRAM NAME TR1_UPTRIG
```

```
BEGIN ATOMIC
INSERT INTO tr2 VALUES(default, o.c1, n.c1);
END
```

6.4.9 TINYINT in CLI

This SQL Call Level Interface (CLI) enhancement allows applications using CLI APIs for binding parameters and output fields for result sets to accept a new bind type named SQL_C_UTINYINT, Unsigned TINYINT. This bind type represents a 1-byte unsigned integer value with values 0 - 255.

For more information, see the IBM i 7.1 SQL CLI documentation at:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/cli/rzadpwhatsnew.htm>

6.4.10 CHGPFM and CHGLFM UNIT support

Change Physical File Member (**CHGPFM**) and Change Logical File Member (**CHGLFM**) commands can now be used to move an individual member to or from an SSD by changing the media preference. One of the main benefits of using these commands is that they do not require a LENR lock. They conflict with another job that has an *SHRNUP, *EXCLRD, or *EXCL lock on the data, however. An exclusive seize is acquired by SLIC DB to actually move the data.

The syntax of these commands is:

- ▶ **CHGPFM t1 UNIT(*SSD)**
- ▶ **CHGLFM v1 UNIT(*SSD)**

If the user is using logical replication, you need the PTFs on the target and the source systems.

6.4.11 SYSTOOLS procedures

SYSTOOLS is a set of DB2 for i supplied examples and tools. SYSTOOLS is the name of a database supplied schema (library). SYSTOOLS differs from other DB2 for i supplied schemas (QSYS, QSYS2, SYSIBM, and SYSIBMADM) in that it is not part of the default system path. As general purpose useful tools or examples that are built by IBM, they are considered for inclusion within SYSTOOLS. SYSTOOLS provides a wider audience with the opportunity to extract value from the tools.

It is the intention of IBM to add content dynamically to SYSTOOLS, either on base releases or through PTFs for field releases. A preferred practice for customers who are interested in such tools is to periodically review the contents of SYSTOOLS.

For more information, see the IBM Information Center at:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/rzajq/rzajqsytools.htm>

6.4.12 Adding a client and server IP address and port number to the QSYS2.TCPIP_INFO view

Client / server identification, logging, and other instrumentation is possible using SQL. The QSYS.TCPIP_INFO view returns detailed information about the current connection.

The SERVER_IP_ADDRESS and SERVER_IP_ADDRESS_TYPE are only available on IBM i 7.1.

Example 6-37 shows the usage of QSYS2.TCPIP_INFO.

Example 6-37 Usage of QSYS2.TCPIP_INFO

```
SELECT * from QSYS2.TCPIP_INFO
```

Figure 6-18 shows the result of Example 6-37.

COLLECTED_TIME	LOCAL_HOST_NAME	CLIENT_IP_ADDRESS_TYPE	CLIENT_IP_ADDRESS	
2011-09-29 11:28:21...	MCV7R1.RCHLAND.IBM.COM	IPV4	9.10.126.136	
CLIENT_PORT_NUMBER	SERVER_IP_ADDRESS_TYPE	SERVER_IP_ADDRESS	SERVER_PORT_NUMBER	HOST-VERSION
1494	IPV4	9.5.168.119	8471	V7R1M0

Figure 6-18 Result of QSYS2.TCPIP_INFO

6.4.13 QSYS2.GROUP_PTF_INFO view

You can use this view to use SQL to retrieve IBM i Group PTF information for the server. The data that is returned is similar to output of the Work with PTF Groups (**WRKPTFGRP**) command.

On IBM i 7.1, the Technology Refresh (TR) level can be determined by using the view definition that is shown in Example 6-38.

Example 6-38 QSYS2.GROUP_PTF_INFO

```
COLLECTED_TIME FOR COLUMN COLLE00001 TIMESTAMP
PTF_GROUP_NAME FOR COLUMN PTF_G00001 VARCHAR(60) ALLOCATE(60)
PTF_GROUP_DESCRIPTION FOR COLUMN PTF_G00002 VARCHAR(100) ALLOCATE(100)
PTF_GROUP_LEVEL FOR COLUMN PTF_G00003 INTEGER DEFAULT NULL
PTF_GROUP_TARGET_RELEASE FOR COLUMN PTF_G00004 VARCHAR(6) ALLOCATE(6)
PTF_GROUP_STATUS FOR COLUMN PTF_G00005 VARCHAR(20) ALLOCATE(20)
```

Run the following command:

```
select * from QSYS2.GROUP_PTF_INFO order by PTF_GROUP_LEVEL DESC;
```

The result of the command is shown in Figure 6-19.

COLLECTED_TIME	PTF_GROUP_NAME	PTF_GROUP_DESCRIPTION	PTF_GROUP_LEVEL
2011-09-30 11:04:19...	SF99710 ...	CUMULATIVE PTF PACKAGE C1116710	11116
2011-09-30 11:04:19...	SF99709 ...	GROUP HIPER	40
2011-09-30 11:04:19...	SF99701 ...	DB2 FOR IBM I	11
2011-09-30 11:04:19...	SF99362 ...	710 BACKUP RECOVERY SOLUTIONS	10
2011-09-30 11:04:19...	SF99368 ...	IBM HTTP SERVER FOR I	10
2011-09-30 11:04:19...	SF99701 ...	DB2 FOR IBM I	10
2011-09-30 11:04:19...	SF99369 ...	IBM I INTEGRATION WITH BLADECENTER AND SYSTEM X	8
2011-09-30 11:04:19...	SF99708 ...	GROUP SECURITY	8
2011-09-30 11:04:19...	SF99363 ...	WEBSPHERE APP SERVER V7.0	6
2011-09-30 11:04:19...	SF99572 ...	JAVA	6
2011-09-30 11:04:19...	SF99617 ...	DB2 WEB QUERY FOR IBM I V1.1.1	6
2011-09-30 11:04:19...	SF99359 ...	WEBSPHERE MQ FOR I5/OS -V7.0.1	5
2011-09-30 11:04:19...	SF99627 ...	7.1 ELECTRONIC SERVICES GROUP PTF	5
2011-09-30 11:04:19...	SF99364 ...	WEBSPHERE APP SERVER V6.1	4

COLLECTED_TIME	PTF_GROUP_NAME	PTF_GROUP_DESCRIPTION	PTF_GROUP_LEVEL
2011-09-30 11:04:19...	SF99710 ...	CUMULATIVE PTF PACKAGE C1116710	11116
2011-09-30 11:04:19...	SF99709 ...	GROUP HIPER	40
2011-09-30 11:04:19...	SF99701 ...	DB2 FOR IBM I	11
2011-09-30 11:04:19...	SF99362 ...	710 BACKUP RECOVERY SOLUTIONS	10
2011-09-30 11:04:19...	SF99368 ...	IBM HTTP SERVER FOR I	10
2011-09-30 11:04:19...	SF99701 ...	DB2 FOR IBM I	10
2011-09-30 11:04:19...	SF99369 ...	IBM I INTEGRATION WITH BLADECENTER AND SYSTEM X	8
2011-09-30 11:04:19...	SF99708 ...	GROUP SECURITY	8
2011-09-30 11:04:19...	SF99363 ...	WEBSPHERE APP SERVER V7.0	6
2011-09-30 11:04:19...	SF99572 ...	JAVA	6
2011-09-30 11:04:19...	SF99617 ...	DB2 WEB QUERY FOR IBM I V1.1.1	6
2011-09-30 11:04:19...	SF99359 ...	WEBSPHERE MQ FOR I5/OS -V7.0.1	5
2011-09-30 11:04:19...	SF99627 ...	7.1 ELECTRONIC SERVICES GROUP PTF	5
2011-09-30 11:04:19...	SF99364 ...	WEBSPHERE APP SERVER V6.1	4

Figure 6-19 Results from QSYS2.GROUP_PTF_INFO call

The PTF group status messages are:

UNKNOWN	The PTF's group status cannot be resolved because a related PTF group is either not found on the system or is in error.
NOT APPLICABLE	All PTFs in the PTF group and related PTF groups are for products that are not installed or supported on this system.
SUPPORTED ONLY	There are no PTFs in the PTF group or related PTF groups that are for installed products on this system. There is at least one PTF that is for a product, release, option, and load identifier that is supported on this system.
NOT INSTALLED	There is at least one PTF that is for an installed product on this system, and not all of the PTFs or their superseding PTFs are temporarily or permanently applied.
INSTALLED	All PTFs for products that are installed on this system are temporarily or permanently applied. If a PTF is superseded, a superseding PTF is either temporarily or permanently applied.
ERROR	The PTF group information is in error. Either delete the PTF group or replace the PTF group information that is on the system.

APPLY AT NEXT IPL	All PTFs for the installed products on the system are either set to be applied at the next IPL or are temporarily or permanently applied.
RELATED GROUP	The PTF group does not have any PTFs for products that are installed or supported on the system. However, it is identified in another PTF group as a related PTF group. Deleting a PTF group in this status causes the other PTF group to have a status of Unknown.
ON ORDER	There is at least one PTF in the group that is on order and has not yet been installed on the system. It is delivered on either physical or virtual media.

6.4.14 QSYS2.DUMP_SQL_CURSORS procedure

It is now possible to capture the list of open cursors for a job in IBM i. The QSYS2.DUMP_SQL_CURSORS are:

- ▶ Job_Name VARCHAR(28),
- ▶ Library_Name CHAR(10),
- ▶ Table_Name CHAR(10),
- ▶ Output_Option integer)

Where:

- ▶ Job_Name is a qualified job name or a special value of '*' to indicate the current job.
- ▶ Library_Name is an optional library name for the procedure output.
- ▶ Table_Name is an optional table name for the procedure output.

Output_Option has these choices:

- ▶ Ignore Library_Name and Table_Name inputs and return a result set.
- ▶ Ignore Library_Name and Table_Name inputs and place the results in table QTEMP/SQL_CURSORS (no result set).
- ▶ Place the results in table in Library_Name and Table_Name (no result set). If the table does not exist, the procedure creates it. If the table does exist, the results are appended to the existing table.
- ▶ Place the results in table in Library_Name and Table_Name (no result set). If the table does not exist, do not create the table.

Example 6-39 shows a possible invocation of the QSYS2.DUMP_SQL_CURSORS procedure.

Example 6-39 Possible invocation

```
-- populate QGPL.SQLCSR1 table with open SQL cursors in this job
call qsys2.DUMP_SQL_CURSORS('*', 'QGPL', 'SQLCSR1', 3);

-- return a result set with open SQL cursors in this job
call qsys2.DUMP_SQL_CURSORS('*', '', '', 1);

-- populate QGPL.SQLCSR1 table with open SQL cursors for a target job
call qsys2.DUMP_SQL_CURSORS('724695/QUSER/QZDASOINIT', '', '', 1);
```

Table/Result Set format:

```
SQL_IDENTITY FOR COLUMN SQL_I00001 INTEGER ,
```

```
DUMPTIME TIMESTAMP ,  
DUMP_BY_USER FOR COLUMN DUMPUSER VARCHAR(18) ,  
CURSOR_NAME FOR COLUMN CSRNAME VARCHAR(128) ,  
PSEUDO_CLOSED FOR COLUMN PSEUDO VARCHAR(3) ,  
STATEMENT_NAME FOR COLUMN STMTNAME VARCHAR(128) ,  
OBJECT_NAME FOR COLUMN OBJNAME CHAR(10) ,  
OBJECT_LIBRARY FOR COLUMN OBJLIB CHAR(10) ,  
OBJECT_TYPE FOR COLUMN OBJTYPE CHAR(10) ,  
JOBNAME CHAR(28)
```

6.4.15 QIBM_SQL_NO_RLA_CANCEL environment variable

The SQL Cancel support includes logic to ensure that DB2 for i programs are active on the stack of the initial thread in the target job for the cancel request. Applications that use Native DB I/O can observe cases where the cancel request is processed and a record level access operation ends with MSGCPF5257 followed by MSGCPF9999.

An environment variable can be used by the customer to direct DB2 for i to avoid canceling RLA access operations. Upon the first cancel request for a specific job, the environment variable QIBM_SQL_NO_RLA_CANCEL is accessed. If the environment variable exists, the cancel request is not honored when RLA is the only database work ongoing within the initial thread at the time the cancel request is received.

The environment variable is the SQL Cancel operational switch. The variable can be created at the job or system level. Creating it once at the system level affects how SQL Cancels are processed for all jobs.

Possible invocations of this variable are:

- ▶ ADDENVVAR ENVVAR(QIBM_SQL_NO_RLA_CANCEL)
- ▶ ADDENVVAR ENVVAR(QIBM_SQL_NO_RLA_CANCEL) LEVEL(*SYS)

6.4.16 QSYS2.CANCEL_SQL and QSYS2.FIND_AND_CANCEL_QSQRVR_SQL procedures

In this section, we describe the QSYS2.FIND_AND_CANCEL_QSQRVR_SQL and QSYS2.CANCEL_SQL procedures.

QSYS2.CANCEL_SQL procedure

The IBM supplied procedure, QSYS2.CANCEL_SQL(), can be called to request the cancellation of an SQL statement for a target job.

SQL Cancel support provides an alternative to end job immediate when you deal with an orphaned or runaway process. End job immediate is like a hammer, where SQL Cancel is more like a tap on the shoulder. Before this improvement, the SQL Cancel support was only available for ODBC, JDBC, and SQL CLI applications. The QSYS2.CANCEL_SQL() procedure extends the SQL Cancel support to all application and interactive SQL environments.

When an SQL Cancel is requested, an asynchronous request is sent to the target job. If the job is processing an interruptible, long-running machine operation, analysis is done within the job to determine whether it is safe to cancel the statement. When it is determined that it is safe to cancel the statement, an SQL0952 escape message is sent, causing the statement to terminate.

If it is not safe to end the SQL statement, or if there is no active SQL statement, the request to cancel is ignored. The caller of the cancel procedure observes a successful return code that indicates that only the caller had the necessary authority to request a cancel and that the target job exists. The caller of the **QSYS2.CANCEL_SQL()** procedure has no programmatic means of determining that the cancel request resulted in a canceled SQL statement.

Procedure definition

The **QSYS2.CANCEL_SQL** procedure is defined as follows:

```
CREATE PROCEDURE QSYS2.CANCEL_SQL (  
  IN VARCHAR(28) )  
LANGUAGE PLI  
SPECIFIC QSYS2.CANCEL_SQL  
NOT DETERMINISTIC  
MODIFIES SQL DATA  
CALLED ON NULL INPUT  
EXTERNAL NAME 'QSYS/QSQSSUDF(CANCEL_SQL) '  
PARAMETER STYLE SQL ;
```

Example 6-40 shows the calling of the procedure.

Example 6-40 CALL QSYS2.CANCEL_SQL procedure

```
CALL QSYS2.CANCEL_SQL('483456/QUSER/QZDASOINIT');
```

Authorization

The **QSYS2.CANCEL_SQL** procedure requires that the authorization ID associated with the statement have *JOBCTL special authority.

Description

The procedure has a single input parameter, that is, the qualified job name of the job that should be canceled. The job name must be uppercase. If that job is running an interruptible SQL statement or query, the statement is canceled. The application most likely receives an SQLCODE = SQL0952 (-952) message. In some cases, the failure that is returned could be SQL0901 or the SQL0952 could contain an incorrect reason code.

This procedure takes advantage of the same cancel technology that is used by the other SQL cancel interfaces:

- ▶ System i Navigator's Run SQL Scripts: Cancel Request button
- ▶ SQL Call Level Interface (CLI): SQLCancel() API
- ▶ JDBC method: Native Statement.cancel() and toolbox com.ibm.as400.access.AS400JDBCStatement.cancel()
- ▶ Extended Dynamic Remote SQL (EDRS): Cancel EDRS Request (QxdaCancelEDRS) API
- ▶ **QSYS2.CANCEL_SQL()** procedure

If the cancel request occurs during the act of committing or rolling back a commitment-control transaction, the request is ignored.

Failures

The procedure fails with a descriptive SQL0443 failure if the target job is not found.

The procedure fails with SQL0443 and SQL0552 if the caller does not have *JOBCTL user special authority.

Commitment control

When the target application is running without commitment control (that is, COMMIT = *NONE or *NC), the canceled SQL statement terminates without rolling back the partial results of the statement. If the canceled statement is a query, the query ends. However, if the canceled statement was a long-running **INSERT**, **UPDATE**, or **DELETE** SQL statement, the changes that are made before cancellation remain intact.

If the target application is using transaction management, the SQL statement is running under the umbrella of a transaction save point level. When those same long running **INSERT**, **UPDATE**, or **DELETE** SQL statements are canceled, the changes that are made before cancellation are rolled back.

In both cases, the application receives back control with an indication that the SQL statement failed. It is up to the application to determine the next action.

Useful tool

The **QSYS2.CANCEL_SQL()** procedure provides a useful tool to database administrators for IBM i machines. After you have the latest DB Group PTF installed, you can start calling this procedure to stop long-running or expensive SQL statements.

QSYS2.FIND_AND_CANCEL_QSQSRVR_SQL procedure

The **QSYS2.FIND_AND_CANCEL_QSQSRVR_SQL()** procedure uses the **QSYS2.FIND_QSQSRVR_JOBS** and **QSYS2.CANCEL_SQL()** procedures that are derived from the set of jobs that has active SQL activity, given a target application job. Each job found is made a target of an SQL Cancel request.

Example 6-41 shows an example of the procedure.

Example 6-41 Using the QSYS2.FIND_AND_CANCEL_QSQSRVR_SQL procedure

```
CALL (QSYS2.FIND_AND_CANCEL_QSQSRVR_SQL('564321/APPUSER/APPJOBNAME'));
```

6.4.17 QSYS2.FIND_QSQSRVR_JOBS() procedure

Anyone responsible for administering, tuning, or explaining the SQL Server Mode (for example, QSQSRVR jobs) activity might find the **QSYS2.FIND_QSQSRVR_JOBS()** procedure a useful tool. This procedure was added to QSYS2 after the application of PTFs. The procedure has a single parameter, which is the qualified job name of an application job. If the target job is active and is set up to use SQL Server Mode, the procedure determines which QSQSRVR jobs are being used by the application, in the form of active SQL Server Mode connections. The procedure collects and returns work management, performance, and SQL information and returns two SQL result sets:

- ▶ Summary information
- ▶ Detailed SQL Server Mode job information

How is this procedure useful? When you have an important application instance (job) that uses QSQSRVR jobs, it can be difficult to determine the “total system impact” of the application. How many SQL Server Mode jobs are in use at that moment? Is this application responsible for a QSQSRVR job that is consuming many processor cycles or holding onto object locks? The **FIND_QSQSRVR_JOBS()** procedure provides some of these answers by tying together the application and its SQL Server Mode job usage.

Example 6-42 shows an invocation of **QSYS2.FIND_QSQRVR_JOBS**.

Example 6-42 Invocation of QSYS2.FIND_QSQRVR_JOBS

```
call QSYS2.FIND_QSQRVR_JOBS('566463/EBERHARD/QPOZSPWP ')
```

Procedure definition

QSYS2.FIND_QSQRVR_JOBS is defined as follows:

```
CREATE PROCEDURE QSYS2.FIND_QSQRVR_JOBS( JOB_NAME VARCHAR(28) )
NOT DETERMINISTIC
MODIFIES SQL DATA
CALLED ON NULL INPUT
DYNAMIC RESULT SETS 2
SPECIFIC FINDSRVR
EXTERNAL NAME 'QSYS/QSQSSUDF(FINDSRVR)'
LANGUAGE C PARAMETER STYLE SQL;
```

Authorization

On IBM i 6.1, to start **QSYS2.FIND_QSQRVR_JOBS**, you need *JOBCTL special authority.

On IBM i 7.1, to start **QSYS2.FIND_QSQRVR_JOBS**, you need *JOBCTL special authority, **QIBM_DB_SQLADM** Function usage, or **QIBM_DB_SYSMON** Function usage. Otherwise, you receive the following message:

```
call QSYS2.FIND_QSQRVR_JOBS('650261/SCOTT/QPOZSPWP')
SQL State: 38501
Vendor Code: -443
Message: [CPF43A4] *JOBCTL special authority, QIBM_DB_SQLADM or QIBM_DB_SYSMON
Function usage is required. Cause.....: The user profile is required to have
*JOBCTL special authority or be authorized to either the QIBM_DB_SQLADM or
QIBM_DB_SYSMON Function through Application Administration in System i Navigator.
The Change Function Usage (CHGFCNUSG) command can also be used to allow or deny
use of the function.
```

Usage

The procedure can be called from any environment. The input parameter is the application qualified job name. When called from within System i Navigator's Run SQL Scripts, two results sets are displayed. When called from Start SQL Interactive Session (**STRSQL**) or elsewhere, you must query the temporary tables to see the data, as shown in Example 6-43.

Example 6-43 Usage for STRSQL

```
select * from qtemp.QSQRVR_DETAIL order by TOTALCPU desc;
select * from qtemp.QSQRVR_SUMMARY;
#Use this query to see the summary information in the same form that is returned
#within the result set:
SELECT SERVER_MODE_JOB,count(*) AS "QSQRVR JOB COUNT",
SERVER_MODE_CONNECTING_JOB, SUM(TOTAL_PROCESSING_TIME) AS "CPU USED
(MILLISECONDS)", SUM(TEMP_MEG_STORAGE) AS "TEMP STORAGE USED (MB)",
SUM(PAGE_FAULTS) AS "PAGE FAULTS", SUM(IO_REQUESTS) AS "I/O REQUESTS" from
SESSION.QSQRVR_SUMMARY GROUP BY GROUPING SETS (SERVER_MODE_JOB ,
SERVER_MODE_CONNECTING_JOB) ORDER BY 1;
```

6.4.18 SQL server mixed mode for batch processing

SQL Server Mode users must be able to fire SQL trigger programs from the Server Mode client job, through native database I/O operations, and have any SQL statements within the trigger program that is run within the client job instead of being rerouted to a QSQSRVR server job.

SQL Server Mode is extended to allow an application to direct the database to run SQL statements within an SQL trigger program within the SQL Server Mode client job instead of rerouting the SQL to a QSQSRVR server job.

The change affects only SQL triggers fired through native database I/O operations.

To enable the new function, an environment variable must exist before any SQL statements are run within the client job. An easy way to deploy the environment variable is to define it at the system level as follows:

```
ADDENVVAR ENVVAR(QIBM_DB2_MIXED_SERVER_MODE) LEVEL(*SYS)
```

Restrictions and usage information

The environment variable just needs to exist; it does not need to be assigned a specific value. After a job chooses to use this mixed mode support, it cannot turn off the choice.

The SQL triggers must not be built to use commitment control.

The SQL triggers must not use statement level isolation level support to run statements using commitment control.

The SQL triggers must not directly or indirectly use Java/JDBC or CLI.

The triggers must not use DRDA.

If the client job is multi-threaded and triggers are fired in parallel over different threads, the mixed-mode server mode solution serializes the execution of the triggers. Only one trigger is allowed to run at a time.

The solution does not apply to native triggers, such as Add Physical File Trigger (ADDPFTRG), built over programs that use SQL. The solution does not include SQL triggers that call procedures, fire user-defined functions, or cause nested triggers to run.

6.4.19 QDBRPLAY() API: Disable or Enable Constraints option

A new option was added to the QDBRPLAY() API that, similar to the Disable triggers option, allows the user of the API to specify a Disable constraints option. This option can improve performance of logical replication products by disabling any constraints on a backup system. A new CHAR(1) option is added to the beginning of the Reserved area.

The disable constraints indicator controls whether constraints that are added or changed as a result of replaying a CT, AC, or GC journal entry should be automatically disabled. The disable constraint indicator does not apply to unique constraints. It has two settings:

0	Do not disable constraints.
1	Disable constraints.

6.4.20 SQL0901 log education

SQL0901 is a message that is issued by DB2 for i when SQLCODE = -901 indicates that an unexpected error was encountered. Sometimes (for example, when the job has no job log), it is difficult to find the cause of such a failure. This situation is why just before the SQL0901 message is issued, DB2 for i creates a record for this incident at the QRECOVERY.QSQ901S table. A unique instance of the job log logs three SQL0901 failures. Any subsequent failures for the job are not logged because in most cases they are uninteresting and do not provide more detail. It is possible to *disable* this logging by setting the environment value QIBM_NO_901_LOGGING.

The QRECOVERY.QSQ901S table has the following definitions:

SERVERNAME - VARCHAR(18)	The server name.
FAILTIME - TIMESTAMP	The time when the failure occurred.
FAILRSN - INTEGER(9)	The unique failure reason that appeared in the SQL0901 message. This reason code is necessary for IBM service and is not documented externally.
CURUSER - VARCHAR(18)	The user who encountered the SQL0901 failure.
JOBNAME - CHAR(28)	The qualified job name, which encountered the SQL0901 failure.
MSGS - VARCHAR(3000)	N/A.
APPLIB - CHAR(10)	The library name of the application.
APPNAME - CHAR(10)	The name of the application.
APPTYPE - CHAR(10)	The type of the application.

The values of APPTYPE are as follows:

- ▶ ***PGM:** APPNAME is a program.
- ▶ ***SRVPGM:** APPNAME is a service program.
- ▶ ***SQLPKG:** APPNAME is an SQL package.
- ▶ ***PROCESS:** APPNAME refers to a process (job).
- ▶ **DYNAMIC:** There is no application name.

There are other functions to help you identify the cause of the failure:

- ▶ `SELECT * FROM qsys2.syscolumns WHERE TABLE_SCHEMA = 'QRECOVERY' and TABLE_NAME = 'QSQ901S' ORDER BY ORDINAL_POSITION;`
- ▶ `SELECT * FROM qsys2.syscolumns2 WHERE TABLE_SCHEMA = 'QRECOVERY' and TABLE_NAME = 'QSQ901S' ORDER BY ORDINAL_POSITION;`

These functions help you get information about the contents of the QRECOVERY.QSQ901S table.

The records in the QRECOVERY.QSQ901S table likely show the internal failures inside DB2 for i and data from this table should be used when you report a problem to IBM, which helps with searching for PTFs for DB2 for i problems.

The SQL0901 logging file can be found in a different library when you use IASPs: QRCYnnnnn/QSQ901S *FILE, where nnnnn is the iASP number.

6.4.21 Retrieving a short name for long schema name (QDBRTVSN)

This enhancement provides a non-SQL based approach to retrieving the library name for an SQL schema.

To retrieve the library name, specify the SQL schema name for the input long object name and blank for the input library name. The first 10 bytes of the output qualified object name contains the short library name and the second 10 bytes are QSYS.

The SQL-based approach is to query the QSYS2/SYSSCHEMAS catalog or start the OBJECT_STATISTICS UDTF. The UDTF performs better, so use it when you know the library or schema name.

In Example 6-44, a schema is created and the methods of accessing the long name with a short name and the short name with a long name using SQL by querying the QSYS2.OBJECT_STATISTICS catalog is shown.

Example 6-44 Using QSYS2.OBJECT_STATISTICS

```
CREATE SCHEMA DANS_LONG_SCHEMA
```

To get the short name given the long name

```
SELECT OBJNAME FROM TABLE(QSYS2.OBJECT_STATISTICS('DANS_LONG_SCHEMA','LIB')) AS A
```

To get the long name of a schema given a short name

```
SELECT OBJLONGNAME FROM TABLE(QSYS2.OBJECT_STATISTICS('DANS_00001','LIB')) AS A
```

Using the SYSSCHEMAS view touches every library object. The following queries are identical:

- ▶

```
SELECT OBJLONGNAME FROM TABLE(QSYS2.OBJECT_STATISTICS('QSYS      ','LIB')) AS A WHERE OBJNAME LIKE 'CATH%';
```
- ▶

```
SELECT SCHEMA_NAME FROM QSYS2.SYSSCHEMAS WHERE SYSTEM_SCHEMA_NAME LIKE 'CATH%';
```

6.4.22 SQL0440 warning on SQL routine build

When a **LANGUAGE SQL** procedure includes an incorrectly coded static **CALL** statement, the build succeeds, but returns a warning:

```
SQLCODE = SQL0440 (+440)  
SQLSTATE = '0168L'
```

Consider the following two **CREATE PROCEDURE** statements:

- ▶

```
CREATE PRODCEDURE ProcPRG1(INOUT parm1 varchar(255)) LANGUAGE SQL PROGRAM TYPE  
SUB BEGIN RETURN 0; END;
```
- ▶

```
CREATE PRODCEDURE ProcPRG2()LANGUAGE SQL PROGRAM TYPE SUB BEGIN CALL ProcPRG1();  
/* incorrect number of parameters?*/ END;
```

The message SQL0440 as in Example 6-45 is generated when you run the previous CREATE PROCEDURE statements.

Example 6-45 SQL0440 message

Message: [SQL0440] Routine PROCPRG1 in not found with specified parameters.
Cause.....: A function or procedure with the specified name and compatible arguments was not found. Recovery: Specify the correct number and type of parameters on the CALL statement or function invocation. Try the request again.

6.4.23 XMLTABLE

XMLTABLE is a table function that evaluates an XQuery expression and returns the result as a relational table.

The **XMLTABLE** built-in table function can be used to retrieve the contents of an XML document as a result set that can be referenced in SQL queries.

The addition of **XMLTABLE** support to DB2 for i users makes it easier for data centers to balance and extract value from a hybrid data model where XML data and relational data coexist.

In Example 6-46, the row-generating expression is the XPath expression \$d/dept/employee. The passing clause refers to the XML column doc of the table emp.

Example 6-46 XMLTABLE example

For example, if the EMP table has an XML column called DOC:

```
SELECT X.*
FROM emp, XMLTABLE ('$d/dept/employee' passing doc as "d"
COLUMNS empID      INTEGER          PATH '@id',
firstname VARCHAR(20)  PATH 'name/first',
lastname  VARCHAR(25)  PATH 'name/last') AS X
```

The output is shown in Figure 6-20.

An XML value in DOC:

<dept bldg="101">
 <employee id="901">
 <name>
 <first>John</first>
 <last>Doe</last>
 </name>
 <office>344</office>
 <salary currency="USD">55000</salary>
 </employee>
 <employee id="902">
 <name>
 <first>Peter</first>
 <last>Pan</last>
 </name>
 <office>216</office>
 <phone>905-416-5004</phone>
 </employee>
</dept>

Results in:

EMPID	FIRSTNAME	LASTNAME
901	John	Doe
902	Peter	Pan

Figure 6-20 XMLTABLE Query output example

The output in Figure 6-21 shows an XML document along with a sample SQL to produce the output in table format.




Figure 6-21 Query the total cost of all items that are purchased on each receipt

The SQL/XML support that is included in DB2 for i 7.1 offers a built-in and standardized solution for working with XML from within SQL. **XMLTABLE** adds significant value by allowing SQL queries to query both XML and relational data in the same query. Working XML data within a relational model is often times a non-trivial task. **XMLTABLE** offers a high degree of flexibility by supporting a wide range of XPath step expressions, predicates, and built in functions that can be used within the row and column generating expressions.

6.4.24 Run SQL (RUNSQL): A new command

The Run SQL (RUNSQL) command runs a single SQL statement without having to construct a source physical file member or write a program. No spool file is generated. Upon failure, the specific SQL failure message is returned to the command caller.

In other SQL interfaces, an SQL statement is limited to 2 MB in length. The limit on this command is 5000 bytes.

The command has many parameters that are similar to the ones used with the **RUNSQLSTM** command. **RUNSQL** runs SQL statements in the invoker's activation group. If **RUNSQL** is included in a compiled CL program, the activation group of the program is used.

Two examples of using **RUNSQL** on the command line are shown in Example 6-47. The example also shows how **RUNSQL** can be used within a CL program.

Example 6-47 Usage of RUNSQL from the command line

```
RUNSQL SQL('INSERT INTO prodLib/work_table VALUES(1, CURRENT TIMESTAMP)')

/* In a CL program, use the Receive File (RCVF) command to read the results of the
query */
RUNSQL SQL('CREATE TABLE QTEMP.WorkTable1 AS
(SELECT * FROM qsys2.systables WHERE table_schema = ''QSYS2'') WITH DATA')
COMMIT(*NONE) NAMING(*SQL)
```

Example 6-48 shows how **RUNSQL** can be used within a CL program. Use the **RCVF** command if you must read the results of the query.

Example 6-48 Usage of RUNSQL within a CL program

```
RUNSQL1: PGM  PARM(&LIB)
              DCL  &LIB TYPE(*CHAR) LEN(10)
              DCL  &SQLSTMT TYPE(*CHAR) LEN(1000)
              CHGVAR  VAR(&SQLSTMT) +
                  VALUE('DELETE FROM QTEMP.WorkTable1 +
                        WHERE table_schema = '' ' || &LIB || '' ' ')
              RUNSQL SQL(&SQLSTMT) COMMIT(*NONE) NAMING(*SQL)
ENDSQL1: ENDPGM
```

6.4.25 Native JDBC driver support added for J2SE 7

With the availability of J2SE 7, the native JDBC driver is enhanced to meet the new JDBC 4.1 specification.

The new support for native JDBC on IBM i 7.1 includes:

- ▶ `java.sql.CallableStatement.getObject(int parameterIndex,java.lang.Class<T>type)`
- ▶ `java.sql.CallableStatement.getObject(java.lang.String parameterName,java.lang.Class<T>type)`
- ▶ `java.sql.Connection.abort(java.util.concurrent.Executor executor)`
- ▶ `java.sql.Connection.setSchema((java.lang.String schema)`
- ▶ `java.sql.Connection.getSchema()`
- ▶ `java.sql.DatabaseMetaData.generatedKeyAlwaysReturned()`
- ▶ `java.sql.DatabaseMetaData.getPseudoColumns(java.lang.String catalog,java.lang.String schemaPattern,`
- ▶ `java.lang.String tableNamePattern,java.lang.String columnNamePattern)`
- ▶ `java.sql.Statement.closeOnCompletion()`
- ▶ `java.sql.Statement.isCloseOnCompletion()`

Because the native JDBC driver does not typically use a network connection, the `Connection.getNetworkTimeout()` and `Connection.setNetworkTimeout()` methods are not implemented.

6.4.26 Adding the ORDERBY parameter to the CPYTOIMPF command

The new (optional) **ORDERBY** parameter specifies the order of the records that the records are inserted in the file.

Here are the **ORDERBY** parameter values:

- ▶ ***NONE**: No specific order requested. This value is the default value for the command.
- ▶ ***ARRIVAL**: The records are inserted in the order they were inserted in to the file.
- ▶ **character-value**: Specifies an SQL **ORDER BY** clause that is used for ordering the records in to the file.

The usage of the **ORDER BY(*ARRIVAL)** parameter and **ORDER BY(character-value)** examples are shown in Example 6-49, where the Stream File is copied and ordered according to the **ORDER BY** parameter value.

Example 6-49 Usage of the ORDERBY clause

```
CPYTOIMPF FROMFILE(CORPDB1/DEPARTMENT) TOSTMF('/iddtdept.file') RCDDL(*LF)
ORDERBY('*ARRIVAL')
```

```
CPYTOIMPF FROMFILE(CORPDB1/DEPARTMENT) TOSTMF('/iddtdept.file') RCDDL(*LF)
ORDERBY(' DEPTNO ASC FETCH FIRST 5 ROWS ONLY ')
```

In addition to the **ORDER BY** parameter, you can use the following parameters:

- ▶ **FETCH FIRST *n* ROWS**
- ▶ **OPTIMIZE FOR *n* ROWS**
- ▶ **FOR UPDATE**
- ▶ **FOR READ ONLY**
- ▶ **WITH <isolation-level>**
- ▶ **SKIP LOCKED DATA**
- ▶ **USE CURRENTLY COMMITED**
- ▶ **WAIT FOR OUTCOME**

6.4.27 System naming convention expanded for permit (/) and (.) qualifiers

Historically, system naming qualification of names required a slash(/) only between the schema-name and the SQL identifier.

Now, when you use system naming, both the slash(/) and dot(.) can be used for object qualification. This change makes it easier to adapt to system naming, as the SQL statement text does not need to be updated.

This enhancement makes it easier to use system naming. Object references can vary and SQL UDFs can now be library qualified with a “.” when you use system naming.

NAMING(*SYS) can be used with (/) and (.), as shown in Example 6-50. However, if the example were used with **NAMING(*SQL)**, it would fail.

*Example 6-50 Using NAMING(*SYS)*

```
SELECT a.ibmreqd, b.ibmreqd FROM sysibm.sysdummy1 a, sysibm/sysdummy1 b
```

6.4.28 Modification of global variables within triggers and functions

SQL global variables are a useful construct to achieve session-specific inter-application communication.

Before this change, triggers and functions were allowed to only *reference* global variables. SQLCODE = -20430 and SQLSTATE = “428GX” were returned to a trigger or function that attempted to *modify* a global variable.

Global variables are modified as shown in Example 6-51, which shows the modification of a global variable.

Example 6-51 Usage of a global variable modification

```
CREATE VARIABLE PRODLIB.MOST_RECENT_ORDER BIGINT DEFAULT 0

CREATE TRIGGER PRODLIB.INSERT_ORDER
BEFORE INSERT ON PRODLIB.ORDER_TABLE
REFERENCING NEW AS N
FOR EACH ROW
MODE DB2ROW
IE:BEGIN ATOMIC
SET PRODLIB.MOST_RECENT_ORDER = NEW_ORDER_VALUE();
SET N.ORDER_VALUE = PRODLIB.MOST_RECENT_ORDER;
END IE
```

6.4.29 New QAQQINI option: SQL_GVAR_BUILD_RULE

This option influences whether global variables must exist when you build SQL procedures, functions, triggers, or running SQL precompiles. Its supported values are:

- ▶ ***DEFAULT:** The default value is set to ***DEFER**.
- ▶ ***DEFER:** Global variables do not need to exist when an SQL routine is created or the SQL precompiler is run. Because global variables are not required to exist, the create does not fail when an incorrect column name or routine variable is encountered because it is not known at create time whether the name is a global variable. Incorrect name usage results in SQL0206 - “Column or global variable &1 not found.” failures when the statement is run.
- ▶ ***EXIST:** Global variables that are referenced by SQL must exist when the SQL routine is created or the SQL precompiler is run. Using this option, an SQL0206 is issued at create time. This choice is similar to the behavior of previous releases, where incorrect name usage will fail at create time with either SQL0312 - “Variable &1 not defined or not usable.” or SQL5001 - “Column qualifier or table &2 undefined”

6.4.30 CPYTOIMPF and CPYFRMIMPF commands include column headings

The **CPYTOIMPF** and **CPYFRMIMPF** commands are enhanced with the following features:

- ▶ **Copy To Import File (CPYTOIMPF):** A new (optional) parameter indicates whether the column names should be included as the first row in the target file.
- ▶ **ADDCOLNAM (*NONE or *SYS or *SQL):**
 - ***NONE:** This option is the default. When it is used, the column names are not included.
 - ***SQL:** The column SQL names are used to populate the first row.
 - ***SYS:** The column System names are used to populate the first row.
- ▶ **Copy From Import File (CPYFRMIMPF):** A new (optional) parameter indicates whether the first row should be skipped when you process the import file.
- ▶ **RMVCOLNAM (*NO or *YES):**
 - ***NO:** This option is the default. When it is used, the first row is not skipped.
 - ***YES:** Remove the column names from the file that is being imported by skipping the first row of data.

6.4.31 Improved performance of the privileges catalogs and the ODBC/JDBC metadata APIs


The performance is improved for the following items.

- ▶ **JDBC:** DatabaseMetaData.getColumnPrivileges() and DatabaseMetaData.getTablePrivileges()
- ▶ **SQL CLI and ODBC:** SQLColumnPrivileges() and SQLTablePrivileges()
- ▶ **DB2 for i Catalogs in SYSIBM:** Variable_Privileges, Usage_Privileges, Routine_Privileges, UDT_Privileges, SQLTablePrivileges, and SQLColPrivileges

The SQL in Example 6-52 produces the output that is shown in Figure 6-22, which clearly shows the privileges by GRANTEE over the selected schema and table.

Example 6-52 Who is allowed to modify this table

```
SELECT GRANTEE, PRIVILEGE FROM SYSIBM.SQLTablePrivileges
WHERE TABLE_SCHEM = 'PRODLIB' AND
TABLE_NAME = 'DEPARTMENT' AND
PRIVILEGE IN ('UPDATE','INSERT','DELETE')
ORDER BY GRANTEE
```



GRANTEE	PRIVILEGE
HRTEAM	UPDATE
SCOTT	INSERT
SCOTT	UPDATE
SCOTT	DELETE
SLFUSER	INSERT
SLFUSER	UPDATE
SLFUSER	DELETE

Figure 6-22 List of users with their relevant privileges

6.4.32 Show statements: Index creates and statistics advised

These two new columns appear when show statements are used against a plan cache snapshot or SQL performance monitor. These two new columns do not require a System i Navigator update. The two new columns can be added to your view by clicking **Columns** in the SQL Plan Cache Snapshots window, as shown in Figure 6-23.




Figure 6-23 Two new columns available for show statements

6.4.33 The RUNSQLSTM command and the OPTION parameter

The **OPTION** parameter of **RUNSQLSTM** provides direct control over the command output file. There is no unnecessary spool file accumulation.

The values for the **OPTION** parameter are:

- ▶ ***LIST** (default): A complete listing with the contents of the source and all errors is generated.
- ▶ ***NOSRC**: A listing with only the messages is generated.
- ▶ ***ERRLIST**: A complete listing is generated only if errors greater than the **ERRLVL** parameter are encountered during processing.
- ▶ ***NOLIST**: No listing is generated. If errors greater than the **ERRLVL** parameter are encountered during processing, the messages are sent to the job log.

Running the statement as shown in Example 6-53 on page 230 obtains a complete listing only if errors greater than the **ERRLVL** parameter are encountered during processing. An example output of ***ERRLIST** is shown in Figure 6-24 on page 230.

Example 6-53 Obtain a complete listing

```
RUNSQLSTM SRCFILE(CLIVEG/QSQLSRC) SRCMBR(STA1) OPTION(*ERRLIST)
```

```
Record *...+... 1 ...+... 2 ...+... 3 ...+... 4 ...+... 5 ...+... 6 ...+... 7
      1 SELECT * FROM GRIFFIC/CUSTOMER;
                                * * * * * E N D   O F   S O U R C E * * * * *
5770SS1 V7R1M0 100423              Run SQL Statements              STA1
Record *...+... 1 ...+... 2 ...+... 3 ...+... 4 ...+... 5 ...+... 6 ...+... 7
MSG ID  SEV  RECORD  TEXT
SQL0084  30          1  Position 1 SQL statement not allowed.
                        Message Summary
      Total      Info  Warning      Error  Severe  Terminal
        1         0        0         0        1         0
30 level severity errors found in source
                                * * * * * E N D   O F   L I S T I N G * * * * *
```

Figure 6-24 *ERRLIST output highlighting error details

6.4.34 Improved NULLID package management

NULLID package management automates the creation of necessary objects (collections and dynamic packages) to more seamlessly allow applications using DB2 family database connectivity methods to access data in DB2 for i.

DB2 for Linux, UNIX, and Windows, IBM DB2 Universal Driver for SQLJ and JDBC, ODBC, CLI, IBM DB2 Connect™, and other application requesters rely upon SQL package (*SQLPKG) objects within the NULLID collection.

The NULLID enhancements include:

- ▶ The NULLID collection is created automatically upon the first attempt by a DRDA client to bind or run using a package within NULLID.
- ▶ The ownership of the NULLID collection is assigned to the QSYS user profile and *PUBLIC is granted *CHANGE authority by default when the system automatically creates the NULLID collection.
- ▶ DB2 for i also automatically creates the dynamic packages that are used by JDBC, ODBC, and CLI. These packages are named SYSzcxy, SYSzcxyy, and SYSzTAT.

Before you use these clients to access data on IBM i, you must create IBM i SQL packages for these application programs.

6.4.35 Java stored procedures and functions: System naming option

The Native JDBC driver is enhanced to include a new method to allow the specification of system naming.

The method can be used to turn system naming mode ON and OFF. To use this method, an appropriate cast of the java.sql.Connection object must be made to the com.ibm.db2.jdbc.DB2Connection class.

Before you exit the method, the method should return the connection to the default value. Failure to do so may cause unexpected behavior in other Java stored procedure and Java user-defined functions.

Example 6-54 shows an example of how system naming is enabled in a Java stored procedure.

Example 6-54 How system naming could be enabled in a Java stored procedure

Parameter style DB2GENERAL:

```
DB2Connection connection = (DB2Connection) getConnection();
connection.setUseSystemNaming(true);
....
.... do work using the connection
....
connection.setUseSystemNaming(false);
```

Parameter style JAVA:

```
DB2Connection connection = (DB2Connection)
DriverManager.getConnection("jdbc:default:connection");
connection.setUseSystemNaming(true);
....
.... do work using the connection
....
connection.setUseSystemNaming(false);
```

Before you exit the method, the method should return the connection to the default value. Failure to do so might cause unexpected behavior in other Java stored procedure and Java user-defined functions

Use ***LIBL** (Library List) in Java procedures and functions.

6.4.36 Named arguments and defaults for parameters: Procedures

There are some enhancements to the named arguments and defaults for parameters on procedures:

- ▶ There are named and default parameters that are supported for SQL and external procedures that make it possible to run more SQL applications on IBM i.
- ▶ You can extend procedures without fear of breaking existing callers and simplify the calling requirements by adding default values.
- ▶ Procedure call invocation has the same type of support as CL commands.

With this enhancement:

1. Parameters may be omitted if the routine was defined with a default value.
2. Parameters may be specified in any order by specifying the name in the call.
3. This enhancement works with **LANGUAGE SQL** and **EXTERNAL** procedures.

Usage of the named arguments and default parameters is shown in Example 6-55. Figure 6-25 shows the **DEFAULT** clause and its options.

Example 6-55 Parameter usage

```
CREATE PROCEDURE p1 (i1 INT, i2 INT DEFAULT 0, i3 INT DEFAULT -1)...
CALL p1(55)
CALL p1(55, i3=>33)
```




Figure 6-25 The default clause

6.4.37 Improved catalog management for procedures and functions

When an SQL or external procedure or function is created, the routine information is stored within the *PGM or *SRVPGM. Previously, when Librarian commands were used to copy, move, or rename the object, the QSYS2/SYSROUTINE, SYSPARMS, and SYSRTNDEP catalogs were left unchanged.

The following commands (and their API counterparts) were changed to keep the catalogs in sync with the executable object for procedures and functions:

- ▶ Create Duplicate Object (**CRTDUPOBJ**): The routine catalog information is duplicated and the SYSROUTINE EXTERNAL_NAME column points to the *newly duplicated* executable object.
- ▶ Copy Library (**CPYLIB**): The routine catalog information is duplicated and the SYSROUTINE EXTERNAL_NAME column points to the newly duplicated executable object.
- ▶ Rename Object (**RNMOBJ**): The routine catalog information is modified with the SYSROUTINE EXTERNAL_NAME column, which points to the *renamed* executable object.
- ▶ Move Object (**MOVOBJ**): The routine catalog information is modified with the SYSROUTINE EXTERNAL_NAME column, which points to the *moved* executable object

There is coverage for Librarian APIs or other operations that are built upon these commands.

The changed behavior can be partially disabled by adding an environment variable. If this environment variable exists, Move Object and Rename Object operations do *not* update the catalogs. The environment variable has no effect on the **CPYLIB** and **CRTDUPOBJ** commands.

Setting the environment variable is shown in Example 6-56.

Example 6-56 Setting the environment variable to partially disable the function

```
ADDENVVAR
ENVVAR(QIBM_SQL_NO_CATALOG_UPDATE)
LEVEL(*SYS)
```

6.4.38 SQE enhancement for Encoded Vector Indexes defined with INCLUDE

Taking advantage of the IBM patented EVI indexing technology, and analytics routines using SQL Constructs such as ROLLUP, CUBE, and Grouping Sets, can produce performance benefits.

SQE is enhanced to make Grouping Set queries aware of EVI **INCLUDE** as an optimization possibility.

Defining an EVI with **INCLUDE** is shown in the Example 6-57, which also shows an example SQL to query the EVI.

Example 6-57 Encoded Vector Index (EVI) INCLUDE example

```
CREATE ENCODED VECTOR INDEX GS_EVI01 ON STAR1G.ITEM_FACT
  (YEAR ASC, QUARTER ASC, MONTH ASC)
INCLUDE (SUM(QUANTITY),
          SUM(REVENUE_WO_TAX),
          COUNT(*) )

SELECT YEAR, QUARTER, MONTH, SUM(QUANTITY) AS TOTQUANTITY, SUM(REVENUE_WO_TAX)
FROM ITEM_FACT WHERE YEAR=2012
GROUP BY GROUPING SETS ((YEAR, QUARTER, MONTH),
                        (YEAR,QUARTER),
                        (YEAR),
                        ( ))
```

Fast Index-Only access is possible for CUBE(), ROLLUP(), and GROUPING SETS().

6.4.39 Navigator: System Name column added to show related and all objects

When you use System i Navigator to observe the objects that are related to a table, only the SQL names of those objects appear. The same situation is true when you use the All Objects view under the Schemas folder.

The System i Navigator 7.1 client is enhanced to include the “System Name” column. This information is useful when long names are used and the system name is not obvious, as shown in Figure 6-26.




Figure 6-26 System name column included

6.4.40 Improved performance for joins over partitioned tables

Before this introduction of this enhancement, data skew (uneven distribution of data values) was not factored in when creating access plans for queries that joined columns in a partitioned table.

This enhancement enables the query statistics engine to detect data skew in the partitioned table if accurate column statistics are available for the partitions. Allowing column level statistics to be automatically collected is the default behavior and a preferred practice.

This enhancement might improve the performance of joins when multiple partitions are queried.

An example of multiple partition joins is shown in Example 6-58.

Example 6-58 Multiple partitions join example

Assume the ITEM table (A) contains the product list and the ORDERS table (B) contains sales order details, partitioned by some date range.

Several products are particularly heavy sellers, making them skewed in the sales order list.

```
SELECT B.ITEM NO, SUM(B.QUANTITY) AS Total_Sold
FROM ITEM A LEFT OUTER JOIN ORDERS B ON
    A.ITEMNO=B.ITEMNO
GROUP BY B.ITEMNO
ORDER BY TOTAL_SOLD DESC
```

6.4.41 Navigator: Table list totals (members, data size, and valid and deleted rows)

New columns under the Schemas folder now show an aggregate set of information for all members and partitions.

You must add these columns by clicking **View** → **Customize this View** → **Columns**. Add the required columns to your view.

The new columns are:

- ▶ Number of rows
- ▶ Number of deleted rows
- ▶ Partitions
- ▶ Size

An example of how to configure your client to show these new columns is shown in Figure 6-27.




Figure 6-27 Configure your client to show the new columns for tables under the Schemas folder

6.5 DB2 database management and recovery enhancements

This section describes the database management and recovery enhancements to DB2 for i.

6.5.1 Preserving the SQL Plan Cache size across IPLs

DB2 for i on IBM i 6.1 and 7.1 are enhanced to preserve the SQL Plan Cache size across IPLs and slip installations of the operating system.

A scratch installation of the OS resets the SQL Plan Cache size to the default size.

You can explicitly increase the size of the SQL Plan Cache to allow more plans to be saved in the plan cache. This action can improve performance for customers that have many unique queries.

During an IPL, the SQL Plan Cache is deleted and re-created. Before this enhancement, when the plan cache was re-created, it was re-created with the default size of 512 MB, even if you explicitly specified a larger plan cache. Now, the size that is specified by you is preserved and used when it is re-created during any subsequent IPL.

After the latest DB2 Group PTFs are installed, you must change the plan cache size one more time (even if it is changed to the same size as its current size) for the size to be persistently saved.

This **CHANGE_PLAN_CACHE_SIZE** procedure can be used to change the size of the plan cache. The procedure accepts a single input parameter, that is, the wanted SQL Plan Cache size in megabytes. If the value passed in is zero, the plan cache is reset to its default value. To use the procedure, run the command that is shown in Example 6-59.

Example 6-59 Usage of QSYS2.CHANGE_PLAN_USAGE procedure

```
CALL qsys2.change_plan_cache_size(1024);
```

It is also possible to get information about Plan Cache properties by using the procedure shown in Example 6-60.

Example 6-60 QSYS2.DUMP_PLAN_CACHE_PROPERTIES

```
call qsys2.dump_plan_cache_properties('QGPL', 'SQECACHE1');
```

```
select heading,value from qgpl.sqecache1  
where heading like '%Time%' ;
```

6.5.2 Prechecking the physical file size during restore

Before this enhancement, when you performed a restore of a large physical file, if the available free space on the target ASP was not sufficient to contain the physical file, the restore used all the available storage and caused a system crash. A precheck is now performed to ensure that enough storage exists on the target ASP.

6.5.3 Preventing index rebuild on cancel during catch up

When a delayed maintenance index is not open, any changes to rows are recorded, but the update of the index binary tree is delayed until the index is next opened. This action improves the performance of row change operations when the index maintenance is delayed.

Before this enhancement, if a job was canceled while it was opening a delayed maintenance index, the entire index was invalidated and had to be rebuilt from scratch. On large indexes, this operation can be a lengthy operation. This enhancement ensures that in this case that the cancel does not cause the entire index to be invalidated.

6.5.4 Adding the QSYS2/SYSDISKSTAT catalog

The QSYS2/SYSDISKSTAT catalog is added to DB2 for i.

The QSYS2.SYSDISKSTAT catalog can be used to quickly and easily understand the DISK information using SQL.

If you query the UNIT_TYPE field, you can identify information about installed SSD media, as shown in Example 6-61, which shows the relevant SQL to query information for all disks or just for the SSDs.

Example 6-61 Usage of the QSYS2/DISKSTAT catalog for analysis of disk usage

Query information for all disks.

```
SELECT * FROM QSYS2.SYSDISKSTAT
```

Query information for all SSD units

```
SELECT * FROM QSYS2.SYSDISKSTAT WHERE unit_type = 1
```

6.5.5 STRDBMON: FETCH statement shows failures and warnings

The database monitor accumulates information about fetch requests. The fetch monitor record is written after the cursor is closed. This enhancement tracks the first warning or failure observed during the fetch requests and uses that information about the monitor fetch record.

Here the columns from the monitor that include this information:

- ▶ QQI8 (SQLCODE)
- ▶ QQC81 (SQLSTATE)

This enhancement enables the use of the **STRDBMON FTRSQLCODE()** pre-filter as a technique to isolate application fetch-time failures or warnings.

6.5.6 STRDBMON: QQI2 result rows for more statements

Database monitor and SQL performance monitor users can see the number of result rows for DML statements by examining the QQI2 column, where QQRID=1000 (INSERT/UPDATE/DELETE/MERGE).

The monitor support is enhanced to return result rows information for other statements.

Here are some considerations about the number of result rows that are returned for QQI2:

- ▶ SQL DML statements (**INSERT**, **UPDATE**, **DELETE**, or **MERGE**) show the total number of rows that are changed.
- ▶ **CREATE TABLE AS** and **DECLARE GLOBAL TEMPORARY TABLE** with the **WITH DATA** parameter shows the total number of rows that are inserted into the table.
- ▶ Any other query-based statement shows the estimated number of rows for the resulting query.
- ▶ All remaining statements show either -1 or 0.

6.5.7 Adding result set information to QUSRJOBI() and Systems i Navigator

Two new fields are returned by **QUSRJOBI()** and SQL Details for jobs:

- ▶ Available result sets: Shows the number of result sets that are currently available to the job.
- ▶ Unconsumed result sets: Shows the cumulative number of result sets in the job that were not consumed by the calling program and were discarded by the system.

The two new result sets returned by **QUSRJOBI()** are displayed in IBM i Navigator within the SQL details for jobs, as shown in Figure 6-28.




Figure 6-28 Two new result sets returned by **QUSRJOBI()** displayed in SQL details for jobs

6.5.8 STRDBMON pre-filtering of QUERY/400 command usage

Query/400 commands Run Query (**RUNQRY**), Work With Queries (**WRKQRY**) and Start Query (**STRQRY**) generate database monitor records. This enhancement can easily identify Query/400 users and the queries run through the Filter by Client Program database monitor pre-filter.

STRQRY and **WRKQRY** activity looks similar to the following output:

```
QQC21  QQ1000
QR     QUERY/400  LIBRARY/FILENAME.MBRNAME
QR     QUERY/400  LIBRARY/FILENAME.MBRNAME
```

RUNQRY activity looks similar to the following output:

```
QQC21  QQ1000
QR     RUNQRY  CLIVEG/QUERY1
QR     RUNQRY  CLIVEG/QUERY2
```

Example 6-62 shows an example of how to perform the filter by client program pre-filter.

*Example 6-62 Examples of using **STRDBMON** with the filter by client to identify **QUERY/400** usage*

IBM i 6.1 Example:

```
STRDBMON  OUTFILE(LIBAARON/QR400mon)
JOB(*ALL)
COMMENT('FTRCLTPGM(RUNQRY)')
```

IBM i 7.1 Example:

```
STRDBMON  OUTFILE(LIBAARON/QR400mon)
```

6.5.9 UNIT SSD supported on DECLARE GLOBAL TEMPORARY TABLE

Support was added for the media preference to the **DECLARE GLOBAL TEMPORARY TABLE** statement. You can place the temporary table directly on to the SSD. The default media preference setting remains UNIT ANY.

The SQL shown in Example 6-63 places the temporary table on the SSD. The existence of this table can be confirmed by using the SQL in Example 6-64. The actual preferences for the table can be identified by the SQL in Example 6-65.

Example 6-63 Placing a global temporary table on to an SSD

```
DECLARE GLOBAL TEMPORARY TABLE SESSION.TEMP_DAN
  (EMPNO CHAR(6) NOT NULL,
   SALARY DECIMAL(9, 2),
   BONUS DECIMAL(9, 2),
   COMM DECIMAL(9, 2))
ON COMMIT PRESERVE ROWS
WITH REPLACE
UNIT SSD
```

The QSYS2.PARTITION_DISKS table can be queried to determine which media the queried table is on.

Example 6-64 Determine whether the table is on the SSD

```
SELECT UNIT_TYPE FROM TABLE(QSYS2.PARTITION_DISKS('QTEMP ','TEMP_DAN ')) AS X

-- indicates the type of disk unit
-- 0 - Not Solid State Disk (SSD)
-- 1 - Solid State Disk (SSD)
```

The QSYS2.PARTITION_STATISTICS table can be queried to determine the media preference of the queried table.

Example 6-65 Determine the media-preference setting for the table

```
SELECT MEDIA_PREFERENCE FROM
TABLE(QSYS2.PARTITION_STATISTICS('QTEMP ','TEMP_DAN ')) AS X

--0 - No media preference
--255 - The table should be allocated on Solid State Disk(SSD), if possible.
```

6.5.10 Adding Maintained Temporary Indexes that are used for a source of statistics within the index advisor and index advisor condenser

The SQL Query Engine is enhanced to use Maintained Temporary Indexes (MTIs) as a source for statistics, which improves the performance of some queries.

The Index Advisor is enhanced to externalize the count (number of times) the MTI is used as a source of statistics and the time stamp of the last time the MTI was used as a source of statistics. This enhancement provides more information that you can use to decide whether to create permanent indexes for MTIs.

Table 6-4 shows the new column definitions for the MTIs.

Table 6-4 MTI column definitions

Column name	Data type	Description
MTI_USED_FOR_STATS	BIGINT	The number of times that this specific MTI is used by the optimizer for statistics.
LAST_MTI_USED_FOR_STATS	TIMESTAMP	The time stamp that represents the last time that this specific MTI was used by the optimizer to obtain statistics for a query.

6.5.11 Adding the QSYS2.REMOVE_PERFORMANCE_MONITOR procedure

DB2 for i Plan Cache snapshots and performance monitors are in physical files and are tracked for System i Navigator users within the QUSRSYS/QAUGDBPMD2 *FILE.

When the **DUMP_PLAN_CACHE** procedure is called to create an SQL Plan Cache snapshot, a procedure interface is needed to later remove the snapshot file object and remove any entries for the file object from the System i Navigator tracking table.

REMOVE_PERFORMANCE_MONITOR() accepts two parameters and removes the SQL table matching that name and associated entries from the QUSRSYS/QAUGDBPMD2 *FILE:

- ▶ **IN FILESCHEMA VARCHAR(10)**
- ▶ **IN FILENAME VARCHAR(10)**

SQL Plan Cache snapshots can be maintained by using the **CALL** statement, as shown in Example 6-66.

Example 6-66 Removing the SQL Plan Cache snapshot and Performance Monitor file

```
CALL QSYS2.DUMP_PLAN_CACHE('CACHELIB','NOV2011')
CALL QSYS2.REMOVE_PERFORMANCE_MONITOR('CACHELIB','MAY2010')
```

6.5.12 STRDBMON: QQI1 fast delete reason code

When the **DELETE** SQL statement is used to remove all records from the table in a performance critical environment, you can now use the database monitor to understand when fast delete was used and why fast delete did not occur.

The QQI1 column contains this reason code when QQRID=1000 and QQC21='DL'.

The QQI1 values are documented in the Database Performance and Optimization book found at <http://pic.dhe.ibm.com/infocenter/iserics/v7r1m0/topic/rzajq/rzajq.pdf> and within the IBM i Technology Updates wiki, which can be found at

<http://www.ibm.com/developerworks/ibmi/techupdates/db2>

In addition to the monitor support, a CPF9898 diagnostic message appears in the job log and QHST, if either:

- ▶ The # of rows in the table is greater than or equal to 100,000.
- ▶ The `SQL_FAST_DELETE_ROW_COUNT` QAQQINI option is set to a non-default value and the # of rows in the table is greater than or equal to that QAQQINI value.

6.5.13 Automatically increasing the SQE Plan Cache size

DB2 for i automatically increases the SQE Plan Cache size unless the plan cache is explicitly set by the user. This enhancement has the following considerations:

- ▶ When processing is initiated to remove plans in the cache because of a size constraint, the efficiency rating of the cache is checked. If the rating is too low, the database automatically increases the plan cache size.
- ▶ The database auto-sizing does not exceed a certain maximum size.
- ▶ The automatic plan cache size decreases the size if the temporary storage on the machine exceeds a certain percentage.
- ▶ The auto-sized value does not survive an IPL. The default plan cache size is used after an IPL and auto sizing begins again.
- ▶ If you explicitly control the plan cache size and want to turn over control of the plan cache size management, set the plan cache size to zero by using the following statement:

```
CALL qsys2.change_plan_cache_size(0);
```

6.5.14 Navigator for i: A new look and no client to manage

You can use the new look for IBM Navigator for i to perform enhanced database management and without client software to manage. Some of the enhanced functions include:

- ▶ A new dynamic navigation area allows for easy navigation to folders and areas in the database.
- ▶ New tab support provides improved object list viewing and interaction.
- ▶ There is improved performance and fewer clicks to reach endpoints.
- ▶ New database capabilities include procedure/function creation/definition and view/edit table contents.

An advanced drill-down using the new look of IBM Navigator for i is shown in Figure 6-29.




Figure 6-29 Enhanced drill-down using the new look of IBM Navigator for i

6.6 DB2 for Web Query for i (5733-QU2, 5733-QU3, and 5733-QU4)

DB2 Web Query for i consists of a foundation product, 5733-QU2, and two more optional products, 5733-QU3, and 5733-QU4. The latest version of these products is Version 1.1.2. IBM introduced a packaging option that includes elements from all three of these products into a new bundle called DB2 Web Query for i Standard Edition.

6.6.1 DB2 Web Query for i (5733-QU2)

This product provides several easy to use, web-based tools for building reports and Business Intelligence applications. Report “authors” can choose from Report and Graph Assistant tools that can web enable Query/400 reports or the new InfoAssist report builder that provides a highly intuitive, drag interface for reporting functions. You can build dashboards, integrate data with new or existing spreadsheets, or choose to deliver reports in an analytical form that you can use to slide and dice through the data interactively. DB2 Web Query for i can either be ordered as modules, or you can choose a Standard Edition for a predefined package of components that provide a robust set of components.

The modules are as follows:

- ▶ **The *base module*** provides the foundation for DB2 Web Query, including the reporting server and the web-based authoring tools. A Query/400 import function allows you to take Query/400 definitions and transform them using the web-based Report Assistant tool. You can use DB2 Web Query to leave your data in DB2 for i and use all security and availability features of the IBM i platform. This base module is priced per processor tier and includes a minimum number of users. Most clients are able to upgrade to the base product at no charge if they own Query/400 and have the latest level of the IBM i software.
- ▶ ***Active Technologies*** can combine the results of queries and create HTML reports that can be made available to users without needing to be connected to the DB2 for i server. Designed for users “on the go”, the reports contain query results, but the data can be viewed in various ways from a browser, including functions to sort, filter the data by different criteria, a calculated field, and chart information for visual impact.
- ▶ The ***DB2 Web Query Developer Workbench*** feature is an open and intuitive environment that you can use for rapid development of more customized web-based reports and metadata. It includes an HTML layout painter for building dashboards, combining multiple reports into a single view. It is a critical component for developing and managing the metadata that is used to shield the complexities of the database from report authors and users.
- ▶ The ***OLAP Module*** provides an interactive visualization of the data that you can use to drill down or slice and dice to find trends or exceptions in an analytical process. A single report can be a starting point for complex data analysis. Setting up dimensional definitions in DB2 Web Query Developer Workbench is a prerequisite to using an OLAP report.
- ▶ Without ***Runtime User Enablement***, each individual user must be licensed to the base product. With the Runtime User Enablement feature, one or more user licenses can now be defined as a group of runtime only users. If you are familiar with Group Profiles, this concept is a similar concept. Each member of the group is able to run reports concurrently, and each group can contain thousands of users, providing an almost unlimited runtime user licensing model. Users that are defined as runtime users cannot create or edit report definitions, but have full functionality in running reports, including parametrized dashboards, OLAP reports, and more.
- ▶ ***DB2 Web Query Spreadsheet Client*** provides enhanced capabilities for users of Microsoft Excel. With the Spreadsheet Client, users can create templates or regularly used spreadsheets that can be repopulated with data from DB2 for i (or Microsoft SQL Server with the following noted adapter feature). Users with the appropriate authority can start the Report Assistant component of the BASE product to build their own query to populate spreadsheet cells. Data computations and totals are brought into Excel as native formulas, and you can add data filtering and style the output to further enhance the data within Excel.
- ▶ ***DB2 Web Query Adapter for Microsoft SQL Server*** provides connectivity from DB2 Web Query to remote SQL server databases. Many IBM i customers have most of their data in DB2 for i, but occasionally want to get real-time access to data in a SQL server database for reporting purposes. The new adapter for SQL server provides access to multiple remote SQL server databases if wanted, and provides seamless access to this data for report authors and users.
- ▶ ***DB2 Web Query Adapter for Oracle JD Edwards*** allows DB2 Web Query to report on data that is stored in World or EnterpriseOne databases within DB2 for i. The adapter provides a level of seamless integration that simplifies authoring of reports. The adapter also preserves data integrity and security during report execution by automatically interfacing to the application’s metadata and security layers.

6.6.2 DB2 Web Query Report Broker (5733-QU3)

This product provides automated report execution and distribution. Use the scheduling facilities to run reports in batch on a daily or weekly basis, on specific dates, or add blackout dates.

Deliver reports in formats such as PDF, spreadsheet, or other PC file formats and automate report distribution through an email distribution list.

6.6.3 DB2 Web Query Software Developer Kit (5733-QU4)

This product is targeted at application developers. The DB2 Web Query SDK provides a set of web services that you can use to integrate DB2 Web Query functions into applications or to customize an interface into DB2 Web Query or DB2 Web Query Report Broker functions.

The web services allow web applications to authenticate users, view domains and folders, determine report parameters, and run DB2 Web Query reports and more. Simplify the programming effort by using the application extension, now part of the SDK. This extension can eliminate the need for programming to the web services and allow you to create a URL interface to report execution that you can embed in an existing or new application.

When you develop using the SDK, the DB2 Web Query BASE product is required and Developer Workbench feature is highly recommended. Deployment (runtime) environments require the BASE product and the Runtime User Enablement feature of DB2 Web Query.

6.6.4 DB2 Web Query for i Standard Edition

DB2 Web Query, Standard Edition simplifies the decision process of which features to order by including the most popular features in a single package. You can order more features, such as the SQL Server or JDE adapter to the Standard Edition if wanted, but the intent is to combine the most commonly chosen functions into a single order. The DB2 Web Query Standard Edition contains:

- ▶ DB2 Web Query for i BASE with the number of users included based on processor group
- ▶ Four more User Licenses (that can be used as individual users or as a group of runtime users)
- ▶ One PC license of Developer Workbench
- ▶ Active Technologies
- ▶ OLAP
- ▶ Runtime User Enablement
- ▶ Spreadsheet Client
- ▶ DB2 Web Query Report Broker
- ▶ DB2 Web Query Software Developer Kit

6.7 OmniFind Text Search Server for DB2 for i (5733-OMF)

The OmniFind Text Search Server for DB2 for i product available for IBM i 7.1 is enhanced to include more SQL programmable interfaces that extend its support beyond traditional DB2 tables. These interfaces allow text indexing and searching of IBM i objects, such as spool files in an output queue or stream file data in the integrated file system.

A text search collection describes one or more sets of system objects that have their associated text data indexed and searched. For example, a collection may contain an object set of all spool files in output queue QUSRSYS/QEZJOBLOG, or an object set for all stream files in the /home/alice/text_data directory.

The text search collection referred to in this documentation should not be confused with a DB2 schema (sometimes also referred to as a collection), or a Lucene collection (part of the internal structure of a DB2 text search index).

When a text search collection is created, several DB2 objects are created on the system:

- ▶ SQL schema with the same name as the collection
- ▶ Catalogs for tracking the collection's configuration
- ▶ Catalogs for tracking the objects that are indexed
- ▶ SQL Stored procedures to administer and search the collection
- ▶ A DB2 text search index for indexing the associated text

Administration of the collection is provided with stored procedures, most of which are created in the schema.

6.7.1 OmniFind for IBM i: Searching Multiple Member source physical files

The OmniFind Text Search Server for DB2 for i product (5733-OMF) for IBM i 7.1 is enhanced to include more SQL programmable interfaces that extend its support beyond traditional DB2 tables.

Multiple Member source physical files are added one at a time to the OmniFind collection. The members from source physical file are retrieved and treated as separate objects.

During the OmniFind update processing, new, changed, or removed members are recognized and processed appropriately.

Two flavors of instrumentation are available:

- ▶ `<collection>.ADD_SRCPF_OBJECT_SET (IN SRCPF_LIB VARCHAR(10) CCSID 1208, IN SRCPF_NAME VARCHAR(10) CCSID 1208, OUT SETID INTEGER) CALL UPDATE`
- ▶ `<collection>.ADD_SRCPF_OBJECT_SET (IN SRCPF_LIB VARCHAR(10) CCSID 1208, IN SRCPF_NAME VARCHAR(10) CCSID 1208) CALL UPDATE`

For more information about this topic and the OmniFind for i product, see the topic "Searching Spool Files and IFS Stream Files" at developerWorks at:

<https://www.ibm.com/developerworks/ibmi/library/i-omnifind/omnifind.html>

You can also find updates for OmniFind at:

<http://www.ibm.com/developerworks/ibmi/techupdates/db2/omf>



Performance tools

The IBM i operating system and its licensed products include various applications for collecting, analyzing, and reporting performance data that are generically called *performance tools* or *performance analysis tooling*. There are basically two components in the product:

- ▶ The collection services that capture data about how the hardware resources of a physical and virtual system are used by the various user and system functions that support the business of a client.
- ▶ The tools for viewing, modeling, aggregating, and analyzing data. Although these tools still support a fixed function terminal (for example, 5250 or a telnet client), most of them are using either a browser-based or a specific client software.

Important: Do not confuse the term *performance tools* with the licensed product 5770-PT1 Performance Tools.

These performance tools actually cover three areas of systems management:

- ▶ Performance and capacity management
- ▶ Diagnostic analysis of processes within a system context
- ▶ Specific diagnostic tests for the IBM support organization

With these functions, you can set up practices for monitoring and managing your system performance to ensure that your IT infrastructure is aligned with the changing demands of your business.

This chapter describes how the Collection Services and the Analysis Tooling changed.

Requirement: To take advantage of all the Performance Tools enhancements that are described in this chapter, the system must have the latest levels of PTFs installed.

This chapter describes the following topics:

- ▶ Introduction to performance tooling
- ▶ Overview of data collection in IBM i 7.1
- ▶ IBM Navigator for i Performance interface
- ▶ IBM iDoctor for IBM i

7.1 Introduction to performance tooling

This chapter starts with a brief review of how the available tools for performance and capacity management are structured.

7.1.1 Gathering data

There are four distinct functions that collect data on IBM i:

- ▶ **Collection Services**

This function provides for the collection of generic system management data. It is the primary collector of system data. You can run this function continuously to know what is happening with your system. Collection Services data is deposited into a management collection object and then converted and put into database files.

The interval data that is collected is specified by either application-defined or user-defined interval data. This Collection Services is part of the IBM i operating system code. This chapter refers to these Collection Services as a separate entity.

- ▶ **IBM i Job Watcher**

This function uses additional instrumentation for the collection of job data for any or all jobs, threads, and tasks on the system. It provides call stacks, SQL statements, objects that are being waited on, JVM statistics, wait statistics, and more, which are used to diagnose job-related performance problems.

- ▶ **IBM i Disk Watcher**

This function provides for the collection of disk performance data to diagnose disk-related performance problems.

- ▶ **Performance Explorer**

This function provides for the collection of detailed data at a program and application level to diagnose problems. It also traces the flow of work in an application and can be used to diagnose difficult performance problems.

Application-defined performance explorer trace points, such as with Domino, NetServer, or WebSphere servers, specify the data that is collected. It is intended to be used as directed by IBM. Performance Explorer data is deposited into a management collection object and then converted and put into database files.

You can use data from all of these collectors and combine it to allow for an in-depth analysis of jobs and processes and how they use system resources.

Operating system functionality: All of the functions that allow the configuration of data collection, to start and end data collection, and to manage the collection objects, are part of the operating system.

7.1.2 Analyzing data

There are several tools and options to view, understand, and analyze data from the collection services. These tools and options are either built into the Performance Tools product, or delivered as an option or a service. This overview excludes any vendor tooling or the analysis you can run using your own developed application:

- ▶ You can use the Display Performance Data GUI to view performance data that is generated by the collection services through a 5250 interface.
- ▶ The Reports function organizes Collection Services performance data and trace data in to a logical and useful format. In this release, there are no changes to this function.
- ▶ You can use the Performance Tools Graphics function to work with performance data in a graphical format. You can display the graphs interactively, or you can print, plot, or save the data to a graphics data format (GDF) file for use by other utilities.
- ▶ You can use the Investigate data function of the IBM Navigator for IBM i to perform deep analyses of the data from within a browser-based interface. For more information, see 18.16, “Performance enhancements” on page 692.
- ▶ System i Navigator provides not only an interface to view and analyze data, but you can use it to define monitoring functions for resource usage by individual, all, or selected jobs.
- ▶ The Job Watcher function of the IBM Navigator for i Performance interface is also included in Performance Tools as an option.
- ▶ The Disk Watcher function of the IBM Navigator for i Performance interface is included in Performance Tools.
- ▶ IBM iDoctor for IBM i provides a separate GUI interface to analyze Collection Services, Job Watcher, Disk Watcher, and Performance Explorer data.
- ▶ IBM Performance Management for Power Systems. The support for the IBM i offering automates the collection, archival, and analysis of system performance data and returns reports to help you manage system resources and capacity.

The Performance Management for Power Systems offering includes the Performance Management Agent (PM Agent). The PM Agent is a function of the operating system that provides automated collection of non-proprietary collection services data, reduces the data, and sends the data to IBM. When you send your data to IBM, you eliminate the need to store all the trending data yourself. IBM stores the data for you and provides you with a series of reports and graphs that show your server's growth and performance. You can access your reports electronically using a traditional browser.

This offering, when used with the IBM Systems Workload Estimator, allows you to understand how your business trends relate to the timing of required hardware upgrades, such as processor or disk. The IBM Systems Workload Estimator can size a systems consolidation or evaluate upgrading a system with logical partitions by having PM Agent send the data for multiple systems or partitions to the IBM Systems Workload Estimator.

7.2 Overview of data collection in IBM i 7.1

This section describes the changes in the collection profiles and in the data that is kept when you collect performance data. These items are adapted to allow an in-depth analysis of the activities on the system to set up the preferred practices to analyze the behavior of applications on a system. In addition, these items provide valuable information about how the workload is built up and how to create a workload management practice to reduce contention and avoid possible conflict situations. This information can also be used to define how to instrument and document capacity management practices of your IT infrastructure.

You can use the IBM i performance tooling to track how the performance and capacity objectives are met in an end-to-end approach. You can use the detailed information that these collection services generate to shape not only the workload management of a system, but also to analyze issues that can be because of many events.

7.2.1 Collection profiles for Collection Services

The existing collection profiles (*MINIMUM, *STANDARD, *STANDARDP and *CUSTOM) stay the same in release 7.1, but there are several categories (except for the *MINIMUM and *ENHPCPLN profiles) added to them:

- ▶ The Standard profile now has the following additional functions:
 - External storage, containing non-standardized data for disk units that are externally attached to an IBM i partition (QAPMXSTGV).
 - System internal data, containing internal data for the system (QAPMSYSINT).
 - Removable storage, containing data about removable storage devices that are connected to the system, more specifically, tape device data.
- ▶ The Standard Plus profile gets information about the logical partition, containing performance data that is collected from eligible partitions if the IBM Director Server (5761-DR1) licensed program is installed on the partition that is running collection services. To collect data from other partitions, the IBM Director Agent (5761-DA1) licensed program must be installed on the other partitions and the server must be authorized to the other partitions.
- ▶ For the Custom profile, already allowing for the customization of the categories, you can now also specify the intervals. This specification means that you can have different categories of data that is collected at different intervals. Figure 7-1 is an example of how you can set up this profile.

Collection profile

☐ Select predefined collection profile
Standard plus protocol

☒ Customize collection profile

Available categories

Select	Category
☐	Memory pool
☐	Memory pool tuning
☐	Hardware configuration
☐	Subsystem
☐	System CPU
☐	System-level data
☐	Jobs (MI tasks and threads)
☐	Jobs (operating system)
☐	SNADS
☐	Disk storage

Page 1 of 3 | 1 | Go | Rows 10 | Total: 29 Selected: 0

Categories to collect

Select	Category	Frequency
☒	System bus	Default interval
☐	IBM HTTP Server (powered by Apache)	

Frequency to collect

☐ 15 seconds
☐ 1 minutes
☒ Use default collection interval

Advanced options

Parameter string:

Figure 7-1 Custom profile settings for the interval

7.2.2 Information repository: Performance tools tables

Performance tools tables are described in the sections that follow.

QAPMBUS

This database file contains data for external system buses.

Support for a particular bus and what metrics are supported for that bus depends on the type of bus, how that bus connects to the system, and whether it is assigned to the partition.

Historically, records were produced for all PCI buses even though data was instrumented only for bus usage within the collecting partition. For that situation now, data is captured only for those buses that have activity within the collecting partition.

For newer technologies, the hardware might provide additional instrumentation. Hardware metrics represent bus usage by all partitions. The collecting partition must be authorized to obtain this data (see the “Allow performance information collection” option within the HMC partition configuration). If the collecting partition is authorized, buses that support hardware metrics are reported independent of partition assignments.

Table 7-1 shows the contents of the QAPMBUS table.

Table 7-1 Contents of the QAPMBUS table

Column name	Description
INTNUM	Interval number: The nth sample database interval that is based on the start time in the Create Performance Data (CRTPFDDTA) command.
DATETIME	Interval date (mmddyy) and time (hhmmss): The date and time of the sample interval.
INTSEC	Elapsed interval seconds: The number of seconds since the last sample interval.
BUIOPB	System bus number: Bus numbering begins with one. Before V5R4, bus numbering began with zero.
BUOPSR	Number of OPSTARTs received: RRCB in server storage.
BUSGLR	Signals received.
BUOPSS	Number of OPSTARTs sent.
BUSGLS	Signals sent.
BURSTQ	Restart queues that are sent.
BUBNAR	Occurrences of BNA received.
BUTPKT	Total packets (sent or received).
BUKBYO	Kilobytes DMAed out.
BUKBYI	Kilobytes DMAed in.
BUNOSR	Normal flow OPSTARTs received.
BUNRDR	Reserved.
BUORQS	OPSTART requests sent.
BUTIMO	Reserved.
BUBNAS	Reserved.

Column name	Description
BUQSAS	Queue space available sent.
BUTYPE	Bus type. Supported values are: ▶ 'S' = SPD bus (no longer supported). ▶ 'P' = PCI bus. ▶ 'V' = Virtual bus.
BUCAT	Bus category. This field indicates if this bus record has some special characteristics, which might require a special interpretation of its performance data. Each bit in this field has an independent meaning: ▶ X'00' = No special category applies. ▶ X'01' = This bus is attached to an I/O hub. (Note: The following might be used depending on what happens with switches.) ▶ X'02' = This record represents a switch. The data that is reported is the sum of all buses under the switch. The bus number that is reported is the first bus under the switch.
BUHUB	Hub number. If this bus is associated with an I/O hub, this number is the number of that hub. (Note: An I/O hub may be embedded in the backplane.)
BUMAXRATE	Maximum byte rate. When available from hardware, this rate is the estimated maximum rate that data might be both sent and received in bytes per second through the hardware port.
BUCBSND	Command bytes sent. When available from hardware, this number is the number of command bytes sent through the hardware port.
BUDBSBD	Data bytes sent. When available from hardware, this number is the number of data bytes sent through the hardware port.
BUCBRCV	Command bytes received. When available from hardware, this number is the number of command bytes received through the hardware port.
BUDBRCV	Data bytes received. When available from hardware, this number is the number of data bytes received through the hardware port.

QAPMBUSINT

This file contains data for internal system buses.

The metrics that are supported depend on the instrumentation within the hardware chips. Support for a particular bus depends on both the type of bus and the chip family.

There might be one or more records for each interval for a reported bus. The number of records and the metrics that are supported depend on both the bus type and chip type.

These metrics are instrumented in the hardware and represent bus usage by all partitions. The collecting partition must be authorized to obtain this data (see "Allow performance information collection" option within the HMC partition configuration).

Table 7-2 shows the contents of the QAPMBUSINT table.

Table 7-2 Contents of the QAPMBUSINT table

Column name	Description
BUNBR	The hardware assigned number that is associated with the bus or hub.
BUTYPE	Bus type. The supported bus types are: ▶ 4 - 12X loop. ▶ 6 - I/O hub (may be embedded in the backplane).

Column name	Description
BUDFMT	Bus data format. This field is provided to help you understand what data is instrumented by the hardware components of the bus.
BUATTR1	Bus attribute 1. The meaning of this field depends on the bus type. One row is present for each bus type (BUTYPE) field: ▶ Type 4: Port identifier. One record is present for each supported port. – 0 = even port – 1 = odd port ▶ Type 6: Category.
BUPKTSND	Packets sent. Not supported for type 6.
BUPKTRCV	Packets received. Not supported for type 6.
BUBYTESND	Data bytes sent. Not supported for type 6.
BUBYTERCV	Data bytes received. Not supported for type 6.
BUMAXRATE	Maximum byte rate. The estimated maximum rate that data can be both sent and received in bytes per second.
BUDATA1	The meaning of this field depends on the type (BUTYPE) field: ▶ Type 4: Reserved. ▶ Type 6: Command bytes sent.
BUDATA2	The meaning of this field depends on the type (BUTYPE) field: ▶ Type 4: Reserved. ▶ Type 6: Command bytes received.

QAPMDISK

You find new entries in this table that detail, per path, the total read and write operations and worldwide node names for external disks. Table 7-3 shows the added columns.

Table 7-3 New columns in the QAPMDISK table

Column name	Description
DSPTROP	The path total read operations reports the number of read requests that are received by internal machine functions, which is not the same as the device read operations reported in the DSDROP field.
DSPTWOP	The path total write operations reports the number of write requests that are received by internal machine functions, which is not the same as the device write operations that are reported in the DSDWOP field.
DSWWNN	The worldwide node name is a unique identifier that represents the external storage subsystem that the disk belongs to. This value is null for non-external disks.

QAPMDISKR

Up to release 6.1, the QAPMDISK table contained a detailed set of data about the performance of the disk unit. This design was kept, but complemented with a new table (QAPMDISKR) that contains only the disk operations per interval. At the same time, it increases the number of bucket definition boundaries reported from 6 to 11, separates the read and write operations in different counters, and reports the bucket definition boundaries in microseconds instead of in milliseconds. These changes apply to all disks, internal or external. Each entry in the QAPMDISKR table contains the number of I/O operations, the response time, and the service time. The associated disk response time boundaries (in microseconds) are reported in the QAPMCONF file in GKEY fields G1–GA, for which there is no interface to change them.

You can find the breakouts for those buckets in Table 7-4. Both QAPMDISK and QAPMDISKRB tables carry the same columns for each row (interval number and device resource name) and can therefore be joined for analysis.

Table 7-4 Boundaries per bucket in the QAPMDISKRB and QAPMDISK tables

QAPMDISKRB (microseconds)			QAPMDISK (milliseconds)		
Bucket	>	<	Bucket	>	<
1	0	15	1	>0	1
2	15	250			
3	250	1,000			
4	1,000	4,000	2	2	16
5	4,000	8,000			
6	8,000	16,000			
7	16,000	64,000	3	16	64
8	64,000	256,000	4	64	256
9	256,000	500,000	5	256	1024
10	500,000	1,024,000			
11	1,024,000		6	1024	

QAPMJOBMI

The QAPMJOBMI table now has information about lock counts on a thread basis, providing details about which locks are held (seizes, process scoped locks, thread scoped locks, process scoped database record locks, and thread scoped database record locks held). It also holds information about the resource affinity status changes of a thread or process.

Cache memory access: On Power Systems, all of the processor cores on any chip can access any of the cache memory in the entire system. The management of the relationship between the processor or “node” where a task runs and the “nodal” location where that task finds its data is called *Memory Resource Affinity*.

Collection Services now includes physical I/O breakdown by SQL activity. These metrics are included in the *JOBMI collection category and reported by Collection Services in the QAPMJOBMI file. An updated template file for QAPMJOBMI is shipped in QSYS with the additional fields. This template file is used only when you create QAPMJOBMI file in libraries where it does not exist. This enhancement enables improved native database versus SQL database performance isolation.

Table 7-5 lists the new columns in the QAPMJOBMI table.

Table 7-5 New columns in the QAPMJOBMI table

Column name	Description
JBNFHN	The identifier of a resource affinity domain that this software thread or task is associated with. A thread or task is associated with the resource affinity domain at a create time, but the operating system can decide to move it to another resource affinity domain later.
JBNFLVL	The resource affinity level specifies the relative strength of the binding between a thread and the internal machine resources with which it has affinity (processors and main storage). The strength is expressed as: ▶ X'00' = Processor normal, main storage normal ▶ X'01' = Processor normal, main storage high ▶ X'10' = Processor high, main storage normal ▶ X'11' = Processor high, main storage high ▶ X'03' = Processor normal, main storage none ▶ X'20' = Processor low, main storage normal
JBNFGRP	The identifier of a resources affinity group or resource affinity domain. This identifier specifies how threads or tasks are related to other threads or tasks in their use of internal machine processing resources, or how they are related to specific resource affinity domains.
JBNFHNC	The amount of processor time that is used by the thread on the resource affinity domain that this thread is associated with. The time is reported in internal model-independent units. This time is called the <i>Local Dispatch Time</i> .
JBNFFNC	The amount of processor time that is used by the thread on resource affinity domains other than the one that this thread is associated with, but within the same group. The time is reported in internal model-independent units. This time is called the <i>Non-Local Dispatch Time</i> .
JBNFHNP	The number of 4-K page frames that are allocated for this thread during this interval from the resource affinity domain that this thread is associated with. These frames are called <i>Local page frames</i> .
JBNFFNP	The number of 4-K page frames that are allocated for this thread during this interval from resource affinity domains other than the one that this thread is associated with, but within the same group. These frames are called <i>Non-local page frames</i> .
JBTNAME	Identifies the name of secondary thread at sample time. The field is blank for primary threads, tasks, and unnamed secondary threads.
JBSLTCNT	If the short lifespan entry count is greater than zero, the entry does not represent a particular task or secondary thread. Instead, it is a special record that is used to report data that is accumulated for tasks and threads whose lifespan was shorter than the reporting threshold that was in effect when the collection started. Short lifespan tasks are reported for the processor node that they were associated with and short lifespan secondary threads are reported for the job to which they belong.
JBSACPU	The accumulated job scaled processor time that is charged (in microseconds). The accumulated scaled interval processor time that is charged for all threads of the job since the job started. This field is provided for primary threads only.
JBINDCPU	The amount of unscaled processor time (in μ s) that represents the work that is done solely within this thread without regard for how server task work is charged.
JBSINDCPU	Thread scaled processor time that is used (in microseconds). The amount of scaled processor (in μ s) time that represents the work that is done solely within this thread without regard for how server task work is charged.
JBCPUWC	The elapsed processor time (in μ s) that a task runs.

Column name	Description
JBVPDLY	The elapsed delay time (in microseconds) because virtualization for a task when it was running. The virtual processor delay time includes virtual processor thread wait event time, virtual processor thread wait ready time, and virtual processor thread dispatch latency.
JBSEIZECNT	The number of seizures held by this thread at the time that the data was sampled.
JBPSLCKCNT	The number of process scoped locks that are held by this thread at the time that the data was sampled.
JBTSLCKCNT	The number of thread scoped locks that are held by this thread at the time that the data was sampled.
JBTSRCDLCK	The number of thread scoped database record locks held by this thread at the time that the data was sampled.
JBNFOGDT	The amount of processor time that is used by the thread in a resource affinity group other than the one that this thread is associated with. The time is reported in internal model-dependent units.
JBNFOGMA	The number of 4-K page frames that are allocated for this thread during this interval from a resource affinity group other than the one that this thread is associated with.
JBFLDR2	Workload capping group delay time (in microseconds). The amount of time that this thread could not be dispatched because of workload capping.
JBFLDR3	Workload capping group. The identifier for the workload capping group that this thread belonged to at the time this data was sampled. A value of zero is reported when no group was assigned.
JBSQLCLK	SQL clock time. The amount of clock time (in microseconds) that this thread spent performing work that is done on behalf of an SQL operation.
JBSQLCPU	The thread unscaled SQL process time that is used. The amount of unscaled processor time (in microseconds) that this thread used performing work that is done on behalf of an SQL operation.
JBSQLSCPU	Thread scaled SQL processor time used. The amount of scaled processor time (in microseconds) that this thread used performing work that is done on behalf of an SQL operation.
JBSQLDBR	SQL synchronous database reads. The total number of physical synchronous database read operations that are done on behalf of an SQL operation.
JBSQLNDBR	SQL synchronous nondatabase reads. The total number of physical synchronous non-database read operations that are done on behalf of an SQL operation.
JBSQLDBW	SQL synchronous database writes. The total number of physical synchronous database write operations that are done on behalf of an SQL operation.
JBSQLNDBW	SQL synchronous nondatabase writes. The total number of physical synchronous non-database write operations that are done on behalf of an SQL operation.
JBSQLADBR	SQL asynchronous database reads. The total number of physical asynchronous database read operations that are done on behalf of an SQL operation.
JBSQLANDBR	SQL asynchronous nondatabase reads. The total number of physical asynchronous non-database read operations that are done on behalf of an SQL operation.
JBSQLADBW	SQL asynchronous database writes. The total number of physical asynchronous database write operations that are done on behalf of an SQL operation.

Column name	Description
JBSQLANDBW	SQL asynchronous nondatabase writes. The total number of physical asynchronous non-database write operations that are done on behalf of an SQL operation.
JBHSQLSTMT	Number of high-level SQL statements. The number of high-level SQL statements that run during the Collection Services time interval. This count includes only initial invocation of the independent SQL statements. It does not include dependent SQL statements started from within another SQL statement. This count also includes initial invocation of independent SQL statements that fail to run successfully.

QAPMJOBSR

This file contains data for jobs that perform save or restore operations. It contains one record per job for each operation type that is performed.

If you click **Collection Services** → **Collection Services Database Files**, and then select **QAPMJOBSR** in Performance Data Investigator (PDI), you see an overview of the data that looks like Figure 7-2.

Select	Interval Number	Interval Date And Time	Elapsed Interval Seconds	Operations Started	Objects Failed	Objects Successful	Operations Completed	Operations Terminated
☐	1	Sep 20, 2011 12:15:00 AM	890	168	0	168	158	10
☐	1	Sep 20, 2011 12:15:00 AM	890	0	0	0	0	0
☐	1	Sep 20, 2011 12:15:00 AM	890	0	0	0	0	0
☐	2	Sep 20, 2011 12:30:00 AM	900	0	0	0	0	0
☐	2	Sep 20, 2011 12:30:00 AM	900	0	0	0	0	0
☐	3	Sep 20, 2011 12:45:00 AM	900	0	0	0	0	0
☐	3	Sep 20, 2011 12:45:00 AM	900	0	0	0	0	0
☐	4	Sep 20, 2011 1:00:00 AM	900	0	0	0	0	0
☐	4	Sep 20, 2011 1:00:00 AM	900	0	0	0	0	0
☐	5	Sep 20, 2011 1:15:00 AM	900	0	0	0	0	0
☐	5	Sep 20, 2011 1:15:00 AM	900	0	0	0	0	0
☐	6	Sep 20, 2011 1:30:00 AM	900	0	0	0	0	0

Figure 7-2 Output from the QAPMJOBSR table

QAPMSHRMP

The QAPMSHRMP table reports shared memory pool data (refer to as Active Memory Sharing in PowerVM). This data is generated only when a partition is defined to use a shared memory pool. Data is reported for both the partition's use of the pool and pool metrics that are the sum of activity that is caused by all partitions using the pool. You must have a POWER6 system and firmware level xx340_075 or later for this data to be available. See Table 7-6 for the data that is kept in this table, which shows the interval number (INTNUM), date and time (DTETIM), and the seconds in the interval (INTSEC) columns.

Table 7-6 Contents of the QAPMSHRM table

Column name	Description
SMPOOLID	Shared memory pool identifier. The identifier of the shared memory pool that this partition is using.
SMWEIGHT	Memory weight. Indicates the variable memory capacity weight that is assigned to the partition. Valid values are hex 0 - 255. The larger the value, the less likely this partition is to lose memory.

Column name	Description
SMREALUSE	Physical real memory used. The amount of shared physical real memory, in bytes, that was being used by the partition memory at the sample time.
SMACCDLY	Real memory access delays. The number of partition processor waits that occurred because of page faults on logical real memory.
SMACCWAIT	Real memory access wait time. The amount of time, in milliseconds, that partition processors waited for real memory page faults to be satisfied.
SMENTIOC	Entitled memory capacity for I/O. The amount of memory, in bytes, assigned to the partition for usage by I/O requests.
SMMINIOC	Minimum entitled memory capacity for I/O. The minimum amount of entitled memory, in bytes, needed to function with the current I/O configuration.
SMOPTIOC	Optimal entitled memory capacity for I/O. The amount of entitled memory, in bytes, that allow the current I/O configuration to function without any I/O memory mapping delays.
SMIOCUSE	Current I/O memory capacity in use. The amount of I/O memory, in bytes, currently mapped by I/O requests.
SMIOCMAX	Maximum I/O memory capacity used. The maximum amount of I/O memory, in bytes, that was mapped by I/O requests since the partition last had an IPL or the value was reset by an explicit request.
SMIOMDLY	I/O memory mapping delays. The cumulative number of delays that occurred because insufficient entitled memory was available to map an I/O request since the partition last underwent an IPL.
MPACCDLY	Pool real memory access delays. The number of virtual partition memory page faults within the shared memory pool for all partitions.
MPACCWAIT	Pool real memory access wait time. The amount of time, in milliseconds, that all partitions processors spent waiting for page faults to be satisfied within the shared memory pool.
MPPHYMEM	Pool physical memory. The total amount of physical memory, in bytes, assigned to the shared memory pool.
MPLOGMEM	Pool logical memory. The summation, in bytes, of the logical real memory of all active partition active partitions that are served by the shared memory pool.
MPENTIOC	Pool entitled I/O memory. The summation, in bytes, of the I/O entitlement of all active partitions that are served by the shared memory pool.
MPIOCUSE	Pool entitled I/O memory in use. The summation, in bytes, of I/O memory that is mapped by I/O requests from all active partitions that are served by the shared memory pool.

QAPMSYSTEM

The QAPMSYSTEM reports system-wide performance data. In IBM i 7.1, columns are added, as shown in Table 7-7.

Table 7-7 New columns in QAPMSYSTEM

Column name	Description
SYPTWAIT	Virtual processor thread wait event time. The elapsed time in microseconds that blocked threads of the partition's virtual processors were waiting for an event that caused them to become ready to run.
SYPTREADY	Virtual processor thread wait ready time. The elapsed time in microseconds that ready to run threads of the partition's virtual processors waited to be dispatched when entitled capacity was exhausted.

Column name	Description
SYPTLATEN	Virtual processor thread dispatch latency. The elapsed time in microseconds that ready to run threads of the partition's virtual processors waited to be dispatched when entitled capacity was not exhausted and a physical processor was not available.
SYPTACT	Virtual processor thread active time. The elapsed time in milliseconds summed for all threads of a virtual processor for the time that the virtual processor is active in the partition. A virtual processor that is active is one that is varied on; a virtual processor that is not active is either varied off or not installed.
SYPTIDLE	Virtual processor thread idle time. The elapsed time in milliseconds summed for all threads of a virtual processor for the time that thread is idle in the partition. A processor thread that is idle is one that is varied on and running the partition's idle loop.
SYPTINTR	Virtual processor thread interrupt time. The elapsed time in milliseconds summed for all threads of a virtual processor for the time that thread is handling interrupts in the partition.
SYFRMCPU	Processor firmware time used (in microseconds). The amount of processor firmware time that is used by this partition.
SYFRMSCPU	Processor scaled firmware time that is used (in microseconds). The amount of scaled processor firmware time that is used by this partition.
SYFRMSCPU	Processor scaled firmware time that is used (in microseconds). The amount of scaled processor firmware time that is used by this partition.
SYPFOLDSW	Identifies the current state of the processor folding ^a switch, where: ▶ <i>Blank</i> means data not available ▶ "0" off ▶ "1" on ▶ "2" system controlled
SYPFOLDST	The current state of processor folding, where: ▶ <i>Blank</i> means data not available ▶ "0" disabled ▶ "1" enabled
SYEMMAJCDE	Energy management major code ^b .
SYEMMINCDE	Energy management minor code ^b .
SYEMATTR	Energy management attributes. Bit 0 identifies the power draw limit type (0 = soft, 1 = hard).
SYEMPWRLMT	Energy management power draw limit in watts.

- a. Processor folding enhances the use of the shared processor pools by minimizing the usage of idle virtual processors (VP). This configuration allows virtual partitions to be configured with more VPs to take better advantage of the shared processor pool. It does so by allowing you to increase the number of VPs on your configuration without a performance impact. It also increases the average VP dispatch cycle. This results in better cache use and reduced workload in the Hypervisor.
- b. Maj 0, Min 0 = Unspecified or unavailable; Maj 0, Min 2 = Disabled (nominal performance); Maj 0, Min 1 = Enabled (maximum performance); Maj 0, Min 3 = Enabled (power saver); Maj 1, Min 0-127 = Enabled (dynamic power optimizer).

QAPMSYSWLC

This database file reports workload capping group data. Data is generated only when one or more workload capping groups were in use during the collection. A record is written for each group that is active. Table 7-8 shows the columns in QAPMSYSWLC.

Table 7-8 Columns in QAPMSYSWLC

Column name	Description
INTNUM	Interval number. The nth sample database interval that is based on the start time that is specified in the Create Performance Data (CRTPFRDTA) command.
DATETIME	Interval date and time. The date and time of the sample interval.
INTSEC	Elapsed interval seconds. The number of seconds since the last sample interval.
SWGROU	Group ID. The identifier for the workload group.
SWGNAME	Group Name. The name that is assigned to the workload group when it is allocated by License Management.
SWPRCASN	Processors assigned. The maximum number of processors that may be used concurrently by all threads of all processes that are associated with the workload group. This value is the value that is associated with the group at the time data was sampled.
SWPRCAVL	Processor time available (in microseconds). The amount of processor time that this group had available to it based on the number of processors that are assigned to the group over time.
SWPRCUSE	Processor unscaled time used (in microseconds). The amount of unscaled processor time that is used within the threads that are assigned to this group. This value does not include the time that charged to a thread by server tasks.
SWSPRCUSE	Processor scaled time that is used (in microseconds). The amount of scaled processor time that is used within threads that are assigned to this group. This value does not include the time that is charged to a thread by server tasks.
SWDELAY	Dispatch latency time. The amount of time ready to run threads could not be dispatched because of the group's maximum concurrent processor limit.
SWPRCADD	Processes added. The number of process instances that became associated with this group during the interval.
SWPRCRMV	Processes removed. The number of process instances that were disassociated from this group during the interval.

QAPMTAPE

The QAPMTAPE table contains the tape device data that is collected in the Removable storage (*RMVSTG) collection category. It contains one record per interval per tape device that is connected to the system. Besides the data about the interval, it contains the columns in Table 7-9.

Table 7-9 New columns in QAPMTAPE

Column name	Description
TPRDS	Number of reads
TPWRTS	Number of writes
TPBRD	Bytes read
TPBWRT	Bytes written

Column name	Description
TPWREQ	Time spent waiting for a request from the client (in milliseconds)
TPWRESP	Time spent waiting for a response from the drive (in milliseconds)
TPSFMCMD	Space by file mark commands
TPFLMRKSPC	File marks spaced
TPSBCMD	Space block commands
TPBLCKSPC	Blocks spaced
TPWFCMD	Write file mark commands
TPFLMRKWRT	File marks written
TPSEODCMD	Space to EOD commands
TPWBCMD	Write buffer commands
TPRELEASES	Release commands
TPREWINDS	Tape rewinds
TPUNLOADS	Tape unloads
TPSTPOSCMD	Set tape position commands
TPRDPOSCMD	Read tape position commands

QAPMXSTGV and QAPMXSTGD

In IBM i 6.1.1, the QAPMXSTGV table, and in IBM i 7.1, the QAPMXSTGD table, were added with performance data of external storage systems (DS8000 and DS6000 storage servers). This data can be analyzed with iDoctor - Collection Services Investigator. The table contains mainly volume and LUN-oriented statistics and can also catch advanced Logsensestats from those storage servers. The support for *EXTSTG is disabled when shipped. For more information, see the *Memo to Users* and APAR SE41825 for PTF information at:

- ▶ <http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzaq9/rzaq9.pdf>
- ▶ http://www-912.ibm.com/n_dir/nas4apar.NSF/c79815e083182fec862564c00079d117/810d72fc51f14ed48625770c004b9964?OpenDocument

For more information, see Chapter 9, “Storage and solid-state drives” on page 367.

7.2.3 Additional enhancements

There are several enhancements that were made to the Collection Services component of Performance Tools, which include:

- ▶ Collection Services was updated to support reporting system-wide usage data for workload capping groups as well as TDE level data to help understand performance issues that are related to capping actions.
 - The *JOBMI data category and QAPMJOBMI file are modified to support additional TDE metrics that identify the group a TDE (thread) was associated with at sample time. These metrics also measure how much time that thread was not able to run because of workload capping constraints.

- The *SYSLVL collection category was modified to collect WLC group data for groups that are in use. A new file named QAPMSYSWLC was created in the target performance database library to contain this data. The QAPMSYSWLC file or member is created only if the source *MGTCOL collection contains data for workload capping groups.
- ▶ A page size flag metric was added to the Dump Main Memory Information (**DMPMEMINF**) command to indicate whether a page is 4 K or 64 K.
- ▶ A new metric was implemented by Collection Services to report the number of full opens that occur within a job and is being reported in the QAPMJ0B0S file using the following fields:
 - JBLBO: The cumulative number of SQL cursors that are fully opened.
 - JBLBS: The cumulative number of SQL cursors that are pseudo-opened (also known as reused SQL cursors).
 - JBNUS: The number of native database (non-SQL) files and SQL cursors that are full opened.

Subtracting the value within field JBLBO from JBNUS yields the number of non-SQL full opens.

7.3 IBM Navigator for i Performance interface

Several enhancements were made to Performance Data Investigator, which can be accessed by selecting the Investigate Data task:

- ▶ A new content package is included, which shows the general health of your partition, and is configurable for user-defined health thresholds.
- ▶ Collection Services collect high-level cross-partition processor performance metrics for all logical partitions on the same single physical server regardless of the operating system. This function is available on POWER6 and later systems, with a minimum firmware level xx340_061. When this data is available, it can be viewed through several perspectives that are found under the Physical System menu.
- ▶ Data can now be exported to an image (charts only), comma delimited, or tab delimited file.
- ▶ Investigate Data can now send data from your current session to Workload Estimator for use in sizing a future system using current performance characteristics.
- ▶ New charts and tables can be developed from within Investigate Data. Adding views, modifying SQL statements, and modifying data series information can all be performed without leaving the page.
- ▶ With the many metrics provided by Investigate Data, sometimes knowing which perspective contains the metric you want to see is not trivial. You can choose your perspective by searching for a metric name. Look for this feature by clicking the **Search** button when you start Investigate Data.
- ▶ A simple Performance Explorer content package is included to start analysis of Performance Explorer data.
- ▶ Many new perspectives were added to Investigate Data, and many new metrics were added to existing perspectives. The new metrics include SAN, Virtual I/O, Energy Management, Virtual Memory, and Communication Data, plus much more.
- ▶ Miscellaneous changes were made to improve the overall capabilities of the performance tasks. Examples within Investigate Data include persistent options, new collection context at the top of each perspective, and a new menu bar that allows for quicker navigation and complete history data.

Remember: Most of these new functions are available for systems that are running IBM i 6.1 with the current IBM HTTP Server for i group PTF (SF99115).

For a full overview of all the functions of the Performance Investigator and Collection management interface, see Chapter 18, “IBM Navigator for i V7.1” on page 639.

7.4 IBM iDoctor for IBM i

This section describes the changes in IBM iDoctor for IBM i for IBM i 7.1. These changes are also available for the 5.4 and 6.1 releases unless otherwise noted.

7.4.1 Components

Job Watcher, PEX Analyzer, Disk Watcher, Plan Cache Analyzer, and Collection Services Investigator are now available for Version 7.1 (and Version 5.4 and Version 6.1).

Heap Analyzer is not available for 7.1.

VIOS Investigator is a new component that uses NMON data to analyze VIOS performance using the power of the DB2 database on System i. This component is available for Version 6.1 and Version 7.1.

A new no additional cost FTP client component is now part of iDoctor. From the connection view, you can choose the connection type (IBM i) or FTP. If you have an IBM i connection and want to try the FTP support, right-click and use the Start FTP session menu. For IBM i FTP support, you have access to both the IFS and libraries. You can download, upload, or cause a transfer of the selected files to another system. The transfer function does not work well if you are moving DB2 files directly from one system to another (the format is corrupted), but it works fine for save files. If you are connecting to AIX, you can transfer the NMON files to the correct IBM i system.

7.4.2 Installation

The following functions were added to the installation process:

- ▶ Validation checks were added for each specified partition to ensure that the partition is able to install iDoctor.
- ▶ The default installation directory is now C:\Program Files\IBM\iDoctor on 32-bit Windows and C:\Program Files (x86)\IBM\iDoctor on 64-bit Windows.
- ▶ A check was added for installing Job Watcher at the 6.1 and 7.1 releases to ensure that the Job Watcher definitions file (QAPYJWDFN) exists in QUSRSYS and contains the IBM-supplied definitions.
- ▶ In the installation, an option on the Component Selection window that is called Create user profile QIDOCTOR (applies to Base support only) with the default cleared was added.

An example video that shows how to install iDoctor is available at:

http://www.youtube.com/watch?v=pURHegItOTQ&feature=channel_video_title

The iDoctor GUI now requires the Visual Studio 2008 SP1 redistributable package to be installed. More information about this requirement can be found at:

https://www-912.ibm.com/i_dir/idoctor.nsf/downloadsClient.html

7.4.3 My Connections View

My Connections View, which is shown in Figure 7-3, provides the following enhancements:

- ▶ Added columns to show access code expiration dates, missing PTFs, system serial number, ASP group name, and relational database name (if the connection uses an independent ASP).
- ▶ New menu options added to Check Expiration Dates or Check PTFs against the wanted partitions. Check PTFs includes checking for the latest Performance Group PTF levels on 5.4 and higher systems.
- ▶ Added menus to Load and Remove all iDoctor Stored Procedures.
- ▶ Added Uninstall iDoctor option.
- ▶ Added an option to Edit an existing connection.
- ▶ Deleted obsolete analysis files for each component.




Figure 7-3 Options of My Connections View

When you sign on to a system, iDoctor uses the configured sign-on setting defined in System i Navigator (you can access this setting by clicking **Properties** and clicking the **Connection** tab for a system). You can use options such as **Use Windows user name and password** to avoid needing to sign on through iDoctor if your Windows password matches the user name and password of the System i to which you are connecting. iDoctor also uses the client access password cache to avoiding prompting for a password unless needed. If you still want to be prompted for a password every time you start iDoctor, set the **Prompt every time** option within System i Navigator.

Support was added to view collections that are stored in libraries that are created in Independent ASPs. Use the Add connection menu or Edit menu from the My Connections View to specify the appropriate ASP group name and relational DB name, as shown in Figure 7-4. These values cause the QZRCRVS and QZDASOINIT jobs that are created by iDoctor to recognize data that is stored in the Independent ASP.




Figure 7-4 Add Connection

7.4.4 Main window

On the main window toolbar, a button was added that enables / disables the display of situational analysis background colors in graphs. A simple click of the button turns it on / off for all graphs (even open ones). Another click of the graph or legend redraws the graph with or without the situations (if found in the data), as shown in Figure 7-5.




Figure 7-5 Button to enable or disable situations analysis background colors

On the main window, which is shown in Figure 7-6, the clock icon can now be used from any component to set the preferred time range interval size. The clock icon now has the following additional time grouping options: one-tenth-second, five-second, fifteen-second, five-minute, four-hour, eight-hour, twelve-hour, and twenty-four-hour. The small groupings are useful in PEX Analyzer and the large groupings are useful in Collection Services. More checking was added to the GUI to ensure that only relevant time grouping options are shown for the current data.




Figure 7-6 Main window

You can save a URL for libraries, collections, call stacks, and so on, in iDoctor using the Copy URL menu option or button. The URL can then be pasted into a web browser or saved for future use. The URL shown in Example 7-1 opens library COMMON within Job Watcher on system Idoc610.

Example 7-1 Opening the COMMON library within Job Watcher on system Idoc610

```
idocitor:///viewinfo1[type=CFolderLib,sys=Idoc610,lib=COMMON,comp=JW]
```

All components now provide several common options (folders) for working with data or parts of the system in various ways:

- ▶ The Libraries folder displays the libraries on the system that contain data for the current component.
- ▶ The SQL Tables folder is a repository in iDoctor for working with the tables created by the iDoctor analyses. Comparison options are available by right-clicking more than one SQL table.
- ▶ A new Browse Collections option was added that provides alternative ways of looking at the collections that are stored on the system. This function is built from a repository that must be updated periodically by the user using the options that are found when you right-click the **Browse Collections** folder.

You can use the Browse Collections option to find collections on the system in several ways:

- Partition Name
- Partition Name and VRM

- Partition and collection type
- Library
- Library and collection type
- Collection type
- Collection type and VRM
- VRM
- Month created
- Month created and collection type

Each of the views gives the total size of all collections (in MBs) in each group and the total number, as shown in Figure 7-7.

- ▶ Added a Super collections folder that you can use to work with the super collections that exist on the system. These collections contain a folder for each collection type that is collected within the super collection.
- ▶ You can use the Saved collections folder to work with any save files that are found on the system that contain iDoctor collections that were saved previously using the iDoctor GUI.
- ▶ You can use the Work Management folder to work with scheduled iDoctor jobs, work with active jobs, or work with the subsystems. Within the subsystems folder are options to work the iDoctor subsystems or all subsystems.

Library name	Collection type	Owner	Total collections	Total size (MBs)
PEXLAB8	PA	QDFTOWN	1	2672.5703
DEMO3	JW	QSECOFR	1	2000
CSILIB1	CS	QDFTOWN	1	1839.9023
JWLABEX5A	JW	QDFTOWN	1	1282
QPFRDATA	CS	QSYS	12	1168.8125
AK247DATA	CS	MCCARGAR	1	1105.0117
PEXLAB11	PA	BSMENGES	5	674.9531
DEMO555	JW	QDFTOWN	1	531
PEXLAB7	PA	QDFTOWN	1	517.0898
JWLABEX4	JW	QDFTOWN	1	515
XSTGD	CS	QCOLSRV	1	338.7969
BRADTEST	JW	BSMENGES	1	334
JWLABEX1	JW	QDFTOWN	1	334
FIXME	JW	BSMENGES	1	285.7979
JWINTERN	JW	QDFTOWN	1	222.9717
DEMO54B	JW	CRAVENS	1	222.9717

Figure 7-7 Browse collections

The iDoctor components now contain two new folders that show the ASPs and disk units that are configured on the current system. You can use the ASPs folder to drill down to view the disk units within an ASP. The Disk Units folder provides WRKDSKSTS type of statistics with updates provided with each refresh (it also includes total I/O and total sizes), as shown in Figure 7-8.

	Disk unit	Disk type	ASP	Size (GB)	% used	% busy	I/Os per second	Avg I/O size (KB)	Avg reads per second	Avg writes per second	Avg read size (KB)	Avg write size (KB)
	49	4326	1	25.76	82.8	5.9	0	0	0	0	0	0
	48	4326	1	30.05	82.8	5.9	0	0	0	0	0	0
	45	4326	1	30.05	82.8	5.9	0	0	0	0	0	0
	44	4326	1	25.76	82.8	5.9	0	0	0	0	0	0
	53	4326	1	30.05	82.8	5.6	.56	101.60	0	.56	0	101.60
	51	4326	1	30.05	82.8	5.6	.11	16	.11	0	16	0
	39	4328	1	120.59	82.8	1.3	.67	85.33	0	.67	0	85.33
	40	4328	1	120.59	82.8	1.2	.56	92	.11	.44	8	113
	33	4328	1	103.37	82.8	1.2	.56	80.80	.11	.44	12	98
	37	4328	1	103.37	82.8	1.1	.67	76.67	0	.67	0	76.67
	38	4328	1	120.59	82.8	1.1	.44	83	0	.44	0	83
	41	4328	1	120.59	82.8	1	.44	81	0	.44	0	81
	55	4326	1	30.05	82.8	0	0	0	0	0	0	0
	56	4326	1	30.05	82.8	0	.11	8	0	.11	0	8
	52	4326	1	25.76	82.8	0	.33	6.67	.11	.22	12	4
	54	4326	1	30.05	82.8	0	.11	68	0	.11	0	68

Figure 7-8 Overview of the disk status in ASP1

Right-click the **Disk Units** or **ASP** folder and click **Reset Statistics** to restart the collection of disk statistics. You can also use the **Select fields** menu option when you right-click the folder to rearrange fields or add more fields. The status bar of the main window shows the times for first disk statistics snapshot, and the last one.

Similarly, you find an Active Jobs (see Figure 7-6 on page 266) folder on the same window, which provides WRKACTJOB-like functionality from the iDoctor client, as shown in Figure 7-9. You can also sort by a statistic and refresh to keep tabs on the top processor users, and so on. There is also a filter option to filter the list by name, user, number, current user, or minimum processor percentage. Click the **Select fields** menu option when you right-click the folder to rearrange fields or add more fields. Expanding a job shows the threads and the thread statistics available for each job. You can start Job Watcher or PEX Analyzer collections or add Job Watcher / PEX definitions using the selected jobs within the Active jobs folder. You can also end the selected jobs or view job logs.

	Job name	Job user	Job number	Threads	Status	Current user	Type	Function	CPU %	CPU time (ms)	Run p...	Disk IO
	QZDASOINIT	QUSER	698951	1	RUN	BABANDY	Interactive - Server		19.5	3543	20	0
	QPADEV0013	BRAU	691313	1	RUN	BRAU	Interactive	Q1PMENU	13.7	2498	20	0
	QZRCSSRV5	QUSER	092109	1	RUN	MCCARGAR	Interactive - Server		.2	38	20	560
	QUMEPVRVGT	QSECOFR	686828	4	DEQW	QSECOFR	Interactive	cimprovagt	.1	34	50	0
	TESTJOB	NICKT	937224	1	DEQW	NICKT	Interactive	TESTRECV	0	0	50	0
	TESTJOB	NICKT	937280	1	DEQW	NICKT	Interactive	TESTRECV	0	0	50	0
	TESTJOB	NICKT	937121	1	DEQW	NICKT	Interactive	TESTRECV	0	0	50	0
	TESTJOB	NICKT	937213	1	DEQW	NICKT	Interactive	TESTRECV	0	0	50	0
	TESTJOB	NICKT	937871	1	DEQW	NICKT	Interactive	TESTRECV	0	0	50	0
	TESTJOB	NICKT	937620	1	DEQW	NICKT	Interactive	TESTRECV	0	0	50	0
	TESTJOB	NICKT	937986	1	DEQW	NICKT	Interactive	TESTRECV	0	0	50	0
	TESTJOB	NICKT	937956	1	DEQW	NICKT	Interactive	TESTRECV	0	0	50	0
	TESTJOB	NICKT	937379	1	DEQW	NICKT	Interactive	TESTRECV	0	0	50	0
	TESTJOB	NICKT	937361	1	DEQW	NICKT	Interactive	TESTRECV	0	0	50	0
	TESTJOB	NICKT	937570	1	DEQW	NICKT	Interactive	TESTRECV	0	0	50	0

Figure 7-9 Overview of Active Jobs (WRKACTJOB)

You can use the objects that are owned by a user folder to manage (view / delete) the objects on the system that are owned by the current user. This function is primarily intended to aid in disk space cleanup and options are provided to view only iDoctor objects, or all objects that are owned by the current user.

7.4.5 Collection options

The Collection menu now contains an Analyses menu for all components. Choosing an option under this menu runs a program that creates SQL tables that are needed for further analysis. In most cases, more reports become available after the analysis completes and the collection is refreshed (by pressing F5.)

The Summarize menu option in CSI and Job Watcher moved to Analyses - Run Collection Summary. Choosing this option now displays a window that you can use to filter the collection data by job name, job user, job number, current user, subsystem, or time range. Filtered results can be viewed under the SQL tables folder. By not filtering the data, the summarized results are accessible using the graphing options that are provided under the collection.

The Create Job Summary menu option in CSI and Job Watcher moved to Analyses - Run Create Job Summary.

There is a new iDoctor Report Generator for all collections (Figure 7-10). To access it, right-click a collection and click the **Generate Reports**. The default web browser is opened to show the HTML report after the reports are captured to JPG files. As reports are running, you can switch to other windows, but before screen captures are taken, the data viewer must be moved to the front of all windows. This action happens automatically, but might look strange the first time you use it.




Figure 7-10 iDoctor Report Generator

With the Save option (see Figure 7-11), you can save multiple collections or monitors. After you use this option, the new Saved collections folder shows a record that identifies the save file, which you can use to restore the data or distribute it.




Figure 7-11 Saving a collection

The Transfer to option now supports all collection types. It was modified as follows:

- ▶ PEX collections, PEX MGTCOL objects, CS collections, CS MGTCOL objects, DW collections, and Job Watcher collections all now support the transfer option.
- ▶ The usage of the **FTPJWCOL** / **FTPPACOL** commands by the GUI is replaced by a combination of **SAVPFCOL** / **FTPFILE**.
- ▶ When you transfer multiple collections, they are saved / sent using the same save file instead of different ones.
- ▶ Monitors (or multiple monitors) can now be transferred.
- ▶ Path / file names increased to 100 chars from 50.
- ▶ You have complete control to set the file name to whatever you want, but the recommended format is given.
- ▶ You now have an option in the action list to FTP the collections to Ecurep.
- ▶ Mixed case passwords, user names, and file names are now supported, which fixes problems when you send data to and from AIX.
- ▶ The **QIDRGUI/FTPFILE** command now supports a GET action instead of the default of PUT. This command allows an iDoctor i user to receive data from another system (such as NMON data from an AIX box).

In all components that support copying a collection, you can now specify the collection name in the target library. You can use this function to copy a collection to a new name in the same library.

7.4.6 Data Viewer

The Data Viewer toolbar has a toolbar that shows all the idle waits (include all buckets) for wait bucket *jobs* and larger grouping graphs in CSI and Job Watcher. This toolbar is a toggle button that you can use to see the idle waits and click again to see the interesting waits.

Previously, the idle waits were not shown for job or higher groupings. Figure 7-12 shows an example.




Figure 7-12 Idle waits toggle button

There is a new menu, Choose Database Members, in the SQL editor that clears the current member selections and opens the member selection window.

In the record quick view, you can see the table alias name before the column name if it is known.

You can now click **Save** on the Data Viewer toolbar to save the current graph and legend as a JPG image.

You can click **File** → **Create Shortcut** to save a Data Viewer graph or table as a shortcut file (*.idr). The file can be sent to other users or saved on the PC to revisit the report later.

The interval grouping option on the Clock icon now has a five-minute and four-hour time interval specification.

The Data Viewer has a button, represented by the sum symbol, to perform math over cells from the selected rows in a flyover as the user moves their mouse pointer over the data. The options are as follows:

- ▶ None (normal flyover behavior)
- ▶ Sum
- ▶ Average
- ▶ Minimum and Maximum
- ▶ Percent of, Delta (current - prior)
- ▶ Delta (prior - current)

There are also changes to the design of graphs and reports.

Side-by-side comparisons

You can use the side-by-side comparisons to sync up the scrolling and Y-axis scaling of any two graphs (or tables) in a Data Viewer.

When two or more graphs or / tables exist in a Data Viewer, the buttons are ready for use. See Figure 7-13.




Figure 7-13 Side by side comparisons

An example video of using side-by-side comparisons can be found at:

http://www.youtube.com/watch?v=0qkg_cmGIAo&feature=channel_video_title

Table filtering options

Right-click the column that you want to filter. Set the wanted operator and value and click **Apply** to refresh immediately. See Figure 7-14.




Figure 7-14 Filter

More options are found by right-clicking a column:

- ▶ Sort
- ▶ Remove filter
- ▶ Hide column
- ▶ Unhide all columns

Graph filter options

You can use new graph filtering options to define a filter over the wanted column that is shown in the legend. Right-click the wanted column description and click **Add Filter**. See Figure 7-15.




Figure 7-15 Example of graph filter options

7.4.7 Collection Services Investigator

The folders available in the Collection Services Investigator (CSI) component changed. Instead of showing libraries that contain CS data, new folders are available, as shown in Figure 7-16:

- ▶ Libraries containing CS database file collections (filterable)
- ▶ Historical summaries that contain broader periods of CS performance data (weeks / months)
- ▶ CS objects for a list of all CS management collections objects on the system

The rest of the folders are described in 7.4.4, “Main window” on page 265.




Figure 7-16 Folders in the CSI component

Historical Summaries

Historical Summaries consist of consolidated or reduced sets of Collection Services data for graphing many days, weeks, or months of data at one time.

Figure 7-17 gives an example of a Historical Summary *Collection Overview Time Signature* graph over 12 days of data.




Figure 7-17 Historical Summaries - Collection Overview Time Signature

This data can be created in two possible ways:

1. Run the **STRCSMON** command (or use the Start Monitor menu option from the Historical Summaries folder), which summarizes new Collection Services data every day at a certain time and adds it to the system's default Historical Summary repository.
2. Right-click one or more libraries and click **Analyses** → **Run Historical Summary**.

Use the 1 hour grouping option when you create the Historical Summary.

Historical Summaries provide a complete set of graphs similar to the graphs provided under normal collections. A full set of “average day” and “average week” graphs are also supplied. More information about Historical Summaries can be found at:

[https://www-912.ibm.com/i_dir/idoctor.nsf/3B3C112F7FBE774C86256F4000757A8F/\\$FILE/iDoctorSep2011.pdf](https://www-912.ibm.com/i_dir/idoctor.nsf/3B3C112F7FBE774C86256F4000757A8F/$FILE/iDoctorSep2011.pdf)

Capacity planning

You can now capture the selected collection's information and import it into the Workload Estimator. Use the Launch Workload Estimator menu on a collection, as shown in Figure 7-18. A window opens with the average processor and disk statistics for the collection.




Figure 7-18 Start the Workload Estimator

When you click **Submit** (Figure 7-19), the data is sent to Workload Estimator for further analysis.

The screenshot shows a web browser window titled "Launch WLE - Mozilla Firefox" with the address bar displaying "file:///C:/temp/WLElauncher.html". The main content area is titled "IDoctor Collection Services Investigator Launch Workload Estimator for upgrade sizing". It contains several sections: "System Name: S659F730", "Operating System: IBM i", "Version: 6.1", "Library Name: ANALYSIS", and "Collection: DS6000". Below this is "CPU Information" showing "Average CPU utilization 29.1187" and "Average interactive utilization 6.0163". The "Disk Information" section contains a table with 8 columns: Disk type, Storage Consumed (GB), Protection, Disk busy, Reads per second, Size per read (bytes), Writes per second, and Size per write (bytes). The table has two rows of data. At the bottom left of the form is a "Submit" button.

Disk type	Storage Consumed (GB)	Protection	Disk busy	Reads per second	Size per read (bytes)	Writes per second	Size per write (bytes)
1750	2003.1017	RAID5	10.620	319.081	39460.0	1075.637	12490.0
1750	58.9203	None	22.250	18.255	34541.0	100.501	5843.0

Figure 7-19 Submit to WLE

Managing collections

On the same drop-down menu that is shown in Figure 7-18 on page 275, you see that a Copy function was added. You can also use the Copy Performance Data (CPYPFRDTA) command to obtain the same result. The Delete function now uses the Delete Performance Data (DLTPFRDTA) command.

The import data to WLE option is accessible from a CSI graph if a time range is selected. The numbers that are provided to WLE are based on the selected time period.

A search function, similar to the one in Job Watcher, is now available in this window, which you can use to generate a report that shows job records that are based on a specific job name, user, number, subsystem, pool ID, or current user profile. From these reports, you can drill down into the graphs for the wanted job over time. You can also search over multiple collections at one time by selecting the wanted collections in the CSI component view's list side and then using the Search menu. After the search results are shown, you can drill down on the same set of collections that are provided for the wanted job or thread.

You can create graphs over multiple collections at one time in the same library. Select the wanted collections from the CSI component view's list side and then right-click and choose the graph of interest. Click **Yes** when prompted if the graph is to be created for all collections selected. From then on, any drill-down you do on rankings and the single object over time graphs apply to this same set of collections.

In CSI and PEX, you now have a Wait Buckets tab that shows the wait buckets or ENUM mapping. Right-click and select **Collection** → **Properties** → **Wait Buckets**, as shown in Figure 7-20.




Figure 7-20 Wait and ENUM display

Situational analysis

To initiate a situational analysis, right-click the collection and click **Analyses** → **Run Situational Analysis**. You can also right-click the collection, click **Analyses** → **Analyze Collection**, and click **Situations** to configure the default situations to be used by CSI.




Figure 7-21 Collection Services Investigator Situations

All the time interval-based Wait graphs include different background colors, each identifying a situation, as shown in Figure 7-22.




Figure 7-22 Situational Analysis example

The current list of situations and the default minimum thresholds are shown in the following list:

- ▶ Interactive feature use high: 100%
- ▶ Write cache overruns: 20%
- ▶ High disk response times: 15 ms
- ▶ High faulting in the machine pool: 10 faults per second
- ▶ High disk storage consumed: 90%
- ▶ Jobs ineligible to run: Three instances for a job per interval
- ▶ Significant changes in pool size: 25% change from one interval to the next

External storage analysis

For an in-depth overview of this instrumentation, see 9.2.5, “External disk storage performance instrumentation” on page 392.

IASP Bandwidth analysis

The purpose of the IASP Bandwidth analysis is to analyze the Collection Services data with the intent to determine whether the system is a good candidate for migrating to Independent ASPs.

Right-click the collection and click **Run IASP Bandwidth** to kick off the analysis. When you run the analysis, you are prompted for several parameters:

- ▶ **Compression rate:** The estimated network (comm line) compression rate (between the system and the IASP). A value of 1 means no compression. 1.5 (the default) means 50% compression. Values less than 1 should not be used.
- ▶ **Full system bandwidth:** Estimated bandwidth that is required by system without IASPs (in megabits per second). Depending on the system or data, you might want to adjust this value to be much higher.
- ▶ **IASP bandwidth:** Estimated bandwidth that is required by system with IASPs implementation (in megabits per second). Depending on the system or data, you might want to adjust this value to be much higher.
- ▶ **ASP filtering:** You can use this option to select which ASPs to include when you run the analysis.

After you run the analysis, a new IASP bandwidth estimations folder is available that contains the generated table and a subfolder with graphs.

The IASP bandwidth estimate table represents the statistics that are generated by the analysis. The statistics that are generated include the number of blocks, writes, database write percentage, and various bandwidth estimates (all in megabits per second).

The IASP bandwidth overview graph displays the database writes for IASPs and the full system writes.

The IASP bandwidth overview graph with lag includes how much bandwidth lag there would be based on the parameter estimation values given when the analysis was created.

SSD Candidate Screening

Available at 7.1 or higher, SSD Candidate Screening helps determine whether SSDs could improve performance, as shown in Figure 7-23. It uses data from QAPMDISKRB. It is similar to the SSD Analyzer tool.




Figure 7-23 SSD candidate screening

If you select **SSD screening summary for ASP** from the right pane, a graph opens in Data Viewer, as shown in Figure 7-24.




Figure 7-24 SSD screening summary for ASP

Physical system graphs

You can now view what was collected by the IBM System Director tool for all partitions on which it was running, as shown in Figure 7-25.




Figure 7-25 Physical system statistics

This tools support all Power Systems as of V5R4.

- ▶ Processor use
- ▶ I/O counts
- ▶ I/O rates
- ▶ Disk space
- ▶ Memory

The information from the hypervisor for systems that are running release 6.1 or higher can now be viewed in the new System Graphs HMC folder (see Figure 7-26).




Figure 7-26 System Graphs (HMC) folder

It includes the following graphs:

- ▶ LPAR CPU time overview: Shows the processor time that is used for all partitions (Figure 7-27).




Figure 7-27 LPAR CPU time overview graph

- ▶ LPAR CPU time rankings: Ranks the partitions by the processor used.
- ▶ LPAR memory overview: Shows the memory consumption for all partitions (Figure 7-28).




Figure 7-28 LPAR memory overview

- ▶ LPAR memory rankings: Ranks the partitions by memory used.
- ▶ LPAR donated time rankings: Ranks the partitions by donated processor time.
- ▶ LPAR uncapped processor time used rankings: Ranks the partitions by uncapped processor time.
- ▶ Virtual shared processor pool utilization.
- ▶ Physical processor utilization overview: Shows the average processor utilization for each physical processor over time (Figure 7-29).
- ▶ LPAR dedicated processor utilization rankings.




Figure 7-29 Physical processor utilization overview

Shared memory graphs

If the QAPMSHRMP file is available, a *shared memory graphs* subfolder is available that contains the following additional graphs:

- Shared memory overview (Figure 7-30)




Figure 7-30 Shared memory overview

► Shared memory pool page faults (Figure 7-31)




Figure 7-31 Shared memory pool page faults

- Memory for I/Os overview (Figure 7-32)




Figure 7-32 Memory for I/Os

Job counts graphs

If you right-click and select **Collection** → **Job counts graphs**, the following options are available:

- Job Counts (Figure 7-33)




Figure 7-33 Job Counts

- Net jobs created
- Net jobs breakdown
- Job created / destroyed

- Job counts rankings (by job grouping; see Figure 7-34)




Figure 7-34 Job Counts by generic job name

- Net jobs breakdown rankings (by job grouping)

Memory pool graphs

If you right-click and select **Collection** → **Memory pool graphs**, you can generate graphs that contain the following information:

- Memory pool consumption (by percentage) (Figure 7-35)




Figure 7-35 Memory pool consumption

- ▶ Memory pool sizes (by percentage)
- ▶ Memory pool sizes
- ▶ Page fault rates
- ▶ Memory pool activity levels
- ▶ Transitions to ineligible rates
- ▶ State transition rates
- ▶ Machine pool sizes and rates
- ▶ DB versus non-DB faulting rates
- ▶ DB versus non-DB paging (pages read/written) rates
- ▶ 64K versus 4K page faults for pool <<JBPOOL>>

You can use the memory pool graphs to right-click the wanted pool and time range to drill down and see the jobs within the wanted pool in various ways.

There is also support to allow multiple collections to be graphed at the same time to compare the evolution in memory use. You can either select multiple collections and right-click and select the wanted memory pool graph or use the Historical Summary analysis to graph multiple collections more easily.

Disk configuration graphs

A new Disk configuration folder under the Collection menu contains information about the ASPs, IOPs, IOAs, and units on the system. This information includes information about the IOAs, including the read / write cache sizes. Three reports are provided where the first report provides a breakdown of disk capacity.

There are two more reports that show the same disk configuration data, where one is a flat table and the other is a tree. The tree provides counts and percentages of the units, IOAs, IOPS, and ASPs within each prior level grouping. To access these reports, right-click and select **Collection** → **Disk configuration**. The window that is shown in Figure 7-36 opens.

Full name	Totals	Disk unit type	Disk type description	ASP number	IOP resource name	IOA resource name	IOA CCIN	IOA write cache	IOA read cache	IOA description (by CCIN)
Total	100% - 70									
ASP 1	100% - 70									
CMB05	25.71% - 18									
DC05	25.71% - 18									
DMP030	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
DMP028	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
DMP026	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
DMP022	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
DMP018	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
DMP016	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
DMP013	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
DMP010	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
DMP052	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
DMP047	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
DMP045	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
DMP044	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
DMP039	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
DMP038	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
DMP033	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
DMP032	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
DMP007	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
DMP005	1.43% - 1	17>	Other	1	CMB05	DC05	280E			4 Gb Single-P>
CMB04	25.71% - 18									
DC04	25.71% - 18									
DMP029	1.43% - 1	17>	Other	1	CMB04	DC04	280E			4 Gb Single-P>
DMP027	1.43% - 1	17>	Other	1	CMB04	DC04	280E			4 Gb Single-P>

Figure 7-36 Disk configuration by tree

Advanced disk graphs

As explained in “QAPMDISKRB” on page 253, there is a new structure for reporting the disk response times in a new set of buckets. These new statistics can be found by right-clicking and selecting **Collection** → **Disk graph** → **Advanced**. An example is shown in Figure 7-37.




Figure 7-37 Advanced disk graphs

12X loops and I/O hubs graphs

The 12x loops and I/O hubs folder (under the Communication graphs folder) contains three styles of graphs.

- ▶ Summarized: Summarizes the loop data over time.
 - Summarized loops/hubs traffic
 - Summarized 12x loops traffic
 - Summarized I/O hub traffic
 - Total traffic breakdown
 - Send and receive traffic breakdown
 - 12x loops traffic breakdown
 - I/O hub traffic breakdown
- ▶ Rankings: Ranks each loop/port by throughput.
- ▶ Advanced: These graphs show every loop/hub (above the filter threshold) over time.

The seven summarized graphs provide drill downs into seven ranking graphs for the wanted time period.

Additional graphs

Changes were applied to a number of graphing capabilities, which supports Version 5.4 and later:

- ▶ The communication graphs folder shows the following information:
 - Average IOP uses
 - Maximum IOP uses
 - SSL authentications
- ▶ Under Disk graphs, new graphs named *I/O size and Ethernet rates* and *Buffer overruns/underruns* are now available.
- ▶ The collection overview wait graphs now show batch and interactive processor usage on the second Y axis.
- ▶ The wait bucket counts are added to the overview graphs for a single thread/job.
- ▶ The IP address family and formatted IP address is added to the end of the job search report.
- ▶ Flattened type graphs provide the capability to customize graph labels, hide values on the Y1/Y2 axis, change scaling, and use the graph definition interface to change the fields that are used, scaling, colors, and sizes.

Starting with Version 6.1, you now have the following information:

- ▶ A new graph for the selected job/user level called “Total pages allocated for <<OBJTYPE>> <<OBJDESC>>,” showing the total pages that were allocated and deallocated for the entire lifespan of the job.
- ▶ A new series of graphs under the I/O graphs folder shows the net pages that are allocated and net page frames that are requested. Net pages that are allocated are shown in megabytes and assume that the page size is 4 KB. Both sets of graphs include the usual rankings graphs to graph the data by thread, job, generic job, user, and so on.

Starting with Version 7.1, seizes and locks graphs were added over the 7.1 lock count fields in the QAPMJOBMI file. Run the Collection Summary analysis and then access the graphs by clicking **Wait graphs** and looking under the Seizes and locks folder.

7.4.8 Job Watcher

The folders available in the Job Watcher component changed. Instead of showing libraries that contain Job Watcher data, new folders are available, as shown in Figure 7-38:

- ▶ Libraries containing Job Watcher database file collections (filterable).
- ▶ A definitions folder provides a list of Job Watcher definitions on the system.
- ▶ The rest of the folders are covered in 7.4, “IBM iDoctor for IBM i” on page 263.




Figure 7-38 Job Watcher folders

Monitors

In the window to start a Job Watcher (or Disk Watcher) monitor, you can specify the maximum collection size (in megabytes) for each collection that is running in the monitor.

The next set of changes applies to the following monitor commands: **STRJWMON**, **STRPAMON**, and **STRDWMON**. These options can also be found in the GUI when you start a monitor.

The Collection Overlap (**OVRLAP**) parameter is no longer used. The monitor now detects that a new collection started before the previous one ended.

The Collection Duration (**MAXSIZE**) parameter can now be specified in minutes with a decimal point (for example, 59.5 minutes).

When you restart a monitor, if the Maximum Historical Collections (**COLNS**) parameter is reduced, there are added checks to delete the extra ones.

Deleting collections in a monitor is now done in a submitted job.

The following changes apply only to **STRJWMON**:

- ▶ Added a Resubmit Collections (**RESUBMIT**) parameter to submit new collections if a collection fails to start or quits early.
- ▶ Added a Max Consecutive Resubmits (**MAXTRIES**) parameter to indicate the number of times collections are resubmitted if the **RESUBMIT** parameter is set to *YES and the current collection ended prematurely.

Create Job Summary analysis

Right-click and select **Collection** → **Analysis** → **Create Job Summary** to produce job totals for the wanted jobs that are based on any filters, as shown in Figure 7-39.

Use this function to query job statistics for the desired collections and produce totals for each job/thread based on the filters provided.

Tip: Leave the filters blank (or at their default values) to include statistics for all jobs.

Collections available:

Library: Ckour1

Collection(s):

Collection name
Ckmon001
Ckmon002

Add >>

Collections to summarize:

Collection name
CKOUR1/CKMON001(540)
CKOUR1/CKMON002(540)

Remove Remove All

Filters (OPTIONAL):

Job name: contains

Job user name:

Job number:

Job current user profile:

Subsystem name contains:

Start time: 2010-02-19-10.13.22

End time: 2010-02-19-10.25.35

Comments:

Creation options:

Library: Ckour1

☒ Job Totals (all collections)

☒ Thread Totals (all collections)

Submit Cancel

Figure 7-39 Create Job Summary analysis

Collection Summary analysis

A collection is summarized by right-clicking and selecting **Collection** → **Analyses** → **Analyze Collection** (for full options) or by clicking **Collection** → **Run Collection Summary**. This new analysis is greatly simplified and many options that were previously on the Summarize Collection window were removed.

The Analyze Collection(s) window now has a new Situations button, which you can use to customize which situational analysis option to run, and the limits that are used for the situations (Figure 7-40). The selected analysis can now run as a batch job. You can use a check box on the Analyze Collection(s) window to indicate if this analysis is done instead of running it in the Remote SQL statement status view (a separate GUI QZDASOINIT job). The same analysis and similar options are found in the CSI component.




Figure 7-40 Analyze Collection(s) window

There are advantages to running in batch:

- ▶ You can run start many analyses simultaneously.
- ▶ You can start the analysis, end your GUI session, and it keeps running.
- ▶ You can start multiple analyses on multiple systems without waiting on the remote SQL statement status view to run them in order.

Within the list of collections, the status field indicates which files are not yet created.

Situational analysis

This option can be found by clicking **Collection** → **Wait graphs**. It has new situations:

- ▶ Concurrent write not enabled
- ▶ Journal caching that are not properly used
- ▶ Jobs ineligible to run
- ▶ Long sync write response times
- ▶ Fixed allocated length setting on a varchar or lob type column is defaulted to 0 or is set too small
- ▶ Contention on DB in use table possibly because of a high number of opens and closes
- ▶ High number of creates and deletes by multiple jobs where all of the objects are owned by the same user profile

- ▶ Potentially large number of locks
- ▶ Deadlock because of DB record locks

Top threads and tasks graphs

These graphs can be found by clicking **Collection** → **Wait graphs** → **Top threads over time**. It displays the threads or tasks that spent the most time in the wanted wait bucket (such as processor) over time.

Objects waited on tab

The Objects waited on tab within the Interval Summary interface now includes the list of jobs that are waiting on an object but did not use processing time in the interval. Previously, only jobs that used processing time in the interval were shown. There is also a check box to show segments that are waited on.

Single interval rankings identifying flags

In Job Watcher, if you drill down into a wait bucket rankings graph on a single (job/thread) interval, you see the new flags field at the end of the job/thread name that contains the following possible values and meanings:

- ▶ W: Has a wait object.
- ▶ H: Holder.
- ▶ B: The current wait bucket is the same as current sort /filter bucket.
- ▶ S: Has an SQL client job (applies to 6.1 and higher only).

A wait object, holder, and a SQL client job are added to the flyover (if one exists).

SQL server mode job information

For Job Watcher 6.1 (with PTFs) or 7.1 only, the interval details property page now includes the SQL server mode client job if found, with the option to drill down and graph the job. See Figure 7-41.




Figure 7-41 SQL client job drill-down options on the Interval Details - Wait Buckets window

Status views

The Remote SQL Statement Status view was updated with new options:

- ▶ Remove/Cancel Selected cancels or removes running SQL statements.
- ▶ Copy Selected SQL Statement(s) to Clipboard.
- ▶ Add SQL Statement defines additional statements to run in the view.

The Remote Command Status view was updated with new options:


- ▶ Remove/Cancel Selected cancels or removes running commands.
- ▶ Copy Selected Commands to Clipboard.
- ▶ Add Command defines additional commands to run in the view.

Additional reporting options

There are several new reporting options for situational analysis, similar to the call stack reports based on the selected situation. From these options, you can double-click a job or thread to get into the call stack or perform further drill down.

You can find JVM statistics on the Java virtual machine interval details tab and J9 call stacks on the Call Stack tab. The J9 Java entries are embedded within the regular Job Watcher call stacks. J9 Java call stack entries are not usable with the call stack reports.

You find a Situations tab in the Collection Properties (Figure 7-42), showing all situation types known to Job Watcher and how many occurred in the collection.



ID	Total	Name	Problem Description	Resolution Description
0	0	Multiple situations occurred		Changing the time interval size to the >
1	0	Seize/lock table	Job(s) are waiting for faulting on the s>	
2	0	Starting/ending commitment control	Job(s) appear to be constantly startin>	Change implementation of the causing>
3	0	Poorly written/performing SQL	Job(s) may be executing poorly imple>	Use Visual Explain to analyze the exe>
4	0	Missed jobs	A high percentage of jobs were NOT >	Shorten the Job Watcher collected int>
5	0	Seize contention due to data forced to disk	Job(s) are using force-end-of-data to f>	Use the N option. This option writes>
6	0	Fixed length of varchar or blob too small	Fixed allocated length setting on a var>	Choose a larger setting for the fixed al>
7	0	High number of opens/closes	Contention on 'DB in use' table possi>	Reduce number of opens
8	0	Contention on user profile	High number of creates and/or deletes>	Reduce the number of creates and del>
9	3	High synchronous write response time		Identify disk performance bottleneck>
10	0	Concurrent write support not enabled	Concurrent write support may not be>	It can be enabled by default if the file >
11	0	Journal cache could help performance	Journal cache may not be utilized.	http://publib-b.boulder.ibm.com/abstr >
12	0	Jobs ineligible to run	Jobs are ineligible to run, activity leve>	Adjust the activity level for the job.

Figure 7-42 Collection situations

In the Interval Details interface, a button was added to go to the primary thread from a secondary thread. An example of this function can be seen at:

http://www.youtube.com/watch?v=A_PZFK9Id18&feature=related

A new Call Stack Summary analysis was added to identify the call stacks, waits, and objects that are associated with the most frequently occurring call stacks that are found in the collection. You can see an example of this function at:

http://www.youtube.com/watch?v=NyycbX6eCL4&feature=more_related

Disk Watcher

The folders available in the Disk Watcher (DW) component changed. Instead of showing libraries that contain Disk Watcher data, new folders are available, as shown in Figure 7-43:

- ▶ Libraries containing DW database file collections (filterable).
- ▶ A definitions folder that provides a list of DW definitions on the system.
- ▶ The rest of the folders are covered in 7.4.4, “Main window” on page 265.




Figure 7-43 Disk Watcher

Collections

In the Start Disk Watcher Wizard (Figure 7-44), you now can collect the hardware resource file, schedule the collections, and check whether there are any PTFs. You can use another parameter in this window to set the maximum collection size (in MB) for each collection.




Figure 7-44 Start a Disk Watcher monitor

A copy function for Disk Watcher collections was added. The **CPYPFRCOL** command (shipped with the OS) may also be used for this purpose.

The Change SQL Parameters interface now has options for changing the library and collection currently being used to display the graph.

Monitors

Disk Watcher monitor server commands were added. These commands are similar to the Job Watcher monitor commands and include **STRDWMON**, **HLDDWMON**, **RLSDWMON**, **ENDDWMON**, and **DLTDWMON**.

Support was added in the GUI to work with or start monitors in Disk Watcher for either Job Watcher or Disk Watcher. The same Monitors folder is also available in Job Watcher, which you can use to work with Disk Watcher monitors from the Job Watcher component.

Definitions

New iDoctor supplied Disk Watcher definitions are available.

- ▶ QFULLO
- ▶ QFULL1MINO
- ▶ QTRCO
- ▶ QTRC1MINO

The reload IBM-supplied definitions option must be used on systems that have definitions so that the new definitions are visible.

Reporting

The graph titles match the naming convention that is used by the trace graphs. The word *pool* was changed to *disk pool*, and *disk unit* to *disk path*.

A new trace Disk Watcher drill-down menu shows the top 25 I/Os rates. The graphs in Disk Watcher trace mode include the following (Figure 7-45) information:

- ▶ I/O counts categorized totals
- ▶ I/O counts categorized writes
- ▶ I/O counts categorized reads
- ▶ I/O time categorized totals
- ▶ I/O time categorized writes
- ▶ I/O time categorized reads




Figure 7-45 Statistical graphs - By disk unit

Figure 7-46 shows an example of these charts.




Figure 7-46 Categorized I/O totals

7.4.9 PEX Analyzer

The folders available in the PEX Analyzer component were changed. Instead of showing libraries that contain PEX Analyzer data, new folders are available, as shown in Figure 7-47:

- Libraries

This folder displays libraries that contain PEX collections or libraries where active PEX collections created with the **STRPACOL** command (or the Start Collection Wizard) are running.

- Active collections

You can use this folder to work with any active PEX sessions on the system. This function is similar to the **ENDPEX** command that lists the active PEX sessions.

- PEX objects

You can use this folder to work with the PEX *MGTCOL objects on the system.

- Definitions

You can use this folder to work with PEX definitions.

- Filters

You can use this folder to work with PEX filters.

The rest of the folders are covered in 7.4.4, “Main window” on page 265.

PEX-Analyzer	Collection name	Collection library	Size (MB)	Created by	Partition collected on	Partition collected on VRM	Descrip...	Owner
Libraries	CKRUN	QUSRSYS	512.1	CKOUR	IDOCEDU	V5R4M0	PEX data	QSYS
Monitors	COL05	QUSRSYS	3.3	IDOCTOR05	IDOCEDU	V5R4M0	PEX data	QSYS
SQL tables	ID10COLL	QUSRSYS	3.3	IDOCTOR10	IDOCEDU	V5R4M0	PEX data	QSYS
Browse collector	ID10TPROF	QUSRSYS	1.8	IDOCTOR10	IDOCEDU	V5R4M0	PEX data	QSYS
Active collector	MBTIPROF	QUSRSYS	3.3	IDOCTOR01	IDOCEDU	V5R4M0	PEX data	QSYS
PEX objects	MYCOLL02	QUSRSYS	3.3	IDOCTOR02	IDOCEDU	V5R4M0	PEX data	QSYS
Super collections	SHIMETPROF	QUSRSYS	1.8	IDOCTOR01	IDOCEDU	V5R4M0	PEX data	QSYS
Definitions	SUPCKOUR	QUSRSYS	513.1	CKOUR	IDOCEDU	V5R4M0	PEX data	QSYS
Filters	SUP01	QUSRSYS	1.8	CKOUR	IDOCEDU	V5R4M0	PEX data	QSYS
Saved collections	TPROF01	QUSRSYS	1.8	IDOCTOR01	IDOCEDU	V5R4M0	PEX data	QSYS
Work management								
ASPs								
Disk units								
Objects owned by								

Figure 7-47 PEX Analyzer

Definitions

The Add PEX Definition Wizard supports defining statistics counters into buckets 5 - 8.

The PEX Analyzer Add/Change PEX Definition interface supports the latest event additions and removals at 6.1/7.1:

- Program events that are removed as of 6.1+: *MIPRECALL, *MIPOSTCALL, *JVAPRECALL, and *JVAPOSTCALL
- Base event *CPUSWT added as of 6.1+
- Base events that are added as of 7.1: *PRCF LDSUSPEND, *PRCF LDRESUME, LPARSUSPEND, and *LPARRESUME
- Storage event *CHGSEGATR added as of 7.1

- ▶ OS *ARMTRC event added as of 6.1
- ▶ Sync event *MTXCLEANUP added as of 6.1

Because collecting DASD start events is no longer necessary for PDIO analysis, the Start PEX Analyzer Collection (**STRPACOL**) command now makes sure that the *PDIO_TIME event type always collects the *READEND, *WRTEND, *RMTWRTSTR, and *RMTWRTEND events.

The **STRPACOL** command (and the Start Collection Wizard) now includes Format 2 events for all MI user problem types (*DB_OPEN, *DB_LDIO, and so on) and the Netsize problem type. Not collecting with Format 2 now requires you to create your own PEX definition.

In PEX Analyzer in the Start Collection Wizard, and when you use one of the iDoctor problem types, the default event format value for PMCO and Taskswitch is now Format 2.

When you create a collection, a QSTATSOPEN problem type collects DB opens into statistics counter #1. It runs concurrently with the QSTATSOPEN filter to ensure that only the user application program opens are counted. You can use this function to determine which programs or procedures caused the most opens by looking at the inline counter 01. The QSTATSOPEN problem type is a PEX definition that is created using ADDPEXDFN by the GUI under the covers before **STRPACOL** is ran.

You can divide a large PEX collection into a more manageable size using the Split option.

Monitors

PEX monitors can be started for systems that are licensed to both Job Watcher and PEX Analyzer. On the server side, there is a new set of commands for PEX monitors in the library QIDRWCH: **STRPAMON**, **ENDPAMON**, **HLDPAMON**, **RLSPAMON**, and **DLTPAMON**.

The Start iDoctor Monitor Wizard supports the creation of PEX monitors into *MGTCOL objects. There is an **ENDPEX** option, on the basic options screen of the wizard, with three possible values: Create DB files, Create *MGTCOL, and Suspend.

Analyses

Several changes were implemented in this menu.

Classic analyses and support for the green panel **QIDRPA/G*** analysis commands was removed and replaced by the SQL-based analyses (SQL stored procedures).

The Analyses menu, found by right-clicking a collection, contains a list of all available analyses (see Figure 7-48). The menu also contains the Analyze Collection option, which allows a user to kick off several analyses at once.




Figure 7-48 Analyses menu for a PEX collection

The TPROF analysis now has the tree table views that display the percentage of processor hits in various ways.

The TPROF analysis folder contains more reports to support MCLI analysis if format 4 PMCO events were collected.

A PDIO analysis is available and provides many graphs and drill downs for analyzing disk performance, including SSD performance. ASP prompting and comparisons are supported when you open a graph that contains data from more than one ASP.

PEX Analyzer has a new analysis called *Hot Sectors*. This SQL-based analysis is only available if the PDIO analysis was run. It allows disk activity to be measured by portions of the disk address of the I/O in megabyte chunks of either 1, 16, 256, or 4096.

A Data Area analysis is available for collections that collected data area events. It provides an SQL-based report similar to the SMTRDTAA file. A similar analysis for data queue events is available.

A CPU Profile by Job analysis is available if PMCO events were collected. It shows the estimated processor consumption during the collection over time and processor thread rankings for the wanted time periods.

The MI user event analyses (LDIO and data area) now resolve the user program if Format 2 events were collected. These analyses allow for MI entry / exit events to be excluded.

A database opens analysis, similar to the database LDIO analysis, provides statistics about the user program that is associated with the DB open events and reports 16 call level stacks, if DBOPEN FMT2 events are collected.

The new IFS analysis is equivalent to the classic version, except it also provides user program names for either MI entry / exit or FMT 2 call stacks, depending on what is available.

There is a new Netsize analysis for 6.1 and higher PEX Analyzer, including several new graphs with drill downs.

A save / restore analysis performs save / restore event parsing in the QAYPEMIUSR table into several reports.

In the Taskswitch analysis, there are added graphs that show what the wait bucket time signature looks like for the wanted thread / task (also known as TDE). See Figure 7-49. More drill-downs and reporting options are also provided.




Figure 7-49 Taskswitch run / wait time signature graph for a single job / thread / task (or TDE)

Reports

The Summarized CPU and I/O by job /pgm / MI instruction report contains two new fields:

- ▶ Inline CPU percent of job-thread total
- ▶ Inline elapsed time percent of job-thread total

The Summarized CPU and I/O by pgm / MI instruction report contains the inline processor percent of total and the inline elapsed time percent of total information.

7.4.10 Plan Cache Analyzer

Plan Cache Analyzer collects and analyzes snapshots of the system's SQL Plan Cache. It complements the features that are available in System i Navigator for analyzing the Plan Cache by providing several graphs and drill-down options that are not available there.

The plan cache is a repository that contains the access plans for queries that were optimized by SQE.

Starting Plan Cache Analyzer

Plan Cache Analyzer is included with a Job Watcher license and is a component of the IBM iDoctor for i suite of tools. iDoctor can be started by clicking **Start** → **Programs** → **IBM iDoctor for IBM i**. After the IBM iDoctor for IBM i application opens, you can start the Plan Cache Analyzer component from the Connection List View by double-clicking the correct system.

A list of available components appears in the next window. Double-click the Plan Cache Analyzer component or select Plan Cache Analyzer and click the **Launch** button to continue, as shown in Figure 7-50.




Figure 7-50 iDoctor Components window

Plan Cache Analyzer: Plan Cache Analyzer is a subcomponent of Job Watcher and is only available if Job Watcher is installed correctly and a valid access code for Job Watcher is applied. This component is included with the Job Watcher license.

For more information about how to use Plan Cache Analyzer, see the IBM iDoctor for IBM i documentation at:

[https://www-912.ibm.com/i_dir/idoctor.nsf/F204DE4F34767E0686256F4000757A90/\\$FILE/iDoctorV7R1.pdf](https://www-912.ibm.com/i_dir/idoctor.nsf/F204DE4F34767E0686256F4000757A90/$FILE/iDoctorV7R1.pdf)

7.4.11 VIOS Investigator

VIOS Investigator combines NMON data and a VIOS to IBM i disk mapping process to help analyze the performance of your VIOS using the power of the DB2 database on IBM i.

You can use VIOS Investigator to import one or more NMON files into the tool. The NMON CSV files are converted and expanded into DB2 SQL tables, which are used to produce graphs with several drill-down options.

Graphing Options:

- ▶ Disk graphs (% busy, counts, sizes, rates, block sizes, service times, and response times)
- ▶ System configuration
- ▶ System graphs (Processor, memory, kernel, paging statistics, and processes)
- ▶ Processor graphs (Processor usage)
- ▶ TOP graphs (Processor usage, paging size, character IO, memory usage, and faults for the top processes)

If a valid disk mapping was created, then you can use the disk graphs to rank the data by disk name, disk unit, disk path, ASP, or disk type. Without the disk mapping, then only rankings by disk name can be performed.

VIOS Investigator can also be used to analyze AIX and Linux systems using NMON data, but the focus is primarily on VIOS analysis with an emphasis on usage by IBM i customers.

VIOS Investigator is a no additional cost tool that is offered *as-is* and does not require an access code. To download VIOS Investigator, you must first accept the license agreement.

NMON

The VIOS Investigator data is created by the NMON or **Topas_NMON** command that is found in AIX.

On AIX V6.1 TL02 and Virtual I/O Server (VIOS) V2.1 (or higher), NMON is installed by default with AIX and the **Topas_NMON** command should be used for collecting data for use with VIOS Investigator.

NMON is the primary/preferred collection tool of AIX performance statistics. NMON is similar in nature to Collection Services on IBM i. Both tools use time intervals and collect high-level statistics for processor usage, disk, memory and much more.

NMON data is collected into CSV (comma-separated values) files, and Collection Services uses DB2 tables. CSV files are difficult to analyze, especially if you want to analyze many of them. VIOS Investigator simplifies this issue by allowing users to import and then analyze multiple NMON files at once.

Disk mappings

Disk mappings (also known as *correlations*) refer to the VIOS to IBM i correlations between hDisks and IBM i disk unit numbers and disk path names assigned to them.

For comparison purposes with the Collection Services Investigator, whenever possible, the disk graphs in VIOS Investigator use the same colors, labels, and field names as the disk graphs in Collection Services Investigator. However, the number of disk metrics that are provided by NMON are far fewer than those disk metrics found in Collection Services (see the QAPMDISK file.)

In most cases, especially if there are any known hardware/system changes, collect the disk mapping before or immediately after you collect the NMON data on your VIOS. This action provides more graphing options (for example, rankings by unit, path, ASP, or disk type) that are otherwise not available.

Disk mappings are collected using a program that is written on the IBM i that interrogates the HMC to determine more disk information that is useful when you perform the analysis that otherwise would not be available with just the NMON data.

Starting VIOS Investigator

VIOS Investigator is a component of the IBM iDoctor for IBM i suite of tools. iDoctor can be started by clicking **Start** → **Programs** → **IBM iDoctor for IBM i**. After the IBM iDoctor for IBM i application opens, the VIOS Investigator component is started from the Connection List View by double-clicking the system.

A list of available components appear in the next window. Double-click the VIOS Investigator component or select **VIOS Investigator** and click the **Launch** button to continue, as shown in Figure 7-51.




Figure 7-51 iDoctor Components window

For more information about how to use VIOS Investigator, see the IBM iDoctor for IBM i documentation at:

[https://www-912.ibm.com/i_dir/idoctor.nsf/F204DE4F34767E0686256F4000757A90/\\$FILE/iDoctorV7R1.pdf](https://www-912.ibm.com/i_dir/idoctor.nsf/F204DE4F34767E0686256F4000757A90/$FILE/iDoctorV7R1.pdf)

7.4.12 iDoctor FTP GUI

A no additional cost GUI exists that provides FTP access to IBM i or other types of FTP servers. This GUI was created primarily for use with VIOS Investigator to ease the sending and receiving of performance data (such as NMON) to and from AIX (or VIOS) systems. It was tested only with connections to IBM i, AIX, or VIOS systems.

FTP connections: The FTP connections are provided through the Windows WININET APIs, which do *not* support any options for FTP using SSL or other secure FTP modes.

The FTP GUI is similar to other iDoctor components and can be accessed through the My Connections view.

From the My Connections view, you can access this option by right-clicking the correct system and selecting **Start FTP session** from the menu, as shown in Figure 7-52.




Figure 7-52 Start FTP Session

Or, you can access the option from the Connection List View by double-clicking the system.

A list of available components appears in the next window. Double-click the iDoctor FTP GUI component or select **iDoctor FTP GUI** and click the **Launch** button to continue, as shown in Figure 7-53.




Figure 7-53 iDoctor Components window

The FTP GUI window opens, as shown in Figure 7-54.




Figure 7-54 iDoctor FTP GUI window

IBM i system: If you connect to an IBM i system, you see subfolders to work with (either the Integrated File System (IFS) or the Libraries on the system). These options are not present for other types of connections.

There are several options available to you when you right-click a file or folder from the FTP GUI, including Upload, Download to PC, and Transfer to another system.

For more information about how to use the iDoctor FTP GUI, see the IBM iDoctor for IBM i documentation at:

[https://www-912.ibm.com/i_dir/idoctor.nsf/F204DE4F34767E0686256F4000757A90/\\$FILE/iDoctorV7R1.pdf](https://www-912.ibm.com/i_dir/idoctor.nsf/F204DE4F34767E0686256F4000757A90/$FILE/iDoctorV7R1.pdf)

7.4.13 More information

For more information about the new features in iDoctor, go to:

http://www-912.ibm.com/i_dir/idoctor.nsf/

Presentations are created every few months with in-depth explanations of the latest features. You can find these presentations at:

http://www-912.ibm.com/i_dir/idoctor.nsf/downloadsDemos.html

There is a YouTube Channel (20+ videos) for IBM iDoctor at:

<http://www.youtube.com/user/IBMiDoctor?feature=mhum>

These videos are also available at <http://www.ibm.com> if your company blocks YouTube. Click the appropriate links from the Video Library page at:

https://www-912.ibm.com/i_dir/idoctor.nsf/videos.html



Virtualization

This chapter describes the following topics:

- ▶ PowerVM enhancements
- ▶ Additional OS level combinations of server and client logical partitions
- ▶ Hardware Management Console virtual device information
- ▶ IBM i hosting IBM i
- ▶ Virtual Partition Manager enhancements
- ▶ Partition suspend and resume
- ▶ HEA Daughter cards
- ▶ 10 Gb FCoE PCIe Dual Port Adapter
- ▶ Live Partition Mobility

8.1 PowerVM enhancements

The following sections summarize recent PowerVM enhancements for IBM i.

8.1.1 Active memory sharing

Active memory sharing (AMS) enables the sharing of a pool of physical memory among IBM i, AIX, and SUSE Linux logical partitions on a single IBM Power Systems server POWER6 or later, helping to increase memory use and drive down system costs. The memory is dynamically allocated among the partitions as needed, to optimize the overall physical memory usage in the pool. Instead of assigning a dedicated amount of physical memory to each logical partition that uses shared memory (referred to as *shared memory partitions*), the hypervisor provides the physical memory from the shared memory pool to the shared memory partitions as needed. The Power hypervisor provides portions of the shared memory pool that are not being used by shared memory partitions to other shared memory partitions that must use the memory.

When a shared memory partition needs more memory than the current amount of unused memory in the shared memory pool, the hypervisor stores a portion of the memory that belongs to the shared memory partition in an auxiliary storage space that is known as a *paging space device*. Access to the paging space device is provided by a Virtual I/O Server (VIOS) logical partition that is known as the *paging service partition*. When the operating system of a shared memory partition accesses data that is in a paging space device, the hypervisor directs the paging service partition to retrieve the data from the paging space device and write it to the shared memory pool so that the operating system can access the data.

For an illustration of these AMS concepts, see Figure 8-1.




Figure 8-1 AMS concepts

The PowerVM Active Memory Sharing technology is available with the PowerVM Enterprise Edition hardware feature, which also includes the license for the VIOS software.

Paging service partitions must be VIOS. Logical partitions that provide virtual I/O resources to other logical partitions can be VIOS or IBM i and must be dedicated memory partitions, but their client partitions are shared memory partitions.

Important: Logical partitions that have dedicated physical resources cannot be shared memory partitions.

In general terms, the setup of ASM includes using the HMC to create a shared memory pool, selecting a paging service partition, selecting a paging space device, and changing the IBM i client partition profile to use shared memory pool. See Figure 8-2.




Figure 8-2 ASM setup using the HMC

You can configure two paging service partitions to access the same, or common, paging space devices. In this configuration, the two paging service partitions provide redundant access to the paging space devices (see Figure 8-3). This function is known as *redundant paging service* partitions. When one paging service partition becomes unavailable, the hypervisor sends a request to the other paging service partition to retrieve the data on the paging space device. For more information about *redundant VIOS partitions* support, see 8.1.8, “Redundant VIOS partitions support” on page 322.




Figure 8-3 Redundant paging service partitions

For IBM i client partitions where the disk storage is virtualized using VIOS partitions and storage area network (SAN) Disk Storage, NPIV and multipath I/O support is available with IBM i 6.1.1 or later. For more information about NPIV, see 8.1.3, “PowerVM Virtualization and I/O enhanced with NPIV” on page 318. For multipath I/O for IBM i client partitions, see 9.2.1, “Multipathing for virtual I/O” on page 386.

Requirement: When you use redundant paging service partitions, common paging space devices must be on SAN Disk Storage to enable symmetrical access from both paging service partitions.

The system requirements for AMS are as follows:

- ▶ IBM Power Systems server or blade with POWER6 processors
- ▶ Virtual I/O Server (VIOS) 2.1.0.1 Fix Pack 21 or later
- ▶ System Firmware level 340_075 or later
- ▶ HMC V7.342 or later
- ▶ IBM i 6.1 plus PTF SI32798 or later
- ▶ AIX V6.1 TL3
- ▶ SUSE Linux Enterprise Server 11

Solid-state disk usage: A solid-state disk (SSD) on VIOS can be used as a shared memory pool paging space device. For more information, see 9.3, “SSD storage management enhancements” on page 398.

For an overview of AMS, go to:

<http://www.ibm.com/systems/power/software/virtualization/whitepapers/ams.html>

For more detailed information about AMS, see *IBM PowerVM Virtualization Active Memory Sharing*, REDP-4470.

8.1.2 Enhanced support for IBM System Storage

On October 2009, IBM announced support for IBM System Storage DS5100 and IBM System Storage DS5300 through native Fibre Channel attachment to POWER6 processor-based servers and IBM i 6.1.1. This announcement also included IBM i 6.1 and 5.4 support for the IBM System Storage DS8700 models 941 and 94E enterprise storage system through native Fibre Channel attachment to POWER5 and POWER6 processor-based servers. In addition, the DS8700 and the IBM System Storage DS5020 are supported by IBM i 6.1 partitions through PowerVM VIOS attached to POWER6 processor-based servers and blades. For a summary of SAN Disk Storage alternatives by IBM i, see Figure 8-4. For more information about storage area networks and IBM i, see Chapter 9, “Storage and solid-state drives” on page 367.

		SVC	DS3000	DS4000	XIV	DS5020	DS5100	DS5300	DS6800	DS8000
Power Systems	IBM i Version	6.1 & POWER6/7	6.1 & POWER6/7 (DS3400)	6.1 & POWER6/7	6.1 & POWER6/7	6.1 & POWER6/7	**6.1 & *6.1.1 POWER 6/7	**6.1 & *6.1.1 POWER 6/7	5.4 & 6.1 POWER 5/6/7	5.4 & 6.1 POWER 5/6/7
	IBM i Attach	VIOS	VIOS	VIOS	VIOS	VIOS	Direct* or VIOS	Direct* or VIOS	Direct	Direct or VIOS
Power Blades	IBM i 6.1 (through VIOS)	Yes (BCH)	Yes BCS DS3200 BCH DS3400 and DS3200	Yes (BCH)	Yes (BCH)	Yes (BCH)	Yes (BCH)	Yes (BCH)	No	Yes (BCH)

Figure 8-4 SAN Disk Storage alternatives for IBM i 5.4, 6.1, and 6.1.1

Support availability: SAN Disk Storage with N_Port ID Virtualization (NPIV) support is available on DS8000 with POWER6 or later (it requires IBM i 6.1.1 or later). Power Blades support the QLogic 8 Gb Blade HBAs to attach DS8100, DS8300, and DS8700 storage systems through NPIV. For more information about NPIV, see 8.1.3, “PowerVM Virtualization and I/O enhanced with NPIV” on page 318.

For an overview of IBM i System Storage solutions, see *IBM i Virtualization and Open Storage Readme First*, found at:

http://www-03.ibm.com/systems/resources/systems_i_Virtualization_Open_Storage.pdf

For more information about available SAN Storage solutions for Power Systems and IBM i, see the *System Storage Interoperation Center* at:

<http://www-03.ibm.com/systems/support/storage/config/ssic/>

For more information about available SAN Storage solutions for BladeCenter and IBM i, see *BladeCenter Interoperability Guide*, found at:

<http://www-947.ibm.com/systems/support/supportsite.wss/docdisplay?ln docid=MIGR-5073016&brandind=5000020>

For more information about IBM i on Power Blades, see the *Read me first* and *Supported Environments* documents, found at:

<http://www-03.ibm.com/systems/power/hardware/blades/ibmi.html>

8.1.3 PowerVM Virtualization and I/O enhanced with NPIV

N_Port ID Virtualization (NPIV) is an industry-standard Fibre Channel (FC) protocol that allows the VIOS to share an NPIV-capable FC adapter among multiple client partitions. For NPIV, the VIOS server acts as an FC pass-through instead of a SCSI emulator, such as when you use Virtual SCSI (see Figure 8-5).




Figure 8-5 Comparing PowerVM storage virtualization with VSCSI and NPIV

With NPIV, a port on the physical FC adapter is mapped to a Virtual Fibre Channel (VFC) server adapter in VIOS, which in turn is mapped to a VFC client adapter in the IBM i client partition, as shown in Figure 8-6.




Figure 8-6 VIOS VFC server adapter and IBM i VFC client adapter

Two unique worldwide port names (WWPNs) are generated for the VFC client adapter, which is available on the SAN so that storage can be mapped to them as you can to any other FC ports. The following issues must be considered when you use NPIV:

- ▶ There is one VFC client adapter per physical port per partition, to avoid a single point of failure.
- ▶ A maximum of 64 active VFC client adapters are permitted per physical port. This number can be less because of other VIOS resource constraints.
- ▶ There can be only 32,000 unique WWPN pairs per system platform.
 - Removing an adapter does not reclaim WWPNs. Can be manually reclaimed through the CLI (`mksyscfg`, `chhwres`, and so on) or through the “virtual_fc_adapters” attribute.
 - If the capacity is exhausted, you must purchase an activation code for more capacity.

Important: Only one of the two WWPN ports is used (port 0). The second WWPN port is not used.

IBM i 6.1.1 supports NPIV, providing direct Fibre Channel connections from i 6.1.1 client partitions to SAN resources. The IBM i clients see the SAN resources with their native device type as though they were natively attached (see Figure 8-7).

Logical Hardware Resources Associated with IOP				
Type options, press Enter.				
2=Change detail		4=Remove	5=Display detail	6=I/O debug
7=Verify		8=Associated packaging resource(s)		
Opt	Description	Type-Model	Status	Resource Name
-	Virtual IOP	6B25-001	Operational	CMB02
■	Virtual Storage IOA	6B25-001	Operational	DC02
-	Disk Unit	2107-A85	Operational	DD004
-	Disk Unit	2107-A85	Operational	DD002
-	Tape Library	3584-032	Operational	TAPMLB02
-	Tape Unit	3580-003	Operational	TAP01
F3=Exit F5=Refresh F6=Print F8=Include non-reporting resources F9=Failed resources F10=Non-reporting resources F11=Display serial/part numbers F12=Cancel				

Figure 8-7 SAN resources as seen by IBM i client partitions when you use NPIV

The 6B25-001 shows a single port (0). The worldwide port name is how the SAN recognizes the Virtual IOA, as shown in Figure 8-8.

Auxiliary Storage Hardware Resource Detail	
Description	Virtual Storage IOA
Type-model	6B25-001
Status	Operational
Serial number	00-00000
Part number	
Resource name	DC02
Port	0
Worldwide port name	C0507600024D0038
SPD bus	
System bus	255
System board	128
System card	9
Storage	
I/O adapter	
I/O bus	127
Controller	
Device	
Bottom	
F3=Exit F5=Refresh F6=Print F9=Change detail F11=Display additional port information F12=Cancel	

Figure 8-8 Virtual Storage IOA 6B25-001 details

With POWER6 or later processor-based servers and Power Blades running IBM i 6.1.1 and PowerVM VIOS 2.1.2 FP22.1, NPIV support is available for the following tape libraries:

- ▶ 3573 (TS3100 and TS3200) with LTO 4 tape drives
- ▶ 3584 (TS3500) with LTO 4 tape drives

With POWER6 or later processor-based servers and Power Blades running IBM i 7.1 or IBM i 6.1.1 with MF47832 and MF48674 and PowerVM VIOS 2.1.3 or later, NPIV support is available for the following tape libraries:

- ▶ 3573 (TS3100 and TS3200) with LTO3 and LTO 4 tape drives
- ▶ 3584 (TS3500) with LTO 3, LTO 4, TS1120, TS1130, and 3592-J1A tape drives
- ▶ 3577 (TS3400) with (TS1120/TS1130) drives
- ▶ 3576 (TS3310) with LTO 3 or LTO 4 drives

8.1.4 Expanded HBA and switch support for NPIV on Power Blades

Power Blades running PowerVM VIOS 2.2.0 with IBM i 7.1 partitions support the QLogic 8 Gb Blade HBAs to attach DS8100, DS8300, and DS8700 storage systems through NPIV. This support allows easy migration from existing DS8100, DS8300, and DS8700 storage to a blade environment. Full PowerHA support is also available with virtual Fibre Channel and the DS8100, DS8300, and DS8700, which includes Metro Mirroring, Global Mirroring, FlashCopy, and LUN level switching.

8.1.5 PowerVM N_Port ID Virtualization attachment of DS5000

IBM i 7.1 partitions on POWER6 or POWER7 rack and tower systems now support N_Port ID Virtualization attachment of DS5100 and DS5300 storage systems. Setting up configurations to share adapters is simpler with NPIV. This support also allows the usage of a Lab Services toolkit to access copy services for the DS5000 storage.

For compatibility information, consult the Storage Systems Interoperability Center at:

<http://www.ibm.com/systems/support/storage/ssic/interoperability.wss>

8.1.6 Enhanced mirroring algorithms

IBM i mirroring algorithms are enhanced to consider any N_Port ID Virtualization (NPIV) attached disks. The locations of the virtual disks are considered when the pairs of mirror disk units are calculated bus level statistics for 12x loops.

8.1.7 PCIe2 Riser Card (Gen2) (#5685) direct support

IBM i now provides direct support for the PCIe Riser Card (Gen2), without the use of VIOS, with IBM i 7.1. This riser card is used in the IBM Power 720 and IBM Power 740 Express to provide greater bandwidth with a smaller number of PCI slots or adapters. It is an optional feature, containing four PCIe Gen2 low-profile (LP) slots for Gen1 or Gen2 adapters. It is physically similar to the PCIe Riser Card (Gen1) (#5610).

For more information about N_Port ID Virtualization (NPIV) for IBM i, see the *IBM i Virtualization and Open Storage Read-me First* topic, found at:

http://www-03.ibm.com/systems/resources/systems_i_Virtualization_Open_Storage.pdf

For more information about SAN Storage solutions for Power Systems and IBM i, see the System Storage Interoperation Center at:

<http://www-03.ibm.com/systems/support/storage/config/ssic/>

For more information about SAN Storage solutions for BladeCenter and IBM i, see the *BladeCenter Interoperability Guide*, found at:

<http://www-947.ibm.com/systems/support/supportsite.wss/docdisplay?ln docid=MIGR-5073016&brandind=5000020>

For more information about IBM i on POWER processor-based blades, see the *Read me first* and *Supported Environments* documents, found at:

<http://www-03.ibm.com/systems/power/hardware/blades/ibmi.html>

8.1.8 Redundant VIOS partitions support

For enhanced availability in a PowerVM VIOS environment, IBM i 6.1.1 or later client partitions can be configured in multipath configurations where one partition uses redundant VIOS partitions to connect to the same IBM System Storage device, as shown in Figure 8-9.

Important: Redundant VIOS support is available on POWER6 or later processor-based servers. It is not supported by BladeCenter and Power Blades.




Figure 8-9 Redundant VIOS using VSCSI

IBM i 6.1.1 or later IBM i clients support Redundant VIOS partitions and N_Port ID Virtualization (NPIV) for attachment to IBM System Storage DS8000 solutions, as shown in Figure 8-10.




Figure 8-10 Redundant VIOS partitions using NPIV

For more information about Redundant VIOS partitions, see the *IBM i Virtualization and Open Storage Read-me First* topic, found at:

http://www-03.ibm.com/systems/resources/systems_i_Virtualization_Open_Storage.pdf

8.1.9 Shared storage pools

Shared storage pools are a new capability that is available with Virtual I/O Server Version 2.2.0.11, Fix Pack 24, Service Pack 1. Shared storage pools provide the following benefits:

- ▶ Simplifies the aggregation of many disks across multiple Virtual I/O Servers.
- ▶ Improves the usage of the available storage.
- ▶ Simplifies administration tasks.

Shared storage pool architecture overview

A shared storage pool is a pool of SAN storage devices that can span multiple Virtual I/O Servers. It is based on a cluster of Virtual I/O Servers and a distributed data object repository with a global namespace. Each Virtual I/O Server that is part of a cluster represents a cluster node.

When you use shared storage pools, the Virtual I/O Server provides storage through logical units that are assigned to client partitions. A logical unit is a file backed storage device that is in the cluster file system in the shared storage pool. It appears as a virtual SCSI disk in the client partition, in the same way as a, for example, virtual SCSI device that is backed by a physical disk or a logical volume.

Prerequisites

A shared storage pool requires the following prerequisites:

- ▶ POWER6 (and above) based servers (including blades).
- ▶ PowerVM Standard Edition or PowerVM Enterprise Edition.
- ▶ Virtual I/O Server requirements:
 - Version 2.2.0.11, Fix Pack 24, Service Pack 1, or later.
 - Processor entitlement of at least one physical processor.
 - At least 4 GB memory.
- ▶ Client partition operating system requirements:
 - IBM AIX 5L™ V5.3 or later.
 - IBM i 6.1.1 or later with the latest PTF.
- ▶ Local or DNS TCP/IP name resolution for all Virtual I/O Servers in the cluster.
- ▶ Minimum storage requirements for the shared storage pool:
 - One Fibre Channel attached disk that acts as a repository, with at least 20 GB of disk space.
 - At least one Fibre Channel attached disk for shared storage pool data. Each disk must have at least 20 GB of disk space.
- ▶ All physical volumes for the repository and the shared storage pool must have redundancy at the storage level.

Virtual I/O Server storage clustering model

The Virtual I/O Servers that are part of the shared storage pool are joined to form a cluster. A Virtual I/O Server that is part of a cluster is also referred to as *cluster node*. Only Virtual I/O Server partitions can be part of a cluster.

The Virtual I/O Server clustering model is based on Cluster Aware AIX (CAA) and RSCT technology. The cluster for the shared storage pool is an RSCT Peer Domain cluster. Therefore, a network connection is needed between all the Virtual I/O servers that are part of the shared storage pool.

Each Virtual I/O Server in the cluster requires at least one physical volume for the repository that is used by the CAA sub system and one or more physical volumes for the storage pool.

All cluster nodes in a cluster can see all the disks. Therefore, the disks must be zoned to all the cluster nodes that are part of the shared storage pools. All nodes can read and write to the shared storage pool. The cluster uses a distributed lock manager to manage access to the storage.

The Virtual I/O Servers in the cluster communicate with each other using Ethernet connections. They share the repository disk and the disks for the storage pool through the SAN.

For an abstract image of a shared storage pool, see Figure 8-11.




Figure 8-11 Abstract image of the clustered Virtual I/O Servers

Shared storage pool layout

The shared storage pool manages logical units as a file. Portions of the logical unit are cached on the client node in a cluster. The logical unit consists of virtual blocks and has a virtual block address space.

The physical volumes in the shared storage pool are managed as an aggregation of physical blocks and user data is stored in these blocks. These physical blocks are managed by a metadata area on the physical volumes. Therefore, the physical volumes in the shared storage pool consist of physical blocks and have a physical block address space.

The translation from a virtual block address to a physical block address is performed by the Virtual Address Translation Lookaside (VATL).

The system reserves a small amount of each physical volume in the shared storage pool to record metadata. The remainder of the shared storage pool capacity is available for client partition user data. Therefore, not all of the space of physical volumes in the shared storage pool can be used for user data.

Thin provisioning

A thin-provisioned device represents a larger image than the actual physical disk space it is using. It is not fully backed by physical storage if the blocks are not in use.

A thin-provisioned logical unit is defined with a user-specified size when it is created. It appears in the client partition as a virtual SCSI disk with that user-specified size. However, on a thin-provisioned logical unit, blocks on the physical disks in the shared storage pool are allocated only when they are used.

Compared to a traditional storage device, which allocates all the disk space when the device is created, this situation can result in significant savings in physical disk space. It also allows overcommitting of the physical disk space.

For an overview of thin provisioning of a shared storage pool, see Figure 8-12.




Figure 8-12 Thin-provisioned devices in the shared storage pool

8.2 Additional OS level combinations of server and client logical partitions

IBM PowerVM continues to enable Power Systems with IBM i to achieve higher resource usage by supporting additional OS level combinations of server and client logical partitions, as shown in Figure 8-13.




Figure 8-13 Additional OS level combinations of server and client for IBM i and VIOS

- ▶ IBM i 6.1 server partition can provide virtual I/O resources to the following elements:
 - IBM i 6.1 and 7.1 or later client partitions
 - AIX V5.2, V5.3, and V6.1, and SLES and Red Hat Linux client partitions
 - iSCSI-attached IBM System x and BladeCenter
- ▶ IBM i 7.1 server partition can provide virtual I/O resources to the following elements:
 - IBM i 6.1 and 7.1 or later client partitions
 - AIX V5.2, V5.3, and V6.1, and SLES and Red Hat Linux client partitions
 - iSCSI attached System x and BladeCenter
- ▶ PowerVM VIOS 2.1.3 server partition can provide virtual I/O resources to the following elements:
 - IBM i 6.1 and 7.1 or later client partitions
 - AIX and Linux client partitions

The following list describes the benefits of using IBM i hosting IBM i:

- ▶ Uses the same technology as IBM i hosting AIX, Linux, and iSCSI x86 servers.
- ▶ Uses the existing hardware investment. You can create new IBM i partitions using only virtual hardware (no IOAs, IOPs, disk units, or I/O slots are necessary for client partitions), but you can also use physical I/O.
- ▶ Rapidly deploy new workloads.
 - You can create a virtual disk with one command or several clicks in IBM Navigator for i.
 - You can deploy new partitions and virtual resources dynamically.
 - You can create test environments without hardware provisioning.
 - Virtual resources allow new test environments of exact sizes to be created and deleted without moving hardware.
 - You can test new applications, tools, and fixes in a virtual test partition.
 - You can test the next release in the client partition.

For more information about PowerVM, see *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940.

For more information about IBM i client partitions, see the IBM i Information Center at:
<http://publib.boulder.ibm.com/iseriess/>

8.3 Hardware Management Console virtual device information

Virtual device information is now available on the Hardware Management Console (HMC) for VIOS logical partitions. The HMC can now display a list of the virtual SCSI adapters for a given VIOS logical partition, as shown in Figure 8-14.




Figure 8-14 Virtual Device Information in the HMC

Tip: To access the Virtual Device Information in the HMC, click **Server** → **Partitions** → **VIOS partition** → **Hardware Information** → **Virtual I/O Adapters** → **SCSI**.

The following list describes the information that is displayed in the HMC:

- ▶ **Virtual Adapter**

This column displays the name of each virtual server SCSI adapter.

- ▶ **Backing Device**

This column displays the name of the storage device whose storage capacity can be used through a virtual SCSI connection to this virtual server SCSI adapter. This storage device is on the same logical partition as the virtual server SCSI adapter.

- ▶ **Remote Partition**

This column displays the partition name and partition ID (in parentheses) of the logical partition to which each virtual server SCSI adapter is set to connect. If this column is blank, then the virtual server SCSI adapter is set to connect to any logical partition.

► Remote Adapter

This column displays the virtual slot ID of the virtual client SCSI adapter to which each virtual server SCSI adapter is set to connect. If this column contains none, then the virtual server SCSI adapter is set to connect to any virtual client SCSI adapter.

► Remote Backing Device

This column displays the name of the virtual disks (or logical volumes) that display on the logical partition with the virtual client SCSI adapter when a virtual SCSI connection exists. The logical partition with the virtual client SCSI adapter can use these virtual disks to store information about the storage device that is owned by the logical partition with the virtual server SCSI adapter. This column contains a value only if the virtual server SCSI adapter is connected to a virtual client SCSI adapter.

Consideration: You can create virtual server SCSI adapters only for Virtual I/O Server and IBM i logical partitions. This window is always blank for AIX and Linux logical partitions.

The following list details the requirements for virtual device information:

- POWER6 or later rack / tower systems
- BladeCenter H
- System firmware level 350_038 or later
- HMC V7.3.5 or later
- VIOS V2.1.2 (FP 22.1) or later
- IBM i 6.1.1 or later (+latest fixes)

Similar information is available by running **lshwres** on the HMC or Integrated Virtualization Manager (IVM) by using the new attribute topology on the **-F** flag.

Example 8-1 lists the Virtual SCSI Adapter attributes in the form of a slash delimited list.

Example 8-1 List Virtual SCSI Adapter attributes

```
#Command:
lshwres -m <system name> -r virtualio --rsubtype scsi --level lpar -F
lpar_name,remote_lpar_name,topology
#Results:
za6bp10,za6bvios2,"OPT01/Active/DC01/vhost6//"/var/vio/VMLibrary/slic611190004AMSTAPE.iso",TAP
01/Active/DC01/vhost6//rmt1"
za6bp11,za6bvios2,"OPT01/Active/DC01/vhost9//"/var/vio/VMLibrary/slic611190004AMSTAPE.iso""
za6bp12,za6bvios2,"OPT01/Active/DC01/vhost10//"/var/vio/VMLibrary/slic611190004AMSTAPE.iso",DP
H001/Active/DC01/vhost10//hdisk28"
za6bp15,za6bvios2,"OPT01/Active/DC01/vhost0//"/var/vio/VMLibrary/WindowsServer2003.iso",DD006/
Active/DC01/vhost0//hdisk29,DD001/Missing/DC01/vhost0//hdisk29"
za6bvios,za6bp4,///vhost0//hdisk2
za6bvios,za6bp4,///vhost1//hdisk3
za6bvios2,za6bp6,///vhost3//"/var/vio/VMLibrary/xpf710_370_B292403.iso""
za6bvios2,za6bp13,///vhost4//,///vhost4//hdisk36
```

Example 8-2 lists the Virtual Fibre Channel Adapters attributes for each logical partition in the form of a slash delimited list.

Example 8-2 Virtual Fibre Channel Adapters attributes

```
#Command:
lshwres -m <system name> -r virtualio --rsubtype fc --level lpar -F
lpar_name,remote_lpar_name,topology
#Results:
za6bp10,za6bvios3,/Active/DC04/vfchost7/fcs0
za6bp10,za6bvios2,"/Active/DC03/vfchost4/fcs4,DD002/Missing/DC03/vfchost4/fcs4"
za6bp10,za6bvios3,"/Active/DC02/vfchost0/fcs1,DD001/Active/DC02/vfchost0/fcs1"
za6bp11,za6bvios2,"/Active/DC03/vfchost7/fcs4,DD001/Active/DC03/vfchost7/fcs4"
za6bp11,za6bvios3,/Active/DC02/vfchost2/fcs1
za6bp12,za6bvios2,"/Active/DC03/vfchost8/fcs4,DD001/Active/DC03/vfchost8/fcs4,DD003/Active/DC03/
vfchost8/fcs4"
za6bp12,za6bvios3,/Active/DC02/vfchost4/fcs1
za6bp13,za6bvios2,unavailable
za6bp13,za6bvios2,unavailable
za6bp15,za6bvios2,"/Active/DC02/vfchost1/fcs3,DD002/Active/DC02/vfchost1/fcs3,DD007/Active/DC02/
vfchost1/fcs3"
za6bvios2,za6bp13,///vfchost5/fcs3
za6bvios2,za6bp10,///vfchost4/fcs4
za6bvios2,za6bp6,///vfchost3/fcs3
za6bvios2,za6bp18,///vfchost2/fcs3
za6bvios2,za6bp13,///vfchost13/fcs4
```

Example 8-3 lists the Virtual Ethernet Adapter attributes.

Example 8-3 Virtual Ethernet Adapter attributes

```
#Command:
lshwres -m <system name> -r virtualio --rsubtype eth --level lpar -F
lpar_name,connect_status,device_name,drc_name,shared_adapter,backing_device
#Results:
za6bp10,active,CMN01,U9406.675.107074C-V10-C2-T1,,
za6bp10,none,CMN02,U9406.675.107074C-V10-C3-T1,,
za6bp12,active,CMN01,U9406.675.107074C-V12-C2-T1,,
za6bp12,active,CMN02,U9406.675.107074C-V12-C3-T1,,
za6bp15,active,CMN03,U9406.675.107074C-V15-C2-T1,,
za6bp15,active,CMN04,U9406.675.107074C-V15-C3-T1,,
za6bvios,active,ent2,U9406.675.107074C-V16-C11-T1,ent4,ent0
za6bvios,active,ent3,U9406.675.107074C-V16-C12-T1,,
za6bvios2,active,ent2,U9406.675.107074C-V17-C11-T1,ent4,ent0
```

For more information about the **lshwres** command, go to the Hardware Information Center at:

<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/iphcg/lshwres.htm>

8.4 IBM i hosting IBM i

An IBM i 6.1 or 7.1 partition can host one or more IBM i partitions, which are known as virtual client partitions. Virtual client partitions typically have no physical I/O hardware that is assigned and instead use virtual I/O resources from the host IBM i partition. The types of hardware resources that can be virtualized by the host partition are disk, tape, optical, and networking. The capability of IBM i to provide virtual I/O resources was used successfully for several years to integrate AIX, Linux, and Microsoft Windows workloads on the same platform. The same virtualization technology, which is part of the IBM i operating system, can now be used to host IBM i partitions, as shown in Figure 8-15.




Figure 8-15 IBM i hosting IBM i components

IBM i hosting IBM i uses an existing function of the system firmware, or IBM Power Hypervisor, which can create VSCSI (VSCSI) and Ethernet adapters in a partition. Virtual adapters are created for each partition in the Hardware Management Console (HMC) or virtual server in the Systems Director Management Console (SDMC). VSCSI adapters are used for storage and optical virtualization; virtual Ethernet adapters are used for network virtualization.

Requirement: POWER6 or later and IBM i 6.1 or later is required to support IBM i client partitions.

Tip: VIOS server partitions can also virtualize a natively attached storage device to IBM i 6.1 or later client partitions. For more information, see *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590.

8.4.1 Disk virtualization

To virtualize integrated disks (SCSI, SAS, or SSD) or LUNs from a SAN system to an IBM i client partition or virtual server, both HMC/SDMC and IBM i objects must be created:

- One VSCSI server adapter on the host partition
- One VSCSI client adapter on the client partition

This VSCSI adapter pair allows the client partition to send read and write I/O operations to the host partition. More than one VSCSI pair can exist for the same client partition in this environment. To minimize the performance impact on the host partition, the VSCSI connection is used to send I/O requests, but not for the actual transfer of data. Using the capability of the POWER Hypervisor for Logical Remote Direct Memory Access (LRDMA), data is transferred directly from the physical adapter that is assigned to the host partition to a buffer in memory of the client partition.

To create Virtual SCSI adapters for the IBM i host partition and IBM i client partition, complete the following steps:

1. Use the managing HMC to create a VSCSI server adapter on the IBM i host partition:
 - a. In the navigation pane, click **Systems Management** → **Servers**, and click the managed system on which the server IBM i host partition is.
 - b. Select the IBM i host partition, click the **Tasks** button, and click **Dynamic Logical Partitioning** → **Virtual Adapters**.
 - c. Click **Actions** and click **Create** → **SCSI Adapter**.
 - d. Use the default VSCSI Adapter number or provide your own number. Write down the VSCSI Adapter number, as you need it in a later step. In Figure 8-16, the number 31 was provided as the Virtual SCSI Adapter number. In the Type of adapter field, select **Server**, and click **OK**.




Figure 8-16 Create VSCSI Server Adapter

- e. Save the current configuration for the IBM i host partition so that the VSCSI adapter continues to exist after you restart the partition.
2. Create the VSCSI client adapter in the IBM i client partition, as shown in Figure 8-17:
 - a. In the navigation pane, click **Systems Management** → **Servers**, and click the managed system on which the IBM i client logical partition is.
 - b. Select the IBM i client partition, click **Tasks**, and click **Dynamic Logical Partitioning** → **Virtual Adapters**.
 - c. Click **Actions** and click **Create** → **SCSI Adapter**.
 - d. Use the default VSCSI adapter number or provide your own number. In this example, the number 31 was provided as this number. In the Type of adapter field, select **Client**.
 - e. Select the IBM i server partition that provides the virtualized optical device as the server partition and specify the **Server adapter ID** from step 1d on page 333, as shown in Figure 8-17, where 31 was used. Click **OK**.
 - f. Create the VSCSI adapter within the partition profile for the IBM i client partition so that the VSCSI adapter continues to exist after you restart the partition.




Figure 8-17 Create VSCSI client adapter

There is no additional configuration that is required in IBM i in the virtual client partition. In the host partition, the minimum required IBM i setup consists of the following requirements:

- ▶ One network server description (NWSDB) object
- ▶ One network server storage space (NWSSTG) object

The NWSD object associates a VSCSI server adapter in IBM i (which in turn is connected to a VSCSI client adapter in the HMC/SDMC) with one or more NWSSTG objects. At least one NWSD object must be created on the host for each client, although more are supported. The NWSSTG objects are the virtual disks that are provided to the client IBM i partition. They are created from available physical storage on the host partition. In the client, they are recognized and managed as standard DDxx disk devices (with a different type and model). The IBM i commands **WRKNWSSTG** and **CRTNWSSTG** can be used to manage or create Network Server Storage Space.

To link an NWSD object with an NWS storage space, complete the following steps:

1. Create a network server description (NWSD). On the IBM i host partition, enter **CRTNWSD** to create a network server description, and press F4 for prompts, then F9 to display all parameters.

Enter the following values:

Network Server Description	Provide a name. CLIENT31 was used in the example in Figure 8-17 on page 334, which corresponds to the name of the IBM i client partition in this example.
Resource Name	Provide the resource name. CTL05 was used in this example.
Network server type	*GUEST.
Server operating system	*OPSYS.
Online at IPL	*YES.
Partition	Specify IBM i client partition's name.
Partition number	Specify IBM i client partition number. The partition and partition number cannot be both specified.
Code page	437.
Restricted service resources	Use *NONE for an NWSD that all optical and tape device resources can use by client partition.
Power Control	Use *NO if the NWSD is only for virtualizing an optical device. Use *Yes if the NWSD is also used to provide virtual disk storage.

Run the command of Create Network Server Desc (**CRTNWS**D), as shown in Figure 8-18.

```

                                Create Network Server Desc (CRTNWS)D

Type choices, press Enter.

Network server description . . . NWS D                > ITSOP2
Resource name . . . . . RSR CNAME                    > CTL05
Network server type:                TYPE
  Server connection . . . . .                > *GUEST
  Server operating system . . .                > *OPSYS
Storage path:                        STGPTH
  Network server host adapter .
  IP security rules:
    Remote interface 1 rule . . .                *DFTSECRULE
    Remote interface 2 rule . . .                *DFTSECRULE
    Remote interface 3 rule . . .                *DFTSECRULE
    Remote interface 4 rule . . .                *DFTSECRULE
                                + for more values
Default IP security rule . . . . DFTSECRULE          *NONE
Multi-path group . . . . . MLTP THGRP                *NONE
                                + for more values

```

Figure 8-18 Specify NWS D name, resource, server type, and OS

2. Create an NWS Storage Space in the IBM i host partition by running the Create NWS Storage Space (**CRTNWS**STG) command, as shown in Figure 8-19.

```

                                Create NWS Storage Space (CRTNWS)STG

Type choices, press Enter.

Network server storage space . . NWSSTG              > ITSOP1
Size . . . . . NWSSIZE                            > 50000
From storage space . . . . . FROMNWSSTG            *NONE
Format . . . . . FORMAT                          > *OPEN
Data offset . . . . . OFFSET                      *FORMAT
Auxiliary storage pool ID . . . ASP                  1
ASP device . . . . . ASPDEV
Text 'description' . . . . . TEXT                  *BLANK

```

Figure 8-19 Create an NWS Storage Space

3. Link the NWSD and NWS storage space. The new NWS storage space cannot be used until it has a link to NWSD. Run the Add Server Storage Link (**ADDNWSSTGL**) command, as shown in Figure 8-20.

Add Server Storage Link (ADDNWSSTGL)		
Type choices, press Enter.		
Network server storage space . .	NWSSTG	> ITS0STG1
Network server description . . .	NWSD	> ITSOP1
Dynamic storage link	DYNAMIC	*NO
Network server type	TYPE	*NWSD
Access	ACCESS	*UPDATE
Drive sequence number	DRVSEQNBR	*CALC
Storage path number	STGPTHNBR	*DFTSTGPTH

Figure 8-20 Add Server Storage Link

Storage spaces for an IBM i client partition do not have to match physical disk sizes; they can be created from 160 MB - 1 TB in size, if there is available storage on the host. The 160 MB minimum size is a requirement from the storage management Licensed Internal Code (LIC) on the client partition. For an IBM i client partition, up to 16 NWSSTGs can be linked to a single NWSD, and therefore to a single VSCSI connection. Up to 32 outstanding I/O operations from the client to each storage space are supported for IBM i clients. Storage spaces can be created in any existing auxiliary storage pool (ASP) on the host, including Independent ASPs. Through the usage of NWSSTGs, any physical storage that is supported in the IBM i host partition on a POWER6 based system can be virtualized to a client partition.

For performance reasons, you might consider creating multiple storage spaces that are associated with multiple NWSDs. The rule of thumb is 6 - 8 storage spaces for each client partition. This setup implies that you are also creating multiple sets of VSCSI adapter pairs between the hosting partition and the client partition. Associate each hosting partition's server VSCSI adapter with a separate NWSD by referencing the VSCSI adapter's resource name in the NWSD, and then link storage spaces to the NWSDs. This action supplies multiple disk arms for the client partition to use.

8.4.2 Optical virtualization

Any optical drive that is supported in the IBM i host partition can be virtualized to an IBM i client partition. An existing VSCSI connection can be used, or a connection can be created explicitly for optical I/O traffic. By default, if a VSCSI connection exists between host and client, all physical and virtual OPTxx optical drives on the host are available to the client, where they are recognized as OPTxx devices. The NWSD parameter **Restricted device resources** can be used to specify which optical devices on the host a client partition *cannot* access.

A virtualized optical drive on the host partition can be used for a D-mode Initial Program Load (IPL) and installation of the client partition, as well as for installing Program Temporary Fixes (PTFs) or applications higher. If the optical drive is writable, the client partition can write to the physical media in the drive.

To locate the optical device in the IBM i client partition, enter the **WRKHDWRSC *STG IBM i** command and complete the following steps:

1. Select option 7 to display the resource details next to each of the CMBxx resources that are listed, as shown in Figure 8-21.

Work with Storage Resources					System: CLIENT031
Type options, press Enter.					
7=Display resource detail 9=Work with resource					
Opt	Resource	Type-model	Status	Text	
7	CMB01	290A-001	Operational	Storage Controller	
	DC01	290A-001	Operational	Storage Controller	
	CMB03	268C-001	Operational	Storage Controller	
	DC02	6B02-001	Operational	Storage Controller	
	CMB06	290A-001	Not detected	Storage Controller	
	DC03	290A-001	Not detected	Storage Controller	
F3=Exit F5=Refresh F6=Print F12=Cancel					Bottom

Figure 8-21 Use Option 7 to locate Virtual Adapter number

2. Look at the last digits for the location code Cxx, where xx corresponds to the virtual adapter number, as shown in Figure 8-22.

Display Resource Detail		System: CLIENT031
Resource name	CMB01	
Text	Storage Controller	
Type-model	290A-001	
Serial number	00-00000	
Part number		
Location : U8233.E8B.100417P-V131-C31		
Logical address:		
SPD bus:		
System bus	255	
System board	128	
Press Enter to continue.		More...
F3=Exit F5=Refresh F6=Print F12=Cancel		

Figure 8-22 Display Resource Detail

- When you find the correct CMBxx resource, look for the DC0xx resource and select option 9 to work with resources, as shown in Figure 8-23.

Work with Storage Resources					System: CLIENT031
Type options, press Enter.					
7=Display resource detail 9=Work with resource					
Opt	Resource	Type-model	Status	Text	
	CMB01	290A-001	Operational	Storage Controller	
9	DC01	290A-001	Operational	Storage Controller	
	CMB03	268C-001	Operational	Storage Controller	
	DC02	6B02-001	Operational	Storage Controller	
	CMB06	290A-001	Not detected	Storage Controller	
	DC03	290A-001	Not detected	Storage Controller	
					Bottom
F3=Exit F5=Refresh F6=Print F12=Cancel					

Figure 8-23 Use option 9 to work with resources

The optical device that is provided by the IBM i server partition is shown in the IBM i client partition, as shown in Figure 8-24.

Work with Storage Controller Resources					System: CLIENT031
Type options, press Enter.					
5=Work with configuration descriptions 7=Display resource detail					
Opt	Resource	Type-model	Status	Text	
	DC01	290A-001	Operational	Storage Controller	
	DD001	6B22-050	Operational		
	DD003	6B22-050	Operational		
	OPT01	632C-002	Inoperative	Optical Storage Unit	
	OPT02	632C-002	Operational	Optical Storage Unit	
	OPT08	632C-002	Inoperative	Optical Storage Unit	
					Bottom
F3=Exit F5=Refresh F6=Print F12=Cancel					

Figure 8-24 Virtualized optical device that is shown on IBM i client partition as type 632C-002

For more information about image catalog, see the IBM i Information Center. Search for “Virtual optical storage.” The IBM i Information Center can be found at:

<http://publib.boulder.ibm.com/iseriess/>

8.4.3 Tape virtualization

Any tape drive that is supported in the IBM i host partition can be virtualized to an IBM i client partition. An existing VSCSI connection can be used, or a connection can be created explicitly for tape I/O traffic. By default, if a VSCSI connection exists between the host and the client, all physical and virtual TAPxx drives on the host are available to the client, where they are recognized as TAPxx devices. The NWSD parameter **Restricted device resources** can be used to specify which optical devices on the host a client partition *cannot* access.

A virtualized tape drive on the host partition can be used for a D-mode initial program load (IPL) and installation of the client partition or applications. The client partition can write to the physical media in the drive.

8.4.4 Networking virtualization

Virtualizing a network adapter and using a virtual LAN (VLAN) for partition-to-partition communication within a system are existing IBM i capabilities. In order for a client to use a host's physical network adapter, a virtual Ethernet adapter must be created in the HMC in both partitions.

To be on the same VLAN, the two virtual Ethernet adapters must have the same Port Virtual LAN ID (PVID). This type of adapter is recognized by IBM i as a communications port (CMNxx) with a different type (268C). In the host partition, the virtual Ethernet adapter is then associated with the physical network adapter through a routing configuration, either Ethernet Level-2 Bridging or network address translation (NAT). This routing configuration allows the client partition to send network packets through the VLAN and the physical adapter to the outside LAN. The physical adapter can be any network adapter that is supported by IBM i 6.1 and later, including Integrated Virtual Ethernet (IVE) ports, also known as Host Ethernet Adapter (HEA) ports.

For more information about Ethernet Level-2 Bridging, see 8.5.1, "Ethernet Layer-2 bridging" on page 342. For more information about network address translation (NAT), see the IBM i 7.1 Information Center, found at:

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp>

8.4.5 Support for embedded media changers

Embedded media changer support extends the automatic media switching capability of virtual optical device type 632B on virtual I/O serving partitions to the client partitions virtual optical device type 632C. One application of this new function is the usage of image catalogs for unattended installations of client partitions. You can use this switching capability to manually switch media in a client virtual optical device without requiring the authority to access the serving partition. You can accomplish this action by using the image catalog interface **WRKIMGCLGE *DEV** command interface on the client partition.

8.4.6 Performance consideration

In IBM i host IBM i environment, disk I/O operations in an IBM i virtual client partition result in I/O requests to the physical disk adapters and drives that are assigned to the host partition. Therefore, the preferred way to ensure good disk performance in the client partition is to create a well-performing disk configuration on the host partition. As the host partition is a standard IBM i partition, all the recommendations in the Performance Capabilities Reference manual (see <http://www.ibm.com/systems/i/solutions/perfmgmt/resource.html>) apply to it. Use the suggestions that are provided in the manual for maximizing IBM i disk performance for the type of physical storage that is used on the host, whether it is integrated disk or SAN.

If only the system ASP exists on the host partition, NWSSTG objects are created on the same physical disk units as all other objects. If the host partition is running production applications in addition to providing virtual storage to client partitions, there is disk I/O contention as both client partitions and IBM i workloads on the host send I/O requests to those disk units. To minimize disk I/O contention, create storage space objects in a separate ASP on the host (Independent ASPs are supported). Performance on the clients then depends on the disk adapter and disk configuration that is used for that ASP. If the host partition is providing virtual storage to more than one client partition, consider using separate ASPs for the storage space objects for each client. This preferred practice should be weighed against the concern of ending up with too few physical disk arms in each ASP to provide good performance.

Disk contention from IBM i workloads on the host partition and virtual client partitions can be eliminated if a separate IBM i partition is used just for hosting client partitions. An additional benefit of this configuration is the fact that an application or OS problem that is stemming from a different workload on the host cannot negatively affect client partitions. These benefits should be weighed against the following items:

- ▶ The license cost that is associated with a separate IBM i partition
- ▶ The maintenance time that is required for another partition, such as applying Program Temporary Fixes (PTFs)
- ▶ The ability to create well-performing physical disk configurations in both partitions that meet the requirements of their workloads

If the host partition runs a heavy-I/O workload and the client partitions also have high disk response requirements, consider using a separate hosting partition, unless separate ASPs on the host are used for storage space objects. If the host partition's workload ranges from light to moderate regarding disk requirements and the client partitions are used mostly for development, test or quality assurance (QA), it is acceptable to use one IBM i partition for both tasks.

8.4.7 Dual hosting

An IBM i client partition has a dependency on its host: If the host partition fails, IBM i on the client loses contact with its disk units. The virtual disks become unavailable if the host partition is brought down to a restricted state or shut down for scheduled maintenance or to apply PTFs. To remove this dependency, two host partitions can be used to simultaneously provide virtual storage to one or more client partitions.

The configuration for two hosts for the same client partition uses the same concepts as for a single host, as described in the 8.4.1, “Disk virtualization” on page 333. In addition, a second VSCSI client adapter exists in the client partition, which is connected to a VSCSI server adapter in the second host partition. The IBM i configuration of the second host mimics that of the first host, with the same number of NWSG and NWSSTG objects, and NWSSTG objects of the same size. As a result, the client partition recognizes a second set of virtual disks of the same number and size. To achieve redundancy, adapter-level mirroring is used between the two sets of storage spaces from the two hosts. Thus, if a host partition fails or is taken down for maintenance, mirroring is suspended, but the client partition continues to operate. When the inactive host is either recovered or restarted, mirroring can be resumed.

8.5 Virtual Partition Manager enhancements

The Virtual Partition Manager (VPM) is a partition management tool that supports the creation of partitions that use only virtual I/O and does not require the Hardware Management Console (HMC), Systems Director Management Console (SDMC), or Integrated Virtualization Manager (IVM). In addition to being able to manage Linux guest partitions, the VPM now supports creation and management of IBM i partitions. The VPM function is available on POWER6 and POWER7 Express Servers that do not have an external management console.

With this enhancement to IBM i 7.1, the ability to create up to four IBM i partitions is enabled in VPM. Client IBM i partitions, which are created with VPM, use virtual I/O to connect back to the IBM i I/O server partition to access the physical disk and network. VPM in the IBM i I/O server partition is used to create the virtual SCSI and virtual Ethernet adapters for the client partitions. You can then use Network Storage Spaces (NWSSTG) and Network Storage Descriptions (NWSG) in the IBM i I/O server partition to define the storage for the client partitions. Tape, disk, and optical can be virtualized to the client partitions. The client IBM i partitions can be IBM i 7.1 or IBM i 6.1 with either 6.1 or 6.1.1 machine code.

8.5.1 Ethernet Layer-2 bridging

IBM i V7R1 has new support for Ethernet Layer-2 bridging between a physical network and the Power Systems virtual Ethernet.

Using Layer-2 bridging, one Ethernet port in an IBM i partition can provide network access for other logical partitions on the same platform. This support is similar in functionality to the Shared Ethernet Adapter (SEA) support provided by a Power Systems Virtual I/O Server (VIOS) partition.

This situation puts two Ethernet adapters (one physical and one virtual) into a mode where they can receive traffic that is not destined for their address, and selectively sends those frames onto the other network according to the IEEE 802.1D standard (“bridging” the frames). Frames that are transmitted by virtual Ethernet adapters on the same VLAN as the bridging virtual Ethernet adapter can be sent to the physical network, and frames from the physical network can be received by adapters on the virtual network.

8.5.2 Preparing for Ethernet Layer-2 bridging

Select a physical Ethernet resource to use for Layer-2 bridging with the following criteria:

- ▶ Any Ethernet resource that supports line speeds of 1 Gbps or greater is supported, except for Host Ethernet Adapter (HEA) resources. (Host Ethernet Adapter supports the ability for multiple partitions to use a single physical port by assigning each partition a logical port.)
- ▶ Layer-2 bridging must not be in use by any varied-on line description, LAN console, or remote support.
- ▶ An aggregate line description can also be used to bridge traffic to the external network.
- ▶ Create a virtual Ethernet resource to use for Layer-2 bridging, and record its resource name.
- ▶ If you are using a Hardware Management Console, create a virtual Ethernet adapter for the wanted VLAN ID, and check the **Access external network** check box to indicate that this virtual Ethernet adapter is used to bridge traffic to the physical network.
- ▶ If you are using the IBM i Virtual Partition Manager, the virtual Ethernet adapter is created automatically and can access the external network.
- ▶ Choose an alphanumeric name (up to 10 characters) for the bridge itself, and make it unique from any existing bridge names.

8.5.3 Preferred practices

Selected Ethernet resources should be used only for Layer-2 bridging (not for the IBM i TCP/IP configuration). There is a processing impact for any host traffic that uses bridged resources.

8.5.4 Configuring Ethernet Layer-2 bridging

Create an Ethernet line description for the physical Ethernet resource, and set its Bridge identifier (BRIDGE) to your chosen bridge name.

Create an Ethernet line description for the selected virtual Ethernet resource, and set its Bridge identifier (BRIDGE) to the same bridge name.

When both line descriptions are varied on, traffic is bridged between the two networks. Any other partitions with virtual Ethernet adapters on the same VLAN as the new virtual Ethernet resource are able to access the same network as the physical Ethernet resource.

8.5.5 Common errors

CHGLINETH cannot be used to change the Bridge identifier of a line description that was created before the latest Technology Refresh. If equivalent behavior is wanted, complete the following steps:

1. Use the **Copy** option on **WRKLIND** to make a temporary copy of the line description.
2. Delete the existing line description.
3. Use the **Copy** option again on **WRKLIND** to replicate the original line description, specifying the wanted bridge identifier.
4. Delete the temporary line description.

No more than one physical Ethernet adapter's line description with a given bridge identifier can be varied on at the same time. Likewise, no more than one virtual Ethernet adapter's line description with a given bridge identifier can be varied on at the same time. An error is returned when you try to vary on any more line descriptions with that bridge identifier, indicating that the configuration is in error. For a bridge, select one physical Ethernet line description and one virtual line description to be bridged. If more than one bridge is required, use a different bridge identifier for each additional bridge.

The selected virtual Ethernet resource must be marked as allowing access to the external network. If an incorrect virtual Ethernet resource is selected, an error is returned when you try to vary on its line description, indicating that the selected resource cannot enable promiscuous mode. Create a virtual Ethernet resource that can be used to access the external network.

8.5.6 Managing Ethernet Layer-2 bridging

While an Ethernet line description is varied off, its bridge identifier (BRIDGE) can be changed to a different name (or to *NONE, indicating that it is not to be used for bridging).

Remember: In IBM i V7R1, an Ethernet line description's bridge identifier is not visible from DSPLIND. Use the **CHGLINETH** command and prompt to see the bridge identifier for an Ethernet line description.

8.6 Partition suspend and resume

PowerVM now includes support for an IBM i 7.1 partition to be suspended, and later resumed. Using Suspend / Resume, clients can perform long-term suspension of partitions, freeing server resources that were in use by that partition, and later resume operation of that partition and its applications on the same server. During the Suspend operation, the partition state (memory, NVRAM, and Virtual Service Processor state) is saved on persistent storage. The Resume operation restores that saved partition state to the server resources. Suspend / Resume can be used to save energy or to allow other partitions to use the resources from the suspended partition.

8.6.1 Requirements for Suspend / Resume

The following items are requirements for Suspend / Resume:

- ▶ All I/O resources must be virtualized using VIOS.
- ▶ All partition storage must be external.
- ▶ Either an HMC or SDMC must be used to manage the partitions.
- ▶ The partition must be resumed on the same server on which it was suspended.
- ▶ POWER7 firmware Ax730_xxx, or later, is required.
- ▶ VIOS 2.2.0.12-FP24 SP02, or later, is required.
- ▶ AIX Version 7.1 Technology Level 0 Service Pack 2 is required.
- ▶ AIX Version 6.1 Technology Level 6 Service Pack 3 is required.
- ▶ For an IBM i logical partition, the logical partition must be running IBM i 7.1 with the latest Technology Refresh.

- ▶ When a logical partition is suspended, the reserved storage device contains the state that is required to resume the logical partition. Therefore, the reserved storage device must be kept persistently associated with the logical partition.
- ▶ The HMC ensures that the reserved storage device pool is configured with at least one active Virtual I/O Server partition available in the pool.
- ▶ You can create or edit the partition profile of a logical partition that can suspend without any restrictions. However, when you activate a logical partition with a specific profile, checks are performed for any of the restrictions that are associated with suspending the logical partition.
- ▶ For NPIV, you must zone both of the WWPNs associated with a Virtual Fibre Channel adapter.

8.6.2 Considerations for Suspend / Resume

The following items are considerations that apply to Suspend / Resume:

- ▶ The logical partition must not have physical I/O adapters that are assigned to the logical partition.
- ▶ The logical partition must not be a full system partition or a Virtual I/O Server partition.
- ▶ The logical partition must not be an alternative error logging partition.
- ▶ The logical partition must not have a barrier-synchronization register (BSR).
- ▶ The logical partition must not have huge pages (applicable only if PowerVM Active Memory Sharing is enabled).
- ▶ The logical partition must not have its rootvg volume group on a logical volume or have any exported optical devices.
- ▶ You cannot suspend an IBM i logical partition while it is active in a cluster.
- ▶ The logical partition must not have a virtual SCSI optical or tape device that is assigned to the logical partition.

The following additional considerations apply for IBM i logical partitions that are enabled for suspension:

- ▶ You cannot activate the logical partition with a partition profile that has a virtual SCSI server adapter.
- ▶ You cannot activate the logical partition with a partition profile that has a virtual SCSI client adapter that is hosted by another IBM i logical partition.
- ▶ You cannot dynamically add any virtual SCSI server adapter.
- ▶ You cannot dynamically add any virtual SCSI client adapter that is hosted by another IBM i logical partition.
- ▶ You cannot dynamically add any physical I/O adapters.
- ▶ You cannot suspend an IBM i logical partition with a varied NPIV attached tape device.
- ▶ All IBM i virtual disks must be backed by physical volumes.

For the latest information about prerequisites, see IBM Prerequisites at:

https://www-912.ibm.com/e_dir/eserverprereq.nsf

For the latest information about the configuration requirements and restrictions for suspending a logical partition, visit the Information Center at:

<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/p7hat/iphathibreqs.htm>

8.7 HEA Daughter cards

POWER7 HEA Daughter cards provide integrated I/O connectors for a CPC enclosure. The connections can be virtualized into the system partitions. All of the connectors are on the rear bulkhead of the CPC enclosure. Choices of features are:

- ▶ #1824 and #1832 for a 1-Gb HEA daughter card with four ports
- ▶ #1825 and #1833 for a Fibre 10-Gb HEA daughter card with two ports
- ▶ #1826 and #1837 for a Copper 10-Gb HEA daughter card with two ports

8.8 10 Gb FCoE PCIe Dual Port Adapter

The 10 Gb FCoE PCIe Dual Port Adapter (#5708) is a high-performance, 10 Gb, dual port, PCIe Converged Network Adapter (CNA) using SR optics. Each port can provide Network Interface Card (NIC) traffic and Fibre Channel functions simultaneously. IBM i supports the usage of this adapter through VIOS.

8.9 Live Partition Mobility

Live Partition Mobility (LPM) is a powerful function that is delivered by PowerVM Enterprise Edition. Using LPM, a running IBM i 7.1 partition can be moved from one POWER7 processor-based server to another server with no application downtime. The migration operation, which takes just a few seconds, maintains complete system transactional integrity. The migration transfers the entire system environment, including processor state, memory, attached virtual devices, and connected users.

IBM Power Systems servers are designed to offer the highest stand-alone availability in the industry. Enterprises must occasionally restructure their infrastructure to meet new IT requirements. By letting you move your running production applications from one physical server to another, LPM allows for nondisruptive maintenance or modification to a system without your users noticing anything. LPM mitigates the impact on partitions and applications that was formerly caused by the occasional need to shut down a system.

Even small IBM Power Systems servers frequently host many logical partitions. As the number of hosted partitions increases, finding a maintenance window acceptable to all becomes increasingly difficult. You can use LPM to move partitions around so that you can perform previously disruptive operations on the machine at your convenience, rather than when it causes the least inconvenience to the users.

Figure 8-25 shows basic hardware infrastructure of LPM.




Figure 8-25 Basic hardware infrastructure for Live Partition Mobility

LPM helps you meet increasingly stringent service-level agreements (SLAs) because you can proactively move running partitions and applications from one server to another server.

The ability to move running partitions from one server to another server means that you can balance workloads and resources. If a key application's resource requirements peak unexpectedly to a point where there is contention for server resources, you might move it to a more powerful server or move other, less critical, partitions to different servers, and use the freed resources to absorb the peak.

LPM may also be used as a mechanism for server consolidation because it provides an easy path to move applications from individual, stand-alone servers to consolidation servers. If you have partitions with workloads that have widely fluctuating resource requirements over time (for example, with a peak workload at the end of the month or the end of the quarter), you can use LPM to consolidate partitions to a single server during the off-peak period so that you can turn off unused servers. Then, you move the partitions to their own, adequately configured servers just before the peak. This approach also offers energy savings by reducing the power to run machines and the power to keep them cool during off-peak periods.

LPM can be automated and incorporated into system management tools and scripts. Support for multiple concurrent migrations allows you to liberate system resources quickly. For single-partition, point-in-time migrations, the Hardware Management Console (HMC) interface offers easy-to-use migration wizards.

LPM contributes to the goal of continuous availability:

- ▶ Reduces planned downtime by dynamically moving applications from one server to another server
- ▶ Responds to changing workloads and business requirements when you move workloads from heavily loaded servers to servers that have spare capacity
- ▶ Reduces energy consumption by allowing you to easily consolidate workloads and turn off unused servers

LPM is the next step in the IBM PowerVM continuum. It can be combined with other virtualization technologies to provide a fully virtualized computing platform that offers the degree of system and infrastructure flexibility that is required by today's production data centers.

8.9.1 Requirements for Live Partition Mobility

The following items are requirements for LPM:

- ▶ The HMC must be used to manage both source and destination systems.
- ▶ HMC Version 7 Release 5 or later is required.
- ▶ Both source and destination systems must be an IBM Power Systems POWER7 technology-based model with firmware service pack 730_51, 740_40, or later.
- ▶ Both source and destination systems must have the PowerVM Enterprise Edition license code installed.
- ▶ All I/O resources of the mobile partition must be virtualized using VIOS, whether they are VSCSI, NPIV, or Virtual Ethernet.
- ▶ VIOS 2.2.1.4 or later is required.
- ▶ Both source and destination systems must use same external storage.
- ▶ An IBM i mobile partition must be running IBM i 7.1 with the latest Technology Refresh.

8.9.2 Considerations for Live Partition Mobility

The following items are considerations for LPM:

- ▶ The mobile partition must have all disks backed by physical volumes.
- ▶ The mobile partition must not be assigned a virtual SCSI optical or tape device or have an NPIV attached tape device that is varied on.
- ▶ The mobile partition cannot be activated with a partition profile that has a virtual SCSI server adapter, and it cannot be host another partition.
- ▶ The mobile partition cannot be activated with a partition profile that has a virtual SCSI client adapter that is hosted by another IBM i logical partition.
- ▶ No virtual SCSI server adapters can be dynamically added to the mobile partition.
- ▶ No virtual SCSI client adapters that are hosted by another IBM i logical partition can be dynamically added to the mobile partition.
- ▶ The mobile partition must not be an alternative error logging partition.

- ▶ An alternative error logging partition is a target from the HMC for error logs.
- ▶ The mobile partition cannot collect physical I/O statistics.
- ▶ The mobile partition must not be a time reference partition. The VIOS synchronizes the time between partitions automatically as part of the migration.

8.9.3 Preparing Live Partition Mobility

To prepare for LPM, complete the following steps:

1. Ensure that both source and destination systems have capabilities for LPM, as shown in Figure 8-26.




Figure 8-26 Power system Capabilities of LPM

2. Set the VIOS logical partitions to **Mover service partition**, as shown in Figure 8-27.




Figure 8-27 Set VIOS as Mover Service Partition

3. Ensure that Resource Monitoring and Control (RMC) connections are established between both the source and destination VIOS logical partitions and HMC.

Sign on HMC with the correct authority and run **lspartition -dlpar**, as shown in Example 8-4, to check the RMC connection between the HMC and VIOS.

Example 8-4 Check the RMC connection between the HMC and VIOS

```
#commands:
lspartition -dlpar
#Results:
<#23> Partition:<5*8205-E6C*06523ER, , 172.16.26.99>
      Active:<1>, OS:<AIX, 6.1, 6100-07-04-1216>, DCaps:<0x4f9f>,
CmdCaps:<0x1b, 0x1b>, PinnedMem:<799>
```

Active:<1> in the results means that the RMC connection is activated.

4. Verify that the same SAN disks that are used as virtual disks by the IBM i mobile partition are assigned to the source and destination VIOS logical partitions, and ensure that the `reserve_policy` attributes of the shared physical volumes are set to `no_reserve`. Sign on both source and destination VIOS logical partitions with the correct authorities and run the commands that are shown on Example 8-5.

Example 8-5 Verify that the physical volumes on external storage are set correctly

List attributes of a physical volume, use the following command:
`#lsdev -dev hdiskX -attr`

If `reserve_policy` is not set to `no_reserve`, use the following command:
`#chdev -dev hdiskX -attr reserve_policy=no_reserve`

5. Ensure that the IBM i mobile partition has no physical I/O resources, and check the **RestrictedIO Partition** check box, as shown in Figure 8-28.




Figure 8-28 Set the IBM i mobile partition as RestrictedIO Partition

8.9.4 Validating Live Partition Mobility

Before you perform a migration, you should complete the validation steps. These steps are optional, but can help eliminate errors. This section shows how to perform the validation steps by using the HMC GUI.

1. In the navigation pane, expand **Systems Management** → **Servers**, and select the source system.
2. In the upper right of the Hardware Management Console (Workplace), select the partition that you will migrate to the destination system.

3. Click the **View popup menu** button and click **Operations** → **Mobility** → **Validate**, as shown in Figure 8-29, to start the validation process.




Figure 8-29 Validate option on HMC

4. Select the destination system, specify **Destination profile name** and **Wait time (in min)**, and then click **Validate**, as shown in Figure 8-30.

Partition Migration Validation - SVRI7740-01-SN06523ER - ibmi_lpm

Fill in the following information to set up a migration of the partition to a different managed system. Click Validate to ensure that all requirements are met for this migration. You cannot migrate until the migration set up has been verified.

Source system : SVRI7740-01-SN06523ER
 Migrating partition: ibmi_lpm
 Remote HMC:
 Remote User:
 Destination system: SVRI7740-02-SN065264R
 Destination profile name: ibmi_lpm
 Destination shared processor pool:
 Source mover service partition:
 Destination mover service partition:
 Wait time (in min): 5
 Override virtual network errors when possible: ☐
 Override virtual storage errors when possible: ☐
 Virtual Storage assignments :

Select	Source Slot ID	Slot Type	Destination VIOS

Figure 8-30 Input information for validation

5. Check for errors or warnings in the Partition Validation Errors/Warnings window, and eliminate any errors.

6. The validation window opens and with information about the destination VIOS logical partition, Virtual Storage, and VLAN, as shown in Figure 8-31. It is ready for migration.

HMC50: Validate - Mozilla Firefox

Partition Migration Validation - SVRI7740-01-SN06523ER - ibmi_lpm

Fill in the following information to set up a migration of the partition to a different managed system. Click Validate to ensure that all requirements are met for this migration. You cannot migrate until the migration set up has been verified.

Source system : SVRI7740-01-SN06523ER
 Migrating partition: ibmi_lpm
 Remote HMC:
 Remote User:
 Destination system: SVRI7740-02-SN065264R
 Destination profile name: ibmi_lpm
 Destination shared processor pool: DefaultPool (0)
 Source mover service partition: i740_01_VIOS1
 Destination mover service partition: i740_02_VIOS1
 Wait time (in min): 5
 Override virtual network errors when possible: ☐
 Override virtual storage errors when possible: ☐
 Virtual Storage assignments :

Select	Source Slot ID	Slot Type	Destination VIOS
☒	3	SCSI	i740_02_VIOS1

Figure 8-31 Validation window with readiness for Live Partition Mobility

8.9.5 Performing Live Partition Mobility

After the validation completed, you can perform LPM. In this scenario, we show how to migrate a IBM i logical partition named IBMi_lpm from the source system (SVRI7740-01-SN06523ER) to the destination system (SVRI7740-01-SN065264R) by using HMC. Complete the following steps:

1. In the navigation pane, expand **Systems Management** → **Servers**, and select the source system.

2. Click the **View popup menu** button and click **Operations** → **Mobility** → **Migrate**, to start the Partition Migration wizard, as shown in Figure 8-32.




Figure 8-32 Migration menu on HMC

3. Check the migration information in the Partition Migration wizard, as shown in Figure 8-33.




Figure 8-33 Migration Information

4. Specify the New destination profile name in the Profile Name window, as shown in Figure 8-34. If you leave the name blank or do not specify a unique profile name, the profile on the destination system is overwritten.




Figure 8-34 Specify destination profile name

5. Enter the Remote HMC network address and Remote User. In our example, we use a single HMC. Click **Next** (Figure 8-35).

The screenshot shows a web-based configuration window titled "Partition Migration - SVRI7740-01-SN06523ER - ibmi_lpm". The window is divided into two main sections. On the left is a vertical navigation pane with a list of steps: "Migration Information", "Profile Name", "Remote HMC" (highlighted with a yellow arrow), "Destination", "Validation Errors/Warnings", "Mover Service Partitions", "VLAN Configuration", "Virtual Storage Adapters", "Shared Processor Pools", "Wait Time", and "Summary". The main content area on the right is titled "Remote HMC" and contains the following text: "Specify remote HMC and user for partition migration. The destination system should have at least one managed system capable of partition migration." Below this text is a checkbox labeled "Remote Migration" which is currently unchecked. Underneath the checkbox are two input fields: "Remote HMC:" and "Remote User:". At the bottom of the window, there is a horizontal bar with four buttons: "< Back", "Next >", "Finish", and "Cancel".

Figure 8-35 Specify the remote HMC IP address and user

6. Specify the destination system for migration, as shown in Figure 8-36. The HMC validates the partition migration environment.




Figure 8-36 Specify destination system for migration

7. Check errors or warnings in the Partition Validation Errors/Warnings window and eliminate any errors. This step was skipped in our example because there were no errors or warnings.

- Specify the mover service partitions to be used for migration in both source and destination systems, as shown in Figure 8-37. In our example, one Virtual I/O Server partition is configured on the destination system, so the wizard window shows only one mover service partition candidate. If you have more than one Virtual I/O Server partition on the source or on the destination system, you can select which mover server partitions to use.




Figure 8-37 Specify MSP partitions

9. Confirm the VLAN configuration, as shown in Figure 8-38.




Figure 8-38 Confirm the VLAN configuration

10. Select the virtual storage adapter assignment, as shown in Figure 8-39. In our example, one Virtual I/O Server partition is configured on each system, so this wizard window shows one candidate only. If you have more than one Virtual I/O Server partition on the destination system, you may choose which Virtual I/O Server to use as the destination.




Figure 8-39 Select virtual storage adapters

11. Select the shared processor pools to be used by the mobile partition in the destination system, as shown in Figure 8-40.




Figure 8-40 Select Shared Processor

12. Specify the wait time in minutes, as shown in Figure 8-41. The wait time value is passed to the commands that are started on the HMC and perform migration-related operations on the relevant partitions using the Remote Monitoring and Control (RMC).

The screenshot shows a web browser window titled "Partition Migration - SVRI7740-01-SN06523ER - Mozilla Firefox". The main content area is titled "Partition Migration - SVRI7740-01-SN06523ER - ibmi_lpm". On the left, a navigation pane lists several steps with checkmarks: "Migration Information", "Profile Name", "Remote HMC", "Destination", "Validation Errors/Warnings", "Mover Service Partitions", "VLAN Configuration", "Virtual Storage Adapters", "Shared Processor Pools", "Wait Time" (highlighted with a yellow arrow), and "Summary". The main panel is titled "Wait Time" and contains the following text: "Wait time is the maximum time, in minutes, that the operating system should wait for applications to acknowledge that a partition migration event is about to happen. You may enter a non-zero value if the default value is unsuitable." Below this text is a text input field labeled "Wait time (in min):" with the value "5" entered. At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

Figure 8-41 Specify Waiting time

13. Check the settings that you specified for this migration in the Summary window, and then click **Finish** to begin the migration, as shown Figure 8-42.




Figure 8-42 Partition Migration Summary

14. The Migration status and Progress are shown in the Partition Migration Status window, as shown in Figure 8-43.




Figure 8-43 Partition Migration Status

15. When the Partition Migration Status window indicates that the migration is 100% complete, you find that the mobile partition is running in the destination system.

8.9.6 Considerations about Live Partition Mobility

This section describes some considerations about IBM i LPM:

- Performance considerations

Active partition migration involves moving the state of a partition from one system to another system while the partition is still running. The partition memory state is tracked while it transfers the memory state to the destination system. Multiple memory transfers are done until enough clean pages are moved.

Providing a high-performance network between the source and destination mover partitions and reducing the partition's memory update activity before migration improves the latency of the state transfer phase of migration. As a preferred practice, use a dedicated network for state transfer, with a nominal bandwidth of at least 1 Gbps.

- Licensing considerations

For a permanent partition migration, the requirements are the same as for a manual migration to a new system. For a temporary partaking migration:

- If software entitlements were acquired on all systems, move partitions where and when they are needed.
- If the software entitlements of the IBM i operating system or standard set of IBM i licensed program products are *not* acquired on the destination systems, you may temporarily move the partition for up to 70 days to a destination system.

You must purchase software entitlements on the source system. The source system must be equal or larger in the processor group than any destination system. After 70 days, you must acquire the entitlements on the destination system, move the entitlements back to the source system, or move the entitlements to another destination system.

- Application considerations

In general, applications and the operating system are unaware that the partition is moved from one system to another. There are some exceptions for applications that recognize system serial number, logical partition ID, and system type/model. When the partition is running on the target system, the Collection Services collector job cycles the collection so that correct hardware information is recorded on the target system.

- Exit program considerations

Two Work Management exits were added: QIBM_QWC_RESUME and QIBM_QWC_SUSPEND. There is no restriction on the number of programs register against these exit points.

- QIBM_QWC_RESUME

The Resume System exit program is called when the system becomes available again after the partition was migrated to another machine or the partition resumes after hibernation.

- QIBM_QWC_SUSPEND

The Suspend System exit program is called before the system becomes temporarily unavailable because the partition is migrating to another machine or the partition is hibernating.



Storage and solid-state drives

This chapter describes the enhancements in recent IBM i releases i 6.1.1 and i 7.1 in the following areas:

- ▶ General and direct-attached storage management enhancements

This section describes enhancements in the areas of disk management, ASP encryption, performance instrumentation, and tape applicable to IBM i internal direct-attached and possibly external storage.

- ▶ SAN storage management enhancements

This section summarizes the enhancements that are related to IBM i storage area network-attached storage systems, such as redundant Virtual I/O Server (VIOS) multipathing support, DS5000 native attachment, changes with IOP-less IOA attached tape libraries, and new DS8000/DS6000 performance metrics.

- ▶ SSD storage management enhancements

This section describes IBM i solid-state drive (SSD) management improvements that are related to DB2 media preference, ASP balancer, and user-defined file systems.

9.1 General and direct-attached storage management enhancements

In this section, the following general storage management enhancements that are applicable to IBM i internal direct-attached storage (DAS) and possibly external storage area network (SAN) storage are described:

- ▶ Concurrent removal of disk units
- ▶ Hot spare for mirroring
- ▶ Dual storage I/O adapters
- ▶ Encrypted ASP enhancements
- ▶ Disk response time buckets enhancements
- ▶ Central processor complex node level mirroring
- ▶ EXP24S SFF Gen2-bay drawer (#5887)
- ▶ Higher capacity 10 K RPM SFF SAS disk drives
- ▶ Tape performance instrumentation
- ▶ Tape library resource name changes for IOP-less IOA attachment
- ▶ Tape library unreadable barcode changes for IOP-less IOA attachment
- ▶ DVD / Tape SAS External Storage Unit for Power 795 CPC Rack
- ▶ RDX support for IBM i

9.1.1 Concurrent removal of disk units

Concurrent removal of disk units with IBM i 7.1 is now supported for SYSBAS (for example, system ASP and user ASP disk units), eliminating the need for downtime to perform an IPL to the DST for removing disk units from the configuration.

Figure 9-1 shows the new Work with Removing Units From Configuration panel in System Service Tools. This panel can be accessed by clicking **System Service Tools** → **Work with disk units** → **Work with disk configuration**.

Work with Removing Units From Configuration

Select one of the following:

1. Display disk configuration

2. Display status of remove operation

3. Remove units from configuration

4. Pause the remove operation

5. Resume the remove operation

6. Cancel the remove operation

7. Cancel the remove operation and balance data in the ASP

Selection

F3=Exit F12=Cancel

Figure 9-1 IBM i SST Work with Removing Units from Configuration panel

Figure 9-2 shows the Remove units from configuration panel with the example of disk unit 11 from ASP1 selected to be removed. This unit can become unconfigured after the removal action.

Remove Units from Configuration

Type options, press Enter.

4=Remove unit from configuration

OPT	Unit	ASP	Serial Number	Type	Model	Resource Name	Status
	2	1	21-DC78C	4328	072	DD006	RAID 5/Active
	3	1	21-DD464	4328	072	DD004	RAID 5/Active
	4	1	21-E72DE	4328	072	DD008	RAID 5/Active
	5	1	21-E7A8D	4328	072	DD005	RAID 5/Active
	6	1	21-E7CB9	4328	072	DD007	RAID 5/Active
	7	1	21-DCA21	4328	072	DD003	RAID 5/Active
	8	1	21-E7B11	4328	072	DD011	RAID 5/Active
	9	1	21-DD3DA	4328	074	DD012	RAID 5/Active
	10	1	21-E7046	4328	074	DD010	RAID 5/Active
4	11	1	21-E7557	4328	074	DD009	RAID 5/Active
	12	2	21-E786C	4328	074	DD002	RAID 5/Active

F3=Exit

F5=Refresh

F11=Display Non-configured units

F12=Cancel

Figure 9-2 IBM i SST Remove Units from Configuration panel

This new disk unit removal function, as with the previously available add disk unit function, works for both SYSBAS and independent ASPs, even if the IASP is varied on.

The remove function does not allow removal if the remaining capacity can result in an exceeded ASP threshold. Media preferences for SSDs are respected by the remove function (for example, DB2 or UDFS media preferences; for more information, see 9.3, “SSD storage management enhancements” on page 398) and are honored if there is remaining capacity on the corresponding media type.

Only one remove operation for one or more disk units of a single system can be started, paused, or canceled at any time. The pause operation prevents further data allocations on the disk units that are selected for removal, similar to the ***ENDALC** option in **STRASPBAL**.

Important: The disk unit remove function in System Service Tools, which supports concurrent disk unit removal with applications using the ASP, does not allow removal of all the disk units from the ASP. An IPL to DST is required to delete the ASP.

9.1.2 Hot spare for mirroring

The usage of hot spare disk units for mirroring is supported by IBM i 6.1.1 or later. The benefit of using hot spares is that a non-degraded array or active mirrored pair state is reached more quickly again after a disk unit failure. No manual intervention is required for resuming drive protection. Although the hot spare selection algorithm selects the hot spare resulting in the highest mirroring protection level, there are situations for which the hot spare configuration does not permit the original protection level. In this case, you might still want to manually replace the newly used hot spare with a replacement drive to reestablish the original protection level, for example, bus level mirroring protection.

For a RAID configuration, the hot spare disk unit is used as a replacement for similar or lower capacity drives. For mirroring, the capacity requirement is more stringent. The hot spare must be the same size or bigger (within 25 GB).

When a disk unit is configured as a hot spare, as shown in Figure 9-3, it is no longer visible as a non-configured or configured disk unit in the **System Service Tools** → **Work with disk units** panels. However, it still shows up in the Hardware Service Manager under the disk IOA as a unique model 51 representing a hot spare disk unit.

Start Hot Spare						
Selecting a disk unit will start using the selected disk unit as a hot spare.						
Type option, press Enter.						
1=Start Hot Spare						
IOA Serial	IOA Resource	Serial	Resource			
Opt Number	Name	Number	Type	Model	Name	
1	1C-3300176	DC06	68-0E7A6	6717	050	DD026
	1C-3300176	DC06	68-0C3CF93	6718	050	DD025
	1C-3300176	DC06	68-0C357DB	6718	050	DD022
	1C-3300176	DC06	68-0C36C54	6718	050	DD027

Figure 9-3 IBM i Start Hot Spare panel

The disk IOA does not control mirror protection, so when a mirror protected disk unit fails, the System Licensed Internal Code (SLIC) detects that the failed drive and completes the following recovery steps (not apparent to the user):

1. SLIC tells IOA to disable the hot spare.
2. The hot spare becomes non-configured.
3. The replace configured unit function is run to replace the failed drive with the now non-configured previous hot spare.
4. The failed drive becomes non-configured for safe physical replacement.

9.1.3 Dual storage I/O adapters

With IBM i 6.1.1 or later, any IBM i serial-attached SCSI (SAS) adapters with write cache that is used for internal disk attachment on POWER6 server or later are supported as dual SAS adapters. Both adapters of the dual SAS adapter pair must have the same size of write cache.

The new dual SAS adapter support provides adapter redundancy with an active and passive I/O path per RAID set, or a mirrored side in a two pair (four adapters) dual SAS adapter configuration with IBM i mirroring. Read and write disk I/O operations are sent by the system only down the active path. The passive path is used only after controller failovers (for example, if the active path fails). Dual SAS adapters are redundantly interconnected through a SAS adapter-to-adapter (AA) cable that connects the top ports of the SAS adapters, and a SAS X cable that attaches to the disk expansion drawer, as illustrated in Figure 9-4.




Figure 9-4 IBM i dual SAS adapter configurations

Both SAS adapters of a dual storage IOA configuration can perform I/O to the attached disk array units. The SAS adapter that is optimized for a configured RAID set is the one driving the I/O to the disk array units. In Figure 9-4, one SAS adapter is optimized for RAID set 0 and the other is optimized for RAID set 1.

Remember: For IBM i mirroring configurations, the disk units that are attached to a dual SAS adapter are each treated as a one-drive parity set.

For a dual SAS adapter pair, there are primary and secondary adapter roles. Only the primary adapter can perform disk management functions (such as creating a RAID array). If the primary adapter becomes unavailable, an automatic failover to the secondary adapter occurs, which becomes the primary adapter. There is no fallback to the original primary adapter when it comes back operational. The current role of a SAS adapter (as the primary or secondary adapter) can be seen by navigating to **System Service Tools** → **Start a service tool** → **Hardware service manager** → **Logical hardware resources** from the panel that shows the details for a dual SAS storage IOA. Select F14=Dual Storage IOA Configuration, as shown in Figure 9-5.




Figure 9-5 IBM i dual storage IOA configuration panel

You can view the disk unit paths for dual SAS adapter connected disk units that are reported as DMPxxx multipath disk unit resources. Navigate to **System Service Tools** → **Work with disk units** → **Display disk unit configuration** → **Display disk unit path status** (Figure 9-6).

Display Disk Path Status						
ASP	Unit	Serial Number	Type	Model	Resource Name	Path Status
1	1	Y2103LM0ACE5	433D	050	DMP147	Active
					DMP148	Passive
1	8	Y2103LN0868T	433C	050	DMP197	Passive
					DMP198	Active
2	2	Y680000FA16A	433C	050	DMP129	Active
					DMP130	Passive
2	3	Y6800024F6C9	433B	050	DMP131	Active
					DMP132	Passive
2	4	Y680000F12FD	433C	050	DMP115	Passive
					DMP116	Active
2	5	Y68000267272	433B	050	DMP135	Passive
					DMP136	Active
2	9	Y68000356821	433B	050	DMP170	Passive
					DMP169	Active
						More...
Press Enter to continue.						
F3=Exit F5=Refresh F9=Display disk unit details						
F11=Display encryption status F12=Cancel						

Figure 9-6 IBM i dual SAS adapter disk unit path status

You can get the best performance without compromising the availability in dual a SAS adapter RAID configuration by balancing the RAID parity sets across both adapters so that each adapter is assigned an equal amount of RAID parity sets with active paths to the disk units. To achieve this availability, the parity optimization method must be set to Performance before you create any RAID 5 or RAID 6 parity sets for dual SAS RAID adapters. See Figure 9-7.

Select Parity Optimization

Select how you want the parity set optimized:

The current parity optimization is: Performance

Type choice, press Enter.

Select parity optimization

1. Availability

2. Balance

3. Capacity

4. Performance

Selection

4

F3=Exit F12=Cancel

Figure 9-7 IBM i parity optimization selection menu

For more information about IBM i dual SAS adapter support, see the “Dual storage IOA configurations” topic in the IBM Systems Hardware Information Center at:

<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/ared5/ared5dualioaconfig.htm>

9.1.4 Encrypted ASP enhancements

The following new functions were implemented in IBM i 7.1 for ASP encryption supported by option 45 “Encrypted ASP Enablement”, which was originally introduced with IBM i 6.1:

- Start or stop ASP encryption for existing user ASPs or IASPs.

The changed encryption mode setting is applied for user ASPs when the system undergoes an IPL past DST and for IASPs when they are varied on. This situation does not mean that an IPL is required, but that the asynchronously run encryption or decryption tasks run only on a system that has undergone an IPL past DST.

Important: For geographic mirroring environments, encryption can be started or stopped only on the production IASP. Data is sent either encrypted or unencrypted to the backup node’s mirror copy IASP. This IASP ensures that the encryption attribute is set.

Figure 9-8 shows the new encryption options available by navigating to **System Service Tools** → **Work with disk units** → **Work with disk configuration** → **Work with encryption**.

Work with Encryption

Select one of the following:

1. Display encryption status

2. Create encrypted ASPs

3. Add units to encrypted ASPs

4. Start encryption on ASPs

5. Stop encryption on ASPs

6. Change data encryption key for basic ASPs

7. Change data encryption key for independent ASPs

Selection

F3=Exit F12=Cancel

Figure 9-8 IBM i work with encryption menu

- Change the data encryption key on existing encrypted user ASPs or IASPs.

The randomly generated 256-bit AES key for user ASP encryption (for example, the ASP master key that is used for IASP encryption) is securely stored in System Licensed Internal Code. This situation is the reason why a SAVSYS is recommended after you start encryption or change the encryption key, as shown in Figure 9-9.

Confirm Change Data Encryption Key for Basic ASPs

Note: This function may take a significant amount of time to complete. During this function, the partition performance may be degraded.

You should perform a Save System (SAVSYS) operation after the data encryption key is changed.

Data will be processed on all encrypted basic ASPs.

Do not change the data encryption key for basic ASPs again until this operation has completed. Do not stop encryption on basic ASPs until this operation has completed.

Press Enter to change the data encryption key.

F3=Exit F12=Cancel

Figure 9-9 IBM i change data encryption key confirmation panel

Requirement: For a clustering environment, an identical ASP master key, which protects the IASP data keys, must be manually created using the same parameter on each cluster node in the device domain to allow the IASP to be varied on.

9.1.5 Disk response time buckets enhancements

For a more granular disk I/O performance analysis, the disk response time buckets that were introduced with IBM i 6.1 and available in IBM i 5.4 since PTF SI23390 (Table 9-1) were extended in IBM i 7.1 from six buckets to 11 buckets (Table 9-2). Although the performance data for the existing six buckets is still stored in the QAPMDISK file, the new buckets are stored separately in the new QAPMDISKRB database file.

Table 9-1 IBM i 6.1 disk response time buckets

IBM i 6.1 disk response time bucket	Range
1	0 < 1 ms
2	1 ms < 16 ms
3	16 ms < 64 ms
4	64 ms < 256 ms
5	256 ms < 1024 ms
6	>= 1024 ms

Table 9-2 IBM i 7.1 disk response time buckets

IBM i 7.1 disk response time bucket	Range
1	0 < 15 us
2	15 us < 250 us
3	250 us < 1000 us
4	1000 us < 4000 us
5	4000 < 8000 us
6	8000 us < 16000 us
7	16000 us < 64000 us
8	64000 us < 256000 us
9	256000 us < 500000 us
10	500000 us < 1024000 us
11	>= 1024000 us

The Performance Data Investigator in IBM Systems Director Navigator for i and the Collection Services Investigator in IBM iDoctor for IBM i are enhanced with new collection services disk response time graphs for the new buckets in IBM i 7.1.

Figure 9-10 shows the disk response time bucket visualization from the IBM Systems Director Navigator for i perspective. You can access this view by navigating to **Collection Services** → **Disk** → **Disk Response Time** → **Detailed** → **Disk I/O Rates Overview – Detailed**.




Figure 9-10 IBM Systems Director Navigator disk response time buckets graph

For more information about the new disk response time buckets in QAPMDISKRB, see the IBM i 7.1 Information Center at:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=%2Frzahx%2Frzahxqapmdiskrb.htm>

9.1.6 Central processor complex node level mirroring

A system that has multiple central processor complex (CPC) nodes now mirrors disk units in a way that allows it to survive a node outage and that allows a concurrent maintenance repair of a CPC node.

When you start mirroring, the operating system considers the CPC node under which the disks are, and attempts to place the two subunits of a mirror protected pair under different CPC nodes. This action allows concurrent maintenance of a CPC node because the two subunits of each mirrored disk unit pair are under a different CPC node, allowing at least one subunit of each mirrored disk unit pair to remain operational during the maintenance operation.

After you install the PTF Group that contains this function, you might want to consider ending and restarting mirroring to recalculate the mirror protected pairs. There is an Advanced Analysis macro named LEVELOFFPROTECTION, which accessible through SST or DST, which you can use to verify the level of protection for each mirrored pair.

The LEVELOFPROTECTION macro is accessed from either the Dedicated Service Tools (DST) menu or the System Service Tools (SST) menu:

- To access the LEVELOFPROTECTION macro, click **System Service Tools** → **Start a service tool** → **Display/Alter/Dump** → **Licensed Internal Code (LIC) data** → **Advanced Analysis**.
- On the Select Advanced Analysis Command panel, there is a blank line at the top. Type a 1 in the Option column to select the blank line, then type LEVELOFPROTECTION, as shown in Figure 9-11. You may also scroll down the list and select the macro.

```

Select Advanced Analysis Command

Output device . . . . . : Display

Type options, press Enter.
1=Select

Option  Command
- - - - -
1      LEVELOFPROTECTION
-      QUEUESPACE
-      RMTTIME
-      SANITIZE
-      SEIZEINFO
-      SEMAPHOREINFO
-      SERVICEDOCS
-      DUMPTRC
-      LEVELOFPROTECTION
-      SMARTCHAIN
-      SPINLOCKTRACE
-      SSDSANITIZE
-      SYNCTOKENINFO

F3=Exit  F12=Cancel

More...
```

Figure 9-11 Selecting the LEVELOFPROTECTION macro

- Press Enter twice on the Select Advanced Analysis Command panel and the help panel is displayed, as shown in Figure 9-12.

```

Display Formatted Data
Page/Line. . . 1 / 1
Columns. . . : 1 - 78
Find . . . . .
.....1.....2.....3.....4.....5.....6.....7.....
DISPLAY/ALTER/DUMP
Running macro: LEVELOFPROTECTION -HELP
This macro generates a hardware config that the mirror code uses
when determining the mirror level of protection

Usage:
  [-UNIT] or [-U]      :Unit number to compute
  [-MP]                :compute muti-path level of protection
  [-ALL] or [-A]       :compute all

F2=Find  F3=Exit  F4=Top  F5=Bottom  F10=Right  F12=Cancel
Bottom

```

Figure 9-12 LEVELOFPROTECTION macro help

- In Figure 9-13, the -UNIT parameter is chosen and disk unit 12 is entered.

```

Specify Advanced Analysis Options
Output device . . . . . : Display
Type options, press Enter.
Command . . . . . : LEVELOFPROTECTION
Options . . . . . : -UNIT 12
F3=Exit  F4=Prompt  F12=Cancel

```

Figure 9-13 Specifying the -UNIT parameter to specify a disk unit

- The macro runs and displays the panel that is shown in Figure 9-14.

```

*
Display Formatted Data
Page/Line. . . 1 / 2
Columns. . . : 1 - 78

Find . . . . .
.....1.....2.....3.....4.....5.....6.....7.....
Running macro: LEVELOFPROTECTION -UNIT 12
UNIT 12
BUSL 'L' 0xD3 'L' 0xD3
CEC 472 476
RING 1670 1930
TOWER 15360 15363
BUSNAME LB44 LB33
BUS 527 379
BOARD 0 0
CARD 0 0
IOP 0 0
IOA -1 -1
IO BUS 0 4
CONTRLR 2 11
DEVICE 0 0
LEVEL PRO 75 CecNodeLevelOfProtection
Bottom
F2=Find F3=Exit F4=Top F5=Bottom F10=Right F12=Cancel

```

Figure 9-14 Level of protection display

The line at the bottom of the display in the box indicates the level of disk protection, which in this case is CecNodeLevelOfProtection.

9.1.7 EXP24S SFF Gen2-bay drawer (#5887)

The EXP24S is a high-density, high-performance SFF drive drawer, holding up to 24 SAS drives in 2U of 19-inch rack space. It has 6 Gb of SAS I/O capability.

The #5887 drawer has double the number of drives than the EXP12S I/O drawer (#5886) and the SFF drives provide significant energy savings compared to the EXP12S 3.5-inch drives.

For more information, see IBM Hardware Announcement letter 111-065 at:

<http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&appname=iSource&supplier=897&letternum=ENUS111-065>

9.1.8 Higher capacity 10 K RPM SFF SAS disk drives

283 GB and 571 GB 10 K RPM SFF disk drives are available, offering a lower cost per gigabyte and more gigabytes per rack space than previous technology. These drives are available in Gen1 and Gen2 features.

For more information, see IBM Hardware Announcement letter 111-065.

9.1.9 Tape performance instrumentation

A new Collection Services *RMVSTG category was introduced in IBM i 7.1 for performance data collection for removable media tape storage devices (for internal and external tape devices).

The tape performance statistics data is stored in the QAPMTAPE structured database file, including physical tape performance I/O statistics counts, such as number of reads and writes, bytes read and written, and number of tape marks and blocks spaced. This data is tracked by the IBM i tape code when you send requests to the tape device driver. Currently, for reviewing the data that is collected in QAPMTAPE, either a user-defined SQL query or a GUI, such as the Systems Director Navigator for i with its Investigate Data function, must be used.

For more information about the structured QAPMTAPE database file performance data, see the IBM i 7.1 Information Center at:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=%2Frzahx%2Frzahxqapmtape.htm>

9.1.10 Tape library resource name changes for IOP-less IOA attachment

When you upgrade a system to IBM i 7.1, a new hardware resource name is assigned to the tape library devices attached through an IOP-less Fibre Channel or SAS adapter. You must ensure that the tape library device description points to the correct hardware resource after the upgrade. Either the hardware resource name must be changed through the System Service Tools Hardware Resource Manager or the device description's resource name field must be updated (for example, by running **CHGDEVMLB**) with the new tape library resource name, which was assigned with IBM i 7.1.

9.1.11 Tape library unreadable barcode changes for IOP-less IOA attachment

Before IBM i 7.1, if, at varyon of the tape library, tape cartridges with unreadable barcodes are found, each of these tape cartridges is loaded in to a drive to read the volume ID. The volume ID is used to generate a corresponding cartridge ID for the unreadable barcode.

This method ensures, for IBM standard labeled (VOL1) tapes, that the volume ID matches the cartridge ID, which is a requirement for IBM i to allow *write* operations to a tape cartridge. The downside of this approach is the time that is required to load and read each cartridge, especially if the library barcode reader itself failed. Also, problems with the barcode label or barcode reader are not made apparent to the user.

With IBM i 7.1 and IOP-less IOA attached tape libraries, if a tape cartridge with an unreadable or missing barcode is manually added, a cartridge ID with a format of UNKXXX is fabricated, with XXX being a sequential decimal number that starts with UNK001. If a cartridge is found in a storage slot with an unreadable barcode, a cartridge ID is fabricated with the format of U@XXXX, with XXXX reflecting the SCSI element address when the tape device driver discovers an unreadable barcode in a slot.

This handling of unreadable barcodes in IBM i 7.1 reveals barcode problems and allows you to read from tapes without barcode labels (which are removed from the library again) quicker, without requiring a tape drive for generating cartridge IDs.

Consideration: With the IBM i 7.1 IOP-less IOA tape library attachment, you should not use cartridges without barcode labels if they are supposed to remain in the library. To write or append to a standard labeled cartridge in a library, a barcode label that matches the volume ID must be stuck on the cartridge.

9.1.12 DVD / Tape SAS External Storage Unit for Power 795 CPC Rack

The #5274 DVD / Tape SAS External Storage Unit for Power 795 CPC Rack is a 1U storage unit that can hold HH DAT160 drives, the #5638 1.5 TB / 3.0 TB LTO-5 SAS Tape Drive, or slimline DVD drives.

For more information, see IBM Hardware Announcement letter 111-065 found at:

<http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&appname=iSource&supplier=897&letternum=ENUS111-065>

9.1.13 RDX support

RDX drives are positioned as an entry tape alternative to VXA-2, VXA-320, DAT72, DAT160, DAT320, or 8 mm, but they are not intended as an alternative to LTO. The offered capacities are 160 GB - 1 TB per removable cartridge.

RDX drives are available through native attach in IBM i 7.1 TR5 with either USB or SATA. SATA support will be added to 6.1.1 through a PTF. The drives are available for POWER7 systems only.

For a USB attach to a Version 7.1 system, you can either use F/C #EU03 (USB) or #EU04 (USB); for a Version 6.1.1 system, you must use F/C #EU07 (SATA).

The RDX dock is available in either 5.25-inch internal (SATA or USB) format or external USB format. The dock supports all RDX cartridges, which have 30+ years of reliability and are rugged.

You can see a comparison matrix between low end tape drives and RDX in Figure 9-15.

	USB Removable Disk Drive (RDX)		Tape Drives
Price drive / dock	Much lower	★	Higher
Maintenance drive/dock	Zero	★	Not zero
Price cartridge	Higher		Lower
Able to add next higher capacity cartridge	Use same docking station	★	Need new drive
Capacity	Up to 1 TB	★	DAT only up to 80 GB uncompressed
Compression	Done by server, not docking station		Most drives have built in compression
Auto loader	No		Not offered on lower end drives
Susceptible to dirt/dust	Very resistant	★	Susceptible
Random access to middle or end of cartridge	Yes (with UDF file system)	★	No
Rugged (drop and contamination resistant)	Yes	★	Less rugged
Encryption	Done by server application, not docking station		Lower end drives use server application, not drive
Head cleaning and cleaning cartridge	No	★	Yes
Backup/Restore time	Somewhat faster than entry	★	Somewhat slower than RDX
External packaging/attachability	Small, mobile, easy USB attach		Larger, SAS/SATA PCI adapter attachment. DAT160 has USB attach option.

Figure 9-15 RDX versus low end tape

Although RDX is a tape replacement, when you configure or operate RDX on IBM i, think of it as a DVD. The devices do not show up as TAP devices. Instead, look for a Removable Mass Storage (RMS) device.

For more information about RDX for IBM i clients, see the section “Technology Release 5 (TR5)” in the Storage Solutions manual, found at:

<http://wilson.boulder.ibm.com/infocenter/iseriess/v7r1m0f/index.jsp?topic=%2Frzam4%2Frzam4rms.htm>

IBM i 7.1 can virtualize RMS devices to other IBM i partitions. VIOS does not virtualize RDX.

9.1.14 POWER7+ 770/780 Native I/O support

For POWER7+ 770 or 780 systems, native I/O attach is supported only for IBM i 7.1. IBM i 6.1 partitions can use native I/O only if virtualized by a VIOS or a Version 7.1 partition, as shown in Figure 9-16.




Figure 9-16 POWER7+ native I/O

Although this virtualization is easily done for disk, SAN, Tape, LAN and optical devices, there is no virtualization support for WAN or crypto cards.

9.2 SAN storage management enhancements

This section describes the following IBM i storage management enhancements specific to IBM i external storage area network (SAN) storage:

- ▶ Multipathing for virtual I/O
- ▶ DS5000 native attachment
- ▶ Level of protection reporting for multipath disk units
- ▶ Library control paths for IOP-less Fibre Channel IOA tape attachment
- ▶ External disk storage performance instrumentation
- ▶ Thin provisioning for DS8700 and DS8800 storage servers and for VIOS shared storage pools

9.2.1 Multipathing for virtual I/O

IBM PowerVM Virtual I/O Server IBM i client support was introduced with IBM i 6.1.

With IBM i 6.1.1 or later, a redundant VIOS configuration (Figure 9-17) is supported by IBM i multipathing across two or more VIOS on the same IBM Power Systems server for protection against VIOS outages because of VIOS updates.




Figure 9-17 IBM i multipathing with a redundant Virtual I/O Server configuration

This new IBM i multipathing support for virtual I/O eliminates the need to use IBM i mirroring for a redundant VIOS configuration, which required duplicate storage capacity.

For further IBM i virtualization enhancements, such as Active Memory Sharing or N_port ID virtualization support, see Chapter 8, “Virtualization” on page 311.

9.2.2 DS5000 native attachment

IBM i SAN storage support is extended with IBM i 6.1.1 to support native attachment of the IBM System Storage DS5100 and DS5300 systems to IBM Power Systems POWER6 or later servers. This new native attached DS5100 and DS5300 storage support provides an easier storage setup (see Figure 9-18) and configuration without needing to deploy the IBM PowerVM VIOS. It is only if some other advanced virtualization functions are being used on IBM i that you must configure a VIOS.




Figure 9-18 IBM i DS5000 native attachment

Figure 9-19 shows how the native attached DS5000 LUNs, created for the IBM i host, report on an IBM i host as device type D818.

Display Non-Configured Units						
Serial Number	Type	Model	Resource Name	Capacity	Status	
Y2103LQ0WGLC	433B	050	DPH001	69793	Non-configured	
Y2103LQ1J064	433B	050	DPH002	69793	Non-configured	
Y2103LQ1J06H	433B	050	DPH003	69793	Non-configured	
Y2103LQ0P0BE	433B	050	DPH004	69793	Non-configured	
Y2103LQ1HV0C	433B	050	DPH005	69793	Non-configured	
Y2103LQ1J6M8	433B	050	DPH006	69793	Non-configured	
Y0C44AC5B4F6	D818	099	DPH007	265333	Non-configured	
Y0C14AC5A32B	D818	099	DPH008	265333	Non-configured	
Press Enter to continue.						
F3=Exit		F5=Refresh		F9=Display disk unit details		
F11=Display device parity status			F12=Cancel			

Figure 9-19 Native attached DS5000 LUNs on IBM i

Storage capacity: Because of the 4 KB page sector conversion from 8 x 520 bytes sectors to 9 x 512 bytes sectors by IBM i SLIC for DS5000 native attachment, the reported usable IBM i capacity is approximately 89% of the configured DS5000 LUN capacity.

The built-in IBM i multipathing in System Licensed Internal Code (SLIC) adheres to the DS5000 active / passive controller concept. Under normal working conditions, I/O is driven across only the active paths to a disk unit (to the controller designated for the LUN as the preferred controller) when the passive paths for a disk unit are used at DS5000 controller failover conditions. Figure 9-20 shows the active and passive path for disk units from a native attached DS5000 after they are added to an ASP. You can access this panel by navigating to **System Service Tools** → **Work with disk units** → **Display disk configuration** → **Display disk path status**.

Display Disk Path Status						
ASP	Unit	Serial Number	Type	Model	Resource Name	Path Status
2	25	Y0C14AC5A32B	D818	099	DMP002	Active
					DMP004	Passive
2	26	Y0C44AC5B4F6	D818	099	DMP001	Passive
					DMP003	Active

Press Enter to continue.

F3=Exit F5=Refresh F9=Display disk unit details
F11=Display encryption status F12=Cancel

Figure 9-20 IBM i active / passive paths for DS5000 disk units

The requirements for IBM i DS5000 native attachment are:

- ▶ IBM i POWER6 or later servers only
- ▶ IBM i 6.1.1 (OS resave RS610-10, SLIC RS611-A, or cumulative PTF C9279610) or later
- ▶ IOP-less Fibre Channel IOA (#5774, #5749, or #5735)
- ▶ DS5100 or DS5300 only
- ▶ DS5000 FW 7.60.28.00 or later (including NVSRAM N1818D51R1060V08 for DS5100 and N1818D53R1060V08 for DS5300 or later)
- ▶ DS5000 Storage Manager 10.60.x5.17 or later
- ▶ DS5000 IBM i Host Kit Feature Code 7735

The following considerations apply for IBM i DS5000 native attachment:

- ▶ The maximum supported LUN size for IBM i is less than 2 TB.
- ▶ The usable IBM i net capacity is 8 / 9 of the configured DS5000 LUN capacity.
- ▶ A maximum of 64 LUNs per IBM i Fibre Channel IOA port is allowed.
- ▶ Unprotected arrays (RAID 0) are not supported for IBM i.
- ▶ IBM i mirroring is not supported for DS5000.
- ▶ Multipathing on a single dual-port Fibre Channel IOA is not supported.

- ▶ DS5000 Dynamic Volume Expansion (DVE) is not supported for IBM i.
- ▶ SSDs in DS5000 are not supported for IBM i.

From an IBM i disk I/O performance perspective, the following preferred practices should be followed:

- ▶ To balance workload across both DS5000 controllers, LUNs should be evenly assigned regarding preferred controller affinity to controller A and B.
- ▶ The LUN size for IBM i IOP-less Fibre Channel of 70 GB applies for DS5000 native attachment.
- ▶ A DS5000 segment size of 128 KB is generally a good compromise for both IBM i transaction and save / restore workload.

For more information about the IBM System Storage DS5000 series, see the following IBM Redbooks publications:

- ▶ *IBM Midrange System Storage Hardware Guide*, SG24-7676
- ▶ *IBM System Storage DS Storage Manager Copy Services Guide*, SG24-7822

For more information about IBM support statements regarding DS5000 Copy Services support with IBM i native attached DS5000, see *IBM i Virtualization and Open Storage Read-me First*, found at:

http://www-03.ibm.com/systems/resources/systems_i_Virtualization_Open_Storage.pdf

IBM STG Lab Services developed a Copy Services Tool Kit offering Advanced Copy Services for PowerHA - DS5000 Edition for DS5000 native-attachment to support IASP storage-based replication solutions with FlashCopy / VolumeCopy and Enhanced Remote Mirroring. For more information about this Copy Services Tool Kit offering for DS5000, see IBM STG Lab Services at:

<http://www-03.ibm.com/systems/services/labservices>

9.2.3 Level of protection reporting for multipath disk units

With IBM i 7.1, the level of protection for multipath attached external disk units is now reported for any multipath disk unit devices from either a natively or VIOS-attached disk storage system. The reported levels of multipath protection reflect the component that can fail without jeopardizing I/O access to the disk units and are the same as the ones used for mirrored protection:

- ▶ Remote Bus
- ▶ Ring (HSL / 12X Loop)
- ▶ Tower
- ▶ Bus
- ▶ IOP
- ▶ IOA
- ▶ IOA-Bus

Figure 9-21 shows the new multipath protection level reporting for the example of DS8000 disk units each attached through three paths. To view this panel, select **System Services Tools → Work with disk units → Display disk unit configuration → Display protection for multiple connection disk units**.

Display Protection for Multiple Connection Disk Units						
ASP	Unit	Serial Number	Type	Model	Resource Name	Protection
7	11	50-70005F0	2107	A04	DMP007	Ring
7	11	50-70005F0	2107	A04	DMP012	Ring
7	11	50-70005F0	2107	A04	DMP009	Ring
7	14	50-53007F0	2107	A04	DMP407	Bus
7	14	50-53007F0	2107	A04	DMP111	Bus
7	14	50-53007F0	2107	A04	DMP208	Bus
Press Enter to continue.						
F3=Exit			F5=Refresh			
F11=Display disk configuration status			F12=Cancel			

Figure 9-21 IBM i protection level reporting for multipath disk units

9.2.4 Library control paths for IOP-less Fibre Channel IOA tape attachment

Tape library devices that are attached to a dual-port Fibre Channel I/O adapter with IBM i 7.1 require at least one control path drive to be attached to each port because the design changed from an adapter-centric to a port-centric control path architecture.

The tape device driver ensures that, from a user perspective, only one library resource per Fibre Channel IOA port is presented for the same logical library, even if multiple control paths are defined. IBM i pools these libraries so all the TAPxx resources for the library are in one TAPMLBxx device description.

Requirement: For IBM i 7.1, a second library control path must be added, preferably before the upgrade to IBM i 7.1, for the second port of a dual-port IOP-less Fibre Channel IOA. Otherwise, the tape drives on the second port can become stand-alone devices without library capability.

Before IBM i 7.1, only one control path drive was required per Fibre Channel IOA for drives in the same logical library. Only one library resource per Fibre Channel IOA is presented for the same logical library, even if multiple control paths are defined.

9.2.5 External disk storage performance instrumentation

New external disk storage performance metrics for IBM System Storage DS8000 and DS6000 series are available with Collection Services in IBM i 7.1. This new data is collected with the new *EXTSTG category and stored in the QAPMXSTGD database file.

Because of a minimum DS8000 Release 4 microcode requirement to support this new external storage performance data collection, the *EXTSTG category is not included in any default collection profile to prevent Product Activity Log (PAL) hardware failure information entries if this DS8000 code requirement is not met. To enable QAPMXSTGD external storage performance data collection with the *STANDARD or *STANDARDP default collection profiles, the following steps as are required to add the *EXTSTG category to these profiles:

1. Run `ENDPFCOL FRCCOLEND(*YES)`.
2. Run `RNMOBJ OBJ(QUSRSYS/QPFCOLDTA) OBJTYPE(*USRSPC) NEWOBJ(QPFCOLDT2)`.
3. Run `CALL QSYS/QYPSCOLDTA PARM('*EXTSTG')`.
4. Run `STRPFCOL`.

For more information about these DS8000 external storage performance data collection requirements, see the *IBM i Memo to Users 7.1* at:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzaq9/rzaq9.pdf>

Requirement: DS8000 Microcode Release 4 or later is required for the QAPMXSTGD external storage performance data collection.

This new QAPMXSTGD database file contains DS8000 or DS6000 external storage subsystem performance data, including Fibre Channel link statistics and rank (RAID array) statistics. The QAPMXSTGV database file that was introduced in IBM i 6.1.1 and part of the *DISK category included in all default collection profiles contains volume level (that is, logical unit (LUN)) cache statistics performance data.

Both the QAPMXSTGD and QAPMXSTGV files store vendor-specific SCSI Log Sense page data in unstructured large data fields. Access to at least a single IBM i LUN on the DS8000 or DS6000 storage system is required to retrieve this log sense data from it when the SCSI Log Sense command is issued against a LUN.

IBM iDoctor for IBM i external storage performance analysis functions

The IBM iDoctor for IBM i suite of analysis tools is recommended for analyzing the external storage performance data. IBM iDoctor for IBM i build C00777 or later is required. This suite has new functions for visualizing and analyzing DS8000 or DS6000 storage performance data.

The new iDoctor Collection Services Investigator functions for analyzing the external storage performance log sense data that is stored in QAPMXSTGV (Log sense page 0x32) and QAPMXSTGD (Log sense pages 0x33 and 0x34) are shown in Figure 9-22.




Figure 9-22 iDoctor analysis functions for external storage

Before you use iDoctor to analyze the external storage performance data, you must click **Analyses** → **Run External Storage** to generate structured SQL tables from the Log Sense data and the new “External storage cache statistics” and “External storage link and rank statistics” report folders. A refresh of the view might be required to display them.

Newly available external storage cache statistics data are shown in Figure 9-23 from the report that is generated by clicking **External storage cache statistics** → **by time interval** → **IO rates totals with cache hits**. The read cache hit% information was available from the QAPMDISK data, but the newly reported write cache hit% from QAPMXSTGV data can check for any potential storage subsystem write cache overruns. These overruns are indicated by write cache hits% < 100%, and might warrant changes in the workload schedule or a cache size upgrade.




Figure 9-23 iDoctor external storage cache statistics

Valuable analysis functions for DS8000 or DS6000 rank and link performance data are available from the External storage link and rank statistics reports.

For example, potential rank overuse issues can easily be visualized and analyzed using a ranking view of the rank IDs based on total I/O by clicking **Rank graphs** → **By rank ID** → **Ranks IO rates totals**. Then, from this view, select one or more ranks with a high I/O rate for a more detailed analysis by clicking **Selected Ranks** → **Ranks IO rates** from the right-click menu, as shown in Figure 9-24.




Figure 9-24 iDoctor rank I/O rates ranking

The read and write rank I/O rates over time for the individually selected rank IDs 5 and 6 from our example are shown in Figure 9-25. If these were RAID 10 instead of RAID 5 ranks, you also could have chosen the option to display the disk I/O rates. This option is not available for RAID 5, because their disk I/O rates cannot be determined accurately from the rank read and write I/O rates.




Figure 9-25 iDoctor rank I/O rates for selected ranks

Similar to the rank I/O performance analysis, reports with graphing functions for host (SCSI) or Peer-to-Peer Remote Copy (PPRC) link performance analysis were added to iDoctor, as shown in Figure 9-26.




Figure 9-26 iDoctor link I/O rates

For more information about the IBM iDoctor for IBM i powerful suite of performance tools, go to the iDoctor website, which offers also a 45-day trial version, at:

https://www-912.ibm.com/i_dir/idoctor.nsf/iDoctor.html

9.2.6 Thin provisioning for DS8700 and DS8800 storage servers and for VIOS shared storage pools

Thin provisioning for DS8700 and DS8800 storage servers, and for VIOS shared storage pools, allows configurations to be set up with a small amount of real disk storage, which can be increased later without changing the partition's view of the storage LUN. Before this enhancement, the full amount of configured storage was allocated at LUN initialization time.

Here are the thin provisioning requirements:

- ▶ **Hardware requirements:**

Thin provisioning enhancement for DS8000 storage servers requires a DS8700 or DS8800 with Release 6.2, available from IBM through FC #1723, or through bundles:

- IBM System Storage DS8700 - level 7.6.2.xx.xx (bundle version 76.20.xxx.xx), or later
- IBM System Storage DS8800 - level 7.6.2.xx.xx (bundle version 86.20.xxx.xx), or later

- ▶ **Software requirements:** IBM i 7.1 with the newest Technology Refresh PTF Group

9.3 SSD storage management enhancements

IBM i with its single-level storage architecture, integrated DB2 database, storage performance analysis, and storage management capabilities is an industry-leading platform for SSD hierarchical storage management.

The integrated hierarchical storage management functions for SSDs in IBM i, such as the DB2 for i and UDFS media preferences or the ASP balancer enhancements for SSDs, allow for an easy and efficient implementation of SSDs on the IBM i platform.

SSDs based on flash memory are considered a revolutionary technology for disk I/O performance and energy efficiency compared to traditional spinning disk drives. SSD I/O response times can be over 200 faster than for spinning disk drives. SSDs are supported in IBM i 5.4.5 or later for IBM i internal storage, and IBM i 6.1.1 plus PTF MF47377 or later if used in IBM System Storage DS8000 series with R4.3 code or later.

For more information about the benefits and usage of SSDs with IBM i, see *Performance Value of Solid State Drives using IBM i*, which is available at the following website:

http://www-03.ibm.com/systems/resources/ssd_ibmi.pdf

The SSD Analyzer Tool for IBM i is a good tool to use for a first analysis about whether SSDs can help improve performance for a particular IBM i system. The tool queries existing Collection Services performance data for retrieving the average system and optional job level disk read I/O response times to characterize whether the workload is a good candidate for SSDs. It can be downloaded as an IBM i save file from the following website:

<http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS3780>

For a reference about the IBM i PTF requirements for SSDs, see the IBM i Software Knowledge Base topic “Requirements for Solid State Drives (SSD)”, which is available at the following website (search for KBS document number 534676318):

http://www-912.ibm.com/s_dir/slkbase.nsf/slkbase

The following subsections describe recent enhancements for management of SSDs in an IBM i environment.

- ▶ DB2 media preference
- ▶ ASP balancer enhancements for SSDs
- ▶ User-defined file system media preference
- ▶ 177 GB SFF SSD with eMLC
- ▶ IBM Disk Sanitizer PRPQ extended to include SSD devices

9.3.1 DB2 media preference

DB2 for i was extended with support for database object placement on SSDs or HDDs. This new function, called *DB2 media preference*, allows the user to have control over which media type selected database files are stored so that DB files that are known to be I/O performance critical can be placed explicitly on high performing SSDs.

Physical and logical DB files (table and indexes) are enhanced with a preferred media attribute that can be set through the **-UNIT** parameter for a certain media preference of either SSDs. You can use **UNIT** parameter values ***SSD** (CL commands) or **SSD** (SQL) for i 6.1 and later, values **255** (CL commands) for i 5.4.5, or HDDs with **UNIT** parameter values ***ANY** (CL commands) or **ANY** (SQL) when create or change these files through the following CL commands or SQL statements:

- ▶ CRTPF, CRTLF, CRTSRCPF, CHGPF, CHGLF, and CHGSRCPF
- ▶ CREATE TABLE, CREATE INDEX, and ALTER TABLE

Parameter usage:

- ▶ The UNIT parameter for the SQL statements is supported by IBM i 6.1 or later.
- ▶ For a partitioned SQL table, the **ALTER TABLE** statement can be used to set a media preference on a partition (member) level.

Figure 9-27 shows the new preferred storage unit parameter (UNIT keyword) for the **CHGPF** command.




Figure 9-27 IBM i CHGPF command

Dynamic data movement

Enhancements were implemented for dynamically changing the DB2 media preference, which starts a dynamic data move done synchronously.

For releases before IBM i 7.1, the following PTFs are required for a dynamic move of physical or logical database files after you change their media preference attribute. Otherwise, a save and restore of those changed database files is required to make the media preference change effective.

- ▶ IBM i 5.4.5 PTFs MF47887, MF47891, and MF47878
- ▶ IBM i 6.1.0 PTFs MF47888, MF47892, and MF47879
- ▶ IBM i 6.1.1 PTFs MF47889, MF47893, and MF47877

DB2 random and sequential reads statistics

To help with SSD media management from a database business logic perspective for determining which database files are good candidates for placement on SSDs, two new fields (RANDOM_READS and SEQUENTIAL_READS) were introduced in IBM i 7.1 for each keyed logical and physical database file. These two 8 byte counters are used to track the amount of random and sequential logical read I/O for each database file, and are continuously updated by the database and reset only at IPL. Because a single logical read I/O can lead to more than one random I/O (for example, because of variable length fields (> 32 KB) or large objects (LOBs), the new RANDOM_READS and SEQUENTIAL_READS usually do not sum up to the reported LOGICAL_READS).

The preferred procedure (after database performance optimization is completed from an application and system perspective and further optimization is warranted at the storage hardware level) to help determine which database files are good candidates for placement on SSDs is as follows:

1. Look at a storage I/O performance critical time period.
2. Compare the RANDOM_READS numbers at the start and end of the time period.
3. Determine the DB files with highest RANDOM_READS I/O count differences, and if these files are critical from a business perspective, they might be good candidates for using DB2 media preference to move them to SSDs.

To query the RANDOM_READS counter for database files, a SQL query against QSYS2/SYSPARTITIONSTAT for physical file statistics or SYSINDEXSTAT for keyed logical file statistics (Example 9-1) or the System i Navigator's Health Center activity tab (Figure 9-28) can be used. Save the query results and use the View History function to compare the results that are retrieved for the start and the end of the critical time period.

Example 9-1 SQL query for physical database file random reads

```
SELECT table_name, logical_reads, random_reads, sequential_reads FROM
QSYS2.SYSPARTITIONSTAT WHERE logical_reads > 0 ORDER BY random_reads DESC
```




Figure 9-28 System i Navigator database health center

Macro for changing storage management SSD allocations

To prevent newly created objects without a media preference from being placed on SSDs in a hybrid ASP with HDDs and less capacity used SSDs, run the **smsetstayoffssd** macro. After you run the macro, storage management tries to stay away from SSDs for non-media preference objects, as shown in Example 9-2. The **smresetstayoffssd** macro resets the storage allocation setting back to the default setting of “Non-media preference will go to best unit”, that is, the lowest percent capacity unit.

Example 9-2 Macro smgetstayoffssd

DISPLAY/ALTER/DUMP

Running macro: SMGETSTAYOFFSSD

Non-media preference will try to stay off SSD.

In Example 9-3, the **smgetstayoffssd** macro is used to reset the storage allocation setting back to the default for a specific independent ASP. For IASPs, the ASP number in hex is required on the **smsetstayoffssd** macro.

Example 9-3 Storage allocation setting

DISPLAY/ALTER/DUMP

Running macro: SMGETSTAYOFFSSD 91

Non-media preference will try to stay off SSD for ASP 0x91.

Statistical view for reporting unit allocations

A new view named SYSPARTITIONDISK, which is used to support reporting of SSD versus HDD space usage for SQL tables and native tables (physical files), is available through PTFs for IBM i 5.4, 6.1, and 7.1.

Figure 9-29 shows a query result example from the new SYSPARTITIONDISK view after you use the ASP balancer to move DB files to SSD and manually move the HSTRY01 table to SSD through DB2 media preference. For each table, the columns SSD_SPACE versus NON_SSD_SPACE show its storage space in bytes allocated on SSDs and non-SSDs (HDDs).

```

select max(table_schema) as table_schema, max(table_name) as table_name,
max(table_partition) as table_partition,
sum(case when disk_type in ('58B0', '58B1') then unit_space_used else null end) as ssd_space,
sum(case when disk_type not in ('58B0', '58B1') then unit_space_used else null end) as non_ssd_space
from qsys2.syspartitiondisk a
where system_table_schema = 'CPW3DB'
group by a.table_schema, a.table_name, table_partition
order by 1,2,3;

```

TABLE_SCHEMA	TABLE_NAME	TABLE_PARTITION	SSD_SPACE	NON_SSD_SPACE
CPW3DB	AREFFIL	AREFFIL	4096	-
CPW3DB	CSTMRPF01	CSTMRPF01	47710208	354942976
CPW3DB	CSTMRPF02	CSTMRPF02	45875200	356777984
CPW3DB	CSTMRPF03	CSTMRPF03	49807360	352845824
CPW3DB	CSTMRPF04	CSTMRPF04	50855936	351797248
CPW3DB	CSTMRPF05	CSTMRPF05	53739520	348913664
CPW3DB	CSTMRPF06	CSTMRPF06	47185920	355467264
CPW3DB	CSTMRPF07	CSTMRPF07	48496640	354156544
CPW3DB	CSTMRPF08	CSTMRPF08	50331648	352321536
CPW3DB	DLVRYLOG	DLVRYLOG	-	24117248
CPW3DB	DLVRYLOGA	DLVRYLOGA	4096	-
CPW3DB	DSTRCT01	DSTRCT01	233472	-
CPW3DB	DSTRCT02	DSTRCT02	233472	-
CPW3DB	DSTRCT03	DSTRCT03	233472	-
CPW3DB	DSTRCT04	DSTRCT04	233472	-
CPW3DB	DSTRCT05	DSTRCT05	233472	-
CPW3DB	DSTRCT06	DSTRCT06	233472	-
CPW3DB	DSTRCT07	DSTRCT07	233472	-
CPW3DB	DSTRCT08	DSTRCT08	229376	4096
CPW3DB	HSTRY01	HSTRY01	48234496	-
CPW3DB	HSTRY02	HSTRY02	2097152	46137344

select max(table_schema) as table_schema, max(table_name) as table_name, max(table_partition) as table_partition, sum(case when disk_type

Figure 9-29 SYSPARTITIONDISK view query result

For more information about the new SYSPARTITIONDISK view and function, see the topic “IBM DB2 for i Statistical View for Solid State Drive Storage Usage Reporting” at:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD105463>

9.3.2 ASP balancer enhancements for SSDs

The ASP balancer function for hierarchical storage management (HSM), which traditionally supports data migration between high performance and low performance (that is, compressed) hard disk drives (HDDs), is extended to support data migration between SSDs and HDDs.

Based on the read I/O count statistics for each 1 MB auxiliary storage extent of an ASP collected during a preceding **TRCASPBAL** run, the ASP balancer enhanced HSM function supports migration of frequently accessed hot extents from HDDs to SSDs and rarely accessed cold extents from SSDs to HDDs. By default, the ASP balancer tries to place all those frequently accessed extents on SSDs, which account for 50% of the total read I/O count.

Typically, the ASP balancer tracing function **TRCASPBAL** is run over a critical I/O workload period, such as a batch processing window, which is optimized for performance with using SSDs. Afterward, the ASP balancer HSM function is started to migrate off both cold data from SSDs and hot data to SSDs. **TRCASPBAL** can be accumulative. Users could clear data at the start of the week, collect the trace across the nightly batch work load window for the week, and balance on the weekend.

Example 9-4 illustrates a typical usage of the ASP balancer tracing and migration functions by clearing the trace statistics first, collecting new trace statistics, starting the migration, and monitoring its completion with the **CHKASPBAL** command.

Example 9-4 ASP balancer tracing and migration

```
TRCASPBAL SET(*CLEAR) ASP(1)
TRCASPBAL SET(*ON) ASP(1) TIMLMT(*NOMAX)
...
TRCASPBAL SET(*OFF) ASP(1)

STRASPBAL TYPE(*HSM) ASP(1) TIMLMT(*NOMAX)
CHKASPBAL
```

The initial ASP balancer accounting only for the extent read I/O counts is enhanced with a more efficient migration algorithm in the weighted ASP balancer version and additional functionality regarding SSD media management.

Weighted ASP balancer

Enhancements were implemented for the HSM function of the ASP balancer for migration of frequently accessed hot data to SSDs and infrequently accessed cold data to HDDs for hybrid ASPs consisting of SSD and HDD disk units.

With IBM i 6.1 plus supersede PTF MF49399, IBM i 6.1.1 plus supersede PTF MF48544, and with IBM i 7.1 base code, the ASP balancer's decision for moving hot or cold data to and from SSDs is now based on a weighted disk read I/O count for the 1 MB auxiliary storage segments to be moved. Not only is the amount of read I/O accesses to a segment counted as before, but its read service time is considered for the migration decision.

This weighted ASP balancer enhancement accounting for the read service times provides more efficient data media placement. For example, frequently accessed data that is derived mainly from read cache hits can no longer be prioritized for migration to SSDs, as it cannot benefit from being placed on SSDs.

ASP balancer migration priority

In IBM i 7.1, the ASP balancer is enhanced with an option that allows the user to specify the migration priority for ***MOVDTA**, ***HSM**, or ***MP** operations at levels of either ***LOW**, ***MEDIUM**, or ***HIGH** (as shown in Figure 9-30). This option influences the number of SLIC internal data moving tasks that are used for the migration. This option is subject to an inherent trade-off between speed of data migration and its effect on disk usage.

Start ASP Balance (STRASPBAL)

Type choices, press Enter.

Balance type		*CAPACITY, *USAGE, *HSM...
ASP number		1-32, *ALL
+ for more values		
ASP device		Name, *ALLAVL
+ for more values		
Storage unit		Number
+ for more values		
Time limit		1-9999 minutes, *NOMAX
Balance priority	*MEDIUM	*LOW, *MEDIUM, *HIGH
Subtype	*CALC	*CALC, *HDD, *SSD

Bottom

F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

Figure 9-30 IBM i ASP balancer migration priority

Also, the **STRASPBAL** command syntax changed in IBM i 7.1 to include a new subtype parameter that, for the ***HSM** balance type, now allows data migration between up to three storage tiers. Tiered storage is the assignment of different categories of data to different types of storage media to reduce total storage cost. You can have the following types of data migration:

- ▶ With subtype ***SSD**, you can have data migration between SSDs and high performance HDDs.
- ▶ With subtype ***HDD**, you can have data migration between high performance HDDs and low performance (compressed) HDDs.

Unless an ASP has disk units from all three storage tiers, the default subtype ***CALC** can be used.

Data migration with the ***HSM** balance type is performed in two phases, with cold data moved off from SSDs first, and then hot data moved to SSDs.

ASP balancer media preference balance type

The *MP balance type is a new ASP balancer function in IBM i 7.1 that helps correct any issues with media preference flagged DB objects or UDFS files not on their preferred media type, which is either SSDs or HDDs. This sweeper function moves objects marked with a media preference attribute to SSDs and non-media preference objects to HDDs when you use the default subtype *CALC. To limit the scope of media preference migration to only one direction, either the *SSD or *HDD subtype can be used for specifying the source media type for the migration.

For earlier releases, this media preference sweeper function is available with the following SST Advanced Analysis interface macros in IBM i 6.1.1 through PTF MF49299 and in IBM i 6.1.0 through PTF MF49371:

- ▶ **movemediapreference asp_num priority [L M H]** (The default is low.)
This macro moves data marked with a media preference attribute to the SSDs and non-media preference data off the SSDs.
- ▶ **movemediapreferencetossd asp_num priority [L M H]** (The default is low.)
This macro moves data marked with a media preference attribute to the SSDs.
- ▶ **movemediapreferenceoffssd asp_num priority [L M H]** (The default is low.)
This macro moves data that does not have the media preference attribute off the SSDs.
- ▶ **movemediapreferencestatus asp_num**
This macro sets the status of the sweeping.
- ▶ **movemediapreferencestop asp_num**
This macro ends the sweeping.

The ASP number in the *asp_num* variable must be specified in hex format.

A scenario for using the media preference sweeper function is after disk units are added to an ASP, then choosing the add and balance option, which does not respect the media preference. It can also be used when disk units are removed from the configuration because of media type capacity constraints within an ASP. The sweeper function can be used to correct these media preference issues after the capacity constraints are solved.

Script for ASP balancer scheduling

The CL script in Example 9-5 is provided to help set up a **TRCASPBAL** and **STRASPBAL** configuration that runs repeatedly to allow for continuous autonomous IBM i hot and cold data migration for SSD and HDD hybrid ASPs.

Example 9-5 CL script for ASP balancer scheduling

```
/* This program runs continuously until the job in which it runs is ended. */
/* Inputs to this program are type of balance that is to be run, the number */
/* of minutes the trace is to run, and the number of minutes the balance is */
/* to run. Once a trace and balance cycle is complete, another trace and */
/* balance cycles is started. */
/* */
/* Parameter declares. The parameters are: */
/* The balance type to run. */
/* The number of minutes the trace is to run. */
/* The number of minutes the balance is to run. */
PGM          PARM(&BALTYPE &TRACEMIN &BALMIN)
DCL          VAR(&BALTYPE) TYPE(*CHAR) LEN(10)
DCL          VAR(&TRACEMIN) TYPE(*CHAR) LEN(4)
```

```

        DCL          VAR(&BALMIN)  TYPE(*CHAR) LEN(4)
/* Declare for a seconds variable for use by the delay DLYJOB command.      */
        DCL          VAR(&SECONDS) TYPE(*DEC) LEN(6 0)
/* Start tracing for ASP 1, wait for the trace to complete and end it.      */
/* An extra 300 seconds is added to the wait to allow for asynchronous      */
/* activity to complete.                                                    */
LABEL1:  TRCASPBAL  SET(*CLEAR) ASP(1)
        TRCASPBAL  SET(*ON) ASP(1) TIMLMT(*NOMAX)
        CHGVAR     VAR(&SECONDS) VALUE(&TRACEMIN)
        CHGVAR     VAR(&SECONDS) VALUE(&SECONDS * 60)
        CHGVAR     VAR(&SECONDS) VALUE(&SECONDS + 300)
        DLYJOB     DLY(&SECONDS)
        TRCASPBAL  SET(*OFF) ASP(1)
/* Start balancing, wait for the balance to complete and end it.          */
/* An extra 300 seconds is added to the wait to allow for asynchronous      */
/* to complete.                                                            */
        STRASPBAL  TYPE(&BALTYPE) ASP(1) TIMLMT(*NOMAX)
        CHGVAR     VAR(&SECONDS) VALUE(&BALMIN)
        CHGVAR     VAR(&SECONDS) VALUE(&SECONDS * 60)
        CHGVAR     VAR(&SECONDS) VALUE(&SECONDS + 300)
        DLYJOB     DLY(&SECONDS)
        ENDASPBAL  ASP(1)
        MONMSG     MSGID(CPF9899) EXEC(GOTO LABEL1)
        DLYJOB     DLY(300)
/* Run another trace and balance cycle.                                     */
        GOTO       CMDLBL(LABEL1)
        ENDPGM

```

For the ASP balancer SSD enhancements, run **TRCASPBAL** for the period of the critical workload, such as a batch window that is to be optimized by using SSDs. The provided CL script might be an alternative if no specific time frame can be identified for optimization.

9.3.3 User-defined file system media preference

You can specify that storage for objects that are created in user-defined file systems (UDFS) should be allocated from SSDs, if available. This support is provided with PTF SI39439 and all of its requisite PTFs. This support includes changes to various commands and APIs.

A new preferred storage unit (**UNIT**) keyword was added to the Create User-Defined FS (**CRTUDFS**) command, as shown in Figure 9-31. The default value is **UNIT(*ANY)**, which indicates that there is no preferred storage media and that storage for objects in the UDFS are allocated from any available storage media. Specifying **UNIT(*SSD)** indicates that storage for objects in the UDFS are allocated from SSD storage media, if available. Online help text for the new keyword is not included in the PTF.

Create User-Defined FS (CRTUDFS)

Type choices, press Enter.

User-defined file system

Public authority for data . . .

Public authority for object . .

+ for more values

Auditing value for objects . . .

Scanning option for objects . .

Restricted rename and unlink . .

Default disk storage option . .

Default main storage option . .

*INDIR

*INDIR

*SYSVAL

*PARENT

*NO

*NORMAL

*NORMAL

Name, *INDIR, *RWX, *RW...

*INDIR, *NONE, *ALL...

*SYSVAL, *NONE, *USRPRF...

*PARENT, *YES, *NO, *CHGONLY

*NO, *YES

*NORMAL, *MINIMIZE, *DYNAMIC

*NORMAL, *MINIMIZE, *DYNAMIC

Additional Parameters

Case sensitivity

Default file format

Preferred storage unit

*MONO

*TYPE2

*ANY

*MIXED, *MONO

*TYPE1, *TYPE2

*ANY, *SSD

More...

F3=Exit

F4=Prompt

F5=Refresh

F12=Cancel

F13=How to use this display

F24=More keys

Figure 9-31 CRTUDFS command

Changes were made to Qp0lGetAttr()--Get Attributes, Perform File System Operation (QP0LFLOP), and statvfs()--Get File System Information, and related APIs to provide support for determining the preferred storage media for a file system. The Retrieve Directory Information (**RTVDIRINF**) CL command also was enhanced to provide this information. For more information about these changes, see the PTF special instructions.

The following considerations apply when you specify a storage media preference for a UDFS:

- Specifying a media preference does not ensure that storage for objects are allocated from the preferred storage media.
- The preferred storage media attribute of a UDFS cannot be changed.
- All objects in a particular UDFS have the same preferred storage media.
- You can display or retrieve only the storage media preference of a user-defined file system, not the individual objects within a file system.
- Objects that are copied or restored into a UDFS are assigned the preferred storage media of the UDFS, regardless of the original object's preferred storage media.
- When you restore a new UDFS to a system, the original storage media preference of the UDFS is retained.

9.3.4 177 GB SFF SSD with eMLC

SAS bay-based SSD options are enhanced with a 177 GB SSD, which provides 2.5 times more capacity per drive than the current 69 GB SSD. The 177 GB drive provides an improved cost per gigabyte and requires a smaller number of SAS bays for the same number of gigabytes.

Enterprise Multi-level Cell technology (eMLC) enables enterprise-level performance and reliability while being more cost-effective than previous technology.

This option is supported on Power 710, 720, 730, 740, 750, 755, 770, 780, and 795 models.

For more information, see *IBM Hardware Announcement letter 111-132* at:

<http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&appname=iSource&supplier=897&letternum=ENUS111-132>

9.3.5 IBM Disk Sanitizer PRPQ extended to include SSD devices

The IBM Disk Sanitizer for i5/OS PRPQ, 5799-SD1 is enhanced to sanitize SSD devices.

The Disk Sanitizer is accessed through a macro interface from either the Dedicated Service Tools (DST) menu or the System Service Tools (SST) menu. To access the Disk Sanitizer, complete the following steps:

1. From DST or SST, select 'Start a service tool'.
2. Select 'Display/Alter/Dump'.
3. Select 1 - 'Display/Alter storage'.
4. Select 2 - 'Licensed Internal Code (LIC) data'.
5. Select 14- 'Advanced Analysis' (you must scroll down to see this option).

- On the Select Advanced Analysis Command panel, there is a blank line at the top. Type a 1 in the Option column to select the blank line, then type SSDSANITIZE, as shown in Figure 9-32. The SSDSANITIZE macro may also be selected from the list of macros.

```

Select Advanced Analysis Command

Output device . . . . . : Display

Type options, press Enter.
1=Select

Option   Command
➔ 1      SSDSANITIZE ←
-        FLIGHTLOG
-        ADDRESSINFO
-        ALTSTACK
-        BATTERYINFO
-        CLUSTERINFO
-        CONDITIONINFO
-        COUNTERINFO
-        DISABLEFLASHSYNC
-        DSTINFO
-        EXCEPTCHAIN
-        FINDFRAMES
-        FINDPTF

F3=Exit  F12=Cancel

More...

```

Figure 9-32 Selecting the SSDSANITIZE macro

- Press the Enter key twice and a help panel is displayed, as shown in Figure 9-33.

```

Display Formatted Data
Page/Line. . . 1 / 1
Columns. . . : 1 - 78

Find . . . . .
.....1.....2.....3.....4.....5.....6.....7.....
DISPLAY/ALTER/DUMP
Running macro: SSDSANITIZE
This macro is used to sanitize Solid State Disk (SSD) units.

An entry will be added to the Product Activity Log (PAL) listing each unit
selected to be sanitized and an indication of whether the unit passed or
failed to sanitize.

The PAL System Reference Code (SRC) is B6000231 for all units on which
sanitize ran successfully. The PAL SRC is B6000232 for all units on which
sanitize did not run successfully.

The control panel will display the SRC C6XX0230 during the sanitize procedure.
The XX will be replaced with the percent complete.

Usage: ssdsanitize <-HELP> <-LIST> <-ALL> <-UNIT RESOURCENAMES> <-IMMED>
More...

F2=Find  F3=Exit  F4=Top  F5=Bottom  F10=Right  F12=Cancel

```

Figure 9-33 SSDSANITIZE macro help panel

The sanitizing SSD units function is nearly identical to sanitizing HDD units from a user interface perspective.

To enable this function, the following PTFs are required:

- ▶ IBM i 7.1: MF52834
- ▶ IBM i 6.1.1: MF50873
- ▶ IBM i 6.1.0: MF50875



Networking enhancements

This chapter describes the following topics that are related to networking enhancements enabled in IBM i 7.1:

- ▶ TCP/IP enhancements summary
- ▶ HTTP server support for PASE FastCGI PHP processing
- ▶ Telnet client support for SSL encryption
- ▶ Sockets programming: New user exits
- ▶ IKEv2 support
- ▶ IPv6 TCP/IP applications and V6R1 enablement PTFs
- ▶ IBM AnyNet support on IBM i 7.1
- ▶ Ethernet link aggregation
- ▶ Ethernet Layer-2 bridging
- ▶ IBM Portable Utilities for i (5733-SC1) supported versions

10.1 TCP/IP enhancements summary

The following sections summarize the key TCP/IP enhancements in IBM i 7.1. These enhancements include the following topics:

- ▶ Additional TCP/IP application enablement for IPv6
- ▶ ISC-based DHCP server
- ▶ Enhancements in HTTP adding support for FastCGI PHP processing
- ▶ TELNET client SSL enablement
- ▶ Addition of SNMP version 3 (SNMPv3) support
- ▶ IKEv2
- ▶ Several new sockets programming user exits

10.1.1 IPv6 support enhancements

Although support for IPv6 was first introduced in IBM i 5.4, it existed primarily at the user application socket level. In IBM i 6.1, many of the standard IBM applications were enhanced to support either Internet Protocol version 4 (IPv4) or version 6 (IPv6). For more information, see 10.6, “IPv6 TCP/IP applications and V6R1 enablement PTFs” on page 429.

IBM i 7.1 extends this support by adding IPv6 for the following applications:

- ▶ DHCP Server
- ▶ DHCP Client
- ▶ SNMP
- ▶ SMTP
- ▶ PPP

10.1.2 ISC-based DHCP server supports IPv6 and failover

With IBM i 7.1, you have the option of using the new DHCP server that is based on the open source code that is provided by the Internet Systems Consortium (ISC). Existing customers can continue to use the old DHCP server that is supplied by IBM in previous releases or replace it with the ISC-based version.

The ISC-based server has several advantages. In addition to supporting IPv4, it also supports IPv6 and DHCP server failover. The DHCP server attributes can be set to run either an IPv4 or IPv6 server or both. There is no GUI support for managing the ISC DHCP server configuration files or for monitoring leases, such as with the old DHCP server. Therefore, by default, the old DHCP server is used.

If you want use the ISC DHCP server, you must add the `QIBM_ISC_DHCP` environment variable, as described in “Using the ISC DHCP IPv6 server on IBM i” on page 415. Then, stop your DHCP server by running the **ENDTCPSVR** command (if it is running) and start the ISC DHCP server with the **STRTCPSVR** command. The IBM i activation code attempts to migrate the old configuration file to the new ISC configuration file the first time DHCP-related code is run (through **CHGDHCPA** or **STRTCPSVR**). The old configuration file is left unchanged after the migration. Any changes that are made to the old configuration file are not migrated to the new one after the initial migration. The new configuration file might require editing to operate properly. The current leases file is also migrated to the ISC leases file. The migration is just a way to get started with the new server. Certain functions that are provided by the old server are not available with the ISC server, so you must weigh the benefits and differences between these two servers and choose which one is best for your environment.

If you want to switch back to the old DHCP server, delete the environment variable, or set the value to 'N', and then stop and restart the DHCP server. If the ISC DHCP server assigned any IP addresses when it was running, those leases are not available to the old DHCP server. There is no backward migration.

Using the ISC DHCP IPv6 server on IBM i

DHCP moved from the base OS to 5770-SS1 Option 31 and requires that 5770-SS1 Option 33 be installed. To use the ISC DHCP IPv6 server on IBM i, complete the following steps:

1. Ensure that IBM i option 31 (Domain Name System (DNS)) and option 33 (Portable Application Solutions Environment (PASE)) are installed on the system.
2. Define an environment variable to tell the operating system to use the ISC DHCP server by running the following command:

```
ADDENVVAR ENVVAR('QIBM_ISC_DHCP') VALUE('Y') LEVEL(*SYS)
```

3. Run the Change DHCP Attributes (**CHGDHCPA**) command.

This command migrates any existing DHCP configuration into the configuration files that are used by the ISC DHCP server. Determine if you want to run an IPv4, IPv6, or both. This setting is managed by setting the DHCP attribute for IPVERSION. The *ALL special value enables support for both IPv4 and IPv6.

```
CHGDHCPA IPVERSION(*IPV6)
```

4. Edit the newly created configuration files.

Access to the ISC DHCP server: A graphical interface is not provided for managing the ISC DHCP server and monitoring the leases that it manages. All associated configuration files must be edited manually.

There are several considerations to make when migrating from the existing IBM i DHCP server to the ISC DHCP server. For example, IBM Navigator for i does not provide an interface for configuring the ISC DHCP server in IBM i 7.1. To configure the ISC DHCP IPv6 server, edit the /QIBM/UserData/OS400/DHCP/ETC/DHCPD6.CONF configuration files manually. Example 10-1 shows an example.

Example 10-1 Edited configuration file

```
authoritative;

subnet6 1ffe:31::/64 {
    default-lease-time 120;
    max-lease-time 86400;
    range6 1ffe:31::d0:ca1 1ffe:31::d0:cef;
}
```

Copy the above into /QIBM/UserData/OS400/DHCP/ETC/DHCPD6.CONF.

Make sure that you have at least one line that is enabled for IPv6 on your system and configured with an IPv6 address, for example something like: *1ffe:31::d0:ccc* so that the line description of the address can be listened and that subnet6 would not be ignored.

Two more files might need to be configured depending on your configuration requirements:

- ▶ /QIBM/UserData/OS400/DHCP/ETC/DHCRELAY6.CONF
- ▶ /QIBM/UserData/OS400/DHCP/ETC/DHCPD6.LEASES

For further information, there are Linux / AIX man pages available for the ISC DHCP server and books such as *The DHCP Handbook*, by Drom, et al, which provides detailed descriptions for the configuration statements available. ISC also has information at the following web page:

<http://www.isc.org/software/dhcp>

10.1.3 DHCPv6 client

The DHCPv6 client is also new in IBM i 7.1. It is not explicitly configured, but is enabled by adding and starting a *IP6SAC interface as follows:

```
ADDTCPIFC *IP6SAC LIND(line-name)
```

The system tries to acquire only IPv6 addresses through DHCPv6 if an IPv6 router on the link tells the system (by turning on the 'M' bit in the Router Advertisement flags) to use the managed configuration to obtain IP addresses. The DHCPv6 client sends multicast messages to find a DHCPv6 server and to request IPv6 address assignment. The DHCPv6 server sends a reply with the addresses assigned. IP addresses obtained from the DHCPv6 server have a preferred and valid lifetime, just like stateless auto configured addresses. Before the preferred lifetime expires, the DHCPv6 client renews the addresses. When the *IP6SAC interface is ended, any DHCP addresses are released.

If the Privacy Extension parameter is enabled on the *IP6SAC interface, you also request temporary addresses from the DHCPv6 server. The request for temporary addresses is sent separately from the request for non-temporary addresses. Temporary addresses are never renewed; when the preferred lifetime is about to be reached, you request new temporary addresses. The old temporary addresses remain until either their valid lifetime is reached or the *IP6SAC interface is ended. The preferred and valid lifetime of DHCP temporary addresses is limited by the IPv6 temporary address valid and preferred lifetimes that are configured through **CHGTCPA**.

To identify itself to the DHCPv6 server, the client uses a DHCP Unique Identifier (DUID). This DUID is generated automatically from a MAC address on the system and a time stamp, and is saved by the TCP/IP configuration. This identifier is a system-wide identifier; the same DUID is used by DHCP on all lines. To identify separate lines, the DHCP message also contains an identity association identifier (IAID), which is a unique value for each separate line (generated and saved by the TCP/IP configuration). The current DUID can be viewed by using the **CHGTCPA** command. The value cannot be changed by the user, but the user can force generation of a new DUID if necessary, by using the ***GEN** option.

As with the DHCPv4 client, additional configuration information can be obtained from the DHCPv6 server beyond just addresses. For DHCPv6, it supports the DNS Server List and Domain Search List options and adds received DNS servers and domains to our configuration when the DHCPv6 client is active.

Additionally, DHCPv6 supports an option to receive configuration information without allocating addresses. This option is automatically selected if the router on the link sends a router advertisement with the O flag (Other configuration) set rather than the M flag. In that case, you request just the DNS Server List and Domain Search List options from the DHCPv6 server.

Support added: IBM i 6.1 added DHCPv4 client support for IPv4 with PTF SI31800.

10.1.4 SNMP

In release 7.1, the IBM i SNMP agent provides basic SNMP version 3 (SNMPv3) support. SNMP version 3 incorporates the use of user-based authentication and data privacy. The IBM i 7.1 SNMP also includes support for IPv6. It is possible to configure SNMP manager IP addresses for both traps and communities through the **CHGSNMPA** and **ADDCOMSNMP** commands. IPv6 support for various MIBs, including RFCs 4022 and 4013, were added.

Enabling the agent to handle SNMPv3 requests

To enable the agent to handle SNMPv3 requests, complete the following steps:

1. If the SNMP server is running, stop it by running **ENDTCPSVR *SNMP**.
2. Change the SNMP server attributes to allow version 3 by running **CHGSNMPA ALWSNMPV3(*YES)**.

Supported functionality: The SNMP agent is still able to receive and handle packets and requests from older versions of SNMP v1 even after you change the SNMP attributes to specify **ALWSNMPV3(*YES)**.

3. Check the engine identifier that is supplied by the SNMP Agent after it is started for the first time after **ALWSNMPV3(*YES)** is set.

In most cases, this engine identifier does not need to be changed. If the generated engine ID must be changed, do so by running **CHGSNMPA**. There are caveats, however. The engine identifier is created using a vendor-specific formula and incorporates the IP address of the agent. Any engine identifier that is consistent with the `snmpEngineID` definition in RFC 3411 and that is also unique within the administrative domain can be specified.

For example, the identifier 80000002010A010203 is a valid engine ID for an IBM i agent with an IP address of 10.1.2.3. The first byte, '80'X, indicates that the engine ID complies with the architecture defined in RFC 3411. The next four bytes, '00000002'X, indicate the private enterprise number for IBM as assigned by the Internet Assigned Numbers Authority (IANA). The next byte, '01'X, indicates that the remaining portion of the engine ID is an IPv4 address. The last four bytes, '0A010203'X, is the hexadecimal representation of the IP address. The **CHGSNMPA SNMPENGID('80000002010A010203')** command is run to specify the engine ID.

Important: An invalid SNMP engine ID can prevent an SNMP manager from communicating with the agent.

Important: Another new SNMPv3 parameter, **SNMPENGB**, was added to the **CHGSNMPA** command, and is the SNMP engine boots counter. Do not manually change this parameter unless you must reset it to a value of zero. This parameter indicates the number of times that the SNMP engine (agent) was started. Each time the **STRTCPSVR *SNMP** command is successfully run, this value increments automatically. Changing the **SNMPENGB** parameter when the agent is active can cause SNMPv3 authentication failures.

4. Add an SNMP user using the **Add User for SNMP** command, with encryption and privacy options that match your SNMP manager.

An SNMP user is not the same as an IBM i user profile. SNMP users must be added and maintained separately. For example, the following command adds an SNMP user who requires authentication using the HMAC-SHA authentication protocol and privacy using the CBC-DES encryption protocol:

```
ADDUSRSNMP USERNAME(testuser) AUTPCL(*HMACSHA) AUTPWD(authpassword)
PVYPCL(*CBCDES) PVYPWD(privpassword)
```

The **USERNAME**, **AUTPWD**, and **PVYPWD** parameters are case-sensitive, so take care when you add SNMP users. The output of the **ADDUSRSNMP** is an entry in the SNMPv3 configuration file.

The configuration file `/QIBM/USERDATA/OS/SNMP/SNMP.CONF` contains the SNMP user information and their generated keys. The passwords that are specified are not stored.

The **SNMP_USER** statements in the configuration file are never edited manually. Instead, the **ADDUSRSNMP**, **CHGUSRSNMP**, and **RMVUSRSNMP** commands are used for maintaining the SNMP users. The **CFGTCPSNMP** command can be run to display and maintain the list of SNMP users.

5. Start the SNMP server by running **STRTCPSVR *SNMP**.
6. Start an SNMPv3 manager application and configure it for the SNMP user that was added in step 4.

Support of functionality: The IBM i 7.1 SNMP manager APIs `snmpGet`, `snmpSet`, and `snmpGetnext` do not support SNMPv3, so a non-native manager such as a PC-based manager must be used. There are a number of these managers available for download, including both no-cost and for-purchase options.

Resolving time synchronization errors

When an SNMPv3 manager first communicates with an SNMPv3 agent, it goes through a discovery process that involves the determination of the agent's SNMP engine ID and SNMP engine boots values. In addition, time synchronization occurs. These steps are necessary before you do any actual SNMP operations, such as `Get` or `GetNext`. If you have problems with these initial steps, the SNMP manager can indicate that a time synchronization error occurred. If this error occurs, complete the following steps:

1. Stop the SNMP agent and manager.
2. Run the following command:

```
ADDENVVAR ENVVAR('QIBM_SNMPV3_AUTH') VALUE('1') LEVEL(*SYS)
```
3. Restart the SNMP agent and manager and try the request again.

10.1.5 SMTP

IPv6 support was added in IBM i 7.1. At this time, there is no IPv6 standard for Real-time Black holes Lists (RBL). The RBL works only for IPv4 addresses. SMTP uses the `getaddrinfo()` API to look up email DNS records. They are looked up first as IPv6 and then as IPv4, which is different from what RFC 3974 recommends. Parts of the DNS resolver were fixed in IBM i 7.1 to be more correct.

SMTP support for RFC 821 and RFC 822 were removed. Starting in IBM i 7.1, Simple Mail Transfer Protocol (SMTP) supports only RFC 2821 and 2822. RFC 2821 and RFC 2822 deprecate many parts of the RFC 821 and RFC 822 email standards. Behavior for SMTP routes, SMTP alias shadowing, and processing mail through Mail Service Framework (MSF) are not compatible with the RFC 2821 and RFC 2822 standards and are to be used on an *as is* basis. The first part of a source route is still honored from RFC 821; other parts of the source route are not contacted. The absolute address is the best way to send email.

The MAILROUTER feature before IBM i 7.1 can, in instances, forward all mail to the mail router even if the email address can be resolved. In IBM i 7.1, MAILROUTER correctly forwards to the mail router only when the email address does not resolve.

The FWDMAILHUB feature was added in IBM i 6.1, and allowed the forwarding of email to a single address. FWDMAILHUB always forwards the email and does not attempt a resolve.

MAILROUTER supports only A and AAAA records, when FWDMAILHUB supports MX, CNAME, AAAA, and A.

The resolve path is shown in Example 10-2.

Example 10-2 Resolve path

```
Forwarding Mail hub(if defined)->
Absolute Address/First part of source route->
  mailrouter(if same domain)->
    mailrouter(different domain) if FIREWALL(*YES).
```

10.1.6 IPv6 support added for PPP connections

Starting in IBM i 7.1, Point-to-Point (PPP) can support both IPv4 and IPv6 addresses. A PPP connection profile can have only IPv4 enabled, only IPv6 enabled, or both IPv4 and IPv6 enabled. By default, both IPv4 and IPv6 are enabled for a PPP connection profile.

Note the following items:

- ▶ IPv6 can be enabled in connection profiles to allow remote workers to use IPv6 to access the company network.
- ▶ If your ISP supports IPv6 addressing, you can also enable IPv6 in the originator profile.
- ▶ If the **Enable IPv6** check box is checked in the TCP/IP IPv6 Settings section of a connection profile, then when the connection profile is activated, IPv6 Stateless Address Auto configuration assigns an IPv6 link-local address to the PPP link. In many cases, this configuration is all that is required to use IPv6 over a PPP link.
- ▶ Additional IPv6 configuration options are available if IP forwarding is enabled in the TCP/IP IPv6 Settings section of a connection profile. Enabling IPv6 datagram forwarding also enables the system to perform router functions for this link, including sending Router Advertisement messages and responding to Router Solicit messages.

Figure 10-1 highlights the configuration changes that are required to enable IPv6 for a connection profile.

Feature availability: PPP configuration enhancements for IPv6 are only available using IBM Navigator for i. It is not available using the PC-based client System i Navigator.




Figure 10-1 Traditional IPv4 PPP remote access configuration

If you want your remote workers to use IPv6 to access the company network, you must enable IPv6 in the connection profile. You do not need to assign a specific IPv6 address. However, if you want the remote workers to have more than the default link-local IPv6 address assigned, you must either configure an IPv6 address prefix or set the appropriate options if a DHCPv6 server is available in the company network.

For this example, assume that you want to advertise an address prefix of 2001:DBA::, a default route, and that a DHCPv6 server in your network can provide IP addresses. A global IPv6 address must be configured in the connection profile to allow the DHCPv6 server to return information to the remote dial-in client. This address can be configured in the Receiver Connection Profile using IBM Navigator for i, as shown in Figure 10-2.

Click **Network** → **Remote Access Servers** → **Receiver Connection Profiles**, and in the right pane, click **Action** → **New Profile**.




Figure 10-2 Create a Receiver Connection Profile

Select **Protocol type**, **Connection type** and **Link configuration** for the new profile, and click **OK**, as shown in Figure 10-3.




Figure 10-3 Select parameters for the new profile

Click **Create Receiver Connection Profile**, and then **TCP/IP IPv6 Settings**. See Figure 10-4.




Figure 10-4 PPP - Create Receiver Connection profile window with the IPv6 option

To advertise an address prefix of 2001:DBA::, a default route, and that a DHCPv6 server in your network can provide IP addresses, configure a global IPv6 address in the connection profile as follows (see Figure 10-5):

1. Select **Enable IPv6**.
2. Specify a global IPv6 address for Fixed local IP address. This address must be compatible with the DHCPv6 server configuration for distributing IPv6 addresses. For this example, click **None**.
3. Click **Generate** for the Interface identifier field.
4. Select **Yes** for the **Allow remote system to access other networks (IP forwarding)** check box.
5. Set the Address prefix to 2001:DBA::.
6. Select **Advertise IPv6 default route**.
7. Select **Advertise DHCPv6** and **Managed address configuration**.
8. Click **OK** to complete the profile.

The screenshot shows the 'PPP IPv6 Settings' dialog box. On the left is a sidebar with tabs: 'General', 'Connection', 'Authentication', 'TCP/IP IPv4 Settings', '*TCP/IP IPv6 Settings' (selected), and 'Other'. The main area is divided into sections: 'Local IP address' with a checked 'Enable IPv6' box and a text field for '*Fixed local IP address' containing '*NONE'; 'Local interface' with an 'Interface identifier' dropdown set to 'Generate'; and 'Routing' with a section titled 'Allow remote system to access other networks (IP forwarding)'. In this section, the 'Yes' radio button is selected, and the 'Address prefix' text field contains '2001:DBA::'. Below this, 'Advertise IPv6 default route' and 'Advertise DHCPv6' are both checked. Under 'Advertise DHCPv6', the 'Managed address configuration' radio button is selected. There is also an 'Accept a default route' checkbox which is checked. At the bottom of the dialog is a checkbox labeled 'Define additional IP addresses based on the user ID of the remote system' which is currently unchecked.

Figure 10-5 Configure PPP IPv6 preferences

10.2 HTTP server support for PASE FastCGI PHP processing

FastCGI is an open standard that extends the Common Gateway Interface (CGI) standard that is supported by many common web servers today. This standard defines how information is exchanged between a web server and FastCGI programs that are isolated in external processes. On IBM i 7.1, these external processes are provided by a FastCGI Apache module that makes external PASE programs (FastCGI programs) available as CGI jobs, which can then be used by the native ILE environment, which results in faster HTTP request processing.

Further implementation details are available at:

http://www-03.ibm.com/systems/resources/systems_i_software_http_docs_pase_config.pdf

Software updates were also included in the HTTP Group PTF packages for January 2010, which enable FastCGI PHP Processing for both IBM i 5.4 and 6.1.

The required components and PTF information for 5.1 and 6.1 are shown in 10.2.1, “IBM i 5.4 required components” on page 424 and 10.2.2, “IBM i 6.1 required components” on page 424.

10.2.1 IBM i 5.4 required components

This section lists the IBM i 5.4 required components:

- ▶ Products
 - 5722SS1 30 Qshell (only needed to create CCSID 819 files)
 - 5722SS1 33 Portable App Solutions Environment
 - 1ZCORE5 *BASE Zend Core for IBM i (Version 2.6.1 or later) only for FastCGI PHP support
 - 5722DG1 *BASE IBM HTTP Server for i
 - 5722DG1 1 IBM HTTP Server for i
 - 5733SC1 *BASE IBM Portable Utilities for i (only for FastCGI PHP support)
- ▶ PTFs
 - SI36004 (PASE) SI36026 (DG1)
 - Group PTFs: SF99114: 540 IBM HTTP Server for i (PTF Group Level: 21 or later)

10.2.2 IBM i 6.1 required components

This section lists the IBM i 6.1 required components:

- ▶ Products
 - 5761SS1 30 Qshell (only needed to create CCSID 819 files)
 - 5761SS1 33 Portable App Solutions Environment
 - 1ZCORE5 *BASE Zend Core for IBM i (Version 2.6.1 or later) (only for FastCGI PHP support)
 - 5761DG1 *BASE IBM HTTP Server for i
 - 5761DG1 1 IBM HTTP Server for i
 - 5733SC1 *BASE IBM Portable Utilities for i (only for FastCGI PHP support)
- ▶ PTFs
 - SI36005 (PASE) SI36027 (DG1)
 - Group PTFs: SF99115: 610 IBM HTTP Server for i (PTF Group Level: 10 or later)

10.3 Telnet client support for SSL encryption

Although the Telnet server for IBM i has long supported SSL for remote clients, IBM i 7.1 can support SSL as a client. This capability was also enabled for IBM i 5.4 and 6.1 through PTFs, and activated using the QIBM_QTV_TELNET_CLIENT environment variable. With this new feature enabled, the **TELNET** command can now connect to a remote telnet server that supports implicit SSL/TLS.

The **PORT** parameter of the **TELNET** command prompt was moved to a new location in the parameter string, and a new parameter, Secure Connection (SSL), was added to the command format. If the environment variable was set up for a secure connection, or the **SSL(*YES)** parameter is selected, the target port number defaults to 992.

The requirements are as follows:

- ▶ You must have the Digital Certificate Manager (5770-SS1 Option 34) installed.
- ▶ You must have certificates that are set up for the client application QIBM_QTV_TELNET_CLIENT. Either a trust list or the remote server certificates must be assigned to this application. The trust list must have the certificate authorities of the wanted remote servers added.

If you want all telnet client users on your system to use SSL, set the QIBM_TELNET_CLIENT_SSL as a system level environment variable.

Encryption is provided using either SSL or Transport Layer Security (TLS) based on negotiation between the Telnet client and the server.

The TELNET client must be assigned an appropriate certificate in the Digital Certificate Manager (DCM) or the connection fails. See Figure 10-6.

```
Start TCP/IP TELNET (TELNET)

Remote system . . . . . RMTSYS
Internet address . . . . . INTNETADR

Port . . . . . PORT          *DFT
Secure connection . . . . . SSL      *ENVVAR
```

Figure 10-6 New parameter of the STRTCPTELN command for SSL enablement

PTF support: This enhancement was made available for V5R4 and V6R1 through the following PTFs:

- ▶ V5R4 - SI32220
- ▶ V6R1 - SI32527

10.4 Sockets programming: New user exits

Functional enhancements to sockets programming include three sockets-related user exit points. These enhancements were added to give a user-defined exit program the ability to control connections that are based on specific runtime characteristics. User-defined exit programs that are registered with the exit points defined in the user registry are able to limit incoming and outgoing connections.

10.4.1 Exit points that are defined in the user registry

User-defined exit programs that are registered with the exit points defined in the user registry are able to limit incoming and outgoing connections. The return codes of the user-defined exit programs indicate whether to allow successful completion to `connect()`, `listen()`, `accept()`, `accept_and_recv()`, or `QsoStartAccept()`. See Table 10-1.

Table 10-1 User exit points

User exit point	Description
QIBM_QSO_ACCEPT	Enables a custom exit program to allow or deny incoming connections that are based on the restrictions that are set by the programs.
QIBM_QSO_CONNECT	Enables a custom exit program to allow or deny outgoing connections that are based on the restrictions that are set by the programs.
QIBM_QSO_LISTEN	Enables a custom exit program to allow or deny a socket the ability to listen for connections that are based on the restrictions that are set by the programs.

Sockets APIs considerations:

- By default, the sockets APIs accepting connections silently ignore rejected connections and wait for the next incoming connection. To give an application the ability to be informed about rejected connections, a socket option is provided. The socket option is enabled by `setsockopt()` with a level of `SOL_SOCKET` and option name `SO_ACCEPTTEPERM`. When the socket option is enabled, sockets APIs accepting connections fail with `EPERM` for each incoming connection that is rejected by the user exit program that is registered for `QIBM_QSO_ACCEPT`.
- Any user that tries to add or remove a sockets-related user exit program is required to have `*IOSYSCFG`, `*ALLOBJ`, and `*SECADM` authority.

10.4.2 Example: User exit program for QIBM_QSO_ACCEPT

This section illustrates an example of the user exit program for the `QIBM_QSO_ACCEPT` user exit point.

The program in Example 10-3 rejects all incoming connections to the Telnet server that come from a particular remote IP address between the hours of 12 a.m. - 4 a.m. The program determines whether the incoming connection can be accepted by the socket API accepting connections or rejected.

Example 10-3 Socket program example using the `QIBM_QSO_ACCEPT` user exit

```
/* **** */
/* System i - Sample User Exit Program for QIBM_QSO_ACCEPT */
/* **** */
```

```

/* Exit Point Name : QIBM_QSO_ACCEPT */
/*
/* Description      : The following ILE C language program */
/*                  will reject all incoming connections to */
/*                  the telnet server (port 23) coming from */
/*                  the remote IP address of '1.2.3.4' between */
/*                  the hours of 12 A.M. and 4 A.M. */
/******
#include stdio.h
#include string.h
#include esoextpt.h /* Exit program formats */
int main(int argc, char *argv[])
{
    Qso_ACPT0100_Format input; /* input format */
    struct in_addr addr;
    char return_code;

    /******
    /* Initialize the address to compare to 1.2.3.4 */
    /******
    addr.s_addr = 0x01020304;

    /******
    /* By default allow the connection. */
    /******
    return_code = '0';

    /******
    /* Copy format parameter to local storage. */
    /******
    memcpy(&input, (Qso_ACPT0100_Format *) argv[1],
           sizeof(Qso_ACPT0100_Format));
    /******
    /* If the local port is the telnet server */
    /******
    if((input.Local_Incoming_Address_Length == sizeof(sockaddr_in) &&
        input.Local_Incoming_Address.sinstruct.sin_port == 23) ||
        (input.Local_Incoming_Address_Length == sizeof(sockaddr_in6) &&
        input.Local_Incoming_Address.sin6struct.sin6_port == 23))
    {
        /******
        /* And the incoming connection is from 1.2.3.4 */
        /******
        if(input.Remote_Address_Length == sizeof(sockaddr_in) &&
            (memcmp(&input.Remote_Address.sinstruct.sin_addr,
                    addr, sizeof(struct in_addr)) == 0))
        {
            /******
            /* And the time is between 12 A.M. and 4 A.M. */
            /* Reject the connection. */
            /******
            if(IsTimeBetweenMidnightAnd4AM())
                return_code = '1';
        }
    }
}

```

```
*argv[2] = return_code;  
return 0;  
}
```

Important: By using the example that is shown in Example 10-3 on page 426, you agree to the terms of the code license and disclaimer information that is available at:

<https://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/rzahg/legalnoticesSW.htm>

10.5 IKEv2 support

You can now use the enhanced IKE version 2 when you perform dynamic key connection. Enhanced Cryptographic Algorithms are also available for use with VPN key exchange policy and data policies. Existing objects were used as much as possible to enable either IKEv1 or IKEv2 exchanges. This design was used to minimize the impacts to the current GUI interface and VPN configuration objects when you enable IKE version 2. See Figure 10-7.

Note the following items:

- ▶ To enable IKEv2, an IKE Version setting is provided on the Dynamic Connection definition.
- ▶ Key Exchange policies can be used for either IKEv1 or IKEv2.
- ▶ There are no more differences to other attributes, such as Key Exchange policy identifiers (all are still supported) and transforms.
- ▶ Main mode / aggressive mode settings are ignored if the Key Exchange Policy is used for IKEv2.
- ▶ Systems IBM Navigator for i is required to configure an IKEv2 connection.




Figure 10-7 Internet Key Exchange V2 enablement for VPN

10.5.1 Enhanced Cryptographic Algorithms

Enhanced Cryptographic Algorithms were added to IBM i 7.1.

Key exchange policy

Note the following items for the key exchange policy:

- ▶ Encryption: AES-CBC – 128 bits
- ▶ Hash/PRF
 - AES-XCBC-MAC (HASH 96 bits; PRF 128 bits)
 - HMAC-SHA-256
- ▶ Diffie-Hellman
 - Group 14
 - Group 24

Data policy

Note the following authentication items for the data policy:

- ▶ AES-XCBC-MAC
- ▶ HMAC-SHA-256

Diffie-Hellman for PFS

Note the following items for Diffie-Hellman for PFS:

- ▶ Group 14
- ▶ Group 24

For more information and configuration details, see the topic “IBM i Security Virtual private networking 7.1”, available at:

<https://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp?topic=/rzaja/rzajagetstart.htm>

10.6 IPv6 TCP/IP applications and V6R1 enablement PTFs

Support for IPv6 is being added incrementally to TCP/IP applications and protocols for the IBM i operating system. The following Information Center entry provides a complete list of which applications and protocols that support IPv6 in IBM i 6.1 and 7.1:

<http://www-03.ibm.com/systems/i/software/tcpip/applications6.html>

This website is updated as PTFs are made available for additional applications or protocols. As of this publication, the following list identifies IBM i 6.1 applications and protocols that support IPv6:

- ▶ IBM Online Help and Eclipse Information Center (IBMHELP) - PTF SI31014
- ▶ INETD - PTF SI29701
- ▶ SNTP - PTF SI30112
- ▶ TFTP - PTF SI30868
- ▶ LPD - PTF SI31015
- ▶ Remote Journal - PTF SI31713
- ▶ Remote Journal - PTF MF44589
- ▶ IPP printer driver - PTF SI31910
- ▶ LPR and Remote output queues - PTF SI31363
- ▶ Enterprise Extender 1 (MSCP) - PTF MF44318
- ▶ Enterprise Extender 2 (HPR) - PTF MF44355

- ▶ Enterprise Extender 3 (HPR) - PTF MF44356
- ▶ Enterprise Extender 4 (DC) - PTF SI31250
- ▶ Enterprise Extender 5 (SW) - PTF SI31223
- ▶ Enterprise Extender 6 (Comm Trace) - PTF SI30790
- ▶ Management Central - PTF SI31888
- ▶ Management Central - PTF SI31892
- ▶ Management Central - PTF SI32720
- ▶ Management Central - PTF SI32721

10.7 IBM AnyNet support on IBM i 7.1

Enterprise Extender was introduced with IBM i 5.4 and was identified at that time as the strategic direction for replacing IBM AnyNet®. Although AnyNet was removed in IBM i 7.1, IBM stated that there are no further enhancements beyond IBM i 6.1 and IBM no longer offers support for it on 7.1.

The IBM i 7.1 Information Center topic “Migrating from IBM AnyNet to Enterprise Extender” provides detailed migration considerations and requirements, and is available at:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/rzajt/rzajtanytoee.htm>

10.8 Ethernet link aggregation

The following sections describe the IBM i Ethernet link aggregation function available in IBM i 7.1:

- ▶ Overview of Ethernet link aggregation
- ▶ Prerequisites for Ethernet link aggregation
- ▶ Configuring Ethernet link aggregation
- ▶ Example: Configuring four aggregated network links

10.8.1 Overview of Ethernet link aggregation

Link aggregation binds several full-duplex Ethernet links that run at the same speed together into one logical link with a single Media Access Control (MAC) address. This aggregation is known by several other names, including IEEE 802.3ad or 802.1ax, Cisco Etherchannel, or the names teaming or trunking.

With the Ethernet link aggregation function available in IBM i 7.1, up to eight Ethernet links can be bound together in a single-line description.

The advantages of this function are:

- ▶ Simplified redundancy and reliability

By binding multiple Ethernet links to a single-line description, if a link fails, the others remain active and the network load is rebalanced across the active links without requiring any system or network administrator actions.

Before this function, if a line failed, IP configuration changes and a manual switch to another line description and link were required.

► Capacity

By aggregating multiple links to a single-line description, outgoing traffic is spread across the links as determined by a user-selected policy and incoming traffic by a policy that is configured at the network switch. This configuration also enables more capacity for a certain IP address.

For example, two aggregated 1 Gbps links can carry up to 2 Gbps for the same IP interface without any additional configuration.

10.8.2 Prerequisites for Ethernet link aggregation

To use Ethernet link aggregation, the environment must have:

- Up to eight Ethernet ports, 10 Gbps or 1 Gbps-capable, on an IBM i partition not in use for other purposes.
- The ports must be on the same network switch and be in a static configuration.
- The newest IBM i 7.1 Technology Refresh PTF and PTF Group.
- A Cisco Catalyst switch with an aggregate in Etherchannel mode that is enabled for static configuration.

Unsupported switches: Other switches that support static aggregation configurations might also work, but they were not tested and are not officially supported.

10.8.3 Configuring Ethernet link aggregation

The following steps are used to configure Ethernet line aggregation. See the example command in Figure 10-8.

```

                                Create Line Desc (Ethernet) (CRTLINETH)

Type choices, press Enter.

Line description . . . . . LIND          > TESTLINE ← a)
Resource name . . . . . RSRNAME         > *AGG
Bridge identifier . . . . . BRIDGE       > *NONE
Online at IPL . . . . . ONLINE          > *YES
Vary on wait . . . . . VRYWAIT          > *NOWAIT
Aggregate policy: AGGPCY
  Standard . . . . .                > *ETHCHL ← b)
  Policy type . . . . .              > *RNDRBN
Aggregate resource list . . . . AGGRSCL  > CMN01 ← c)
                                         > + for more values > CMN02
Local adapter address . . . . . ADPTADR  > *ADPT
Exchange identifier . . . . . EXCHID     > *SYSGEN
Ethernet standard . . . . . ETHSTD       > *ETHV2
Line speed . . . . . LINESPEED          > 1G
Duplex . . . . . DUPLEX                  > *FULL
Serviceability options . . . . . SRVOPT  > *NONE

More...

F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys
```

Figure 10-8 Creating an aggregated Ethernet line description

1. Create a line description with the resource name (**RSRCNAME**) parameter set to a new special value ***AGG**, which enables new aggregation-specific parameters.
2. Choose which aggregate policy to use for spreading outgoing frames across the aggregated links using the Aggregate Policy (**AGGPCY**) parameter.
 - a. The first element specifies the standard aggregation technology to be used. Only the Etherchannel technology (***ETHCHL**) value is supported.
 - b. The second element specifies the Policy type. Here are the allowed special values and their meanings:
 - ***DFT**
The adapter selection algorithm uses the last byte of the Destination IP address (for TCP/IP traffic) or MAC address (for ARP and other non-IP traffic). This mode is typically the best initial choice for a server with many clients.
 - ***SRCPRT**
The adapter selection algorithm uses the source TCP/IP or UDP port value.
 - ***DESTPORT**
The outgoing adapter path is selected through an algorithm using the destination TCP/IP or UDP port value.
 - ***SRCDESTP**
The outgoing adapter path is selected through an algorithm using the combined source and destination TCP or UDP port values.
 - ***RNDRBN**
Outgoing traffic is spread evenly across all the adapter ports in the Etherchannel. This mode is the typical choice for two hosts connected back-to-back (that is, without an intervening switch).
3. Choose a list of the communication resources that should be aggregated and specify them in the Aggregated resource list parameter.
4. You must select a single speed for all of the adapters as specified in the Line speed parameter.
5. You must set the **DUPLEX** parameter to full-duplex (***FULL**). Then, create the aggregated line description.
6. The corresponding ports on the switch must be bound together into an aggregate according to the switch's configuration manual. The configuration must indicate that all of the ports are always aggregated (as opposed to being negotiated according to some protocol).

10.8.4 Example: Configuring four aggregated network links

A logical view of another example is shown in Figure 10-9.




Figure 10-9 Logical view of example with four aggregated links

In the example, four links and IBM i communication resources (CMN14, CMN17, CMN08, and CMN11) are aggregated together with one line description named ETHLINE.

The command that is shown in Figure 10-10 creates the line description for the aggregated configuration.

```
CRTLINETH LIND(ETHLINE)
          RSRCNAME(*AGG)
          AGGPCY(*ETHCHL *RNDRBN)
          AGGRSCL(CMN14 CMN17 CMN08 CMN11)
          LINESPEED(1G)
          DUPLEX(*FULL)
          TEXT('Four link aggregated line')
```

Figure 10-10 Example CRTLINETH command for four aggregated links

For more information about configuring Ethernet resources and link aggregation, see the IBM i Information Center at the following website. For Ethernet requirements, see the hardware requirements section.

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

10.9 Ethernet Layer-2 bridging

The following sections describe the sharing of physical Ethernet connections through Ethernet Layer-2 bridging using IBM i 7.1:

- ▶ Introduction to Ethernet Layer-2 bridging
- ▶ How Ethernet Layer-2 bridging works on IBM i
- ▶ IBM i prerequisites for Ethernet Layer-2 bridging
- ▶ Configuring a shared network connection through Ethernet Layer-2 bridging

10.9.1 Introduction to Ethernet Layer-2 bridging

Logical partitions in a Power Systems system typically need access to an IP network, usually through Ethernet. However, it is not always possible or cost-effective to assign a physical Ethernet adapter to every logical partition in a Power Systems system.

One answer to this dilemma is the new Ethernet Layer-2 bridging function in IBM i 7.1. Although similar in concept to the Shared Ethernet Adapter (SEA) support provided by a Power Systems Virtual I/O Server (VIOS) partition, this IBM i function enables a single physical LAN connection to be shared by multiple logical partitions on a physical system without using Virtual IO Server (VIOS).

With IBM i 7.1, an IBM i partition can bridge a physical Ethernet port to the virtual LAN. This function reduces costs in the following ways:

- ▶ Sharing an Ethernet port means fewer Ethernet cards on the server.
- ▶ Fewer ports are needed at the network switch and fewer cables are required.
- ▶ There might be reduced administration costs because there are fewer physical resources to manage.
- ▶ Complexity might be reduced because no Virtual I/O Server partition is needed to manage the port sharing.

10.9.2 How Ethernet Layer-2 bridging works on IBM i

Ethernet Layer-2 bridging on IBM i works as follows:

1. A single partition is defined to have a physical Ethernet adapter.
2. A bridge using the IEEE 802.1D standard is configured to link a virtual Ethernet LAN to the physical Ethernet adapter.
3. Frames that are transmitted by virtual Ethernet adapters on the same VLAN as the bridging virtual Ethernet adapter can be sent to the physical network.
4. Frames that are sent from the physical network can be received by adapters on the virtual network.
5. After the bridge is in place, other partitions can access the physical network by using virtual Ethernet adapters on the bridged virtual LAN.

A logical view of the Layer-2 bridging as implemented on IBM i is shown in Figure 10-11.




Figure 10-11 Ethernet Layer-2 bridging logical view on IBM i

10.9.3 IBM i prerequisites for Ethernet Layer-2 bridging

To use Ethernet Layer-2 bridging, you must have the following prerequisites in place:

- ▶ A partition with the newest IBM i 7.1 Technology Refresh PTF and PTF Group installed
- ▶ An unused 10 Gbps or 1 Gbps-capable Ethernet adapter in the IBM i partition (excluding Host Ethernet Adapter logical ports)
- ▶ Access to the management console for the system, which can be either:
 - The Hardware Management Console
 - The IBM i Virtual Partition Manager

Tip: Use the selected Ethernet resources only for Layer-2 bridging and not for IBM i TCP/IP configuration, as there is a significant increase in processor usage for any host traffic that uses bridged resources.

10.9.4 Configuring a shared network connection through Ethernet Layer-2 bridging

To configure Ethernet Layer-2 bridging to share a network card, complete the following steps:

1. As the administrator, use the management console to create a virtual Ethernet adapter in the IBM i partition, indicating that the adapter is used for external access.
 - Configuring a virtual adapter using a Hardware Management Console

If you are using a Hardware Management Console, click **Systems Management** → **Servers** → **Configuration** → **Manage Profiles** → **Edit profile** → **Virtual Adapters** to reach the window that is shown in Figure 10-12.




Figure 10-12 Select Create Virtual Adapter - Ethernet Adapter on an HMC

When you create the virtual Ethernet adapter, select the **Access external network** check box to indicate that this virtual Ethernet adapter is used to bridge traffic to the physical network, as shown in Figure 10-13.




Figure 10-13 Creating the virtual Ethernet adapter on an HMC

- Configuring a virtual adapter using IBM i Virtual Partition Manager

If you are using the IBM i Virtual Partition Manager, the virtual Ethernet adapter is automatically created with the ability to access the external network. To create the adapter, navigate to the Change Partition Configuration panel by clicking **STRSST** → **Work with system partitions** → **Work with partition configuration** → **Change**. Create the virtual adapter by changing one or more of the virtual Ethernet identifiers to 1 (Yes), as shown in Figure 10-14.

System: RCHASRS4

Change Partition Configuration

Type changes, press Enter.

Partition identifier and name 1 10-5E050

Number of available system processors . . . : 0

Number of partition processors 1

Minimum / maximum number of processors . . . 1 / 1

Use shared processor pool 1 1=Yes, 2=No

Shared processor pool units 0 . 75

Minimum / maximum processor pool units . . . 0 . 10 / 1 . 00

Uncapped processing 2 1=Yes, 2=No

Uncapped processing weight 0 0, 64, 128, 255

Size of available memory (MB) 0

Size of partition memory (MB) 15680

Minimum / maximum size of memory (MB) . . . 320 / 16384

Enable workload management 2 1=Yes, 2=No

Virtual Ethernet Identifiers (1=Yes, 2=No)

1	2	3	4
<u>1</u>	<u>2</u>	<u>2</u>	<u>2</u>

F3=Exit F12=Cancel

Figure 10-14 Virtual Partition Manager with virtual Ethernet ID1 activated

2. On the IBM i partition with the physical adapter, create two Ethernet line descriptions:
 - a. Create a one line description for the Ethernet link (physical communications resource CMN09) connected to the physical network, as shown in Figure 10-15.

```
CRTLINETH LIND(ETHLINEP) RSRNAME(CMN09) BRIDGE(COVERED) LINESPEED(1G)
DUPLEX(*FULL) TEXT('Line for physical Ethernet link')
```

Figure 10-15 Command to create line description for the physical link

- b. Create a one line description for the new virtual Ethernet adapter (virtual resource CMN14), as shown in Figure 10-16.

```
CRTLINETH LIND(ETHLINEB) RSRNAME(CMN14) BRIDGE(COVERED) LINESPEED(1G)
DUPLEX(*FULL) MAXFRAME(8996) TEXT('Line for virtual Ethernet bridge')
```

Figure 10-16 Command to create the virtual line description for the bridge link

The resource name for a virtual adapter is found by selecting a CMNnn resource with type of 268C. Communications resources can be displayed through the Work with Hardware Resources (**WRKHDWRSC**) command by specifying the **TYPE(*CMN)** parameter.

3. To establish the bridge, the user gives the two line descriptions the same bridge name, which is a new parameter on the **CRTLINETH** and **CHGLINETH** commands for the purposes of this support. In the example commands above, the bridge name is “COVERED”.
4. When both line descriptions are active, traffic is bridged between the physical network and the virtual networks.
5. On any IBM i partition that uses the bridged connection, a line description must be created by specifying a virtual communications resource and the bridge name of “COVERED”. The command to do this task is shown in Figure 10-17.

```
CRTLINETH LIND(ETHLINVRT) RSRNAME(CMNxx) BRIDGE(COVERED) LINESPEED(1G)
DUPLEX(*FULL) MAXFRAME(8996) TEXT('Line for virtual Ethernet')
```

Figure 10-17 Command to create a virtual line description on another partition to use the bridge

For more information about configuring Ethernet resources and Layer-2 bridging, see the IBM i Information Center at the following website address. For Ethernet requirements, see the hardware requirements section.

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

10.10 IBM Portable Utilities for i (5733-SC1) supported versions

License Program Offering 5733-SC1 - IBM Portable Utilities for i contains the OpenSSH, OpenSSL, and zlib open source packages that are ported to IBM i using the PASE for i runtime environment.

For IBM i 7.1 and IBM i 6.1, the most current versions are listed along with their respective PTFs in Table 10-2.

Table 10-2 IBM Portable Utilities function, version, and PTFs

Function	Version	Version 7.1 PTFs	Version 6.1 PTFs
OpenSSH	4.7p1	SI38685, SI39965	SI40092
OpenSSL	0.9.8	SI41724	SI36892
zlib	1.2.3	None, in base	None, in base



Printing enhancements

This chapter describes the following topics that are related to printing and output that were enhanced in IBM i 7.1:

- ▶ Printing overview
- ▶ Overview of IBM i 6.1 print enhancements
- ▶ Enhancements to base print functions
- ▶ Print Services Facility for IBM i enhancements
- ▶ Transform Services enhancements
- ▶ IBM Navigator for i print enhancements
- ▶ IBM i Access for web print enhancements
- ▶ Host Print Transform enhancements

11.1 Printing overview

IBM i has powerful printing and output functions. You can present information using overlays, bar codes, graphics, images, and more. IBM i supports various industrial-type printing and presentation solutions.

Choosing and implementing a printing and presentation solution requires you to be familiar with both your organization's requirements and resources, and the capabilities that are provided by IBM i.

IBM i has both Basic Printing and Advanced Function Presentation (AFP). AFP is an architecture-based system of hardware and software for creating, formatting, viewing, retrieving, printing, and distributing information using a wide variety of printer and display devices. AFP is the original, integrated data stream on IBM i for generating fully composed pages of data.

The following list offers a high-level overview of the IBM i printing process:

1. The printing process starts when an application program runs. The application program creates output data. The output data is based on the application program and information that is contained in the printer file.
2. If print spooling is selected, the output data is placed in a spooled file and the spooled file is placed in an output queue. If direct printing is selected, the output data is sent directly to the printer.
3. The destination of the output data is based on values that are stored in several printing elements, such as job description, user profile, workstation description, printer file, and system values. Output queues are used to manage spooled files.
4. Spooled files in output queues can be used in the following ways:
 - Printed
 - Kept as records
 - Used as input to other applications
 - Transferred to other output queues
 - Sent as email
 - Used to create PDF files
5. The printer writer program interacts between the output queue and the printer and can be used to convert the printer data stream.
6. The printer writer program included in IBM i supports various printer data streams. IBM Print Services Facility™ for IBM i provides an additional function that provides support for the Advanced Function Presentation (AFP) Intelligent Printer Data Stream (IPDS).

Each printer must have a printer device description. The printer device description contains a configuration description of the printer. Printers can be attached by various attachment methods.

7. You can use a remote writer to route spooled files from an output queue on your system to another system.

Figure 11-1 shows the IBM i printing process.




Figure 11-1 IBM i printing process

11.2 Overview of IBM i 6.1 print enhancements

Major enhancements to print were made in IBM i 6.1:

- ▶ Date and time stamps in the spooling function were changed to Coordinated Universal Time. This change also required enhancements in data queue formats and addition of a third format that supports the Coordinated Universal Time time / date format.
- ▶ System i Navigator and IBM Navigator for i were enhanced with the addition of an Add Printer wizard.
- ▶ The ability to generate PDF documents directly from IBM i applications with only a minor change or override to your existing printer files was added. This function includes the following elements:
 - New IBM Transform Services for i5/OS 5761TS1 - option 1, which provides the AFP to PDF transform.
 - New printer file parameter enhancements to direct output to IFS stream files and to request that data be transformed.

For more information about the IBM i 6.1 print enhancements, see the *IBM i 6.1 Technical Overview*, SG24-7713.

11.3 Enhancements to base print functions

The sections that follow describe the following enhancements, which were made to the system base print functions:

- ▶ CPYSPLF command enhancements for copying to PDF or TIFF stream files
- ▶ New QIBM_QSP_SECURITY exit point and formats
- ▶ Encryption of PDF output (Infoprint Server LP)

11.3.1 CPYSPLF command enhancements for copying to PDF or TIFF stream files

The Copy Spoolfile (**CPYSPLF**) command was enhanced to greatly simplify copying of spooled files to the Integrated File System (IFS) as PDF or TIFF stream files.

The command interface changes are:

- ▶ Added a special value ***TOSTMF** to the To database file (**TOFILE**) parameter, which specifies the command output is to go to a stream file.
- ▶ Added the To stream file (**TOSTMF**) parameter, which is used to specify the stream file to which the output is directed.
- ▶ Added the Workstation customizing object (**WSCST**) parameter, which specifies the Workstation customizing object, which specifies the conversion to be done.
- ▶ Added the Stream file option (**STMFOPT**) parameter, which specifies if the target stream file is to be replaced.

The detailed syntax for the new parameters follows:

- ▶ To stream file (**TOSTMF**)

Specifies the stream file where the output data is written. All directories in the path name must exist. New directories are not created. This parameter must specify a value other than ***NONE** if the To database file (**TOFILE**) parameter is ***TOSTMF**.

– ***NONE**

The output is written to a user-defined physical file. This value is only valid if the To database file (**TOFILE**) parameter specifies a user-defined physical database file.

– **path-name**

Specify the path name for the stream file where the output data is to be written. This value is only valid if the To database file (**TOFILE**) parameter specifies ***TOSTMF**.

If the stream file exists, the CCSID associated with the stream file is not changed.

- ▶ Workstation customizing object (**WSCST**)

Specifies the workstation customizing object to use to transform the spooled file output to its final form before it is written to a stream file. If the To database file (**TOFILE**) parameter specifies a physical database file, the **WSCST** parameter is ignored. To convert a spooled file using the (**WSCST**) parameter, the device type (**DEVTYPE**) of the spooled file must be ***SCS** or ***AFPDS**.

– ***NONE**

Specifies that no workstation customizing object is used.

If the To database file (**TOFILE**) parameter specifies *TOSTMF and the device type of the spooled file is *AFPDS or *USERASCII, the spooled file data is copied directly to the stream file. If the stream file does not exist, the associated CCSID of the stream file is set to 65535.

For other types of spooled files, the spooled file data is copied to the stream file using the Control character (**CTLCHAR**) parameter to format the data. Lines end with carriage return and line feed controls to indicate record boundaries. If the stream file does not exist, a CCSID obtained from the spooled file attributes is associated with the stream file.

If the spooled file has a **CHRID** attribute other than *DEV D, the **CHRID** attribute is used to select the CCSID to be associated with the stream file. If the spooled file has a **CHRID** attribute of *DEV D, the CCSID of the job that created the spooled file is used.

- *PDF

The output is transformed to Portable Document Format (PDF) before it is written into a stream file. If the stream file does not exist, the CCSID associated with the stream file is set to 65535.

- Qualifier 1: Workstation customizing object name

Specify the name of the customizing object. When a named customizing object is used and the stream file does not exist, the CCSID associated with the stream file is set to 65535.

- Qualifier 2: Library

- *LIBL

All libraries in the library list for the current thread are searched until the first match is found.

- *CURLIB

The current library for the job is used to locate the customizing object. If no current library entry exists in the library list, QGPL is used.

- name

Specify the name of the library where the customizing object is located.

- Stream file option (**STMFOPT**)

Specifies whether the copy operation replaces or fails to copy the records to the stream file if a stream file with the specified name exists. If the stream file does not exist, it is created.

- *NONE

No records are copied and the operation fails.

- *REPLACE

The records replace the existing stream file records.

Because this function is implemented by PTF, there is no online or prompter help for the new parameters and the new *TOSTMF value.

In Figure 11-2, the CPYSPLF command prompt is displayed and F11 was pressed to display the parameter names. The *TOSTMF value is specified for the TOFILE parameter, as shown in the red box. This setting directs the copy function to copy a spooled file to a stream file.

Copy Spooled File (CPYSPLF)		
Type choices, press Enter.		
Spooled file	FILE	> QPJOBLOG
To data base file	TOFILE	> *TOSTMF
Library		_____
Job name	JOB	*_____
User		_____
Number		_____
Spooled file number	SPLNBR	> *LAST
Job system name	JOBSYSNAME	*ONLY
Spooled file created:	CRTDATE	
Creation date		*ONLY
Creation time		_____
To member	TOMBR	*FIRST
Replace or add records	MBROPT	*REPLACE
Control character	CTLCHAR	*NONE
More...		
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display		
F24=More keys		

Figure 11-2 Copy spooled file command prompt - TOFILE parameter with *TOSTMF value

In Figure 11-3, the next panel of the **CPYSPLF** command is shown. In the red box, the specification of the new **TOSTMF**, **WSCST**, and **STMFOPT** parameters is shown.

- ▶ The **TOSTMF** parameter directs to copy output to the /pdfoutput file.
- ▶ The **WSCST** parameter specifies a conversion to a PDF file.
- ▶ The **STMFOPT** parameter specifies the replacement of the current stream file.

```

Copy Spooled File (CPYSPLF)

Type choices, press Enter.

Channel values:                CHLVAL      -
Channel . . . . .             - *NORMAL
Line . . . . .                -
+ for more values
To stream file . . . . . TOSTMF      - /pdfoutput
-----
Workstation customizing object  WSCST    > *PDF
Library . . . . .
Stream file option . . . . . STMFOPT   > *REPLACE

F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys
Bottom

```

Figure 11-3 Copy spooled file command prompt - New **TOSTMF**, **WSCST**, and **STMFOPT** parameters

PTFs and licensed programs: This function requires PTF SI43471 for IBM i 7.1 and the 5770TS1 IBM Transform Services for i and Transforms – AFP to PDF Transform licensed program.

11.3.2 New **QIBM_QSP_SECURITY** exit point and formats

Spooled file security is primarily controlled through the output queue that contains the spool file. In V6R1, there were four main ways in which a user can be authorized to control a spooled file:

- ▶ The user is assigned spool control authority (SPCAUT(*SPLCTL)) in the user profile.
- ▶ The user is assigned job control authority (SPCAUT(*JOBCTL)) in the user profile, the output queue is operator-controlled (OPRCTL(*YES)), and the user has *EXECUTE authority to the library that the output queue is in.
- ▶ The user has the required object authority for the output queue.
- ▶ The user is always allowed to control the spooled files that are created by that user.

Spooled file security is enhanced through the addition of a spooled file security exit point. This exit point can be used with a spooled file security exit program to allow more granular access to individual spooled files based on the operation to be performed.

This new exit point changes the access criteria in the following ways:

- ▶ The user is granted authority by the use of a spooled file security exit program. A user can be granted or denied access to any spooled file for one or several operations by the exit program.
- ▶ The last two access methods that are mentioned in the previous list can be overridden by the spooled file security exit program.

This feature is another way to access spool files for the Copy Spooled File (**CPYSPLF**), Display Spooled File (**DSPSPLF**), and Send Network Spooled File (**SNDNETSPLF**) commands. If DSPDTA(*YES) is specified when the output queue is created, any user with *USE authority to the output queue is allowed to copy, display, send, or move spooled files.

If the user is authorized to control the file by one of the ways that are already listed, using DSPDTA(*NO) when you create the output queue does not restrict the user from displaying, copying, or sending the file. DSPDTA authority is checked only if the user is not otherwise authorized to the file. All of the previous access methods override DSPDTA(*NO).

Exit Point QIBM_QSP_SECURITY allows registered exit programs to control access to spooled files on a file-by-file basis. The exit programs are called at the beginning of each IBM i spool command or API, except under any of the following conditions:

- ▶ The job or thread has spool control (*SPLCTL) special authority. The special authority can originate from the user profile, group profile, or adopted authority.
- ▶ The job or thread has job control (*JOBCTL) special authority and the spooled file is on an output queue with OPRCTL(*YES). The special authority can originate from the user profile, group profile, or adopted authority.
- ▶ The command or API is run in a system job (including SCPF), a subsystem monitor job, or any job that is running under one of the system user profiles that are listed in Table 11-1.

Table 11-1 System user profiles

QAUTPROF	QCLUMGT	QCOLSRV	QDBSHR	QDBSHRDO
QDFTOWN	QDIRSRV	QDLFM	QDOC	QDSNX
QFNC	QGATE	QLPAUTO	QLPINSTALL	QMSF
QNETSPLF	QNFSANON	QNTP	QPEX	QPM400
QRJE	QSNADS	QSPL	QSPLJOB	QSRVAGT
QSYS	QTCP	QTFTP	QTSTRQS	

The following commands use the exit point

- ▶ **CHGSPLFA**: Change Spooled File Attributes
- ▶ **CPYSPLF**: Copy Spooled File
- ▶ **DLTSPLF**: Delete Spooled File
- ▶ **DSPSPLF**: Display Spooled File
- ▶ **EXPORT**: System i Navigator API to export an EBCDIC spooled file to an ASCII desktop
- ▶ **HLDSPLF**: Hold Spooled File
- ▶ **QGSRLRSC**: List Spooled File AFPDS Resources API
- ▶ **QSPMOVSP**: Move Spooled File API
- ▶ **QSPOPNSP**: Open Spooled File API
- ▶ **QUSRSPLA**: Retrieve Spooled File Attributes API
- ▶ **RLSSPLF**: Release Spooled File
- ▶ **RST**: Restore Object, Restore Library, or QSRRSTO API
- ▶ **SAV**: Save Object, Save Library, or QSRSAVO API

- ▶ **SNDNETSPLF**: Send Network Spooled File
- ▶ **SNDTCPSPLF**: Send TCP/IP Spooled File
- ▶ **WRKPRTSTS**: Work with Printing Status
- ▶ **WRKSPLFA**: Work with Spooled File Attributes

More details about the exit program format names, formats, and parameters are available in the IBM i 7.1 Information Center at:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp>

An example of creating an exit program and using the QIBM_SP_SECURITY exit point is available in IBM Software Technical Document 560810071 - "QIBM_QSP_SECURITY Exit Point: Let's See How it Works" at:

http://www-912.ibm.com/s_dir/SLKBase.nsf/1ac66549a21402188625680b0002037e/4dce2d7df8415e9c862577230076acdd?OpenDocument

11.3.3 Encryption of PDF output (Infoprint Server LP)

In IBM i 7.1, when you transform an Intelligent Printer Data Stream (IPDS) file to PDF, you can now encrypt the PDF stream or spooled file.

PDF encryption: PDF encryption is a feature of the Infoprint Server license program.

There is a new **PDFENCRYPT** parameter for the user-defined data **USRDFNDA** parameter, which is used to specify whether to encrypt an output PDF stream file or spooled file and whether to send it as email. There are several ways to specify the **USRDFNDA** parameter with the **PDFENCRYPT** parameter:

- ▶ It can be specified for an existing spool file using the Change Spooled File Attributes (**CHGSPLFA**) command.
- ▶ It can be specified in a printer file using the Create Printer File (**CRTPRTF**) command.
- ▶ It can be specified by using the Override Printer File (**OVRPRTF**) command.

Values for the **PDFENCRYPT** parameter are as follows:

- ▶ ***NONE**

Encryption and email distribution options are not specified at the spooled file level. The encryption values specified in a PDF map object are used. Specifying **PDFENCRYPT(*NONE)** is the same as not specifying **PDFENCRYPT**. A value of ***NONE** cannot not be specified with any other values. If other values are specified with ***NONE**, they are ignored.

- ▶ ***NOMAIL**

The email is not sent. If this value is not specified, the email is sent. Use ***NOMAIL** to encrypt a stream file or spooled file without sending an email.

- ▶ ***STMF**

The generated PDF file that is placed in a stream file is encrypted. If this value is not specified and a stream file is generated, the stream file is not encrypted. If the stream file distribution option is not specified, this value is ignored.

- ▶ ***SPLF**

The PDF file that is placed in a spooled file is encrypted. If this value is not specified and a spooled file is generated, the spooled file is not encrypted. If the spooled file distribution option is not specified, this value is ignored.

Generating encrypted PDF output

To encrypt the output PDF file, complete the following steps:

1. Specify the email distribution option in a PDF map object entry or segment.
2. Specify encryption values for the email distribution option.
3. Specify the name of the PDF map object with the **PDFMAP** parameter in the Print Services Facility (PSF) configuration object.
4. Specify the default IBM-supplied user program with the **PDFMAPPGM** parameter in the PS configuration object.
5. Specify the **PDFENCRYPT** parameter on the printer file's or spooled file's **USRDFNDDTA** parameter.

The following list details examples of this procedure:

- ▶ To convert a spooled file to an encrypted PDF, send it as email, and write it as an encrypted PDF file to the integrated file system, complete the following steps:
 - a. Specify the stream file distribution option in the PDF map object entry or segment.
 - b. Specify this parameter on the printer file or spooled file's attributes:
`USRDFNDDTA('PDFENCRYPT(*STMF)')`
- ▶ To convert a spooled file to an encrypted PDF file and spool it as an encrypted PDF file without sending it as email, complete the following steps:
 - a. Specify the spooled file distribution option in the PDF map object entry or segment.
 - b. Specify this parameter on the printer file or spooled file's attributes:
`USRDFNDDTA('PDFENCRYPT(*SPLF *NOMAIL)')`
- ▶ To convert a spooled file to an encrypted PDF file, spool it as an encrypted PDF file, and write it as an encrypted PDF file to the integrated file system without sending it as email, complete the following steps:
 - a. Specify the stream file distribution option in the PDF map object entry or segment.
 - b. Specify the spooled file distribution option in the PDF map object entry or segment.
 - c. Specify this parameter on the printer file or spooled file's attributes:
`USRDFNDDTA('PDFENCRYPT(*NOMAIL *STMF *SPLF)')`
- ▶ To convert a spooled file to an encrypted PDF file, spool it as a PDF file without encryption, and write it as an encrypted PDF file to the integrated file system without sending it as email, complete the following steps:
 - a. Specify the stream file distribution option in the PDF map object entry or segment.
 - b. Specify the spooled file distribution option in the PDF map object entry or segment.
 - c. Specify this parameter on the printer file or spooled file's attributes:
`USRDFNDDTA('PDFENCRYPT(*STMF *NOMAIL)')`

11.4 Print Services Facility for IBM i enhancements

This section describes the Print Services Facility for IBM i (PSF) enhancements in IBM i 7.1. They include:

- ▶ Disabling offset stacking and edge mark printing.
- ▶ Specifying public data authority for directories that are created by PSF.
- ▶ Improved PSF debugging capabilities.

11.4.1 Disabling offset stacking and edge mark printing

PSF enables disabling offset stacking on cut-sheet printers and edge mark printing on continuous forms printers through a new PSF Defined Option (**PSDFNOPT**) parameter named **OFFSTACK**:

▶ **PSDFNOPT (OFFSTACK(*YES))**

This command is the default. It specifies that there is no change to offset stacking and edge marking. Offset stacking and edge marking occurs between each spooled file. All offset stacking and edge mark changes in the form definition are honored.

▶ **PSDFNOPT (OFFSTACK(*NO))**

This command specifies that no offset stacking or edge marking is done, which includes offset stacking and edge marking done between spooled files and through the form definition.

This parameter is specified within a PSF configuration object that is either created through the Create PSF Configuration (**CRTPSFCFG**) command or specified by using the Change PSF Configuration (**CHGPSFCFG**) command.

11.4.2 Specifying public data authority for directories that are created by PSF

PSF is enhanced to enable specification of the public data authority for any directories that PSF creates when a mapping program specifies that PDF files are to be written to the Integrated File System. This enhancement applies to customer-written PDF mapping programs and to map objects. This function is enabled through a new PSF Defined Option (**PSDFNOPT**) parameter named PDF Data Authority (**PDFDTAAUT**).

The various **PDFDTAAUT** values and their functions are as follows:

▶ **PSDFNOPT (PDFDTAAUT(*INDIR))**

The authority for the directory to be created is determined by the directory in which it is created. The directory immediately preceding the new directory determines the authority.

A directory that is created in the root (/), QOpenSys, or user-defined file system is assigned the same public, private, and primary group authority, authorization list, and primary group as the directory in which it is created.

A directory that is created in QDLS for a folder defaults to *EXCLUDE for a first level folder. If created in the second level or higher, the authority of the previous level is used.

The QOpenSys and root (/) file systems use the parent directory IFS Data Authority value. If the value *INDIR is specified, PSF specifies the value *INDIR for the *PUBLIC object authority.

► **PSFDFNOPT (PDFDTAAUT(*RWX))**

You can change the object and perform basic functions on the object except those functions that are limited to the owner or controlled by object existence (*OBJEXIST), object management (*OBJMGT), object alter (*OBJALTER), and object reference (*OBJREF) authorities. Read, write, execute (*RWX) authority provides object operational (*OBJOPR) and all data authorities.

► **PSFDFNOPT (PDFDTAAUT(*RW))**

You can view and change the contents of an object. Read, write (*RW) authority provides *OBJOPR and data read (*READ), add (*ADD), update (*UPD), and delete (*DLT) authorities.

► **PSFDFNOPT (PDFDTAAUT(*RX))**

You can perform basic operations on the object, such as run a program or display the contents of a file. You cannot change the object. Read, execute (*RX) authority provides *OBJOPR and data *READ and execute (*EXECUTE) authorities.

► **PSFDFNOPT (PDFDTAAUT(*WX))**

You can change the contents of an object and run a program or search a library or directory. Write, execute (*WX) authority provides *OBJOPR and data *ADD, *UPD, *DLT, and *EXECUTE authorities.

► **PSFDFNOPT (PDFDTAAUT(*R))**

You can view the contents of an object. Read (*R) authority provides *OBJOPR and data *READ authorities.

► **PSFDFNOPT (PDFDTAAUT(*W))**

You can change the contents of an object. Write (*W) authority provides *OBJOPR and data *ADD, *UPD, and *DLT authorities.

► **PSFDFNOPT (PDFDTAAUT(*X))**

You can run a program or search a library or directory. Execute (*X) authority provides *OBJOPR and data *EXECUTE authorities.

► **PSFDFNOPT (PDFDTAAUT(*EXCLUDE))**

You cannot access the object. The OBJAUT value must be *NONE, if this special value is used. The value *NONE is not supported for **PDDTAAUT**.

If you specify a value for **PDDTAAUT** that is not supported, PSF issues PQT0038 with reason code 5 and terminates. Message PQT0038 is: Printer writer ended because of an error. Reason code 5 is: Value not recognized.

This parameter is specified within a PSF configuration object that is either created through the Create PSF Configuration (**CRTPSFCFG**) command or specified by using the Change PSF Configuration (**CHGPSFCFG**) command.

11.4.3 Improved PSF debugging capabilities

PSF added diagnostic trace files and new reason codes to aid in troubleshooting mapping problems. The new PSFTRACE diagnostic spool file contains in-depth information about your mapping program and can be used with the reason codes issued with messages PQT4140 and PQT4151 (Incorrect data was returned by mapping program) to diagnose problems with your mapping program.

Example 11-1 Command to create a data area to force a PSFTRACE spool file

```
CRDTAARA DTAARA(library/prINTER_DEVICE_NAME)
TYPE(*CHAR)
LEN(40)
AUT(*ALL)
VALUE(X'E6E6D7C4E3D9C8D98000000000000080000000000000000320000000000000000000000000000000000000000')

```

For more information about PSFTRACE and interpreting the data within it, see the “Troubleshooting Mapping Problems” section of the *Advanced Function Presentation* PDF at the IBM i 7.1 Information Center, which can be found at:

The following error codes were added to the PQT4151 (Incorrect data was returned by mapping program) error message to support new function:

- ▶ 43: Value for PDF Email Comma Delimiter must be '0' or 'X'00' when SNDDST is the mail server.
- ▶ 44: Encryption of stream file or spooled file requested but encryption settings not specified.
- ▶ 45: Value for Encrypt PDF stream file must be '0' or '1'.
- ▶ 46: Value for Encrypt PDF spooled file must be '0' or '1'.

The following Transform Services functional enhancements are described in this section:

- ### 11.5.1 Generating PDFs from existing spooled files

In IBM i 7.1, the user can now generate PDF output from a spooled file. If the job ends after the spooled file is generated and closed, the user can generate the PDF from the spooled file, regardless of whether the spool file was closed before the job ended.

Chapter 11. Printing enhancements **453**

11.5.2 PDF transform enhancements

The following sections describe enhancements to the PDF transform.

Additional barcode support

Barcode support of the PDF transform was enhanced to render all barcode types that are listed for the DDS BARCODE keyword, including intelligent USPS barcodes and two-dimensional codes.

A CPD6DF0 diagnostic message, Barcode data did not print correctly due to errors, is logged if invalid data or parameters are specified.

AFP font support improvements

The PDF transform now converts Advanced Function Presentation fonts to PDF Type 1 or Type 3 fonts and embeds them to preserve text appearance and text content.

The transform continues to revert to the PDF standard font references if font resources are not available in font libraries and the library list. The text is limited to ANSI characters.

Improved globalization

Eastern European languages require embedded fonts to display all the characters.

Non-Latin1 character identifiers (CHRIDs) are now automatically mapped to the appropriate AFP font resources.

Where possible, font attributes such font size, bold fonts, italic fonts, and so on, are honored. Font mapping can be customized through a workstation customization (WSCST) object.

For these languages and character sets, the following products might be required:

- ▶ 5648-B45 AFP Font Collection for i V3.1
- ▶ 5648-E77 InfoPrint Fonts for Multi-platform

11.6 IBM Navigator for i print enhancements

IBM Navigator for i has the following enhancements that are related to printing:

- ▶ The View as PDF task was added to enable viewing spooled files as PDFs, which include AFPDS and SCS printer output files. You can use this function to open and view the contents of a printer output file using Acrobat Reader.
- ▶ The Export task was replaced with an “Export as” menu. The menu adds the ability to export spooled files as PDF files. You can use the new export of PDF files function to save the contents of a printer output file to your client desktop, to the integrated file system (IFS), to an output queue, or to an email.

Both the “View as PDF” and “Export as PDF to client” desktop tasks use Transform Services for the AFPDS and SCS conversions to PDF.

Requirement: The “Export as PDF” to IFS, to an output queue, and to an email require the 5722IP1 Infoprint Server for iSeries licensed program.

11.6.1 Viewing printer output in PDF format

To view printer output in PDF format, you must first navigate to the printer output.

Figure 11-4 shows the navigation to access the Printer Output function. In the IBM i Management list, the first arrow points to the Basic Operations link. When that link is selected, the Basic Operations menu opens.




Figure 11-4 Navigating to the printer output list in IBM Navigator for i

The second arrow in Figure 11-4 on page 455 points to the Printer Output menu item on the Basic Operations menu. When selected, a list of printer output is displayed, as shown in Figure 11-5.




Figure 11-5 Printer Output list with menu in IBM Navigator for i

Select a file and right-click, or click the **Actions** drop-down menu, and click **View as PDF**.

When the View PDF task is selected, you see the output as a PDF, as shown in Figure 11-6.




Figure 11-6 PDF displayed from View PDF in IBM Navigator for i

11.6.2 Exporting the printer output in PDF format

Print functions of IBM Navigator for i now use Transform Services and Infoprint Server to export spooled files in PDF format. The output can be processed as one of the following options:

- ▶ Export to the client's file system
- ▶ Export to an output queue
- ▶ Export to the Integrated File System (IFS)
- ▶ Export to email

Requirement: For the latter three options, the Infoprint Server licensed program (5722-IP1) is required. Users can use the native IBM Transform Services for i (5770-TS1) licensed program to export to the IFS, but they must map a network drive to the IFS and then select the first option.

The option to use the Infoprint Server licensed program to convert spooled files to PDF remains.

The navigation to a specific printer output file is identical to what is shown in Figure 11-4 on page 455. Select a file and right click, or click the **Actions** drop-down menu. The menu opens, as shown in Figure 11-7.




Figure 11-7 Printer output list with Export PDF options shown

Click **Export as**. A menu with PDF options opens, as shown in Figure 11-7. Click **PDF using Transform Services** or **PDF using Infoprint Server**. In the next menu that opens, click the appropriate export option.

11.6.3 System i Navigator and exporting printer output in PDF format

To export printer output in PDF format using Transform Services, complete the following steps:

1. In the navigation pane, click **Basic Operations** → **Printer Output**. Select a file, right-click, and select **Convert to PDF**, as shown in Figure 11-8.




Figure 11-8 Convert to PDF using System i Navigator

2. The Convert Printer Output to PDF wizard window (Figure 11-9) opens.




Figure 11-9 Convert Printer Output to PDF wizard

The arrow points to the Store in Stream File radio button, which is consistent with saving the output in the IFS.

3. Click **Next** to have the wizard request a printer, as shown in Figure 11-10.

Because the system has no active printers capable of PDF conversion, the printer selection is disabled and the **Create new printer** radio button is automatically selected.




Figure 11-10 Convert Printer output to PDF - select printer

4. Click **Next**. You are prompted for a printer name and port, as shown in Figure 11-11. The Port defaults as shown.




Figure 11-11 Creating a PDF printer

5. Set the printer name to PDFPRINTER, and click **Next**.
6. Specify advanced configuration parameters in the printer configuration window, as shown in Figure 11-12. Click **Next**.




Figure 11-12 Specify advanced configuration parameters

The advanced parameters window (Figure 11-13) opens.




Figure 11-13 Advanced PDF printer configuration parameters

The arrow in Figure 11-13 points to an important function that minimizes the PDF size. Transform Services embed the PDF fonts in the PDF to preserve text appearance and text content. This action increases the size of the PDF file. This option directs transforms to not embed the PDF fonts.

7. Click **Next**. Another advanced parameters window (Figure 11-14) opens. Accept the defaults and click **Next**.




Figure 11-14 Advanced parameters

Enter the IFS path where the PDF is to be stored (Figure 11-15).




Figure 11-15 Setting the destination in the IFS

8. Create and secure the directories in the IFS, according to the following rules:
 - The directories must exist.
 - The QSPLJOB user (or *PUBLIC) must have *RWX (read / write / execute) authority to the root (/) directory.
 - The QSPLJOB user must have a minimum of *X (execute) authority to the directories in the path.
 - The QSPLJOB user must have *RWX (read / write / execute) authority to the directory where the files are stored.
9. Click **Next** to continue, and click **Finish** in the confirmation window to print.

11.7 IBM i Access for web print enhancements

IBM i Access now uses Transform Services to view spooled files in PDF format.

Support was added for the AFP to PDF Transform (option 1) of the IBM Transform Services for i (5770-TS1) licensed program when you view printer spool output as a PDF document.

Output can be viewed in a browser or placed in the IBM i integrated file system (IFS).

For more information, see the “IBM i Access for Web” topic in the IBM i 7.1 Information Center, or the IBM i Access for web PDF at:

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzamm/rzamm.pdf>

Change in requirements: System i Navigator and IBM i Access for Web previously required the 5722-IP1 IBM Infoprint Server for iSeries product to view output as PDF. This option is still usable for users that have the software, but it is no longer required.

11.8 Host Print Transform enhancements

Host Print Transform now uses Transform Services. The following sections explain the changes with examples.

11.8.1 The Host Print Transform API and Transform Services

The Host Print Transform (QwpzHostPrintTransform) API now uses Transform Services to convert existing *SCS and *AFPDS spooled files to PDF.

The function behavior is directed by a Workstation Customization Object (WSCST), which must specify the CTXFORM attribute. Three such WSCST objects with CTXFORM attribute are shipped with the system. One of these objects can be used, or you can create your own using the Create WSCST (**CRTWSCST**) command. The three shipped WSCST objects are:

- QSYS/QCTXPDPF

This WSCST object includes tags to map all of the CJK World Type fonts included in IBM i option 43 to corresponding registered CMaps and character collections. This object allows ideographic characters to be rendered without embedding these large fonts within the document. The smaller World Type fonts are not mapped by this object and are embedded within the output document to allow non-Latin1 SBCS languages to be presented.

- QSYS/QCTXPDPFWT

This WSCST object contains tags to map all of the World Type fonts included with IBM i option 43 to corresponding Type 1 fonts or registered CMaps and character collections. This situation generally results in the smallest possible PDF file size, but only reliably presents Latin and CJK languages.

- QSYS/QCTXPDPFMBD

This WSCST is the minimal PDF object. It omits all font tags, so that the default font mapping is used for all font references. The default behavior is to embed all TrueType font references.

When the API is started with a WSCST object with the CTXFORM attribute, the job reads the input data stream from the spooled file that is specified in the API. Transform Services is called to generate the PDF output from the input spooled file data. Transform Services returns the PDF output in the output buffer that is provided on the API. For more information, see the API documentation in the IBM i 7.1 Information Center at:

<http://pic.dhe.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=%2Frzahg%2Ficmain.htm>

11.8.2 Example: Host Print Transform API with Transform Services

This sample flow illustrates the usage of the API to transform a spooled file to a PDF using a workstation customization object of type CTXFORM:

1. The Host Print Transform API is called with the initialize (10) process option.

The Host Print Transform API sets any initial information and returns to the caller. The first eligible spooled file is selected for processing.

2. The Host Print Transform API is called with the process file (20) process option.

The name and identifier of the spooled file is passed in the input information. A workstation customization object of type *CTXFORM is passed in.

The Host Print Transform API checks to see that the iCTT transform service is available and determines whether it transforms the spooled file. It returns to the caller an indication whether it transforms the spooled file. It also passes back an indication that it reads the spooled file directly and the caller does not pass it any data. The end file (40) process option must be used before another process file (20) process option is allowed.

3. The Host Print Transform API is called with the transform data (30) process option.

The Host Print Transform API reads the spooled file data, transforms the data, and passes back the transformed data to the caller. If the buffer used by the caller cannot hold all of the transformed data, the Host Print Transform API indicates to the caller that it is not done transforming the file. The caller must call the API again with another transform data (30) process option. This step happens repeatedly until the entire transformed file is passed back to the caller, at which point the Host Print Transform API indicates that it is done transforming the file.

4. The Host Print Transform API is called with the end file (40) process option.

The Host Print Transform API returns to the caller any remaining data to be sent to the printer.

5. The caller selects the next eligible spooled file to be transformed.

Steps that start with the process file (20) process option are repeated.

6. All transformation is completed.

The Host Print Transform API is called with the terminate (50) process option. The Host Print Transform API cleans up any workspaces that it created and returns to the caller.

11.9 References

The following references have more information about IBM i printing:

- ▶ *IBM i Printing Basic Printing*, found at:
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzalu/rzalu.pdf>
- ▶ *IBM i Printing Advanced Function Presentation*, found at:
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzau6/rzau6.pdf>
- ▶ *IBM Advanced Function Printing Utilities for iSeries: User's Guide*, found at:
http://publib.boulder.ibm.com/infocenter/iseriess/v6r1m0/topic/books_web/s5445349.pdf
- ▶ *IBM i Files and File Systems Spooled Files*, found at:
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzata/rzata.pdf>
- ▶ *System i Programming - DDS for Printer Files*, found at:
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzakd/rzakd.pdf>
- ▶ *iSeries Guide to Output*, found at:
<http://publib.boulder.ibm.com/infocenter/iseriess/v6r1m0/topic/rzalu/s5445319.pdf>
- ▶ InfoPrint AFP Font Collection, found at:
http://www-03.ibm.com/systems/i/software/print/afpfonthome_m_ww.html

- ▶ InfoPrint AFP Print Utilities for System i, found at:
http://www-03.ibm.com/systems/i/software/print/afputilhome_i_ww.html
- ▶ Print Services Facility for IBM i for i5/OS, found at:
http://www-03.ibm.com/systems/i/software/print/ipfontscmp_m_ww.html



Integration with IBM BladeCenter and IBM System x

This chapter describes the new enhancements to the integration with BladeCenter and System x with IBM i 7.1

This chapter describes the following topics:

- ▶ iSCSI software targets
- ▶ Defining iSCSI software target support
- ▶ Service Processor Manager function
- ▶ VMware support changes
- ▶ New planning worksheets
- ▶ IBM Navigator for i
- ▶ New IBM i CL commands
- ▶ IBM i changed CL commands
- ▶ Create NWS Configuration (CRTNWSCFG) and Change NWS Configuration (CHGNWSCFG) CL commands
- ▶ Install Linux Server (INSLNXSVR) CL command
- ▶ No new integrated Linux servers
- ▶ Fewer IBM i licensed programs are required
- ▶ Changes to IBM i integration with BladeCenter and System x documentation

12.1 iSCSI software targets

With i 7.1, IBM i now supports iSCSI software targets using standard Ethernet Network Interface Cards (NICs), as shown in Figure 12-1. Software targets provide additional flexibility for the IBM iSCSI target solution.




Figure 12-1 Attaching servers to IBM i using iSCSI

Within the IBM i Integrated Server Environment, you are limited to 1 Gb connectivity if you use physical iSCSI Target HBAs.

With the new software target solution, you can now use dedicated Ethernet ports with 1 Gb or 10 Gb connectivity. It is now possible to intermix hardware and software target adapter environments.

However, if you are using an iSCSI software initiator in combination with an iSCSI software target, you have full 10 Gb connectivity.

12.1.1 IBM i Integrated server object model with a hardware target

With a physical iSCSI Target HBA, the NWSH object identifies the hardware resource and configures the IP address information (and other attributes), as shown in Figure 12-2.




Figure 12-2 NWSH object for a physical iSCSI target HBA

12.1.2 IBM i Integrated server object model with a software target

In IBM i 7.1, with a software iSCSI Target, the NWSH object, along with a line description and TCP interface, identify the hardware resource and configure the IP address information (and other attributes), as shown in Figure 12-3.




Figure 12-3 NWSH object for a software target

The physical iSCSI HBA is replaced by an Ethernet NIC, which is defined in the Ethernet line description.

Within the TCP/IP interface, configure the IP address for this Ethernet NIC.

This same IP address is also used in the local SCSI interface parameter for the NWSH configuration object.

12.1.3 Direct connect software targets

IBM i Integrated Server Support is enhanced to provide support for iSCSI direct connect when you use an IBM i software target (Ethernet NIC). This enhancement eliminates the requirement for a switch between an IBM i iSCSI software target and a BladeCenter blade or System x iSCSI initiator. This support is provided in IBM i 7.1 with IBM i integration with BladeCenter and System x Group PTF SF99369 Level 12 or later.

12.2 Defining iSCSI software target support

In IBM i 7.1, it is possible to define iSCSI software target support in a text-based interface and by using IBM Navigator for i.

12.2.1 CRTDEVNWSH CL command interface

To define a iSCSI software target (Figure 12-4), you must specify *VRT for the resource parameter for the Create Device Description for a Network Server Host (CRTDEVNWSH) command.

```

                                Create Device Desc (NWSH) (CRTDEVNWSH)

Type choices, press Enter.

Device description . . . . . DEVD          > SWTARGET
Resource name . . . . . RSRNAME          > *VRT
Local (target) interface: LCLIFC
  Subnet mask . . . . .                  > *NONE
  Port speed . . . . .                   > *AUTO
  Duplex . . . . .                       > *AUTO
Local SCSI interface:
  Internet address . . . . .              > 172.16.211.100
  Gateway address . . . . .              > *NONE
  SCSI TCP port . . . . .                > 3260
Local LAN interface:
  Internet address . . . . .              > *IPSCSI
  Gateway address . . . . .              > *NONE
  Virtual Ethernet base UDP port         > 8801
  Cable connection . . . . .             > *NETWORK
Online at IPL . . . . . ONLINE           *NO

More...
F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
F13=How to use this display   F24=More keys
```

Figure 12-4 CRTDEVNWSH command

For the **LCLIFC** parameter, specify the ***IPSCSI** option, which indicates that the local LAN interface IP address is the same as the local SCSI interface IP address.

12.2.2 IBM Navigator for i changes for iSCSI software target support

It is possible to define the iSCSI software target support by using IBM Navigator for i. You can use the New Network Server Host Adapter option in the Integrated Server Administration window, as shown in Figure 12-5.




Figure 12-5 Network server host adapter - virtual resource

Here you can specify the Virtual for the Hardware resource parameter to create the Network server host adapter device description for the iSCSI software target.

Within IBM Navigator for i, it is possible to create a TCP/IP interface and a corresponding line description when you create an NWSH configuration object. You can do this task by clicking **New**, as shown in Figure 12-6.




Figure 12-6 Create a TCP/IP interface for an NWSH configuration object

12.3 Service Processor Manager function

With IBM i 7.1, the Service Processor Manager function of IBM i Integrated Server Support is now used for integrated server management connections and power control.

Before IBM i 7.1, this function was provided by IBM Director (5722-DR1), which is no longer used for this purpose.

12.4 VMware support changes

With IBM i 7.1, VMware ESX 4, ESXi 4, and ESXi 5 are now supported on iSCSI-attached integrated servers. ESXi 4 or later also includes support for iSCSI software initiators.

For the specific OS versions that are supported by each IBM i release, see the “Integrated Server Operating System (Server OS) Versions” section of the *IBM i iSCSI Solution Guide* at: http://www.ibm.com/systems/i/advantages/integratedserver/iscsi/solution_guide.html

12.4.1 New NWSD types

New network server description (NWSD) types are provided for VMware ESX servers. The new types eliminate the requirement for an “install” drive (the second drive) on ESX servers.

Figure 12-7 shows the new *ESX Server operating system Network server type for the CRTNWSD command.

Create Network Server Desc (CRTNWSD)

Type choices, press Enter.

Network server description . . .	> VMWESX	Name
Resource name	> *NONE	Name, *NONE, *AUTO
Network server type:		
Server connection	> *ISCSI	*IXSVR, *ISCSI, *GUEST...
Server operating system . . .	> *ESX	*WIN32, *AIXPPC, *ESX...
Storage path:		
Network server host adapter .		Name, *NONE
IP security rules:		
Remote interface 1 rule . . .	*DFTSECRULE	1-16, *DFTSECRULE, *NONE
Remote interface 2 rule . . .	*DFTSECRULE	1-16, *DFTSECRULE, *NONE
Remote interface 3 rule . . .	*DFTSECRULE	1-16, *DFTSECRULE, *NONE
Remote interface 4 rule . . .	*DFTSECRULE	1-16, *DFTSECRULE, *NONE
	+ for more values	
Default IP security rule	*NONE	1-16, *NONE
Multi-path group	*NONE	1-4, *NONE
	+ for more values	
More...		
F3=Exit	F4=Prompt	F5=Refresh
F12=Cancel	F13=How to use this display	
F24=More keys		

Figure 12-7 CRTNWSD command for the *ESX Server operating system

For VMware ESXi embedded servers, the system drive (the first drive) is no longer required.

Requirement: VMware ESX servers that were installed on prior IBM i releases must be migrated to the new NWSD type after you install IBM i 7.1.

12.4.2 VMware ESX server management

With VMware ESX Server, the IBM i Integrated Server Support software (including the administration functions, such as shutdown) does not run directly on the VMware ESX server. Instead, an iSCSI-attached integrated Windows server serves as a management server for the VMware ESX server.

An integrated Windows server can serve as the management server for any number of integrated VMware ESX servers in the same IBM i logical partition. At least one integrated Windows server is required in each IBM i logical partition that hosts integrated VMware ESX servers.

12.4.3 SWA storage spaces for VMware ESX servers

With IBM i 7.1, save while active (SWA) support is provided for integrated VMware ESX servers. Storage spaces for VMware ESX servers can be saved from IBM i while the ESX server is active. This setup allows a concurrent save of ESX data without requiring the ESX server to be shut down or applications ended.

For more information, see the *IBM i iSCSI Solution Guide*, found at:

http://www.ibm.com/systems/i/advantages/integratedserver/iscsi/solution_guide.html

12.5 New planning worksheets

New planning worksheets were added to *IBM i iSCSI Solution Work Sheets* PDF:

- ▶ IBM i TCP/IP interface worksheet
This worksheet is used when you plan iSCSI software targets.
- ▶ IBM i line description worksheet
This worksheet is used when you plan iSCSI software targets.
- ▶ Integrated server installation worksheet
This worksheet replaces the Install Windows Server (**INSWNTSVR**) command worksheet.
- ▶ VMware ESX post-installation worksheet
This worksheet is used for VMware ESX server post-installation tasks.

In addition, these worksheets are enhanced to allow them to be completed and saved as softcopies.

The instructions for filling out these worksheets are in the *IBM i iSCSI Solution Guide* PDF. Both PDFs are available at:

http://www.ibm.com/systems/i/advantages/integratedserver/iscsi/solution_guide.html

The instructions and worksheets were previously part of the *iSCSI Network Planning Guide* topic in the Information Center.

12.6 IBM Navigator for i

The IBM Navigator for i web GUI is now the preferred user interface for managing integrated servers. Therefore, most integrated server management tasks are documented using the web GUI.

GUI tasks: The System i Navigator GUI is still available in IBM i 7.1 and works for many tasks. However, the new GUI tasks that are listed in the following paragraphs and support for IBM i 7.1 enhancements are not available in the System i Navigator GUI.

New GUI tasks are available within the IBM Navigator for i web GUI, and are described in the following sections:

- ▶ Create Server task
- ▶ Clone Integrated Windows Server task
- ▶ Delete Server task
- ▶ Launch Web Console

12.6.1 Create Server task

This task, as shown in Figure 12-8, creates an iSCSI-attached integrated server.




Figure 12-8 Create Server option in the IBM Navigator for i web GUI

To create a server in an IBM i Integrated Server environment, use IBM Navigator for i for a walk-through of a server installation.

You must configure the following objects on the IBM i side:

- ▶ A network server host adapter (NWSH) configuration object
- ▶ A remote system configuration for iSCSI attached servers
- ▶ A service processor configuration for iSCSI attached servers

In the window that is shown in Figure 12-8 on page 475, click **Create Server**. This action starts the Create Server wizard that is shown in Figure 12-9. This wizard is helpful for creating a server that uses IBM i virtual storage.




Figure 12-9 Create Server wizard

The wizard guides you through the following tasks:

- ▶ Selecting a supported operating system
- ▶ Allocating IBM i virtual storage for the server operating system, if necessary
- ▶ Selecting the hardware for the server
- ▶ Selecting an IBM i iSCSI target to use for I/O
- ▶ Selecting an IBM i memory pool to use for I/O
- ▶ Providing a name to identify the server configuration on IBM i

Tip: Review the “Server installation roadmap and checklist” chapter of the *IBM i iSCSI Solution Guide* before you use this wizard. This guide can be found at:

http://www-03.ibm.com/systems/resources/systems_power_ibmi_iscsi_solution_guide.pdf

For more information, go to:

http://www.ibm.com/systems/i/advantages/integratedserver/iscsi/solution_guide.html

12.6.2 Clone Integrated Windows Server task

The New Based On...(cloning) task, which is shown in figure Figure 12-10, creates an iSCSI-attached integrated Windows server that is based on one that was previously installed.




Figure 12-10 New Based On...(cloning) option in the IBM Navigator for i web GUI

In Figure 12-10 on page 477, click **New Based On....** This action starts the Create Server Based On (cloning) wizard that is shown in Figure 12-11.




Figure 12-11 Create Server Based On (Cloning) wizard

The cloning wizard guides you through the following tasks:

- ▶ Providing a name to identify the clone server configuration on IBM i
- ▶ Selecting virtual storage to copy from the base server to the clone server
- ▶ Selecting the hardware for the clone server

The server cloning process is provided for integrated servers that are running supported Windows Server editions. The cloning process requires that you prepare the base server for cloning before you use the cloning task. Additional configuration is required after the server is cloned.

Review Chapter 5, “Server cloning roadmap and checklist”, of the *IBM i iSCSI Solution Guide* PDF before you use this wizard. It can be found at:

http://www.ibm.com/systems/i/advantages/integratedserver/iscsi/solution_guide.html

12.6.3 Delete Server task

This new task deletes an integrated server configuration, as shown in Figure 12-12.




Figure 12-12 Delete Server option in the IBM Navigator for i web GUI

This option is only available when the server is not active or starting.

12.6.4 Launch Web Console

This new task launches the service processor web console for an iSCSI-attached BladeCenter blade or System x server. An IBM BladeCenter can launch the Advanced Management Module web interface, as shown in Figure 12-13.




Figure 12-13 Launching the web console from the IBM Navigator for i web GUI

12.6.5 Simplified Windows File Level Backup (FLB) from IBM i

The IBM Navigator for i web GUI is enhanced to simplify the task to enable which Windows share names under the IBM i /QNTC/servername directory can be saved from IBM i. The web GUI now provides a File Level Backup tab on the integrated server properties page, as shown in Figure 12-14.




Figure 12-14 File Level Backup tab on the integrated server properties

This new tab provides a way to select share names to enable for backup and eliminates the need to manually add a member to the QAZLCSAVL file in QUSRSYS and then manually edit the file member to list the share names to enable for backup.

For more information about this topic, see section 7.7.2, "Enabling Windows share names for file level backup from IBM i" in the *IBM i iSCSI Solution Guide* PDF. This guide can be found at:

http://www.ibm.com/systems/i/advantages/integratedserver/iscsi/solution_guide.html

12.7 New IBM i CL commands

The following new IBM i control language (CL) commands are available for integrated servers:

- ▶ Install Integrated Server (**INSINTSVR**)
- ▶ Delete Integrated Server (**DLTINTSVR**)

12.7.1 Install Integrated Server (INSINTSVR)

The **INSINTSVR** command installs an iSCSI-attached integrated Windows Server 2008 or VMware ESX server, as shown in Figure 12-15.

```

                                Install Integrated Server (INSINTSVR)

Type choices, press Enter.

Network server description . . . NWSD
Operating system type . . . . . OSTYPE
Remote system NWSCFG . . . . . RMTNWSCFG
Storage path:                      STGPTH
  Network server host adapter .
Pool identifier . . . . . POOL          *BASE
Server storage space sizes:        SVRSTGSIZE
  System size . . . . .                *CALC
Storage space ASP:                 SVRSTGASP
  System ASP . . . . .                1
Server storage ASP device:         STGASPDEV
  System ASP device . . . . .
Text 'description' . . . . . TEXT      *BLANK

More...

F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys
```

Figure 12-15 *INSINTSVR* command

For more information about this topic, go to:

<http://publib.boulder.ibm.com/infocenter/iseres/v7r1m0/topic/cl/insintsvr.htm>

12.7.2 Delete Integrated Server (DLTINTSVR) command

The **DLTINTSVR** command, which is shown in Figure 12-16, deletes an integrated server configuration.

```

                                Delete Integrated Server (DLTINTSVR)

Type choices, press Enter.

Network server description . . . NWSD

Bottom

F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys
```

Figure 12-16 *DLTINTSVR* command

12.8 IBM i changed CL commands

The following new IBM i control language (CL) commands are changed for integrated servers:

- ▶ Install Windows Server (**INSWNTSVR**) CL command
- ▶ Create NWS Configuration (**CRTNWSCFG**) and Change NWS Configuration (**CHGNWSCFG**) CL commands

12.8.1 Install Windows Server (INSWNTSVR) CL command

The Install Windows Server (**INSWNTSVR**) command has a number of parameter changes that might require you to recompile any existing CL programs that use this command:

- ▶ The Windows server version (**WNTVER**) parameter no longer supports the installation of Windows 2000 Server. Related to this change is the removal of the *TSENABLE special value for element 3 of the License mode (**LICMODE**) parameter.
- ▶ The Windows server version (**WNTVER**) parameter no longer supports the installation of Windows Server 2008. To install Windows Server 2008, use the new Install Integrated Server (**INSINTSVR**) command.
- ▶ The Installation type (**INSTYPE**) parameter is obsolete and was removed. ServerGuide assisted installations are no longer supported.
- ▶ The Install option (**OPTION**) parameter is obsolete and was removed. There are no supported upgrades that can be processed using the Install Windows Server (**INSWNTSVR**) command.
- ▶ The Enable unicast (**ENBUNICAST**) parameter is obsolete and was removed. Communication with the service processor is configured by specifying a value for Service processor name (**SPNAME**) or SP Internet Address (**SPINTNETA**) parameter.
- ▶ The Boot device ID (**BOOTDEVID**) parameter is obsolete and was removed.
- ▶ The Remote (initiator) interfaces (**RMTIFC**) parameter, SCSI, and LAN Gateway addresses are not supported. The iSCSI attached servers do not support bridged networks that require the gateway address. These elements were removed.

12.8.2 Create NWS Configuration (CRTNWSCFG) and Change NWS Configuration (CHGNWSCFG) CL commands

The Enable Unicast (**ENBUNICAST**) parameter was removed from the Create NWS Configuration (**CRTNWSCFG**) and Change NWS Configuration (**CHGNWSCFG**) commands. This parameter change might require you to recompile any existing CL programs using these commands.

In IBM i 7.1, iSCSI-attached integrated servers no longer support the multicast discovery method for the remote server service processor. Instead, unicast discovery of the remote server service processor must be used. Existing network server configurations of type *SRVPRC that have Enable Unicast (**ENBUNICAST**) configured to *N0 must use the Change NWS Configuration (**CHGNWSCFG**) command to specify either the Service Processor Name (**SPNAME**) or Service Processor IP Address (**SPINTNETA**) parameter.

iSCSI-attached network server descriptions cannot vary on until the network server configurations of type *SRVPRC with **ENBUNICAST** configured to *N0 are changed.

12.8.3 Install Linux Server (INSLNXSVR) CL command

The Install Linux Server (**INSLNXSVR**) CL command is no longer supported in IBM i 7.1. The **INSLNXSVR** command was used for Linux and VMware ESX server installations on prior IBM i releases.

Tip: There are no alternatives available for Linux server installs. For VMware ESX server installs, use the Create Server web GUI task or the **INSINTSVR** command.

12.8.4 No new integrated Linux servers

New integrated Linux servers cannot be installed on IBM i 7.1. Integrated Linux servers that were installed on prior IBM i releases and migrated forward to IBM i 7.1 can continue to run as is, but without service support. The suggested migration path for these servers is to install an integrated VMware ESX server and run the Linux server as a virtual machine under VMware ESX.

12.9 Fewer IBM i licensed programs are required

The following IBM i licensed programs are no longer needed for integrated server functions:

- ▶ IBM Extended Integrated Server Support for i5/OS (5761-LSV)
- ▶ IBM Director (5722-DR1)
- ▶ Qshell (5770-SS1 option 30)

12.10 Changes to IBM i integration with BladeCenter and System x documentation

The following sections list the changes that were implemented in various supporting documentation that is available.

12.10.1 A new IBM i iSCSI Solution Guide PDF

A new *IBM i iSCSI Solution Guide* PDF and associated planning worksheets are now available at:

http://www.ibm.com/systems/i/advantages/integratedserver/iscsi/solution_guide.html

This guide provides the information that you need to plan for and install an IBM BladeCenter blade or System x server that is attached to IBM i using an iSCSI network (iSCSI). This guide contains the following information:

- ▶ iSCSI solution support matrixes: See the capabilities that the solution provides, which IBM BladeCenter and System x server models and operating systems are supported, and much more.
- ▶ Concepts: Learn about how the solution works.
- ▶ Server installation roadmap and checklist: Required information to install a server that is integrated with IBM i.

- ▶ Server cloning roadmap and checklist: Required information to clone a Windows server that is integrated with IBM i.
- ▶ BladeCenter and System x configuration: iSCSI configuration tasks for BladeCenter blade and System x servers.
- ▶ Additional topics: Other topics that are related to the iSCSI solution.

This guide consolidates and replaces the following information:

- ▶ iSCSI installation roadmap PDFs for i 7.1, i 6.1, and i 5.4
- ▶ *iSCSI Initiator Hardware Configuration* PDF
- ▶ iSCSI Install Read Me First website
- ▶ BladeCenter and System x models that are supported by the iSCSI website
- ▶ iSCSI target and iSCSI initiator tables that were formerly on the iSCSI solution (iSCSI) website
- ▶ Ethernet switches for the iSCSI website
- ▶ Ordering channels for the iSCSI website
- ▶ Various websites that are related to tape and optical device support
- ▶ Various websites that are related to migration

12.10.2 IBM i 7.1 Information Center

The IBM i integration with BladeCenter and System x topic in the IBM i 7.1 Information Center was updated at:

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp?topic=/rzahq/rza hqntspo.htm>

Here are the most significant updates to this topic:

- ▶ Information that is related to various IBM i 7.1 enhancements was added.
- ▶ The following documentation was migrated to the *IBM i iSCSI Solution Guide* PDF:
 - The Integrated server installation roadmap chapter
 - BladeCenter and System x hardware installation and configuration information
 - Microsoft Windows Server and VMware ESX Server installation and configuration information
- ▶ The iSCSI network planning worksheets were migrated to the *IBM i iSCSI Solution Work Sheets* PDF.
- ▶ Integrated Windows servers that run on an Integrated IBM eServer xSeries® Server (IXS) or a server that is attached using an Integrated xSeries Adapter (IXA) are no longer documented.
- ▶ Because the new Create Server task is now available within IBM Navigator for i, the Windows server installation advisor is no longer provided.
- ▶ Integrated Linux servers are no longer documented
- ▶ Network server description configuration files are no longer documented.

12.10.3 New IBM i integration with BladeCenter and System x group on IBM developerWorks

This IBM developerWorks group is the place where developers and IT professionals who work with the IBM i iSCSI solution congregate to connect, share, and collaborate. The website can be found at:

<http://www.ibm.com/developerworks/groups/IBMiIntegratedServer>

This group provides:

- ▶ A message board
- ▶ Bookmarks for important websites and reference material
- ▶ A set of Wiki pages, including:
 - Documentation
 - Learning Resources
 - Service and support (including required group PTFs for each IBM i release)

12.10.4 New IBM i Technology Updates page on developerWorks

The new Integration with BladeCentre and System x page on the IBM i Technology Updates Wiki lists the latest integrated server enhancements, including the required group PTF levels. It can be found at:

<http://www.ibm.com/developerworks/ibmi/techupdates/IBMiIntegratedServer>

12.10.5 IBM i integration with BladeCenter and System x Marketing website

This Marketing website was streamlined and can be found at:

<http://www.ibm.com/systems/i/advantages/integratedserver>

Most of the technical content that was previously on this website (for example, the iSCSI Install Read Me First web page) was migrated to the *IBM i iSCSI Solution Guide* or to developerWorks. For more information, see the previous sections.



IBM Advanced Job Scheduler for i enhancements

This chapter describes the new functions that are available in IBM i 7.1, which are related to job scheduling, through the 5770-JS1 IBM Advanced Job Scheduler for i licensed program.

This chapter describes the following topics:

- ▶ Advanced Job Scheduler capabilities
- ▶ IBM Systems Director Navigator for i AJS support
- ▶ Other AJS enhancements
- ▶ References

13.1 Advanced Job Scheduler capabilities

The IBM Advanced Job Scheduler for i (AJS) can be used to schedule complex batch job flow, distribute spooled files as an email attachment, and send job status notifications. Jobs can run based on a schedule or an event. Jobs can run on local or remote IBM i systems and commands can be processed on non- IBM i platforms.

13.1.1 Scheduling a job

Jobs can be scheduled with a simple schedule (such as every week Monday through Friday) or more complex schedules (such as the seventh working day every month). If the standard scheduling conventions within the AJS are not enough, you can create a scheduling calendar that contains all the dates that a job is to run. Holiday calendars can be used for holidays or exceptions to the normal scheduling. These dates are the dates that you do not want a job to run. These calendars can be used in multiple jobs.

13.1.2 Job groups

Job groups are made up of AJS jobs that are grouped to run consecutively in the order that is specified in the group sequence field. This dependency is one form of dependency in the AJS. A normal completion is required before the next job in the group is submitted.

13.1.3 Notification and report distribution

You can use the notification feature in the AJS to notify people of the status of jobs by email. Important messages can be sent with an escalation list attached that specifies the length of time that passes before the next person in the list is notified. Escalation continues until someone stops escalation. You can use notification, which is coupled with report distribution, to distribute spooled files to one or more email addresses or to one or more output queues automatically.

13.1.4 Remote command jobs

Commands can run on non- IBM i platforms such as PCs. They can be processed on a single system or a group of systems. The Run Remote (**RUNRMTCMD**) command is used to run the commands on the remote systems. You can run **RUNRMTCMD** to run a command on a remote system that is running the target portion of this function. The target portion of this function can be a Remote Executing Daemon (REXECD). Use the incoming remote command (IRC) service of the IBM i Access for Windows on the remote systems to process the commands.

13.1.5 Multiple scheduling environments

The AJS manages multiple scheduling environments. A scheduling environment is all the objects in the QUSRIJS data library that is duplicated into another library. QUSRIJS is the main data library that is created during the installation of AJS. It contains the physical files that store all the information about the scheduled jobs. Each scheduling environment is its own entity and can be accessed and managed using the AJS interfaces. When you define a scheduling environment within the AJS, a monitor switch is available to designate whether the scheduling environment can be active. A scheduling environment must be active to submit jobs automatically.

13.2 IBM Systems Director Navigator for i AJS support

IBM Systems Director Navigator for i (formerly known as IBM Systems Director Navigator for i5/OS) is a web console interface for IBM i administration where you can work with the web-enabled actions of the System i Navigator client.

IBM Systems Director Navigator for i includes a number of welcome pages that enable the user to find the action that they want to perform quickly. Most functions found in IBM Systems Director Navigator for i are also found in IBM Systems Director, which handles multiple IBM i systems and non- IBM i system platforms.

In IBM i 6.1, the AJS function in IBM Systems Director for i5/OS was limited to viewing. Viewing options include the following elements:

- ▶ Activity logs for the system, for a scheduled job, and for a specific execution of a job
- ▶ Configured jobs, their properties, and their status
- ▶ Configured groups, their properties, and their status
- ▶ Scheduled jobs and their status

The major limitation of the IBM i 6.1 function was that it could not change anything, and it could not add, change, hold, or remove scheduled jobs.

The new IBM i 7.1 AJS functions in the Systems Director interface now includes most of the same functions that are found in the System i Navigator Windows client, with the advantage that the interfaces are web-based, not client-based.

The rest of this section walks through the AJS web pages and describes the new functions during the walkthrough.

13.2.1 Navigating to the AJS menu

The AJS menu is found as part of the Work management advanced function menu. It is not shown on the basic function menu, so navigation instructions follow.

IBM Systems Director Navigator for i navigation

After you log on to the IBM Systems Director Navigator for i web page, the Welcome window opens, as shown in Figure 13-1. On initial access, the IBM i Management option, pointed to by the top arrow, is minimized. When this option is expanded, the list of IBM management functions is shown. One of these functions is Work Management, to which the second arrow points.




Figure 13-1 Welcome window with IBM i Management item expanded

When the Work Management function is selected, the Work Management main menu that is shown in Figure 13-2 opens.




Figure 13-2 Work Management main menu

The Work Management menu includes many of the operations level actions of IBM i work management. None of the job scheduler actions display on this window. To access the job scheduler functions, click **Show All Work Management Tasks**.

To make the subsequent window larger, select the icon to hide the navigation area, and then click **Show All Work Management Tasks**.

The Work Management Advanced menu is shown in Figure 13-3.




Figure 13-3 Work Management all tasks menu

The first time the Work Management advanced menu opens, the AJS menu is minimized. When you select the icon to which the arrow points (**Advanced Job Scheduler**), the menu is expanded, showing the job scheduler menu items as highlighted by the vertical bar.

To the right of each item is an icon that, when selected, displays a menu of actions. Figure 13-3 shows the menu box for the AJS menu item.

Following each of the menu items are icons that indicate that menus are available for that item. The AJS item's menu has vital actions that are described in 13.2.2, "AJS menu" on page 496.

IBM Systems Director navigation

After you log on to the IBM Systems Director, you must navigate to an IBM i system. Complete the following steps:

1. Click the **Navigate Resources** task in the Navigation window.
2. Select the **All Systems** systems group from the Navigate Resources window, as shown in Figure 13-4.





Figure 13-4 Selecting Navigate Resources and the All Systems resource group

3. Select an IBM i system from the All Systems group list window.

In our example, we selected the **All Systems** group in Figure 13-4 on page 493 and the All Systems group systems window opened. We then scrolled down in search of a specific IBM i system.

Figure 13-5 shows the All Systems group list, with the check box selected for the wanted IBM i system.



The screenshot shows a web-based interface titled 'Navigate Resources'. Below the title is a breadcrumb 'Groups > All Systems (View Members)'. There is a search bar with the text 'Search the table...' and a 'Search' button. Below this is a table with columns: Select, Name, Type, Access, Problems, Compliance, IP Addresses, and OS Type. The first row is selected, indicated by a checked checkbox in the 'Select' column. The 'OS Type' for this row is 'IBM i', which is highlighted with a red box. A red arrow points to the 'Actions' dropdown menu in the top left of the table area. The table contains four rows of data, all of which are 'Operating System' type. The first row has 'OK' status in all other columns. The second row has a 'Minor' problem. The third and fourth rows have 'OK' status. The bottom of the window shows a pagination bar with 'Page 2 of 2', a '2' in a box, and 'Selected: 1 Total: 19 Filtered: 19'.

Select	Name	Type	Access	Problems	Compliance	IP Addresses	OS Type
☒	x0716p3.rchland.ibm.com	Operating System	OK	OK	OK	9.5.53.254, 127.0.0.1	IBM i
☐	x0716p7.rchland.ibm.com	Operating System	OK	Minor	OK	9.5.52.190, 200.1.1.1	Linux
☐	z1014p26.rchland.ibm.com	Operating System	OK	OK	OK	9.5.103.39	IBM i
☐	z1014p28.rchland.ibm.com	Operating System	Offline	Information	OK	9.5.105.200	IBM i

Figure 13-5 Selecting an IBM i system from the All Systems group list

4. Click **Actions** after the system is selected. A drop-down menu opens.

5. Select the **IBM i Management** item from the Actions drop-down menu. A second drop-down menu opens. On this second menu, click **Work Management**, as shown in Figure 13-6.




Figure 13-6 Selecting Work Management from the Actions drop-down menu

When you click **Work Management**, a menu for basic Work Management tasks opens, as shown in Figure 13-7.




Figure 13-7 IBM i Work Management basic menu that is accessed from Systems Director

The Work Management menu that opens is nearly identical in function to the Work Management menu shown by IBM Systems Director Navigator for i, including the Show All Work Management Tasks button.

6. Click **Show All Work Management Tasks**. An advanced Work Management window opens.

When it first opens, the AJS task menu is minimized. Upon expanding the menu by selecting the icon to which the arrow points, the window that is shown in Figure 13-8 opens.




Figure 13-8 Systems Director - IBM i Work Management all tasks menu

13.2.2 AJS menu

In "IBM Systems Director Navigator for i navigation" on page 490, the Advanced Jobs Scheduler menu was shown in Figure 13-3 on page 492.

The AJS menu contains items of system-wide scope, rather than of job scope. This menu is the menu where the job scheduler functions themselves are configured and maintained. In the following sections, the Advanced Jobs Scheduler menu actions are reviewed.

13.2.3 Start Scheduler action

Completing the following steps starts the AJS monitor:

1. In the Work Management Advanced menu, select the icon to the right of Advanced Jobs Scheduler option (Figure 13-3 on page 492).
2. Select the **Start Scheduler** option. The confirmation window that is shown in Figure 13-9 opens.




Figure 13-9 Start AJS monitor confirmation window

Clicking **Yes** starts the AJS monitor. Clicking **No** returns you to the window shown in Figure 13-3 on page 492.

13.2.4 Stop Scheduler action

Completing the following steps stops the AJS monitor:

1. In the Work Management Advanced menu, select the icon to the right of Advanced Jobs Scheduler option (Figure 13-3 on page 492).
2. Select the **Stop Scheduler** option. The confirmation window that is shown in Figure 13-10 opens.




Figure 13-10 Stop AJS monitor confirmation window

Clicking **Yes** stops the AJS monitor. Clicking **No** returns you to the previous window.

13.2.5 AJS properties menu

The AJS properties are divided into six tabs in the left pane of Figure 13-11. The following sections describe each tab.




Figure 13-11 AJS properties - General tab

General tab

The General tab that is shown in Figure 13-11 on page 497 is used to view and set the general properties of the job scheduler. You can specify the following options:

- ▶ How long to retain activity
- ▶ How long the log remains
- ▶ The period during which scheduled jobs are not allowed to run
- ▶ The working days that jobs are allowed to process
- ▶ The notification command that sends a notification whenever a job completes successfully or fails

Schedules tab

The Schedules tab that is shown in Figure 13-12 provides a display of the existing schedules on your system. The buttons at the right of the page provide the following functions

- ▶ New: Enables you to create a schedule.
- ▶ New Based On: Enables you to create a schedule that is based on an existing one.
- ▶ Remove: Removes an existing schedule.
- ▶ Properties: Enables modification of an existing schedule.




Figure 13-12 AJS Properties Schedules tab

Clicking the **New** button in Figure 13-12 on page 498 opens the New Schedules window that is shown in Figure 13-13.

New Schedule

Name:

Description:

April 2010

Sun	Mon	Tue	Wed	Thu	Fri	Sat
28	29	30	31	1	2	3
4	5	6	7	8	9	10
11	12	13	14	15	16	17
18	19	20	21	22	23	24
25	26	27	28	29	30	1
2	3	4	5	6	7	8

3/30/2010 -- 12/31/2039

Dates to run

Frequency:

☐ On selected dates

☒ Weekly Skip count:

☐ Monthly ☐ Select as working days

☐ Yearly

Details

☒ Sunday

☒ Monday

☒ Tuesday

☒ Wednesday

☒ Thursday

☒ Friday

☒ Saturday

Figure 13-13 AJS Properties - New Schedule window

A schedule is a specification of days on which a job can run. There are no times that are listed. With scheduled jobs and group jobs, you can select a schedule rather than make one when you created the jobs.

Data Libraries tab

The Data Libraries tab displays all job scheduler data libraries. As shown in Figure 13-14, you can add, remove, and modify a data library, and can start and end the job scheduler monitor job for a specific data library.




Figure 13-14 AJS Properties - Data Libraries window

Clicking the **Add** button opens the New Data Library window that is shown in Figure 13-15.




Figure 13-15 AJS New Data Library window

From this window, you can accomplish the following tasks:

- ▶ Create a jobs scheduler data library.
- ▶ Specify a job scheduler monitor.
- ▶ Specify a monitor job's name.
- ▶ Start the job scheduler monitor automatically.

A system can have multiple job scheduler data libraries, and each library can have a monitor job that is running simultaneously. You can switch from one job scheduler that is running a production environment to the other job scheduler library that is running a test environment. This capability is covered in more detail in 13.3.2, “Multiple Scheduling Environments function” on page 528.

Systems & Groups tab

You can use the Systems & Groups tab (Figure 13-16) to add, remove, and change existing IP addresses and other information for IBM i and non- IBM i systems that are used by the AJS.

Select	Name	Description	Type	IP Address	Communication Type
☐	RCHASA20		IBM i	9.10.229.21	Use product TCP/IP server
☐	RCHASA21		IBM i	9.10.229.43	Use product TCP/IP server
☐	RCHASA22		IBM i	9.10.229.70	Use product TCP/IP server
☒	RCHASA23		IBM i	9.10.229.81	Use product TCP/IP server

☒ Automatically start product TCP/IP server

Figure 13-16 Job Scheduler properties - Systems and Groups tab

Users tab

The Users tab (Figure 13-17) enables maintenance of a list of job scheduler users that are associated with a job scheduler data library. The Add button adds users, the Properties button changes the properties of a user, and the Remove button removes a user.




Figure 13-17 AJS Properties - Users tab

Suppose that there is a system with multiple applications and each application's personnel are not allowed to access the scheduled jobs of other applications. The system administrator can set up each application with their own job scheduler data library. The system administrator uses the Users tab function to assign each application's personnel to their own job scheduler data library. Because a user can access only one job scheduler data library, the administrator effectively locks the users to their own application's job scheduler while locking access to the others.

Calendars tab

The buttons in the right pane of Figure 13-18 provide the following functions:

- ▶ New: Create a calendar.
- ▶ New Based On: Create a calendar that is based on an existing one.
- ▶ Remove: Remove calendars.
- ▶ Properties: Maintain calendars.

Scheduling calendars and holiday calendars use different windows because they have different parameters.




Figure 13-18 AJS properties - Calendars tab

The Holiday Calendar Properties window is shown in Figure 13-19.

Holiday Calendar Properties

*Name: OMITCAL

Description:

Reference calendar: Use entry from below

Date: 3/30/10

Alternate day to run

☐ Every year ☒ Special: Do not run ☐ Specific:

Details:

Select	Date	Alternate day to run
☐	March 11, 2007	Do not run

Holiday days

Alternate day to run:

☐ Sunday Do not run

☐ Monday Do not run

☐ Tuesday Do not run

☐ Wednesday Do not run

☐ Thursday Do not run

☐ Friday Do not run

☐ Saturday Do not run

Figure 13-19 Holiday Calendar Properties window

13.2.6 Scheduled Jobs menu

This section reviews the various actions of the Scheduled Jobs menu. This menu is accessed through the Work Management Advanced menu (Figure 13-3 on page 492) by expanding the **Advanced Job Scheduler** option.

The Scheduled Jobs Menu is shown in Figure 13-20. The menu offers the following options:

- ▶ **Scheduled Jobs**

This action displays a table of scheduled jobs that you can use to perform various actions against the jobs. You can customize which columns appear in the table, sort on any column, and filter the table. Usage of the scheduled job table is described in 13.2.8, “Using the Scheduled Jobs table” on page 508.

- ▶ **New Scheduled Job**

You can use this action to create a scheduled job.

- ▶ **Properties**

You can use this action to change scheduled jobs properties for all scheduled jobs. The properties include the following ones:

- Application for Scheduled Job
- Base Periodic frequency on start time
- Reset held jobs
- Reset after change
- Start time of day

- ▶ **Reset Scheduled Jobs**

You can use this action to reset all scheduled jobs.




Figure 13-20 Scheduled Jobs menu

13.2.7 Using the Scheduled Jobs action

In Figure 13-20 on page 505, the Scheduled Jobs menu is highlighted. This section provides an overview of what this menu can do.

When you select the **Scheduled Jobs** action on the Scheduled Job menu, a table of scheduled jobs on the system is displayed, as shown in Figure 13-21.

The screenshot shows a web browser window titled "Scheduled Jobs - Rchasa23.rchland.ibm.com". It features a "Refresh" button, a toolbar with icons for various actions, and a table of scheduled jobs. A context menu is open for the "Bkupalt" job, listing actions such as "Job Dependencies", "Activity", "Activity Log", "New Based On...", "Add Selected Job to Group...", "Permissions", "Hold", "Skip Next Run", "Run Now...", "Reset Scheduled Job", "Rename...", "Delete", and "Properties...".

Select	Job Name	Status	Next Run	Schedule
☐	Autoagelog	Ready to run	5/16/10 12:00 PM	*NUMDAY
☐	Autoezpage	Ready to run	4/26/10 12:00 PM	*NUMDAY
☐	Autogtmon	Ready to run	3/29/10 5:05 PM	Weekly
☐	Autojmsgq	Ready to run	3/29/10 6:20 PM	Weekly
☐	Autonumcpu	Ready to run	4/5/10 9:00 AM	Weekly
☐	Autorsmlog	Ready to run	4/27/10 1:05 PM	Fourth
☐	Bksavsys	Completed		Once
☐	Bkupalt	Ready to run		
☐	Bkupalt			
☐	Bkupd		3/29/10 5:00 PM	Weekly
☐	Bkupw		3/31/10 5:00 PM	Weekly
☐	Bkup1		4/3/10 1:00 PM	Weekly
☐	Fallback			Once
☐	Hthamt		4/5/10 6:10 AM	Weekly
☐	Hthjobds		3/31/10 6:20 AM	Last day of each month

Figure 13-21 Scheduled Jobs list table

You can select multiple jobs by selecting their check boxes. There are options that enable selection of all jobs. For now, you should review what you can do to a specific job.

To illustrate the many tasks you can accomplish in a specific job, see the menu for the BKUPD job that is shown in Figure 13-21. Although it is beyond the scope of this book to describe each of these actions in detail, a brief description of each action follows:

- **Job Dependencies**

This action enables the display and updating of job dependencies, including predecessors and successors, and whether all or just one dependency must be met.

- **Activity**

This action enables listing of the job activity (history) for a specific job scheduler entry. A specific job can be selected from the job activity, which has another menu of actions.

- **Activity Log**

This action enables listing of entries in the job scheduler log for a specific job scheduler entry.

- **New Based on**

This action creates a new scheduled job that is based on the selected job. You can override any of the based on job's parameters.

- **Add Selected Job to Group**

You can use this action to add the selected job to a job group.

► **Permissions**

You can use this action to set the permissions and authorities to the scheduled job, including which users can submit, manage, and set permissions.

► **Hold**

This action allows a scheduled job to be held. This action is shown only for jobs that are not already held.

► **Release**

This action allows a held scheduled job to be released. This action is shown only for jobs that are held.

► **Skip Next Run**

You can use this action to skip (omit) the next scheduled run for the job.

► **Run Now**

You can use this action to run a scheduled job now. You can also use it to run a job at a specified date and time and at a certain number of minutes from now. The action can also override the first command to be run and the last command to be run. You can also check and update job dependencies.

► **Reset Scheduled Job**

This action resets the selected scheduled job.

► **Rename**

This action renames the selected scheduled job.

► **Delete**

This action deletes the selected scheduled job.

► **Properties**

You can use this action to view and change virtually every parameter and value for a scheduled job. The parameters are divided into nine pages of parameters:

– **General**

Specifies the type of scheduled job, job name, job (text) description, and the commands and sequence to be run.

– **Schedule**

Specifies when the job is to be run.

– **Batch Information**

Specifies the information that is required to submit the job, such as run on a single system or system group, basic submission parameters, such as the job queue, job description, job user, and advanced submission parameters, such as job accounting code, job and run priorities, routing data, and command logging.

– **Notification**

Specifies options and parameters for the notification of the job status.

– **Documentation**

Provides a field for the job documentation.

– **Problem Recovery**

Specifies the maximum run time, alternative job in case of failure, action to take when a job cannot be started at the specified time, handling of inquiry messages, and message logging parameters to be used when job fails.

- Communication
Specifies communications parameters, such as local device, local location name, remote device, and remote location name.
- Local Data Area
Specifies what the scheduled job places in the local data area when run.
- Last Run
Shows the last run details.

The Scheduled Jobs properties and display windows are described in 13.2.9, “Adding a new scheduled job” on page 514.

13.2.8 Using the Scheduled Jobs table

This section reviews the Scheduled Jobs table (including the sorting, changing columns, and using filtering tasks).

The Select Action drop-down menu at the top of the table enables options that were not previously described. With no scheduled jobs selected, the menu has the following options, as shown in Figure 13-22:

- New: You can select this option to create a scheduled job.
- Reset Scheduled Jobs: You can reset all scheduled jobs.
- Show find toolbar: This option displays a small toolbar that is used for searching the table.
- Table Actions: This option opens another menu with more actions for manipulating the table.




Figure 13-22 Scheduled Jobs Table Action drop-down menu

If you select one or more jobs in the table, the Select Action menu adds job-specific actions to the list options. These options are equivalent to the actions listed in the scheduled job's menu that is shown in Figure 13-21 on page 506.

If the Table Actions menu item is selected, a second drop-down menu is displayed, as shown in Figure 13-23. The actions of this menu are the following ones:

- ▶ **Select All**
Selects all scheduled jobs.
- ▶ **Deselect All**
Clears all scheduled jobs.
- ▶ **Show Filter Row**
Shows a row at the top of the table that is used to configure and select criteria to determine which scheduled jobs are shown.
- ▶ **Clear All Filters**
Removes all filters.
- ▶ **Collapse Table**
Does not show any table data. This option can be helpful in speeding changes to the table, as the data does not have to be refreshed on the window between multiple changes.
- ▶ **Configure Columns**
This action is used to change the columns in the table.
- ▶ **Restore Defaults**
Removes the changes that are made to the table by sorting and filtering column changes.




Figure 13-23 Scheduled Jobs - Table Actions

Resizing columns


Columns can be resized by dragging the column separator in the table column description (header) line.

Sorting the table

Any of the columns that have a ^ (circumflex) symbol are sortable. The first sort is in ascending order, but the sort can be reversed by selecting the sort symbol a second time. After the sorting, the ^ (circumflex) symbol is replaced by a triangle symbol that indicates the sort direction.

In Figure 13-24, the Next Run column was sorted in ascending order, then resorted in descending order. The triangle in the Next Run header is pointing down, indicating a sort in descending order.

To return the data back to its original sort order, click the **Clear All Sorts** action from the Table Actions menu.



Select	Job Name ^	Status ^	Next Run ▼	Schedule
☐	Autoezpage	Ready to run	5/8/10 12:00 PM	*NUMDAY
☐	Autoagelog	Ready to run	5/21/10 12:00 PM	*NUMDAY
☐	Jbadasd	Ready to run	4/6/10 5:30 AM	*NUMDAY
☐	Autonumcpu	Ready to run	4/5/10 9:00 AM	Weekly
☐	Pgmagedev	Ready to run	4/5/10 7:00 AM	Weekly
☐	Jbacpu	Ready to run	4/5/10 6:15 AM	Weekly
☐	Hthamt	Ready to run	4/5/10 6:10 AM	Weekly
☐	Seclabdwk3	Ready to run	4/5/10 6:00 AM	Weekly
☐	Seclabdwk4	Ready to run	4/5/10 6:00 AM	Weekly
☐	Seclabdwk	Ready to run	4/5/10 6:00 AM	Weekly
☐	Seclabdwk1	Ready to run	4/5/10 6:00 AM	Weekly
☐	Hthusracc	Ready to run	4/5/10 6:00 AM	Weekly
☐	Seclabdwk2	Ready to run	4/5/10 6:00 AM	Weekly
☐	Newjob	Held from running	4/4/10 6:00 AM	Weekly
☐	Pspprtjobd	Ready to run	4/4/10 6:00 AM	Weekly

Figure 13-24 Descending sort of scheduled job on Next Run column

Configuring table columns

The table in its initial format includes all columns. However, you can remove columns and change the order of the columns by using the **Configure Columns** option of the Table actions menu. When **Configure Columns** is clicked, a light shaded column section area is displayed, as shown in Figure 13-25.




Figure 13-25 Table column selection and reordering

The vertical bar points to a window that contains a list of columns with a check box in front of each one for selection purposes. The window has a slide, which is used to navigate the list as pointed to by the lower arrow.

If you want a selected column to not appear in the table, select the check box.

The two scroll buttons pointed to by the upper arrow are used to move a selected column up and down (where up appears first and down appears last in a left to right order). To select the column for moving, click its description.

Clicking **OK** puts the column changes into effect and the table is displayed again.

To return the table back to its original format, close the page and reopen it.

Using Filters

To use filtering, click **Show Filter Row** from the Table Actions menu. A light blue filter row is added to the top of the scheduled job table, as shown in Figure 13-26. The arrow points to the filter row.




Figure 13-26 Scheduled Jobs list filtering

To add a filter to a column, click **Filter** in the filter row for that column. Figure 13-26 shows what is displayed when the Filter tag for the Status column is selected.

The red box shows the filtering data. The Condition value has the following options:

- ▶ Contains (as shown)
- ▶ Starts with
- ▶ Ends with

The Text value is the text that the filter is supposed to match.

In Figure 13-26, the filter shows scheduled jobs with status values containing Held (all held scheduled jobs).

To disable the filter, select the filter's check box, as shown by the top arrow in Figure 13-27. The filter remains, but it does not filter. This option is used when you want to retain a filter for future use.




Figure 13-27 Scheduled Jobs table filter options

There are two other actions that can be selected from the Table Actions menu that impact the filters shown in Figure 13-27:

- Hide Filter Row
This option hides the filter row, but leaves the filters in their current state.
- Clear All filters
This option clears all filters and deletes them.

The other tables in the job scheduler functions, such as Job Groups, Scheduled Job Activity, and Activity Log, use similar operations.

13.2.9 Adding a new scheduled job

This section describes the functions and windows that are used to add a new scheduled job.

To add a new scheduled job, complete the following steps:

1. Click the **New Scheduled Job** action from the Scheduled Job menu, as shown in Figure 13-28.




Figure 13-28 Selecting the New Scheduled Job action

The General window of the New Scheduled Job menu (Figure 13-29) opens.

Figure 13-29 New Scheduled Job General window

In Figure 13-29, the job name is set to NEWJOB. The job is a scheduled job (not an alternative job or non-scheduled job), and the job (text) description is entered.

2. Add one or more commands by clicking **Add** in the Commands section. The Command Properties window (Figure 13-30) opens.

Figure 13-30 Command Properties window

3. Specify the command. In this example, RTVDSKINF is entered into the Command box in Figure 13-30 on page 515.

In our example, we have no messages to monitor, so we leave those fields blank.

4. Check the **RTVDSKINF** command's parameters by clicking **Prompt**.

There is one field, ASP Device, with a default value of *SYSBAS. This window shows only the basic parameters. See Figure 13-31.




Figure 13-31 Basic parameters window

5. Check the advanced parameters by clicking **Advanced**. The window refreshes with another **RTVDSKINF** command parameter displayed with a default value, as shown in Figure 13-32.




Figure 13-32 Advanced parameters window

6. Click **View**. The View menu opens, as shown in Figure 13-33.




Figure 13-33 View menu

7. Click the **Keywords** option. If you want to see the keywords rather than the text descriptions for the fields, or see all parameters, or see the command string, select those options.

The window refreshes and the window that is shown in Figure 13-34 opens.




Figure 13-34 Keywords option window

The last thing of note about command prompting is the command help. Figure 13-35 shows the Help menu.




Figure 13-35 Help menu

Click **Help** and the command help window opens, as shown in Figure 13-36.




Figure 13-36 RTVDSKINF Command help

Now that the correct command is specified, click **OK** to close the Help window. Click **OK** on the **RTVDSKINF** command prompt display to accept the **RTVDSKINF** command. Click **OK** on the Command Properties window to accept the command properties for **RTVDSKINF** usage. The General window of the New Scheduled Job function opens.

In this example, two more commands were added to the scheduled job (not shown), but the second command is out-of-order; it must be the first command to be run. Click the out-of-order second command, and then click the **Move Up** button to move the second command to the first run position, as shown in Figure 13-37.

The screenshot shows the 'New Scheduled Job' dialog box with the following details:

- Job name:** NEWJOB
- Type of job:** Scheduled (selected)
- Description:** This is a new sample job for the IBM i Redbook
- Commands:**

Select	Enabled	Commands
☐	☒	RTVDSKINF
☒	☒	CHGJOB OUTQ(QUSRSYS/DISKINF)
☐	☒	PRTDSKINF RPTTYPE(*OBJ) OBJ(*ALL) MINSIZE(20000)
- Buttons:** Add, Add Based On, Remove, Move Up, Move Down, Properties
- Options:**
 - ☐ Run as remote command
 - Application: [dropdown]
 - Report distribution list: [dropdown]
 - ☐ Use start time of day
- Footer:** OK, Cancel

Figure 13-37 Moving a command to change the run order

When the command order is correct, click **OK**.

New scheduled job: Schedule window

The Schedule window is shown in Figure 13-38. It is used to specify when and how often the scheduled job runs. There are many scheduling options that are beyond the scope of this book.

Scheduled Job Properties

General

Schedule

Batch Information

Notification

Documentation

Problem Recovery

Communication

Local Data Area

March 2010 Refresh

Sun Mon Tue Wed Thu Fri Sat

28 1 2 3 4 5 6

7 8 9 10 11 12 13

14 15 16 17 18 19 20

21 22 23 24 25 26 27

28 29 30 31 1 2 3

4 5 6 7 8 9 10

3/29/2010 -- 12/31/2039

Additional Calendars...

Schedule name: [Go]

Times to run

☒ On specific times Example: 12:30 PM Add Remove

6:00 AM

☐ Periodically Frequency: 30 hours minutes

Start time: 12:00:00 AM Example: 12:30:00 PM

End time: 11:59:00 PM Example: 12:30:00 PM

Dates to run

Frequency:

☐ On selected dates

☒ Weekly Skip count: 0

☐ Monthly Select as working days

☐ Yearly

Details

☒ Sunday

☐ Monday

☐ Tuesday

☐ Wednesday

☐ Thursday

☐ Friday

☐ Saturday

OK Cancel

Figure 13-38 New Scheduled Job Schedule window

In the upper left corner of the window, there is a calendar, with which you can select a date range, multiple dates, or one date, as shown in Figure 13-39.

Scheduled Job Properties

General

Schedule

Batch Information

Notification

Documentation

Problem Recovery

Communication

Local Data Area

March 2010 Refresh

Sun Mon Tue Wed Thu Fri Sat

28 1 2 3 4 5 6

7 8 9 10 11 12 13

14 15 16 17 18 19 20

21 22 23 24 25 26 27

28 29 30 31 1 2 3

4 5 6 7 8 9 10

3/29/2010 -- 12/31/2039

Additional Calendars...

Figure 13-39 Calendar section of the Schedule Job Schedule window

In the upper right corner of the window, which is shown in Figure 13-40, you can specify to repeat a run periodically during a time range or apply a Schedule name, which uses a previously created schedule.

Schedule name: ▼ Go

Times to run

☒ On specific times Example: 12:30 PM Add

6:00 AM Remove

☐ Periodically Frequency: 30 hours ☐ minutes ☒

Start time: 12:00:00 AM Example: 12:30:00 PM

End time: 11:59:00 PM Example: 12:30:00 PM

Figure 13-40 Times to run section of the Schedule Job Schedule window

The “Dates to run” section, which is shown in Figure 13-41, has a set of radio buttons (circled) to specify the frequency at which the job is to run. The Details section varies depending on the selected Frequency button. In this example, we selected options to run the scheduled job at 6:00 a.m., to run weekly, but on Sundays only.

Dates to run

Frequency:

☐ On selected dates

☒ Weekly Skip count: 0

☐ Monthly ☐ Select as working days

☐ Yearly

Details

☒ Sunday

☐ Monday

☐ Tuesday

☐ Wednesday

☐ Thursday

☐ Friday

☐ Saturday

Figure 13-41 Dates to run section of the Schedule Job Schedule window

New Scheduled Job: Batch Information window

Typically, when a job is submitted, there is a host of parameters that must be specified for the job to run. Use the Batch Information window that is shown in Figure 13-42 to set these parameters for the new scheduled job.

Scheduled Job Properties

General
Schedule
***Batch Information**
Notification
Documentation
Problem Recovery
Communication
Local Data Area

System to run job

☒ Single system Use local system
☐ System group

Submission Information

Job description: Use user's job description Output queue: Use job description value
Library: Use job control value
Use user's job description
Use entry from below Library: Use entry from below
Job queue: Use job description value User: MS773
Library: Use entry from below Current library: Use job control value
Library list: Use job description value

Advanced Batch Information

Accounting code: Use job control value Print text: Use system value
*Run priority: No change required Routing data: Use job description value
*Job priority: Use job description value Log CL commands: Yes
*Output priority: Use job description value Hold on job queue: No
*Printer device: Use job description value Initial ASP group: Use job description value

Figure 13-42 New Scheduled Job Batch Information window

In the Batch Information window, the fields have drop-down options and, in most cases, include a blank field to enter the correct values manually. Parameters such as “Job description” and “User” have Browse buttons that open a list window from which a selection can be made.

After you set these parameters and clicking **OK**, the user can open the Notification page. In many situations, the scheduled job runs as expected.

New Scheduled Job: Notification window

You can use the Notification window that is shown in Figure 13-43 to set up various methods of notification for various job conditions.

Figure 13-43 New Scheduled Job Notification window

There are five notification options in the Notification window. Each option has a check box for enabling or disabling the notification.

- Send completion message

For this function, add the QSYSOPR value by selecting the **Use entry from below** drop-down option and entering the value into the blank field. You can select the message queue by clicking **Browse** and selecting a queue from a list of queues.

- Notify if job completes successfully

In this example, select **Use job control value** from the drop-down menu and **Use completion message** from the Message drop-down menu.

- Notify if Job Fails

In this example, select the Recipient value **Use job control value** from the drop-down menu, and enter the message manually.

- Notify if job does not start within limit

In this example, the default values are used.

- Notify with error messages

In this example, you choose to send the job log.

The drop-down menus prevent you from entering incorrect options.

New Scheduled Job: Documentation window

The Documentation window, which is shown in Figure 13-44, has a large text field that can be used for job documentation (such as who owns the job, what the job actually does, and what an operator does when it fails).




Figure 13-44 New Scheduled Job Documentation window

New Scheduled Job: Problem Recovery window

The Problem Recovery window is shown in Figure 13-45. On this window, the user specifies problem handling parameters and the message logging parameters for the scheduled job.

The screenshot shows the 'Scheduled Job Properties' window with the 'Problem Recovery' tab selected. The left sidebar contains tabs for General, Schedule, Batch Information, Notification, Documentation, Problem Recovery (selected), Communication, and Local Data Area. The main area contains the following settings:

- Maximum run time: No maximum (dropdown) with a Go button.
- Maximum run time action: End the job (dropdown).
- ☐ Use an alternate job if job fails.
- Job name: [text box] with a Browse... button.
- Group: [text box].
- Sequence: [text box] 0 0 - 99.
- Action if job cannot start when scheduled: Do not start (dropdown).
- Reply to inquiry message: Use system reply list (dropdown).
- Message logging** section:
 - Level: Use job description value (dropdown).
 - Severity: Use job description value (dropdown).
 - Text: A dropdown menu is open showing options: Use job description value, Use job control value, Use job description value, Use message text, Use message text and message help (highlighted), and None, if completes normally.

OK and Cancel buttons are at the bottom right.

Figure 13-45 New Scheduled Job Problem Recovery window

New Scheduled Job: Communication window

Figure 13-46 shows the New Scheduled Job Communications window.

The screenshot shows the 'Scheduled Job Properties' window with the 'Communication' tab selected. The left sidebar contains tabs for General, Schedule, Batch Information, Notification, Documentation, Problem Recovery, Communication (selected), and Local Data Area. The main area contains the following settings:

- Local device: Use communications device (dropdown).
- Location name: Use network attributes (dropdown).
- Remote device: Use job control value (dropdown).
- Remote location name: Use job control value (dropdown).

OK and Cancel buttons are at the bottom right.

Figure 13-46 New Scheduled Job Communication window

New Scheduled Job: Local Data Area window

Figure 13-47 shows the New Scheduled Job Local Data Area window.




Figure 13-47 New Scheduled Job Local Data Area window

Similar windows are shown when you change an existing scheduled job.

13.3 Other AJS enhancements

This section covers other IBM Advanced Job Scheduler for i enhancements.

13.3.1 Notification enhancements

Notification capabilities are enhanced by the ability to send a spooled file to an IFS path name with the option to replace the file if it exists. You can use replacement variables to make the file and the path unique. The path name is specified in a recipient file, as shown in Figure 13-48.

The window in the background shows the New Recipient menu and its General window, as accessed through IBM Systems Director Navigator for i. The window in the foreground shows the New Recipient Path window and the path parameter. The path value shows the use of substitution variables. This function is also available through IBM i Access client, but is not available through the character-based interface.




Figure 13-48 The Path window of the Create Recipient menu

13.3.2 Multiple Scheduling Environments function

With the Multiple Scheduling Environments (MSE) function, you can set up any number of scheduling environments that run at the same time. The environments use different job scheduler data libraries. This function is helpful for having a test and production set of scheduled jobs or when you replicate live data from a production system to a backup system.

In the character-based interface, use the Change Data Library using JS (**CHGDLJS**) command to start more environments, as shown in Figure 13-49.

Change Data Library using JS (CHGDLJS)

Type choices, press Enter.

Data library	> <u>AJSTEST</u>	Name
Monitor	> <u>*YES</u>	*SAME, *YES, *NO
Monitor job name	> <u>TESTJSMON</u>	Name, *SAME
Automatically start monitor . .	> <u>*YES</u>	*SAME, *YES, *NO
Text	> <u>'Test Job Scheduler Environment.'</u>	

Bottom

F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

Figure 13-49 Starting an addition job schedule environment using the CHGDLJS command

You can also specify which users can use which job scheduler environment by running the Set Data Library (SETDLJS) command, as shown in Figure 13-50.




Figure 13-50 Setting your MSRJJ's access to a job scheduler data library

Consideration: A user profile can be associated with only one job scheduler data library at a time. If you are not careful, you can be prevented from accessing the default QUSRIJS data library.

The equivalent multiple scheduling environments function is also available through the System i Access (GUI) and IBM Systems Director Navigator for i (web) interfaces.

13.3.3 Scheduling and running jobs

This section describes the enhancements for scheduling and running jobs.

Predefined schedules

Jobs in a group with a sequence number greater than 1 (jobs other than the first one in a group) can now use a predefined schedule. This option is helpful when you have a group of jobs where you want a subset of the jobs in the group to run on another schedule than the rest. For example, if you have a group of jobs that run on a daily schedule, but one in the group must run on Fridays only, a schedule can be used for the Friday job that overrules the daily schedule of the group. The schedule can also be a holiday calendar. This function adds flexibility for configuring which jobs in a group run on different days and dates without breaking up the group.

Start Group job enhancements

Starting group jobs was enhanced to add new “Based on” parameters to enable the user to specify an override capability of which jobs are run in a group.

The Start Group using JS (**STRGRPJS**) command was changed to add “Based on” parameters, as shown in Figure 13-51.

Start Group using JS (STRGRPJS)

Type choices, press Enter.

Job entry:

Job	_____	Name
Group	_____	Name
Group sequence	_____	1-99
Submit time	*SCHED	*SCHED, *IMMED
Based on day of week	*CURRENT	*CURRENT, *DATE, *MON...
Based on date	*CURRENT	Date, *CURRENT
Parameters: _____		
Parameter name	*NONE	Name, *NONE
Parameter data	_____	

+ for more values _

Bottom

F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

Figure 13-51 New Based on parameters of STRGRPJS command

The “Based on” parameters enable you to use the day of week, the date, or both to override which jobs are run in the group job.

If *FRI (Friday) is specified for the “Based on day of week” parameter, the group jobs, other than the first one, run as though the day were a Friday. Jobs in the group that list *FRI as a day to run and jobs in the group that use a predefined schedule specifying to run on *FRI run even if the **STRGRPJS** command was issued on a Wednesday.

The “Based on date” parameter works similarly. If the specified date is December 1, 2009, the job scheduler determines which of the jobs in the group can run on that date and runs them when the **STRGRPJS** command is run.

The new “Based on” function is also found in the System i Access (GUI) and IBM Systems Director Navigator for i (web) interfaces.

New time offset added to the Submit Scheduled Job function

When you submit a scheduled job manually, you can now specify a submit time offset. Suppose that it is 6 p.m., and you want to submit a scheduled job to run at 9 p.m., but do not want to change the job's schedule, as this job is a one-time need. You can specify a time offset of 180 minutes, causing the job to run at 9 p.m.

The Submit Job using Job Scheduler (**SBMJOBJS**) command in Figure 13-52 has a new "Submit time offset" parameter.

Submit Job using Job Scheduler (SBMJOBJS)

Type choices, press Enter.

Job entry:

Job	> <u>PSPQDFTOWN</u>	Name
Group	<u>*NONE</u>	Name, *NONE
Group sequence	<u>*NONE</u>	1-99, *NONE
Submit time	<u>*SCHED</u>	Time, *SCHED, *IMMED, *OFFSET
Submit time offset	<u>180</u>	Minutes, 1-720
Submit date	<u>*CURRENT</u>	Date, *CURRENT
Starting sequence	<u>*FIRST</u>	Number, *FIRST
Ending sequence	<u>*LAST</u>	Number, *LAST
Check for dependencies	<u>*YES</u>	*YES, *NO
Update dependencies	<u>*NO</u>	*NO, *YES, *GRP

More...

F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

Figure 13-52 Submit Job using Job Scheduler with new Submit time offset parameter

The user can specify the amount of time in minutes 1 - 720 to determine the submit time.

A similar function is available as the “Minutes from new” parameter through IBM Systems Director Navigator for i, as shown in Figure 13-53.




Figure 13-53 Run scheduled job minutes from now

This new function is also available through the System i Access (GUI).

Job dependency enhancements

Jobs can now depend on the contents of a data area or a value that is returned from a program. When you specify a resource dependency for an object type data area, you can specify a value that must exist in the data area before the associated job runs. A starting position and a length can be used in case the value to be checked is smaller than the data in the object. The value can be up to 50 alphanumeric characters. This new function is found in the System i Access (GUI), IBM Systems Director Navigator for i (web) interfaces, and character-based interfaces.

13.4 References

For more information about the topics that are covered in this chapter, see the following resources:

- ▶ IBM AJS for i What's New website, found at:
http://www-03.ibm.com/systems/i/software/jscheduler/v7r1_feat.html
- ▶ IBM AJS for i PDF, found at:
<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzaks/rzaksads.pdf>
- ▶ Connecting to IBM i IBM Systems Director Navigator for i PDF, found at:
<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzatg/rzatgdirect.pdf>



Web serving

This chapter provides an overview of the changes that are made to the IBM HTTP Server for i (5770-DG1).

This chapter describes the following topics:

- ▶ Triggered cache manager removed
- ▶ Web server search engine and web crawler support removed
- ▶ Plug-ins and LoadModule directives
- ▶ HTTP Server for i with HA IPv6 support

14.1 Triggered cache manager removed

Triggered cache manager support (option 1 of DG1) was removed from the licensed products list in IBM i 7.1.

14.2 Web server search engine and web crawler support removed

The web server search engine and crawler support was removed in IBM i 7.1.

The following commands were removed:

- ▶ Configure HTTP Search (**CFGHTTPSCH**)
- ▶ Start HTTP Crawler (**STRHTTPCRL**)
- ▶ End HTTP Crawler (**ENDHTTPCRL**)
- ▶ Resume HTTP Crawler (**RSMHTTPCRL**)

Any existing CL programs that use these commands might need to be modified. Any web interfaces that use the search support must be modified so it no longer uses this removed option.

14.3 Plug-ins and LoadModule directives

These directives apply to all IBM HTTP servers that are associated with WebSphere Application Server Version 6.1 (5733-W61) or Version 7 (5733-W70) when you upgrade to IBM i 7.1.

All Application Server service programs that implement the HTTP plug-ins might need to be updated before you start the HTTP servers on IBM i 7.1.

For more information related to WebSphere Application Server in IBM i 7.1, see 15.1, “IBM Integrated Web Services for i” on page 542.

In IBM i 7.1, the LoadModule directives that are used by external HTTP servers that are associated with WebSphere Application Server Version 6.1 and WebSphere Application Server Version 7 changed.

For HTTP servers that are associated with WebSphere Application Server Version 6.1 or Version 7, the LoadModule directive must be changed to match the following format:

```
LoadModule was_ap20_module /QSYS.LIB/<product_library>.LIB/QSVTAP22.SRVPGM
```

Where <product_library> is the product library for the Application Server installation.

The product library for each WebSphere Application Server installation on your system contains the program and service program objects for the installed product:

- ▶ For WebSphere Application Server V6.1:
 - The product library name for Version 6.1 is QWAS61x (where x is A, B, C, and so on).
 - The product library for the first WebSphere Application Server V6.1 product that is installed on the system is QWAS61A.

- For WebSphere Application Server V7:
 - The product library name for Version 7 is QWAS7x (where x is A, B, C, and so on).
 - The product library for the first WebSphere Application Server V7 product that is installed on the system is QWAS7A.

The product library can be determined by examining the `<profile_root>/properties/.instance.properties` file, where `<profile_root>` is the root directory of the associated WebSphere Application Server profile.

In the `.instance.properties` file, `<product_library>` is specified by the `was.install.library` property.

The `LoadModule` directive can be modified from the IBM Web Administration for i window.

We now describe an example where we assume that we created a WebSphere Application Server 7 server instance WAS70TOMV on IBM i 6.1:

1. Start the IBM Systems Director Navigator for i and click the **IBM i Tasks Page** link on the Welcome window, as shown in Figure 14-1.




Figure 14-1 IBM Systems Director Navigator for i - Welcome window

2. The window that is shown in Figure 14-2 opens. Click **IBM Web Administration for i**.




Figure 14-2 IBM Web Administration for i

The IBM Web Administration for i window opens, as shown in Figure 14-3.




Figure 14-3 IBM Web Administration for i - Manage HTTP Servers

Complete the following steps:

1. From the Manage tab, click **HTTP Servers**.
2. Select your server from the Server drop-down menu.

3. Click **Edit Configuration file** in the navigation frame.

Figure 14-4 shows the configuration file of the corresponding HTTP server instance WEB_TOMV that was configured.

```
Display Configuration File
HTTP server:   WEB_TOMV
Selected file:  /www/web_tomv/conf/httpd.conf

1  WebSpherePluginConfig
/QIBM/UserData/WebSphere/AppServer/V7/Base/profiles/WAS70TOMV/config/cells/M
ERCURE_WAS70TOMV/nodes/MERCURE.BE.IBM.COM-node/servers/IHS_WEB_TOMV/plugin-c
fg.xml
2  LoadModule was_ap20_module /QSYS.LIB/QHTTPSVR.LIB/QSVT2070.SRVPGM
3  # HTTP server (powered by Apache) configuration
4  DocumentRoot /www/web_tomv/htdocs
5  ServerRoot /www/web_tomv
6  Options -ExecCGI -FollowSymLinks -SymLinksIfOwnerMatch -Includes
-IncludesNoExec -Indexes -MultiViews
7  Listen *:10000
8  LogFormat "%h %T %l %u %t \"%r\" %>s %b \"%{Referer}i\"
\"%{User-Agent}i\"" combined
9  LogFormat "%{Cookie}n \"%r\" %t" cookie
10 LogFormat "%{User-agent}i" agent
11 LogFormat "%{Referer}i -> %U" referer
12 LogFormat "%h %l %u %t \"%r\" %>s %b" common
13 CustomLog logs/access_log combined
14 SetEnvIf "User-Agent" "Mozilla/2" nokeepalive
15 SetEnvIf "User-Agent" "JDK/1\0" force-response-1.0
16 SetEnvIf "User-Agent" "Java/1\0" force-response-1.0
17 SetEnvIf "User-Agent" "RealPlayer 4\0" force-response-1.0
18 SetEnvIf "User-Agent" "MSIE 4\0b2;" nokeepalive
19 SetEnvIf "User-Agent" "MSIE 4\0b2;" force-response-1.0
20 <Directory />
21     Order Deny,Allow
22     Deny From all
23 </Directory>
24 <Directory /www/web_tomv/htdocs>
25     Order Allow,Deny
26     Allow From all
27 </Directory>
```

Figure 14-4 WEB_TOMV HTTP server configuration file

4. Look for the library name that is associated with the WebSphere Application Server instance within the .instance.properties file at following path within the IFS, \QIBM\UserData\WebSphere\AppServer\V7\Base\profiles\WAS70TOMV\properties, as shown in Figure 14-5.

```
instance.name=WAS70TOMV
instance.type=appserver
instance.creating.product=BASE
instance.use.j9=false
instance.j9.path=$(j9path)
instance.j9.version=classic
default.server.name=WAS70TOMV
was.install.library=QWAS7A
was.install.path=/QIBM/ProdData/WebSphere/AppServer/V7/Base
```

Figure 14-5 was.install.library property within the.instance.properties file

5. Update the LoadModule directive by changing QHTTPSVR to QWAS7A and change QSVT2070 to QSVTAP22, as shown in Figure 14-6.

```
WebSpherePluginConfig
/QIBM/UserData/WebSphere/AppServer/V7/Base/profiles/WAS70TOMV/config/cells/M
ERCURE_WAS70TOMV/nodes/MERCURE.BE.IBM.COM-node/servers/IHS_WEB_TOMV/plugin-c
fg.xml
LoadModule was_ap20_module /QSYS.LIB/QWAS7A.LIB/QSVTAP22.SRVPGM
# HTTP server (powered by Apache) configuration
DocumentRoot /www/web_tomv/htdocs
ServerRoot /www/web_tomv
Options -ExecCGI -FollowSymLinks -SymLinksIfOwnerMatch -Includes
-IncludesNoExec -Indexes -MultiViews
Listen *:10000
LogFormat "%h %T %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%{User-Agent}i\""
combined
LogFormat "%{Cookie}n \"%r\" %t" cookie
LogFormat "%{User-agent}i" agent
LogFormat "%{Referer}i -> %U" referer
LogFormat "%h %l %u %t \"%r\" %>s %b" common
CustomLog logs/access_log combined
SetEnvIf "User-Agent" "Mozilla/2" nokeepalive
SetEnvIf "User-Agent" "JDK/1\0" force-response-1.0
SetEnvIf "User-Agent" "Java/1\0" force-response-1.0
SetEnvIf "User-Agent" "RealPlayer 4\0" force-response-1.0
SetEnvIf "User-Agent" "MSIE 4\0b2;" nokeepalive
SetEnvIf "User-Agent" "MSIE 4\0b2;" force-response-1.0
<Directory />
    Order Deny,Allow
    Deny From all
</Directory>
<Directory /www/web_tomv/htdocs>
    Order Allow,Deny
    Allow From all
</Directory>
```

Figure 14-6 LoadModule Directive

6. On the IBM Web Administration for i window, which is shown in Figure 14-7, click **Apply** to change the configuration file.




Figure 14-7 IBM Web Administration for i - Applying changes to the configuration file

7. You now get a message that the configuration was successfully changed. Do not restart the server now. Click **OK** and from now on you can start the upgrade from IBM i 6.1 towards IBM i 7.1. After the upgrade, the HTTP server can be successfully started on IBM i 7.1.

14.4 HTTP Server for i with HA IPv6 support

Highly available HTTP servers take advantage of IBM i clustering technology and make it possible to build a highly available website. High availability improves the availability of business-critical web applications that are built with static Hypertext Markup Language (HTML) pages or Common Gateway Interface (CGI) programs.

Before IBM i 7.1, only IPv4 addresses were supported by PowerHA. Now PowerHA for i on IBM i 7.1 fully supports IPv6 address (including all HA-related APIs, commands, and GUIs), so HA IPv6 support was added to HTTP Server for i on IBM i 7.1. You can use IPv6 addresses to configure all web servers in the cluster and access your web applications that are running in the highly available web server environment.

Network requirement: The IPv6 network among the clients and the cluster must be set up already and available for access. Every client should be able to **ping** through the clustered IPv6 address.

HTTP Server for i with HA IPv6 support has the following requirements:

- ▶ Software requirements:
 - 5770SS1 40 HA Switchable Resources
 - 5770DG1 *BASE HTTP Server for i
 - 5770HAS *BASE IBM PowerHA for i Standard Edition
 - 5770HAS 1 PowerHA for i Enterprise Edition
- ▶ Required PTFs: Current Group PTF for 5770DG1 SF99368 (minimum level 10)

14.5 IBM HTTP Server plug-in module

The IBM HTTP Server plug-in module was enhanced to handle files larger than 2 GB. For clients that still require functionality from plug-ins that were added before the December PTF groups, they might need to recompile those plug-ins to enable the new support. For clients using IBM WebSphere Application Server on IBM i, installing the latest fix packs provides the corrected HTTP server plug-in. For more information, see the IBM i Technology Updates wiki on IBM developerWorks:

<http://www.ibm.com/developerworks/ibmi/techupdates/web>



Web enablers

This chapter describes the following topics:

- ▶ IBM Integrated Web Services for i
- ▶ Changes in IBM WebSphere Application Server
- ▶ IBM Web Administration for i

15.1 IBM Integrated Web Services for i

Integrated Web Services for i enables integrated Language Environment (ILE) applications to work in the web services and service-oriented architecture (SOA) arena with little effort, knowledge, and resources on your part. The convergence of web service and IBM i technologies can help you liberate these core business assets by making it easier to enrich, modernize, extend, and reuse them well beyond their original scope of design.

The IBM i operating system integrates software technologies that support the externalization of an ILE program object as a web service and the consumption of a web service by an ILE program object. These technologies are the integrated web services server and the integrated web services client for ILE.

The following sections describe enhancements that were made to the integrated web services support on IBM i.

For the latest news about Integrated Web Services for i support, go to:

<http://www.ibm.com/Systems/i/software/iws/>

15.1.1 Integrated Web Services Client for ILE

The Integrated Web Services Client for ILE provides a set of libraries and Java tools that enable you to build web service client applications from existing Web Service Description Language (WSDL) files. You use the tools to generate code that can be called from ILE applications to use a web service.

In the past, the tools supported only the generation of C and C++ code, and you had to manually generate the program or service program that contained the generated code. The tools were enhanced to generate RPG code in addition to automatically creating the service program that contains the generated code.

To use this support, ensure that you have the latest PTFs or the latest replacement PTFs. At the time of publication, the latest PTFs are SI44364 and SI44363.

More information can be found in the *Web Services Client for ILE Programming Guide* at:

<http://www.ibm.com/Systems/i/software/iws/documentation.html>

15.2 Changes in IBM WebSphere Application Server

The following sections describe enhancements that were made to WebSphere Application Server.

15.2.1 Packaging for WebSphere Application Server

In IBM i 6.1, IBM Web Enablement for IBM i includes the following products:

- ▶ IBM WebSphere Application Server - Express V6.0 for OS/400
- ▶ IBM WebSphere Application Server - Express V6.1 for i5/OS
- ▶ IBM WebSphere Application Server - Express V7.0 for IBM i
- ▶ IBM WebSphere Application Server - Express V8.0 for IBM i

In IBM i 7.1, IBM Web Enablement for IBM i includes the following products:

- ▶ IBM WebSphere Application Server - Express V6.1
- ▶ IBM WebSphere Application Server - Express V7.0
- ▶ BM WebSphere Application Server - Express V8.0

IBM WebSphere Application Server - Express V6.0 is not included, as it is not supported and is not operational on IBM i 7.1.

15.2.2 Java for IBM i 7.1

Java Developer Kit 1.4, 5.0 and 6 (5761JV1 options 6, 7, and 10), which are referred to as classic Java, are no longer supported in IBM i 7.1 and were replaced by IBM Technology for Java. If your applications are still using classic Java, you must upgrade to IBM Technology for Java, but before you do, be aware of the following information (see also Table 15-1):

- ▶ Classic Java virtual machine (JVM) is a 64-bit virtual machine and migrating to the 32-bit IBM Technology for Java (default JVM) reduces the Java object heap to no larger than 3 GB, which is approximately 1000 threads. If you require more than 1000 threads or a Java object heap larger than 3 GB, use the 64-bit version of the IBM Technology for Java.
- ▶ If you have ILE programs that use Java Native Interface functions, you must compile these programs with teraspace storage enabled.
- ▶ Adopted authority for Java program is not supported by IBM Technology for Java virtual machine.
- ▶ PASE for i now enforces stack execution disable protection.
- ▶ You must install the latest Group PTF for Java SF99572.

Table 15-1 Classic Java levels and the suggested IBM Technology for Java replacement

Current product classic Java	Option	Replacements of IBM Technology for Java	Option
Java Developer Kit 1.4 - 5761JV1	6	Java SE 6 32 bit - 5761JV1 ^a	11
		Java SE 6 64 bit - 5761JV1 ^a	12
		J2SE 5.0 32 bit - 5761JV1	8
		J2SE 5.0 64 bit - 5761JV1	9
		J2SE 1.4 64 bit - 5761JV1	13
Java Developer Kit 5.0 - 5761JV1	7	Java SE 6 32 bit - 5761JV1 ^a	11
		Java SE 6 64 bit - 5761JV1 ^a	12
		J2SE 5.0 32 bit - 5761JV1	8
		J2SE 5.0 64 bit - 5761JV1	9
Java Developer Kit 6 - 5761JV1	10	Java SE 6 64 bit - 5761JV1	11
		Java SE 6 64 bit - 5761JV1	12

a. Use Java SE 6 when migrating from Java Developer Kit 1.4 or 5.0.

For more information about what's new in Java for IBM i 7.1, see the Information Center for the product at:

<http://publib.boulder.ibm.com/eserver/ibmi.html>

Supported options: 5761JV1 options 6, 7, and 10 are the only options that are not supported in IBM i 7.1

15.2.3 Installation

Installing IBM WebSphere Application Server V6.1 on IBM i 7.1 requires a refreshed installation version of the product. The refreshed installation version of IBM WebSphere Application Server V6.1 is based on fix level 6.1.0.29. For ordering information, go to:

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

For IBM WebSphere Application Server V7.0, apply Fix Pack 7 (7.0.0.7) or later after you complete the installation.

15.2.4 WebSphere Application Server Installation Management Supports on IBM Web Admin for i

WebSphere Application Server V8.0 and higher is installed and managed by the IBM Installation Manager (IM) tool. Although this tool is available for IBM i, there is no GUI support on the platform. Therefore, WebSphere Application Server V8.0 and later can be installed only through the command line.

To address this problem, IBM Web Administration for i now provides a GUI interface for WebSphere Application Server to provide the following actions:

- ▶ Install product.
- ▶ Uninstall product.
- ▶ Update the installation with interim fixes or fix packs.
- ▶ View interim fixes on the installation.
- ▶ Delete interim fixes.

To use the installation management support, your IBM HTTP Server for i DG1 must be at the following minimum Group PTF levels:

- ▶ For IBM i 7.1: SF99368 level 16
- ▶ For IBM i 6.1: SF99115 level 28

Ensure that your user profile has at least the *ALLOBJ, *IOSYSCFG, and *SECADM special authorities.

Verify that your HTTP Admin server is active by running **WRKACTJOB SBS(QHTTPSVR)**. The displayed jobs should look as shown in Figure 15-1.




Figure 15-1 QHTTPSVR subsystem

To access your web admin GUI, enter `http://<your_system>:2001/HTTPAdmin` in your browser, and then click **Manage** → **Installations** to open the Manage Installations functions window that is shown in Figure 15-2.




Figure 15-2 Manage installations

In this window, you can install, update, uninstall, and review the fix level of WebSphere Application Server V8.0 and later.

You can use the Install button to install a new WebSphere Application Server product on your system. If the installation manager (IM) tool is not on your system, the wizard helps you install IM or update IM to the required level.

Complete the following steps:

1. Start the wizard by clicking Install WebSphere Application Server in the left navigation pane or by clicking the **Install** button, as shown in Figure 15-3.




Figure 15-3 Start the installation

2. After the welcome window opens, specify the IM installation packages location, as shown in Figure 15-4.



Figure 15-4 IM location

The installation package location can either be in your IFS or on a network file system (NFS) share. You can obtain the IM packages from the physical installation media or download them from the IM website found at:

<http://www-01.ibm.com/support/docview.wss?uid=swg24032358>

3. Specify the location of your WebSphere installation files, as shown in Figure 15-5. The product packages can be on your local system or a remote server. To access the packages on the local system, enter the location of the packages into the Location field or click the **Browse** button to find the location. If accessing the packages' location requires authentication, select **Specify authentication to access the install location** and complete the User and Password fields. Click **Next**.

Install WebSphere Application Server
Specify Product Install Packages Location - Step 3 of 6

To install a new WebSphere Application Server product installation, the product install packages are required. Specify the path of the product install packages. The product install packages can be located on the local system or a remote system. Specify the user authentication if required to access the remote install location. ?

Product Install Packages

Location: **Browse**

e.g. /builds/WAS/, http://www.ibm.com/software/repositorymanager/com.ibm.websphere.ND.v80/repository.config

☒ Specify authentication to access the install location

User:

Password:

Back **Next** **Cancel**

Figure 15-5 Product install Packages

4. Figure 15-6 shows the Liberty profile. The Liberty profile is a new feature in WebSphere Application Server V8.5. You can select to install it or not.

Install WebSphere Application Server
Select to Install Liberty Profile - Step 4 of 6

Liberty profile is a new lightweight application foundation provided since WebSphere Application Server V8.5. Select whether to install the Liberty profile on this installation.

☒ Install Liberty profile on this installation

Figure 15-6 Liberty Profile

- The next two steps (5 and 6) are just verification steps that are not described here. In the final window, click **Finish** to start the installation, as shown in Figure 15-7.




Figure 15-7 Finish installation

To uninstall a WebSphere product, select the version that you want to uninstall and click **Uninstall**, as shown in Figure 15-8.




Figure 15-8 Uninstall

The product and all profiles that are based on it are deleted from the system.

To update a WebSphere Application Server installation with interim fixes or fix packs, click the **Update** button in Figure 15-8 on page 548 instead. You must specify a location for your fixes, as shown in Figure 15-9.

Update Websphere Application Server

Specify the Location of the Fix Packages - Step 1 of 3

This wizard helps you update an existing WebSphere Application Server installation with a fix pack or interim fixes. Product fix packs contain bundled service to bring the WebSphere Application Server up to a new product level. Interim fixes provide corrective service for specific known problems.

Specify the fix pack or interim fixes package locations to update the WebSphere Application Server installation. The packages can be located on the local system or a remote system. Specify the user authentication if required to access the remote packages location. If more than one remote package is specified, make sure the packages are located on the same remote system.

Specify the Location of the Fix Packages: ?

☒ Fix Pack

Location:

e.g. /QIBM/WAS/WASFixPacks/WAS/FIXPACK/, ftp://hostname/WAS/FixPack/

☒ Interim Fix

Interim Fix Package Location	
Example	/home/WAS/iFixes/fix_name
Example	http://hostname/WAS/iFixes/fix_name
☒	Browse

☒ Specify authentication to access the package location

User:

Password:

Figure 15-9 Fix location

Click **Finish** to complete the installation, as shown in Figure 15-10.




Figure 15-10 Finish updating

To see the installed fix levels, select the installation and click **View Fix**, as shown in Figure 15-11.




Figure 15-11 View fixes

You can also delete fixes through the View Fix function. This topic is not described in this book.

For more information about the topics in this section, use the following resources:

- ▶ Web Administration website, found at:
http://www-03.ibm.com/systems/i/software/http/examples/getting_started.html
- ▶ WebSphere Application Server V8.5 Information Center, found at:
<http://pic.dhe.ibm.com/infocenter/wasinfo/v8r5/index.jsp>
- ▶ IM Information Center, found at:
<http://pic.dhe.ibm.com/infocenter/install/v1r5/index.jsp>

15.2.5 Upgrading to IBM i 7.1

If you are using IBM WebSphere Application Server V6.0 and upgrading to IBM i 7.1, you must migrate to the WebSphere Application Server V6.1.0.29 or V7.0.0.7 products. Migration can occur before or after you upgrade to IBM i 7.1.

Before you upgrade to IBM i 7.1, ensure that all WebSphere Application Server installations meet the minimum required fix levels. The version identifier is contained in the `<app_server_root>/properties/version/WAS.product` file, where `<app_server_root>` is the root directory of the IBM WebSphere Application Server installation. The version is also displayed on the IBM Web Administration GUI. It is listed on the introduction page under the Properties link. For WebSphere Application Server V6.1, apply Fix Pack 29 (6.1.0.29) or later if needed. For WebSphere Application Server V7.0, apply Fix Pack 7 (7.0.0.7) or later if needed.

Tip: Switch any WebSphere servers that are running classic to J9 before the OS upgrade. You can accomplish this task by running the following command:

```
/qibm/proddata/websphere/appserver/<version>/<edition>/bin/enablejvm -jvm  
std32'
```

When you upgrade to IBM i 7.1, enable WebSphere Application Server to use IBM Technology for Java virtual machine. The classic Java virtual machine is not available for IBM i 7.1. It is not operational until it is enabled to use IBM Technology for Java virtual machine.

After you upgrade to IBM i 7.1, if you upgraded from V5R4 or earlier, update the WebSphere Application Server service programs for IBM i 7.1. To update the programs, complete the following steps:

1. Start the Qshell interpreter.
2. Change the directory to `<app_server_root>/bin`.
3. Run **export OSVER=V7R1M0** to export the OSVER environment variable to the Qshell environment.

4. Run the `_postfpexit` script as follows:

```
_postfpexit<app_server_root><product_library>.
```

`<product_library>` is the product library for the WebSphere Application Server installation that is listed in `<app_server_root>/properties/product.properties`.

5. Upload the LoadModule directives. LoadModule directives are used by external HTTP servers that are associated with WebSphere Application Server V6.1 and WebSphere Application Server V7. For more information about updating the LoadModule directives, see 14.3, “Plug-ins and LoadModule directives” on page 534.

15.2.6 IBM Installation Manager for WebSphere Application Server V8

WebSphere Application Server V8 is the first full version to be installed by Installation Manager rather than by the programs that are used to install, update, and uninstall previous versions. Installation Manager is a single installation program that can use remote or local software repositories to install, modify, or update new WebSphere Application Server products. It determines available packages (including products, fix packs, interim fixes, and so on) checks prerequisites and interdependencies, and installs the selected packages. You also use Installation Manager to uninstall the packages that it installed.

Considerations

Note the following items:

- ▶ Do not use the same response files that are used with WebSphere Application Server V7 or earlier to install or uninstall Version 8 and later. Use response files that are based on Installation Manager to install, update, or uninstall Version 8 and later.
- ▶ The Installation Manager GUI is not available on IBM i; all interaction with Installation Manager on IBM i is done through the command line or response files.

Overview of Installation Manager

Installation Manager is a general-purpose software installation and update tool that runs on a range of computer systems. Installation Manager can be started through a command-line interface. You can also create response files in XML and use them to direct the performance of Installation Manager tasks in silent mode.

For more information about Installation Manager, go to the IBM Installation Manager Version 1.4 Information Center at:

<http://publib.boulder.ibm.com/infocenter/install/v1r4/index.jsp>

Packages and package groups

Each software product that can be installed with Installation Manager is referred to as a package. An installed package has a product level and an installation location. A package group consists of all of the products that are installed at a single location.

How many Installation Managers do you need

You must run Installation Manager only on those systems on which you install or update product code. You normally need only one Installation Manager on a system because one Installation Manager can track any number of product installations.

Creating an Installation Manager

When the installation kit is available on your system, you can create an Installation Manager. An Installation Manager consists of a set of binary files that are copied from the installation kit and a set of runtime data that describe the products that are installed by this particular Installation Manager. Before you created an Installation Manager, you must decide in which mode the Installation Manager runs, as well as where the binary files and runtime data, called “agent data” or “appdata”, are. Then, issue the Installation Manager installation command from the appropriate user ID to create the Installation Manager.

Accessing product repositories

All software materials that are installed with Installation Manager are stored in repositories. Each repository contains program objects and metadata for one or more packages, that is, software products at a particular level. Repositories can also contain product maintenance, such as fix packs and interim fixes. Whenever you install a new product, you can choose from any of the available product levels in any accessible repository.

Installing the product

After you create an Installation Manager and have access to all necessary product repositories, you can use Installation Manager command-line commands or response files to perform the actual product installations. When you install a product, you provide the package name, optionally the product level to be installed, the product location, and any other optional properties. For example, some products have optional features that you can select at installation time or a list of optional supported language packs from which you can select.

Working with installed products

You can use Installation Manager commands to list installed products and product levels. You can also obtain this information for installed copies of WebSphere Application Server V8 products by issuing **versionInfo** from the product file system. You can use Installation Manager commands or response files to install a new product level, roll back to a previous level, or modify the product by adding or removing optional features or language packs.

15.2.7 Enabling IBM Technology for Java virtual machine

IBM Technology for Java virtual machine is available in both 32-bit and 64-bit versions and is included in the licensed program 5761-JV1 with the system CDs. To install the IBM Technology for Java virtual machine option, complete the following steps:

1. Enter the Go Licensed Program (**GO LICPGM**) command and select Option 10.
2. If you do not see the program listed, then complete the following steps:
 - a. Enter the **GO LICPGM** command on the command line.
 - b. Select Option 11 (Install licensed program).
 - c. Choose option 1 (Install) for licensed program 5761-JV1 *BASE and select the option that you want to install.
3. Load the latest Java PTF group.

4. Set the JAVA_HOME environment variable to the home directory of the Java Development Kit that you want to use. At a command line, enter one of the following commands:
 - **ADDENVVAR ENVVAR(JAVA_HOME) VALUE(' /QOpenSys/QIBM/ProdData/JavaVM/jdk14/64 bit')**
 - **ADDENVVAR ENVVAR(JAVA_HOME) VALUE(' /QOpenSys/QIBM/ProdData/JavaVM/jdk50/32 bit')**
 - **ADDENVVAR ENVVAR(JAVA_HOME) VALUE(' /QOpenSys/QIBM/ProdData/JavaVM/jdk50/64 bit')**
 - **ADDENVVAR ENVVAR(JAVA_HOME) VALUE(' /QOpenSys/QIBM/ProdData/JavaVM/jdk60/32 bit')**
 - **ADDENVVAR ENVVAR(JAVA_HOME) VALUE(' /QOpenSys/QIBM/ProdData/JavaVM/jdk60/64 bit')**

Suggestion: Upgrade to IBM Technology for Java before you migrate to IBM i 7.1.

15.3 IBM Web Administration for i

IBM Web Administration for i is an integrated management interface for creating many types of servers, including web servers and application servers, such as WebSphere Application Server for i, WebSphere Portal Server, IBM Integrated Web Application Server for i, and IBM Integrated Web Services Server for i.

IBM Web Administration for i has several wizards that guide you through a series of advanced steps to accomplish a task. With a few clicks of a button, you can have a web server or web application server running in no time at all.

Enhancements were made to IBM Web Administration for i to include a web log monitor and permissions.

15.3.1 Web Log Monitor

In today's web application environment, notices and error messages that appear in the web servers logging files can be easily lost or missed. This situation leads to you to missing important information of which you might need to be aware. This support is included as part of the IBM Web Administration for i interface.

You can use Web Log Monitor to monitor the contents of log files for any web-related server, such as Integrated Web Application Server, Integrated Web Services Server, WebSphere Application Server, WebSphere Portal, and IBM HTTP Server.

Users can set rules that Web Log Monitor inspects. If a match is found, a notification is sent to one of the following sources:

- ▶ The *QSYSOPR system message queue
- ▶ One or more email addresses
- ▶ Both the *QSYSOPR system message queue and email addresses

Web Log Monitor ensures that important messages are not lost or missed.

Remember: Web Log Monitor inspects the log files if Web Administration for i is started. The minimum OS supported is IBM i 6.1 PTF SF99115 level 12 or higher.

Configuring Web Log Monitor

To activate Web Log Monitor, complete the following steps:

1. Start the IBM Systems Director Navigator for i by accessing the following URL from a web browser, where *your_system* is your IBM i server host name:
`http://your_system:2001`
2. From the IBM Systems Director Navigator for i welcome window (Figure 15-12), click the **IBM i Tasks Page** link in the right navigation pane.




Figure 15-12 IBM Systems Director Navigator for i

3. Click the **IBM Web Administration for i** link (Figure 15-13).




Figure 15-13 IBM Web Administration for i

4. Click **Manage** (Figure 15-14).




Figure 15-14 Manage tab

5. Click **Application Servers** (Figure 15-15).

IBM Web Administration for i

Setup **Manage** Advanced | Related Links

All Servers | HTTP Servers **Application Servers**

Running [Stop] [Refresh] Server: WSERVICE - V1.3 (web services)

WSERVICE

Manage Web Services Server

Server: **WSERVICE**

Web services server WSERVICE, created by the Create Web Services Server wizard.

The Web services server provides a convenient way to externalize existing programs running on IBM i, such as RPG and COBOL programs, as Web services. Web service clients can then interact with these IBM i program based services from the Internet or intranet using Web service based industry standard communication protocols such as SOAP. The clients can be implemented using a variety of platforms and programming languages such as C, C++, Java and .NET. An easy to use wizard is provided to configure the Web services server and the services for IBM i program objects. Other management functions such as starting, stopping and deleting services are also provided. For more information, please visit: <http://www-03.ibm.com/systems/i/software/iws/>

Manage Deployed Services

Server: "WSERVICE"
There are no entries for this table.

Note: To update the status, click Refresh

Figure 15-15 Application Servers

6. Click **Web Log Monitor** under the Problem Determination link in the left navigation window (Figure 15-16).




Figure 15-16 Web Log Monitor

7. Click **Enable Web Log Monitor** (Figure 15-17).




Figure 15-17 Enable Web Log Monitor

8. Specify the log file to monitor. Click **Browse** to select the log file (Figure 15-18). Only log files that you are authorized to use are shown in the browser window. Select the log file and click **Next**.



Figure 15-18 Specify Log File to monitor

When the log file is selected, you can use the Basic tab (Figure 15-19) or the Advanced tab to configure the rule.




Figure 15-19 Basic tab

You can use keywords to filter the content of specified log file. To specify more than one keyword, use a comma or semicolon to separate each key word. Three modes are provided:

- Match any keyword

For example, if the monitored file contains a line such as JVMDUMP0061 Processing Dump Event gpf.detail- Please wait and the keywords that are specified here are Dump, Failed, and Error, then the line is considered a match.

- Match all keywords

For example, if the monitored file contains a line such as JVMDUMP0061 Processing Dump Event gpf.detail- Please wait and if the keywords specified here are Dump, Event, and Wait, then this line is not considered a match, as the white space or blank character in front of Wait is also treated as part of the keyword. If the specified keywords are Dump, Event, and Wait, then this line is considered a match, as all three specified keys are in the line.

- Keyword A+any string+Keyword B

For example, if the monitored file contains a line such as JVMDUMP0061 Processing Dump Event gpf.detail- Please wait and the keywords that are specified here are Dump and detail, then this line is considered a match.

The Advanced Tab (Figure 15-20) specifies the complex patterns that are used in the rule:

- Exclude the following keywords
This option specifies whether to ignore a line that contains any of the keywords that follow.
- Configure regular expressions
You can use this option to customize the rule by using a regular expression.

The screenshot shows the IBM Web Administration console interface. The top navigation bar includes 'Setup', 'Manage', 'Advanced', and 'Related Links'. Below this, there are tabs for 'All Servers', 'HTTP Servers', and 'Application Servers'. The 'Application Servers' tab is selected, and a dropdown menu shows 'Server: IWAIAS - V7.1 (int app svr)'. On the left, a sidebar contains a tree view with categories like 'Common Tasks and Wizards', 'Application Server Wizards', 'Server Properties', 'Applications', 'Resource Configuration', 'Web Performance', and 'Problem Determination'. The 'Web Log Monitor' option is highlighted under 'Problem Determination'. The main content area is titled 'Configure Web Log Monitor' and 'Configure Rule - Step 2 of 3'. It contains a text area for defining the filter rule, with 'Basic' and 'Advanced' tabs. The 'Advanced' tab is active, showing options to 'Exclude the following keywords' and 'Configure regular expression'. A text input field contains the regular expression '(?!).*\\(Open\\E).*'. Below the input field, an example is provided: 'e.g. *Failed opening*on server*, ^[0-9]*[1-9][0-9]*\$', ^s*\\s*\$'. At the bottom, there are 'Back', 'Next', and 'Cancel' buttons.

Figure 15-20 Advanced tab

- If you have multiple log files to monitor, click **Add** in the File and Rule Tab, as shown in Figure 15-21.

IBM Web Administration for i

Setup **Manage** Advanced | Related Links

All Servers | HTTP Servers **Application Servers**

Running [Stop] [Start] [Refresh] Server: IWAIAAS - V7.1 (int app svr)

Common Tasks and Wizards

- Create Web Services Server
- Create HTTP Server
- Create Application Server
- Create WebSphere Portal

Application Server Wizards

- Create Database Connection
- Install New Application

Server Properties

- Properties
- Server Tracing
- View HTTP Servers

Applications

- Manage Installed Applications

Resource Configuration

- Manage Database Connections

Web Performance

- Web Performance Advisor

Problem Determination

- View Logs
- Web Log Monitor
- View Create Summary

Web Log Monitor provides users the ability to monitor the contents of log files for HTTP and application servers. You can define rules to describe what contents in a log file are to be monitored for. When a defined rule is matched in the specified log file, the *QSYSOPR system message queue and/or a specified email address.

Disable Web Log Monitor

File and Rule Notification Miscellaneous

File and rule configuration:

Log File	Configuration				
☐ /www/iwaiaas/lwi/logs/error-log-0.xml	Associated rule and notification text : <table border="1"> <thead> <tr> <th>Rule</th> <th>Notification Text</th> </tr> </thead> <tbody> <tr> <td>[Include]"open"</td> <td>File /www/iwaiaas/lwi/logs/error-log-0.xml has something</td> </tr> </tbody> </table>	Rule	Notification Text	[Include]"open"	File /www/iwaiaas/lwi/logs/error-log-0.xml has something
Rule	Notification Text				
[Include]"open"	File /www/iwaiaas/lwi/logs/error-log-0.xml has something				
☐ /www/iwaiaas/lwi/logs/error-log-2.xml	Associated rule and notification text : <table border="1"> <thead> <tr> <th>Rule</th> <th>Notification Text</th> </tr> </thead> <tbody> <tr> <td>[Include]"exception"</td> <td>File /www/iwaiaas/lwi/logs/error-log-2.xml has something</td> </tr> </tbody> </table>	Rule	Notification Text	[Include]"exception"	File /www/iwaiaas/lwi/logs/error-log-2.xml has something
Rule	Notification Text				
[Include]"exception"	File /www/iwaiaas/lwi/logs/error-log-2.xml has something				

Add Modify Remove

OK Apply Cancel

Figure 15-21 Add multiple log files

10. When the rules are set, define where the notification text is sent when a match is found by clicking the **Notification** tab (Figure 15-22). Here are the choices that you can use:

- Send message to QSYSOPR message queue
When selected, the notification text is sent to the “QSYSOPR system message queue when a match is found for a specified log file.
- Send email to
A notification is sent to the specified email address. More than one email address can be specified. Use a comma to separate multiple email addresses.
- Sender’s email address
Specifies the email address that is used to send the notification.
- SMTP server
Specifies the Simple Mail Transfer Protocol (SMTP) server that is used to send the mail notification.
- Account
Specifies the account that is used to log on to the SMTP server.
- Password
Specifies the password of the account that is used to log on to the SMTP server.

Figure 15-22 Notification tab

11. On the Miscellaneous tab, specify the interval in minutes and the maximum number of notifications to be sent per hour to prevent notification flood (Figure 15-23):

- Monitor interval

This option specifies the frequency of checking the log file. Set the interval to a small value when log files are updated frequently and a large value when the specified log files are updated infrequently.

- Notification flood prevention

This option specifies the maximum permitted number of notifications to be sent per hour to prevent notification flood. If the specified maximum number is exceeded, no notification is sent even when a match is found.




Figure 15-23 Miscellaneous tab

To disable Web Log Monitor, click **Disable Web Log Monitor** on the Web Log Monitor introduction window.

Requirement: Only users who have Developer or higher authority can configure Web Log Monitor.

15.3.2 Permissions

By default, only users with *IOSYSCFG and *ALLOBJ special authority can manage and create web-related servers on the system through IBM Web Administration for i. To get to the Permission tab, go to the IBM i Task window and click **IBM Web Administration for i** → **Advanced** → **Permissions**.

The Manage Permissions form allows administrators to give permissions to users without *IOSYSCFG and *ALLOBJ. There are two roles that can be granted:

- ▶ Operators
- ▶ Developers

A permission is the ability to perform an operation on a server. The ability for a user to perform operations on a server is determined by the role they are assigned for the server. The Web Administration for i roles are defined with the permissions listed in Table 15-2.

Table 15-2 Permissions corresponding to each role

Permissions	Administrator	Developer	Operator
Start/Stop server	x	x	x
Delete server	x	x	
Install/Remove applications ^a	x	x	
Install/Remove web services ^a	x	x	
Start/Stop applications	x	x	x
Start/Stop web services ^a	x	x	x
Modify server attributes	x	x	
Modify application attributes	x	x	
Create database connections	x	x	
Delete database connections	x	x	
Modify server tracing	x	x	x
Use Web Performance Advisor	x	x	
Use Web Performance Monitor	x	x	
Use Web Log Monitor	x	x	
Create Server ^b	x		

a. Web services that are deployed with integrated web services servers.

b. An administrator that grants permissions to a user profile must explicitly grant the create-server permission.

A new feature, *group profile support*, adds the ability to grant or revoke permissions to a group of users all at once. Otherwise, these users must be granted or revoked permissions separately, which is time consuming and error prone. When you use this feature, when a user has one or more supplemental groups, you combine the permissions that the individual has and the ones from its groups. The cumulative and highest permissions are treated when the user needs appropriate permissions to perform an operation through the Web Administration Interface.

The group profile support applies only to IBM i 6.1 and later.

Adding permissions

The Add Permissions wizard considers all aspects of a selected server. If a selected application server is associated with an HTTP Server, the Add Permissions wizard takes this situation into account and recommends that permissions are specified correctly for that web environment. This specification ensures that the specified user can successfully manage the server that is based on granted permissions.

Users who are granted permission to servers can be given a different role for each server. When a user is granted permission to create new servers, any server that they do create is automatically updated to give them developer permission to that newly created server.

To add permissions for a user ID, complete the following steps:

1. Click **Add Permissions** (Figure 15-24). Select the user ID to which you want to give permission. Click >> and click **Next**.




Figure 15-24 Add Permissions - step 1

2. Click **Server** and select the permissions that you want to grant to the user. Select the role and select the **Create server** check box if you want to grant the user the ability to create servers (see Figure 15-25). The options are explained as follows:
 - Server: Specifies how the server’s category is to be handled. There are three values:
 - *ALL: The users that are specified in step 1 on page 567 are granted permission to all the servers for this category, including all existing servers and any other server that is created in the future.
 - *NONE: The users that are specified in step 1 on page 567 have no permission to the servers for this category.
 - Specify: When this option is selected, an additional interface is displayed when the **Next** button is clicked. The Specify Specific Servers interface allows the Administrator to specify a list of existing servers that the specified users can work with.
 - Role: Lists the permissions that are granted to a user for a server. There are three values:
 - Operator: Basic permission to the server. See Table 15-2 on page 566 for details.
 - Developer: All permissions to the server. See Table 15-2 on page 566 for details.
 - Create Server: The specified users can be granted permission to create servers for this category.

IBM Web Administration for i

Setup | Manage | **Advanced** | Related Links

Settings | Internet Users and Groups | **Permissions**

▼ Common Tasks and Wizards

- Create Web Services Server
- Create HTTP Server
- Create Application Server
- Create WebSphere Portal

▼ Permissions

- Add Permissions
- Manage Permissions

Add Permissions

Specify Permissions for Servers - Step 2 of 3

Servers managed: ?

Server Type	Servers	Role	Create server
Integrated Web Application Server	*ALL	*ALL	☒
Integrated Web Services Server	*ALL	Developer	☒
WebSphere Application Server	*ALL	Developer	☒
IBM HTTP Server	*ALL	Operator	☒

Back Next Cancel

Figure 15-25 Add Permissions - step 2

- When you have defined the correct authority and role, click **Next**. Check what is displayed (Figure 15-26). If you must change anything, click **Back**; otherwise, click **Finish**.

IBM Web Administration for i

Setup | Manage **Advanced** | Related Links

Settings | Internet Users and Groups | **Permissions**

▼ Common Tasks and Wizards

- Create Web Services Server
- Create HTTP Server
- Create Application Server
- Create WebSphere Portal

▼ Permissions

- Add Permissions
- Manage Permissions

Add Permissions

Summary - Step 3 of 3

The following users are being granted permission:

User ID
CHUA

The following server permission roles are being granted to the specified users:

Server Type	Role/Servers
Integrated Web Application Server Create server <i>Enabled</i>	Developer *ALL Operator *NONE
Integrated Web Services Server Create server <i>Enabled</i>	Developer *ALL Operator *NONE
WebSphere Application Server Create server <i>Enabled</i>	Developer *ALL Operator *NONE
IBM HTTP Server Create server <i>Enabled</i>	Developer *ALL

Back Finish Cancel

Figure 15-26 Add Permissions - step 3

Modifying permissions

The Modify Permissions wizard allows an administrator to modify permissions for a specified server or user. The Modify Permissions wizard guides the administrator through this process.

The Modify Permissions wizard considers all aspects of a selected server. If a selected application server is associated with an HTTP Server, the Modify Permissions wizard considers this situation and recommends that permissions are specified correctly for that entire web environment. Either add or remove the permissions for all servers within that specified web environment. This action ensures that the specified user can either successfully manage the server based on the granted permissions or no longer successfully manage the server.

There are two ways to modify the permissions:

- Modify the permissions for a specific server (Figure 15-27). This action is started from the By Server view on the Manage Permissions window when you choose a specific server.

The screenshot shows the IBM Web Administration console interface. The top navigation bar includes 'Setup | Manage | Advanced | Related Links'. The left sidebar contains a tree view with 'Common Tasks and Wizards' (Create Web Services Server, Create HTTP Server, Create Application Server, Create WebSphere Portal) and 'Permissions' (Add Permissions, Manage Permissions). The main content area is titled 'Modify Permissions' and 'Specify Permissions for a Server - Step 1 of 2'. It contains instructions: 'Modify the user permissions for specified server. Allow users without the necessary default special authorities *ALLOBJ & *IO ability to create and manage servers or server groups they have been granted permission to.' The 'Specified Server' is set to '*ALL'. There are two sections for specifying user IDs: one for the 'Developer' role with a text box containing 'CHUA' and a 'Specify' button, and another for the 'Operator' role with an empty text box and a 'Specify' button. A note below the Operator section says 'Click the Specify button to choose servers for this permission.' At the bottom are 'Back', 'Next', and 'Cancel' buttons.

IBM Web Administration for i

Setup | Manage | **Advanced** | Related Links

Settings | Internet Users and Groups | **Permissions**

▼ Common Tasks and Wizards

- Create Web Services Server
- Create HTTP Server
- Create Application Server
- Create WebSphere Portal

▼ Permissions

- Add Permissions
- Manage Permissions

Modify Permissions

Specify Permissions for a Server - Step 1 of 2

Modify the user permissions for specified server. Allow users without the necessary default special authorities *ALLOBJ & *IO ability to create and manage servers or server groups they have been granted permission to. ?

Specified Server: *ALL

Specify the user IDs for **Developer** role:

User ID

CHUA

Specify

Specify the user IDs for **Operator** role:

User ID

Click the Specify button to choose servers for this permission.

Specify

Back Next Cancel

Figure 15-27 Modify permissions by Server view

- Modify the permissions for a specific user (Figure 15-28). This action is started from the By User view on the Manage Permissions window when you choose a specific user.

IBM Web Administration for i

Setup | Manage | **Advanced** | Related Links

Settings | Internet Users and Groups | **Permissions**

Modify Permissions
Specify Permissions for a User - Step 1 of 2

Modify the permissions for the specified user. The permissions specified allow this user the ability to work with these servers within IBM Web Administration for i with out the special authorities *ALLOBJ & *IOSYSCFG. This includes the ability to create and manage specific servers or server groups.

Specified user: CHUA

Servers managed: ?

Server Type	Servers	Role	Create server
Integrated Web Application Server	ALL	Developer	☐
Integrated Web Services Server	*ALL	Operator	☒
WebSphere Application Server	*ALL	Developer	☒
IBM HTTP Server	*ALL	Developer	☒

Back Next Cancel

Figure 15-28 Modify permissions by User view

Removing permissions

The Remove Permissions wizard provides an administrator with the ability to remove the permissions for a specified server or user. The removing of permissions removes the ability of the specified user to work with and manage a server within the IBM Web Administration for i interface.

The Remove Permissions wizard considers all aspects of a selected server. If a selected application server is associated with an HTTP Server, the Remove Permissions wizard takes this situation into account and also removes the permissions for all servers within that specified web environment. This action ensures that the specified user no longer successfully manages the server.

There are two ways to remove the permissions:

- Remove all permissions for a specific server and the other servers from the same web environment (Figure 15-29). This action is performed from the By Server view on the Manage Permissions window when you choose a specific server.




Figure 15-29 Remove permissions by Server view

- Remove all the permissions for a specific user (Figure 15-30). This action is performed from the By User view on Manage Permissions window when you choose a specific user.




Figure 15-30 Remove Permissions by User view

Click **Remove** to remove the permissions. This operation cannot be undone.

15.3.3 Web Performance Advisor

Web Performance Advisor (Figure 15-31) provides a consolidated place to review, evaluate, and modify the performance attributes for your web environments on IBM i. This tool is a great first stop to ensure that you are running at a reasonable set of performance-related settings. It is updated to include the latest performance settings for IBM i 7.1 as it relates to each of the supported web-based servers. WebSphere Portal Server V6.1.5 is now supported on Web Performance Advisor.

If you are using WebSphere Application Server V6 and upgrading to IBM i 7.1, you must update the web performance profile. The classic Java virtual machine is not available for IBM i 7.1. If your WebSphere Application Server installation is enabled to use classic Java, it is not operational until it is enabled to use IBM Technology for Java virtual machine. For more information, see 15.2.2, “Java for IBM i 7.1” on page 543.

IBM Web Administration for i

Setup **Manage** Advanced | Related Links

All Servers | HTTP Servers **Application Servers**

Stopped [Start] [Stop] [Refresh] Server: IWAIAS - V7.1 (int app svr)

Web Performance Advisor ?

The overall Web performance evaluation cannot be determined at this time. See the system and Web environment sections for additional details.

System Performance Attribute Information

Possible performance improvements may be realized by updating system Web performance attributes to acceptable values.

Host name:	CTCV71	Memory:	11.85 GB
System model:	MMA	Disk units:	12
Processor feature:	5462	Total disk storage:	1.30 TB
System CPW:	76900		

Manage system attributes

Web Environment: IWAIAS - LWISVR7.1

Evaluation for this Web environment is unknown at this time; Web environment servers must be running.

Select	Name	Type	Evaluation
☒	IWAIAS	V7.1 (int app svr)	? Unknown - Web environment is not running
☐	IWAIAS	Apache-HTTP/Apache/2.2.11 (i5)	! Improvements possible

Manage attributes

Export performance profile **Import performance profile** **Cancel** **Generate**

Figure 15-31 Web Performance Advisor



Application development

This chapter describes the following topics:

- ▶ High-level programming languages
- ▶ Control Language (CL)
- ▶ PHP
- ▶ XMLSERVICE
- ▶ IBM Lotus products for IBM i
- ▶ Native archive and unarchive API support
- ▶ IBM Toolbox for Java JDBC
- ▶ JTLite Java Classes Library for mobile application development
- ▶ Application Runtime Expert for i

16.1 High-level programming languages

IBM i has a suite of IBM Integrated Language Environment® (ILE) compilers that include RPG, COBOL, C, C++, and CL. IBM Rational Development Studio for i V7.1 delivers significant enhancements to the compilers.

For more information about ILE Compilers, see 17.4, “IBM Rational Development Studio for i” on page 618.

16.2 Control Language (CL)

This section covers the following topics:

- ▶ New workload capping commands
- ▶ Retrieving CL source support for ILE CL
- ▶ Longer Integer CL variables for ILE CL
- ▶ Showing DO and SELECT nesting levels in a compiler listing
- ▶ Encrypting the debug listing view for ILE CL
- ▶ Nested INCLUDE support

16.2.1 New workload capping commands

New workload capping commands were added to the IBM i commands. You can use workload capping to set a usage limit for a licensed program by restricting the number of processor cores that are available to be used by the licensed program.

Here are the new workload capping commands:

- ▶ Add Workload Capping Group (**ADDWLCGRP**)
- ▶ Add Workload Capping Product Entry (**ADDWLCPRDE**)
- ▶ Change Workload Capping Group (**CHGWLCGRP**)
- ▶ Display Workload Capping Group (**DSPWLCGRP**)
- ▶ Remove Workload Capping Group (**RMVWLCGRP**)
- ▶ Remove Workload Capping Product Entry (**RMVWLCPRDE**)

For more information about these commands, see the following website:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=%2Frbam6%2Frbam6whatsnew.htm>

16.2.2 Retrieving CL source support for ILE CL

Developers can move to IBM Integrated Language Environment (ILE) without fear of not being able to retrieve their source statements. IBM i 7.1 included this enhancement with the ILE module. For completeness, even the Declare Processing Options (**DCLPRCOPT**) command was updated to reflect that the **ALWRTVSRC** keyword now applies to ILE program as well.

The Retrieve CL Source (**RTVCLSRC**) command can be used to retrieve control language CL source statements from an Integrated Language Environment ILE module. The module must be created with the Create CL Module (**CRTCLMOD**) command or the Create Bound CL program (**CRTBNDCL**) command by specifying *YES for the **ALWRTVSRC** parameter. The module that contains the CL source to be retrieved can be a module (*MODULE) object or a module within an ILE program (*PGM) or service program (*SRVPGM) object.

The ability to retrieve the CL source from a CL program or CL module can make it easier to diagnose and fix problems in CL code on systems where the original CL source code is not available.

The following command retrieves the CL source from module MOD1 in the ILE program MYCLPGM:

```
RTVCLSRC PGM(MYCLPGM) MODULE(MOD1) SRCFILE(MYLIB/QCLSRC)
```

The retrieved CL source is stored in member MOD1 of the source physical file QCLSRC in library MYLIB. The shipped default value for the **ALWRTVSRC** parameter is *YES.

16.2.3 Longer Integer CL variables for ILE CL

All variables must be declared to the CL program or procedure before they can be used by the program or procedure.

The Declare CL variable (**DCL**) command supports a value of 8 for the **LEN** parameter for signed integer *INT and unsigned integer *UINT variables if the CL source is compiled using the **CRTCLMOD** or the **CRTBNDCL** commands. This capability is useful when you call API programs and API procedures that define 8 byte integer fields in input or output structures.

Important: **LEN(8)** can be specified only if the CL source is compiled with the Create CL Module (**CRTCLMOD**) or the Create Bound CL Program (**CRTBNDCL**) command.

16.2.4 Showing DO and SELECT nesting levels in a compiler listing

You can specify a CL compiler option to see the nesting level for all types of **DO** and **SELECT** commands.

CL source programs contain **DO** and **SELECT** commands where these commands are nested several levels deep. For example, between a **DO** command and the corresponding **ENDDO** command can be a **DOFOR** and another **ENDDO** command. The CL compiler supports up to 25 levels of nesting for **DO** commands and **SELECT** commands.

You can specify **OPTION(*DOSLTLVL)** and the Create CL Program (**CRTCLPGM**) command or the **CRTCLMOD** or the **CRTBNDCL** commands.

This compiler option adds a column to the compiler listing, which shows the nesting levels for the following elements:

- ▶ Do DO
- ▶ Do For DOFOR
- ▶ Do Until DOUNTIL
- ▶ Do While DOWHILE
- ▶ SELECT

If you do not want to see this nesting level information, you can specify *NODOSLTLVL for the **OPTION** parameter.

16.2.5 Encrypting the debug listing view for ILE CL

You can create a compiler listing view of the ILE CL procedure in a CL module by specifying *LIST for the Debugging view **DBGVIEW** parameter on the **CRTCLMOD** or the **CRTBNDCL** command.

The listing view can be seen by anyone with sufficient authority to the program or service program object that contains the CL module.

In IBM i 7.1, a new keyword parameter Debug encryption key **DBGENCKEY** was added to the **CRTCLMOD** and the **CRTBNDCL** commands. Specifying an encryption key value for the **DBGENCKEY** parameter and also specifying ***LIST** for the **DBGVIEW** parameter causes the debug listing data to be encrypted before it is stored with the module ***MODULE** or ILE program ***PGM** object. To see the listing view during debugging, you must provide the same encryption key value.

When you start the debug session, you are prompted for the encryption key value. If the same value is not specified for the debug session that was specified when the CL module was created, no listing view is shown.

16.2.6 Nested **INCLUDE** support

You can use the Include CL Source (**INCLUDE**) command to split your CL source code so that the CL source code can be compiled across multiple source file members.

The CL source to be embedded can be in another member of the same source file that is identified on the Source file **SRCFILE** parameter of the CL compiler commands or another source file. The CL compiler commands include **CRTCLPGM**, **CRTCLMOD**, and the **CRTBNDCL** Program.

You can run the **RTVCLSRC** command later to retrieve either the original CL source (which contains just the **INCLUDE** commands) or the expanded CL source (which contains the embedded CL source commands).

16.3 PHP

PHP stands for PHP Hypertext Preprocessor. PHP is an open source scripting language that is designed for web application development and enables simple scripting.

PHP applications are easily integrated with data in IBM DB2 for i, RPG, COBOL, and other business applications that are running on IBM i.

PHP is used for content management, customer relationship management, database access, e-commerce, forums, blogs, wikis, and other web-based applications.

Zend and IBM worked together to deliver Zend Solutions for IBM i, a complete PHP development and production environment solution for the IBM i platform.

Here are the Zend Solutions for IBM i:

- ▶ Zend Server Community Edition for IBM i
- ▶ Zend Server for IBM i
- ▶ Zend Studio for IBM i

Zend Solutions for IBM i can be downloaded from the Zend website at:

<http://www.zend.com/en/solutions/modernize-ibm-i/ibm-i-products>

The following sections summarize features and enhancements of Zend products for IBM i.

16.3.1 Zend Server Community Edition for IBM i

Zend Server Community Edition (CE) is a fully tested and enhanced version of the open source PHP. It provides the PHP run time and is packaged to make the software installation easier and faster with the instant PHP setup. It is enhanced to take advantage of IBM i specific resources and capabilities.

Zend Server CE for IBM i is a lightweight version of Zend Server, and replaces Zend Core. It offers the following features:

- ▶ Preinstalled on IBM i 6.1 and IBM i 7.1 starting April 2010.
- ▶ Includes extensions and a toolkit that provides the following functions:
 - Enables PHP application to easily access DB2 for i data
 - Takes advantage of RPG and COBOL applications in IBM i
 - Supports for Program call, Procedure call, Data Area, Data Queue, Message Queue, Commands, and System values
- ▶ Simple to install and to use, and provides basic performance optimization.
- ▶ Available at no charge to for using in development or in production, and it comes with an initial year of Silver Support that is provided by IBM.
- ▶ Zend Server is the only Zend certified and supported version of PHP for IBM i.

16.3.2 Zend Server for IBM i

Zend Server is a robust PHP production environment that helps ensure that applications written in PHP run smoothly at all time. It is designed for IT personnel and businesses that require commercial grade web applications in highly reliable production environments.

Zend Server replaces the Zend platform. It offers all the features that are provided in Zend Server CE for IBM i and the following additional features:

- ▶ High performance and scalability to provide customers with an improved web experience and response time.
- ▶ Delivers application uptime and reliability through enhanced PHP monitoring and immediate problem resolution.
- ▶ Includes the Java Bridge for integrating PHP application with Java application.
- ▶ Includes 5250 bridge for integrating 5250 applications with PHP applications. The 5250 bridge allows running interactive IBM i based applications from a web browser.

The following website shows a comparison between the features that are offered in Zend Server CE for IBM i and Zend Server for IBM i:

<http://www.zend.com/en/products/server/editions-ibm-i>

16.3.3 Zend Studio for IBM i

Zend Studio for IBM i is an industry-leading PHP integrated development environment (IDE) designed for professional developers. It includes all the development components necessary for the full PHP application lifecycle and simplifies complex projects.

Zend Studio for IBM i includes the following features and enhancements:

- ▶ Enhanced to work with the integration toolkit provided with Zend Server and Zend Server CE for IBM i
- ▶ Includes a comprehensive set of editing, debugging, analysis, optimization, database, and testing tools
- ▶ Toolkit support for easy integration with earlier IBM i applications and data
- ▶ Customizable and context-sensitive templates for IBM i Toolkit functions
- ▶ Create / generate PHP Toolkit code quickly
- ▶ Call an RPG/CL/COBOL program, run a CL command, retrieve Spooled file entries, and access Data Area, Data Queue, and User Space.

For more information about Zend products for IBM i, see the following websites:

- ▶ Zend and IBM i:
<http://www-03.ibm.com/systems/i/software/php/index.html>
- ▶ Zend Products for IBM i:
<http://www.zend.com/en/solutions/modernize-ibm-i/ibm-i-products>

16.4 XMLSERVICE for IBM i

XMLSERVICE is open source RPG code that enables web XML scripting calls of IBM System i resources, such as PGM, SRVPGM, CMD, PASE utilities, and SQL DB2, over various driver transports (DRDA, ODBC, and REST) through 1-tier (IBM i) or 2-tier clients (Linux/Windows).

XMLSERVICE can be accessed from PHP, Java, Ruby, RPG, Groovy, or other languages that can work with XML.

All the information, including several demonstrations, can be found at:

<http://www.youngipprofessionals.com/wiki/XMLSERVICE>

The runtime version is available through the following PTFs:

- ▶ IBM i 7.1: 5770-SS1 SF99368 HTTP PTF Group Level 16, or later
- ▶ IBM i 6.1: 5761-SS1 SF99115 HTTP PTF Group Level 28, or later

For more information about XML or open source programming languages, go to:

<http://www.iprodeveloper.com/article/associate/unleash-your-ibm-i-with-xmlservice-65781>

16.4.1 XMLSERVICE highlights

Here are some highlights of XMLSERVICE:

- ▶ The code is developed in RPG.
- ▶ XMLSERVICE can be licensed to be used at no additional cost (XMLSERVICE library download).
- ▶ XMLSERVICE can be deployed as either PC (2-tier) or IBM i (1-tier).
- ▶ All communication is through XML documents.

- ▶ The called RPG programs can maintain state, commit boundaries, and so on.
- ▶ XMLSERVICE provides an *interprocess communication* (IPC) mechanism that synchronizes multiple web client sources.
- ▶ There is traditional DB2 connection security 1–2 tier (when you use the provided DB2 stored procedures).
- ▶ XMLSERVICE handles the conversions of XML string text to and from the actual program required parameter type (packed decimal 12p2, zoned decimal 7s2, integer 10i0/5i0, float 4f2/8f4, or hex binary 400b), and character data is converted automatically to ASCII/EBCDIC (1024A/10a varying='on').
- ▶ Arrays of data structures are allowed both as parameters and return, including nested data structures or arrays of data structures.

16.4.2 XMLSERVICE start

There are several options to start XMLSERVICE. It can be called from a web server, any application server that handles HTML or REST web services, through the ODBC Driver, or through the PHP toolkit.

To start XMLSERVICE, complete the following steps:

1. Install the XMLSERVICE programs.
2. Define the communication method to use:
 - 1 Tier: Apache Web Server or PHP (Zend Server)
 - 2 Tier: ODBC Driver or Web Application Server / PHP Server in the middle tier

Installing XMLSERVICE

The XMLSERVICE Library is on your IBM i systems if you apply the latest HTTP Server PTF Group:

- ▶ IBM i 7.1: 5770-SS1 SF99368 HTTP PTF Group Level 16, or later
- ▶ IBM i 6.1: 5761-SS1 SF99115 HTTP PTF Group Level 28, or later

XMLSERVICE can be downloaded from the Young IBM i Professionals website at:

<http://www.youngiprofessionals.com/wiki/XMLSERVICE>

After you have the XMLSERVICE package, upload it in to your IBM i System by running the commands that are shown in Example 16-1.

Example 16-1 Uploading the XMLSERVICE.zip file

```
#IBM i CL Command:
CRTSAVF FILE(QGPL/XMLSERVICE)
#PC:
unzip xmlservice-rpg-1.7.3.zip
ftp IBMiHostName ...
bin
cd QGPL
put /path/xmlservice.savf XMLSERVICE
```

The commands that are shown in Example 16-2 finish the installation process.

Example 16-2 CL Commands to install XMLSERVICE and test programs

```
#IBM i CL Commands:
RSTLIB SAVLIB(XMLSERVICE) DEV(*SAVF) SAVF(QGPL/XMLSERVICE)
ADDLIB XMLSERVICE
CRTCLPGM PGM(XMLSERVICE/CRTXML) SRCFILE(XMLSERVICE/QCLSRC)
CALL CRTXML
#If you want to create some optional programs for testing and demo, execute the
#following commands:

CRTCLPGM PGM(XMLSERVICE/CRTTEST) SRCFILE(XMLSERVICE/QCLSRC)
CALL CRTTEST
CRTCLPGM PGM(XMLSERVICE/CRTTEST6) SRCFILE(XMLSERVICE/QCLSRC)
CALL CRTTEST6

#Do not forget to provide the corresponding authorization:
CHGAUT OBJ('/qsys.lib/XMLSERVICE.lib') USER(QTMHHTTP) DTAAUT(*RWX) OBJAUT(*ALL)
SUBTREE(*ALL)
CHGAUT OBJ('/qsys.lib/XMLSERVICE.lib') USER(QTMHHTTP1) DTAAUT(*RWX) OBJAUT(*ALL)
SUBTREE(*ALL)
```

CCSIDs: If you are running an IBM system with CCSID 65535 (hex), you must set the Apache web server to a valid CCSID (such as 37 or 284) by running the following command:

```
DefaultFsCCSID 37    CGIJobCCSID 37
```

Configuring the HTTP Server

To configure the XMLSERVICE REST, open the HTTP administrative console session by using the URL `http://hostname:2001/HTTPAdmin` and add the interface by using the RPG CGI (`xml cgi .pgm`), as shown in Figure 16-1.




Figure 16-1 HTTP Server configuration to add XMLSERVICE

Testing XMLSERVICE

You can create a few HTML files, including XML tags to call IBM i resources. Figure 16-2 shows a program call example using XMLSERVICE.




Figure 16-2 HTML file with XML tags to call an IBM i program

User ID and password: Depending on your security setup, you might need to change the session parameters “uid” and “pwd” to a valid user ID and password.

To start the program through the web browser, enter the IBM i URL address or host name and the previous HTML file name:

`http://HostName/yourhtmlfilename.html`

The resulting XML document is shown in Figure 16-3. It is a basic sample of a Program call. However, there is a rich set of XML tags that you can use with other IBM i resources, such as <cmd>, <sh>, <pgm>, and <sql>.

```

- <script>
- <pgm name="ZZCALL" lib="XMLSERVICE">
- <parm>
  <data type="1A">C</data>
</parm>
- <parm>
  <data type="1A">D</data>
</parm>
- <parm>
  <data type="7p4">321.1234</data>
</parm>
- <parm>
  <data type="12p2">1234567890.12</data>
</parm>
- <parm>
- <ds>
  <data type="1A">E</data>
  <data type="1A">F</data>
  <data type="7p4">333.3330</data>
  <data type="12p2">4444444444.44</data>
</ds>
</parm>
- <return>
  <data type="10i0">0</data>
</return>
</pgm>
</script>

```

Figure 16-3 Returned XML file by XMLSERVICE

For more information about XMLSERVICE commands, parameters, and reserved words, go to:

<http://www.youngiprofessionals.com/wiki/XMLSERVICE>

16.5 IBM Lotus products for IBM i

Lotus software delivers robust collaboration software that empowers people to connect, collaborate, and innovate when optimizing the way they work. With Lotus, you can achieve better business outcomes through collaboration. The following Lotus Software products are supported on IBM i 7.1:

- ▶ Lotus Domino V8.5.1 or later
- ▶ Lotus Enterprise Integrator (LEI) V8.5.2 or later
- ▶ Lotus Sametime V8.5.1 or later
- ▶ Lotus VQuickr 8.5 or later
- ▶ IBM Forms Server V4.0 or later
- ▶ Lotus Forms Server V3.5.1 FP2 or later
- ▶ Lotus Workflow V7.0 or later
- ▶ IBM Integrated Domino Fax for i5/OS V4R5

Product withdrawal: Domino Fax for i5/OS V4R5 was withdrawn from worldwide marketing on January 14, 2011 and end of support was on April 30, 2012.

Only specific releases of these products are supported on IBM i 7.1. Before you upgrade to IBM i 7.1, check the most current details about the product releases supported at the following website:

http://www-03.ibm.com/systems/resources/systems_power_ibmi_lotus_releasesupport.pdf

16.6 Native archive and unarchive API support

IBM i 7.1 now supports both native creation and extraction of archive files.

This support includes the following native APIs and a service program to create archive files:

- ▶ QZIPZIP API
- ▶ QZIPUNZIP API
- ▶ QZIPUTIL Service program

16.6.1 QZIPZIP API

Multiple files and directories within IBM i can be compressed and packaged into a single archive file using the QZIPZIP() API.

The QZIPZIP syntax is as follows:

```
#include <qziputil.H>
void QzipZip(
    Qlg_Path_Name_T * fileToCompress,
    Qlg_Path_Name_T * compressedFileName,
    char * formatName,
    char * zipOptions,
    char * errorStruct)
```

Parameters

Table 16-1 shows the list of QZIPZIP API parameters.

Table 16-1 Parameters for QZIPZIP API

Name	Type	Description
fileToCompress	Input	The name of the file or directory that is to be compressed into an archive file. The path name must be in Qlg_Path_Name_T structure.
CompressedFileName	Input	The name of the compressed archive file. This file is created by the API. The path name must be in Qlg_Path_Name_T structure.
formatName	Input	The format name to pass the user's options for compressing a file or a directory to an archive file. For more information, see "ZIP00100 format description" on page 587.
zipOptions	Input	This pointer passes in the user's options to the QZIPZIP API in ZIP00100 format.
errorStruct	Input/output	This item is a pointer to an error code structure to receive error information.

Authorities and locks

The user must have at least the following authorities to be able to use the API:

- *R data authority to the file that is to be compressed
- *R data authority to each directory in the path name that precedes that file
- *W data authority to the directory where the compressed file is written

During the time, this API reads a file for compressing it. The file is locked and shared with reading only mode. The API releases the lock on the file after it reads the file. If the file that is to be compressed is locked, then an error message is sent. Further compression is stopped.

ZIP00100 format description

Table 16-2 shows the format for passing the user's options to compress files or directories.

Table 16-2 Format description

Offset		Type	Field
Dec	Hex		
0	0	CHAR(10)	Verbose option
10	A	CHAR(6)	Subtree option
16	10	CHAR(512)	Comment
528	210	BINARY(4), UNSIGNED	Length of the comment

16.6.2 QZIPUNZIP API

The contents of the (.zip) archive file can be extracted by running QZIPUNZIP on the target IBM i system.

The QZIPUNZIP syntax is as follows:

```
#include <qziputil.H>
void QzipUnzip(
    Qlg_Path_Name_T * compressedFileName,
    Qlg_Path_Name_T * dirToPlaceDecompFiles,
    char * formatName,
    char * unzipOptions,
    char * errorStruct)
```

Parameters

Table 16-3 shows the list of QZIPUNZIP API parameters.

Table 16-3 Parameters for QZIPZIP API

Name	Type	Description
CompressedFileName	Input	The name of the archive file that is to be extracted. The path name must be in the Qlg_Path_Name_T structure.
dirToPlaceDecompFiles	Input	The directory in which you place the contents of the archive file. The path name must be in the Qlg_Path_Name_T structure.
formatName	Input	The format name to pass the user's options for extracting an archive file. See "UNZIP100 format description" for a description of this format.

Name	Type	Description
unzipOptions	Input	This item is a pointer that passes in the user's options to the QZIPUNZIP API in UNZIP100 format.
errorStruct	Input/output	This item is a pointer to an error code structure to receive error information.

Authorities and locks

The user must have at least the following authorities to be able to use the API:

- ▶ *R data authority to the file that is to be decompressed
- ▶ *R data authority to each directory in the path name that precedes that file
- ▶ *W data authority to the directory where the decompressed file is written

During the time, this API reads a compressed archive file for extracting it, and the file is locked and shared for reading only. The API releases the lock on the file after it reads the file completely. If the file that is to be extracted is locked, then an error message is sent. Further extracting is stopped.

UNZIP100 format description

Table 16-4 shows the format for passing the user's options to compress files or directories.

Table 16-4 Format description

Offset		Type	Field
Dec	Hex		
0	0	CHAR(10)	Verbose option
10	A	CHAR(6)	Subtree option

16.6.3 QZIPUTIL Service program

The service program QZIPUTIL has entry points that can be called by any other ILE program to create and extract archive files. It is a system state user domain service program that adopts *USER authority.

The QZIPUTIL service program exported QZIPZIP and QZIPUNZIP APIs.

16.7 IBM Toolbox for Java JDBC

JDBC is an API included in the Java platform that enables Java programs to connect to a wide range of databases.

You can use the IBM Toolbox for Java JDBC driver to use JDBC API interfaces to issue Structured Query Language (SQL) statements to and process results from databases on the server.

The following sections describe the enhancements that are done to IBM Toolbox for Java JDBC support for IBM i 7.1.

16.7.1 JDBC 4.1

Java JDBC interface now supports several features of the latest JDBC 4.1 definitions with DB2 i. Java developers can use this enhancement to continue to use the latest defined options for JDBC.

16.7.2 XML Data type support

The JDBC 4.0 interface specification adds new methods and classes for XML Data type support. IBM Toolbox for Java implements XML support in its JDBC 4.0 driver.

This enhancement allows JDBC clients easy access to IBM i 7.1 XML support.

16.7.3 Database metadata updates

Database metadata is obtained by calling methods of the `AS400JDBCDatabaseMetaData` class. Starting with IBM i 7.1, the default behavior of IBM Toolbox for Java JDBC is to obtain this metadata from a set of standard system stored procedures.

This enhancement brings IBM Toolbox for Java into alignment with IBM i native JDBC support in addition to JDBC drivers on other platforms.

To provide compatibility with earlier versions of the metadata functionality, a new connection property, “metadata source”, can be used to force IBM Toolbox for Java JDBC to use the old method of retrieval of database metadata.

16.7.4 Currently committed support

Lock timeouts and deadlocks can occur under the isolation levels that perform row-level locking, especially with applications that are not designed to prevent such problems. Some high throughput database applications cannot tolerate waiting on locks that are issued during transaction processing, and some applications cannot tolerate processing uncommitted data, but still require non-blocking behavior for read transactions.

Under the new currently committed semantics, if currently committed is enabled, then only committed data is returned, as was the case previously, but now readers do not wait for writers to release row locks. Instead, the data that is returned to readers is based on the currently committed version, that is, data before the start of the write operation.

This feature also implements a way to direct the database manager to wait for the outcome when it encounters data while being updated.

16.7.5 Array type support

IBM Toolbox for Java supports the IBM i 7.1 SQL array data type in stored procedure parameters. Arrays of all of the various DB2 types are supported, except for data that is returned in a locator.

IBM Toolbox for Java JDBC adds support for arrays as **IN**, **OUT**, and **INOUT** parameters to stored procedures. However, `ResultSet`s returned from stored procedures or other queries that contain arrays is not supported.

JDBC supports the calling of stored procedures in the `java.sql.CallableStatement` class, which IBM Toolbox for Java implements in `AS400JDBCCallableStatement`.

16.7.6 Long schema name support

IBM i 7.1 DBMS added support for 128 byte schema names. IBM Toolbox for Java JDBC is adding support for long schema names.

16.8 JTLite/JTOpenLite Java Classes Library

JTLite is an alternative to the JT400 and JTOpen Java toolkits, and provides a small footprint (about 420 KB) and good performance that is focused on application development for mobile devices.

The Lite version of the Java toolbox for IBM i provides a new smaller foot print toolkit for accessing IBM i native objects from Java running on mobile devices. Table 16-5 shows the difference between both Java Libraries.

Table 16-5 Differences between JTOpen and JTOpen Lite

JTOpen and JT400	JTOpenLite and JTLite
Database - JDBC (SQL) and record-level access (DDM)	Database - JDBC (SQL) and record-level access (DDM)
Integrated File System	Integrated File System - Open, Read, Write, Delete
Program calls (RPG, COBOL, Service Programs, and so on)	Program calls (RPG, COBOL, Service Programs, and so on)
Commands	Commands
Data Queues	
Data Areas	
Print/Spool resources	
Product and PTF information	
Job and Joblogs	Job information
Messages, messages queue, message files	
Data Queues	
Users and groups	User and Group information
User spaces	
System values	
System status	

For access to the project, go to the following website:

<http://jtlite.sourceforge.net/>

16.9 Application Runtime Expert for i

IBM Application Runtime Expert for i (ARE) is a new product that has the potential to revolutionize how you do application service and support. ARE can help you ensure consistent performance and deployment for any workload that is running on your system. You can use ARE to build and maintain an application knowledge base that can be used to automatically apply its knowledge to verify a system.

ARE provides a GUI that you can use to collect and verify a customized set of information, system settings, and attributes about the following items:

- ▶ Applications
- ▶ An IBM i system
- ▶ A runtime environment

ARE collects the needed information and builds it into a template. This template can then be used in verifying the application and its environment, on the same system where the template was built, or any other IBM i system.

When you run a template against an IBM i system, the system is verified against the information that is stored in the template, and the results are documented in a set of reports. These reports give a clear and concise view of what changed, or what is different.

Templates can be updated easily to include new information and settings from the ARE GUI.

16.9.1 Deployment template

A deployment template represents the expected attributes and state of a deployment and its environment. Typically, a deployment consists of application attributes and attributes of the environment the application runs in. Users can customize different plug-ins by defining what information to include in the deployment template using the *Deployment Template Editor*. Each plug-in can verify a different aspect of a deployment.

Here are common examples of a deployment:

- ▶ IBM i products
- ▶ Applications on IBM i:
 - ISV application
 - WebSphere Application Server
 - Integrated WebSphere Application Server
 - RPG application
 - Directory tree in IFS
- ▶ Custom selection of software products, system, and environment information

The deployment template is used as an input to the ARE Core. ARE Core uses the deployment template as the basis for comparison for the attributes and state of the deployment on the system that is being verified.

Deployment Template Editor

You can use the Deployment Template Editor to create, edit, import, and export deployment templates. Users can use the editor to customize their templates. Templates that are created or edited can be used to verify systems using the console.

16.9.2 Customizing plug-ins for templates

IBM Application Runtime Expert for i has a wide array of application attributes and settings it can collect, along with many of the system settings and values that the most critical for your applications. After the template customization is complete, you can build it by using the Build template button.

The following sections describe some of the possible attributes and values that can be collected by ARE:

- ▶ Files and directories
- ▶ Software requirements
- ▶ Network
- ▶ System Environment
- ▶ Advanced

Files and directories

You can customize the template to verify various files and directories settings:

- ▶ File and Directory Attributes
- ▶ File and Directory Authorities
- ▶ Configuration Files
- ▶ Resource Collector

File and Directory Attributes plug-in

The File and Directory Attributes plug-in verifies attributes such as existence, creation date, last modification date, size, and CCSID. Attributes can be verified for files and directories in IFS, and objects in the Library file system. The precise file and directory attributes that are verified by this plug-in are fully customizable.

This plug-in can also verify that a symbolic link is truly a symbolic link and not a real file. This plug-in is useful to detect cases where a symbolic link was replaced by an actual copy of the file that it is supposed to be a link to.

File and Directory Authorities plug-in

The File and Directory Authorities plug-in verifies authority attributes, such as owner, authorization list, primary group, and private data and object authorities. Authority attributes can be verified for files and directories in IFS, and objects in the Library file system. The precise authority attributes verified by this plug-in are fully customizable.

Configuration Files plug-in

The Configuration Files plug-in verifies that the contents of configuration files are correct.

The following configuration file types are supported:

- ▶ XML configuration files
- ▶ Property files
- ▶ Apache HTTP configuration files

The items to check in a configuration file are fully customizable, as are the types of checks performed (equals, does not equal, contains, exists, exclude, and more). The expected value for configuration items can also be customized, as shown in Figure 16-4.




Figure 16-4 Edit Expected Value

Resource Collector plug-in

You can use the Resource Collector plug-in to collect and package files and objects from the target system into an archive file for more review. This feature allow Application Runtime Expert to verify specific files and attributes and to collect the data that is needed to review and possibly debug an issue on a system.

The plug-in allows specific files and directories from IFS to be collected and packaged into a single archive file. You can also use it to gather native IBM i objects and libraries by saving these native object and libraries into SAVF objects that can then be included in the specified archive file.

Software requirements

You can customize the template to verify various software requirements and prerequisites:

- ▶ IBM i Products
- ▶ PTFs

IBM i products

You can use this feature to select specific IBM i products from the list of all IBM i products. The selected products are added to the list of products to verify in the template.

You can use the Check Product Option (**CHKPRDOPT**) CL command during product verification to provide useful, product-specific information. Select the **CHKPRODOPT** check box to enable this feature, as shown in Figure 16-5.

IBM Application Runtime Expert for i

Home > Plugin Selection and Customization > IBM i Products

IBM i Products

Select the IBM i Products to Verify

Products [Learn more...](#)

Licensed program	Option	Release	Status
5770PT1	3	V7R1M0	*INSTALLED
5770QU1	*BASE	V7R1M0	*INSTALLED
5770RD1	*BASE	V7R1M0	*INSTALLED
5770RD1	10	V7R1M0	*INSTALLED
5770RD1	11	V7R1M0	*INSTALLED
5770RD1	12	V7R1M0	*INSTALLED
5770SM1	*BASE	V7R1M0	*INSTALLED
5770SS1	*BASE	V7R1M0	*INSTALLED
5770SS1	1	V7R1M0	*INSTALLED
5770SS1	2	V7R1M0	*INSTALLED
5770SS1	5	V7R1M0	*INSTALLED
5770SS1	6	V7R1M0	*INSTALLED
5770SS1	7	V7R1M0	*INSTALLED
5770SS1	8	V7R1M0	*INSTALLED

Products to verify [Learn more...](#)

Licensed program	Option	Release	CHKPRDOPT	Severity
5761JV1	11	V6R1M0	☒	Error
5770DG1	*BASE	V7R1M0	☒	Error
5770SS1	3	V7R1M0	☐	Error
5770SS1	12	V7R1M0	☒	Error
5770SS1	30	V7R1M0	☐	Error
5770SS1	33	V7R1M0	☐	Error

Previous Next Filter 7/11

Position to:

Collection name:

OK Cancel

Figure 16-5 Verify IBM i products

When a problem is found during product verification, the problem is added to the IBM i product verification section of the report. The severity of the problem, which determines how the problem is recorded in the report, can be customized for each product verified.

To select an IBM i product that is not installed on the system, click the **Filter** button and select the **Show all products supported by IBM i** option, as shown in Figure 16-6.



Figure 16-6 Filter IBM i products

PTFs

The PTFs plug-in verifies the status of individual and group PTFs on an IBM i system. It also lists all group PTF levels on the system, which is equivalent to the information displayed by the **WRKPTFGRP** CL command. The specific PTFs and group PTFs that are verified are fully customizable. This plug-in can verify that a group PTF level meets both a minimum and recommended level.

Network

You can use this feature to verify various TCP/IP configuration settings, network characteristics, and ports status:

- ▶ Network Configuration
- ▶ Port validation

Network Configuration plug-in

The Network Configuration plug-in verifies various TCP/IP configuration settings and network characteristics. Much of the TCP/IP configuration that is verified are the configuration settings viewable by running the **CFGTCP** CL command.

The primary purpose of this plug-in is to attempt to verify that the system's network configuration is such that applications that require frequent access to the network, such as web applications, can do so in a fast, reliable, and repeatable manner.

Port validation plug-in

The Port validation plug-in verifies if one or more specified ports are active or free on the system. The plug-in can verify if one or more specified ports are restricted.

Ports on the system can be restricted by running **CFGTCP** and selecting option 4 - Work with TCP/IP port restrictions. If an application requires a port that is already in use or is restricted, the application might fail to start or not work correctly.

System Environment

You can use this feature to verify various system environments settings, including network attributes, environment variables, and user IDs:

- ▶ System Values and Network Attributes
- ▶ Environment Variables
- ▶ User IDs
- ▶ Scripts and Commands
- ▶ SQL Query Verifier

System Values and Network Attributes plug-in

The System Values and Network Attributes plug-in verifies system values and that network attributes are configured as a deployment expects. The expected value for a system value or network attribute is customizable, including the capability to specify a single value, list of possible values, a range of possible values, and more.

There is also the option to list, but not check, a value in the report, which is a useful mechanism to automate the collection of system configuration information.

Environment Variables plug-in

The Environment Variables plug-in collects and optionally verifies system-wide environment variables. System-wide environment variables can affect the runtime attributes of any job on the system, so their existence, and their values, are an important part of the environment on which any application runs.

User IDs plug-in

The User IDs plug-in verifies attributes of one or more IBM i user profiles. Over 20 different attributes can be verified, such as special authorities, group profile, supplemental group, profile enabled, and more. The exact attributes to check are customizable, as shown in Figure 16-7.

Home > Plugin Selection and Customization > User IDs > Edit User ID Collection

Edit User ID Collection

Select User IDs to Verify

All user IDs on the system [? Learn more...](#)

QBRMS

QCLUMGT

QCLUSTER

QCOLSRV

QDBSHR

QDBSHRDO

QDESADM

QDESUSR

QDFTOWN

QDIRSRV

>

<

User IDs to verify [? Learn more...](#)

QDESADM

Previous

Next

7/15

Position to:

Collection name:

Attributes to verify:

- CCSID
- Character ID
- Country ID
- Group Authority
- Group Authority Type
- Group Profile
- Job Description
- Language ID
- Limit Capabilities
- Locale
- Message Queue
- Output Queue
- Owner
- Password is *NONE
- Special Authority
- Special Environment
- Status
- Supplemental Group
- User Class

Customize

OK

Cancel

Figure 16-7 Verify User IDs

Scripts and Commands plug-in

You can use the Scripts and Commands plug-in to verify the results of Qshell scripts or CL commands on a remote system. This plug-in provides a powerful and simple way to extend the base verification function that is provided by Application Runtime Expert for i. The results of each verification are customizable, including the capability to specify the source of a result, expected value, a range of expected values, and more.

SQL Query Verifier plug-in

The SQL Query Verifier plug-in can verify database information using SQL statements. An SQL statement can be specified to be run on the target system. The results of this SQL query can be verified. The verification of specific column data or the number of records that are returned can be verified. The complete query results can also be returned as part of the ARE report.

Advanced

ARE offers some advanced features for plug-in customization:

- ▶ Custom plug-ins
- ▶ Other Resources
- ▶ Plug-in Configuration

Custom plug-ins

You can use this feature to manage custom plug-ins that are included in the deployment template. Custom plug-ins are Java classes that can augment the verification that is performed by a deployment template. Anyone can write a custom plug-in, and after the plug-in is created, it can be added to any deployment template.

A custom plug-in plugs into the IBM Application Runtime Expert for i environment, and is run along with the other plug-ins selected and customized using the deployment template editor.

Other Resources plug-in

You can use this feature to manage other resources, which are any items that are not a custom plug-in and are included in the deployment template. Other resources are commonly needed with a custom plug-in. For example, a custom plug-in might require several other Java classes for it to run. These additional Java classes can be added to the deployment template using the Other Resources window.

Plug-in Configuration

There are several advanced plug-in features that can be edited through the Plug-in Configuration window. Additionally, any plug-in that can be used by a template, including plug-ins that are not configured elsewhere in the GUI, can be configured and added to a template using the Plug-in Configuration window.

16.9.3 Application Runtime Expert console

After a template is built, there are two ways it can be used to verify a system:

- ▶ Use the console web user interface.
- ▶ Use a script that can be run from Qshell.

This section describes using the console to verify IBM i systems.

The console is a web user interface that enables a system, or group of systems, to be verified using a deployment template that was created or imported using the deployment template editor.

Systems are verified by the following sequence of events:

- ▶ The console remotely logs in to the system.
- ▶ The console starts the ARE to perform the verification, using the specified deployment template.
- ▶ The results of the verification go back to the console so that they can be reviewed.

Requirement: A valid, enabled user profile name and password for the target system must be provided. The user profile must have *ALLOBJ special authority on the target system because the verification of the target system might involve the inspection of many different files, authorities, PTFs, user profiles, and other items.

The Groups feature provides a way to group one or more systems into a discrete entity. Groups are useful for creating different sets of systems to verify, especially when each set of systems needs a different deployment template for verification, as shown in Figure 16-8.

Console

The IBM Application Runtime Expert for i is capable of verifying multiple IBM i systems using a deployment template created or imported using the deployment template editor. [? Learn more...](#)

Groups:

Default

Production_Systems

Test_Systems

Operations ▾

Systems to verify:

	System name	User ID	Password	Template
☒	9.5.168.119	amonsour	●●●●●●●●	sample ▾
☒	9.172.23.77	amansour	●●●●●●●●	DB2_Extender_Template ▾
☒	9.5.36.131	cstore	●●●●●●●●	sample ▾

Add

Remove

Select all

Deselect all

Runtime properties

Note: Only selected systems will be verified.

Verify systems

View previous results

Launch deployment template editor

Import and view report

Figure 16-8 IBM Application Runtime Expert console

Verifying systems

After the system verification is complete, a Complete icon is shown in the status column for that system and a brief summary of its verification is shown in the result column. If the console failed to perform the verification on a target system, a red icon is shown followed by a brief reason description of why the verification could not be performed, as shown in Figure 16-9.

[Console](#) > [Result](#)

Console

System verification status: 3/3 complete [Refresh](#)

System name	Template	Status	Result
9.5.168.119	DB2_Extender_Template	Complete	No problems found
9.172.23.77	sample	Network connection error	View log
9.5.168.119	sample	Complete	4 error, 1 warning, 1 info

[Summary Report](#)[Detailed Report](#)[XML Report](#)[Download Archive](#)

[Runtime Log](#)[Remote runtime log](#)[Remote stdout](#)[Remote stderr](#)

[Export all report files](#)[View console server logs](#)

[Back](#) [Stop all](#)

Figure 16-9 System Verification Status

More details about the failure can be found by clicking the **View log** link in the result column.

Reports

For systems that were verified, the results column contains a brief summary of the verification results. The summary text in the result column is also a link. This link provides access to the verification reports generated by the ARE core during the verification of the target system, as shown in Figure 16-9.

All three ARE reports (summary, detailed, and XML) are available through the link in the result column. You can also download all the reports in an archive file by clicking the **Download Archive** link.

Summary report

The summary report contains a summary of all problems that are detected during the verification of a system. Each row in the summary table contains the results for a specific plug-in, such as the “Authority Verifier” or “System Value Verifier” plug-ins. The icon directly before the plug-in name indicates the highest severity problem that was found by that plug-in. The other items in each row indicate the number of attributes that are verified by the plug-in, and the number of problems that are found at each severity level (Error, Warning, and Info). The final item in each row, the “Fix actions” column, indicates how many of the detected problems can be fixed directly from the console web interface.

Figure 16-10 and Figure 16-11 show examples of two summary reports.

Console > Result > Summary Report

Summary Report

System name: 9.5.168.119

Remote core version: 1.6.8

Template: DB2_Extender_Template

User ID: AMONSOUR

Summary details [? Learn more...](#)

Plugin name	Attributes verified	Error	Warning	Info	Fix actions
✓ Authority Verifier	596	0	0	0	None
✓ File Attribute Verifier	327	0	0	0	None
✓ Product Verifier	8	0	0	0	None

Back [Summary report \(plain text\)](#)

Figure 16-10 Summary report 1

Summary details [? Learn more...](#)

Plugin name	Attributes verified	Error	Warning	Info	Fix actions
✓ Hello plugin	0	0	0	0	None
✗ SampleAutoFix2	0	1	0	0	Fix actions (1)
✗ SampleFixAction	0	1	0	0	Fix actions (1)
✓ SampleLoadResource	0	0	0	0	None
✓ SampleParameter	0	0	0	0	None
✗ SampleReport	0	2	1	1	None

Figure 16-11 Summary report 2

Detailed report

The detailed report is a text report that contains every status and problem message reported during verification. This report is a complete record of everything that was checked during verification and the result of each check, even if the check did not detect a problem.

XML report

The XML report is an XML formatted report that contains every status and problem message reported during verification. This report is a complete record of everything that was checked during verification and the result of each check, even if the check did not detect a problem. In this regard, the XML report is exactly like the detailed report, except in an XML format instead of plain text.

XML report details: XML reports include information about how to fix detected problems. ARE Core uses an XML report as a guide for automatically fixing detected problems.

Automatic fixes

IBM Application Runtime Expert for i offers an additional important feature, which is the ability to automatically fix problems that are detected by ARE.

The “Fix actions” column in the summary report allows certain problems that were detected during verification to be fixed directly from the console web interface. You can use this convenient feature to solve certain problems without logging in to a different system.

The following list shows some of the problems that ARE can automatically fix:

- ▶ Authority: This category includes ownership, primary group, authorization list, and private authority.
- ▶ User Profiles: Some, but not all, user profile attributes can be fixed.
- ▶ Symbolic Links: If ARE detects that a symbolic link is missing, it can re-create the link.

Important: It is important to understand that only certain types of detected problems can be fixed directly from the console.

Figure 16-12 shows a summary report in which the “Authority Verifier” plug-in has a fix action available.

Summary details ? Learn more...					
Plugin name	Attributes verified	Error	Warning	Info	Fix actions
✖ Authority Verifier	596	1	0	0	Fix actions (1)
✔ Configuration Attribute Verifier	104	0	0	0	None
✔ File Attribute Verifier	327	0	0	0	None
✔ Product Verifier	8	0	0	0	None
✔ User Profile Verifier	20	0	0	0	None

Figure 16-12 Summary report fix action

Clicking the **Fix action** link opens a new window that summarizes all of the problems that are detected by that plug-in that can be fixed directly from the console. You can select which problems to fix by selecting the check box that precedes each problem description, and then clicking the **Fix** button, as shown in Figure 16-13.

Console > Result > Summary Report > Fix Problems

Fix Problems

System name: 9.5.168.119
Problems found by: Authority Verifier

Select	Description
☒	✖ User QDESADM data authorities to object /QSYS.LIB/QDB2TX.LIB/DB2TX.PGM is not what we expect Expected: *RWXDA Actual: *WXDA

▼ Processed problems

Result	Description
This table is empty.	

Figure 16-13 Fix plug-in problem

The console begins the process of fixing the selected problems on the target system. After the problems are fixed, a second table is displayed that shows the fix results.

Other than the console web user interface, ARE can also fix detected problems using a script. For more information, see the “areFix.sh” script section in the document found at:

http://www-03.ibm.com/systems/resources/systems_i_are_script_interfaces.pdf

16.9.4 Application Runtime Expert requirements

The IBM Application Runtime Expert for i product (5733ARE) requires several IBM i products and PTFs to be installed on the system before you install ARE.

Prerequisite products

Here is the list of required software products on IBM i 7.1:

- ▶ 5770SS1 option 3 - Extended Base Directory Support
- ▶ 5770SS1 option 12 - Host Servers
- ▶ 5770SS1 option 30 - Qshell
- ▶ 5770SS1 option 33 - PASE
- ▶ 5761JV1 option 11 - J2SE 6 32 bit
- ▶ 5770DG1 - IBM HTTP Server for i

PTF requirements

The latest Group PTF level must be installed on the system before you install ARE. For up-to-date PTF requirements, go to:

<http://www-03.ibm.com/systems/power/software/i/are/gettingstarted.html>

To benefit from the latest ARE enhancements, you must install the latest PTF. For the latest PTFs, see the ARE support that is found at:

<http://www-03.ibm.com/systems/power/software/i/are/support.html>

16.9.5 More information

For more information about Application Runtime Expert, see *Uncovering Application Runtime Expert*, REDP-4805.



IBM Rational products

This chapter describes the latest Rational products for Power Systems Software and IBM i, including the following topics:

- ▶ IBM Rational Developer for Power Systems Software
- ▶ IBM Rational Team Concert
- ▶ IBM Rational Developer for i for SOA Construction
- ▶ IBM Rational Development Studio for i
- ▶ IBM Rational Open Access: RPG Edition
- ▶ IBM Rational Application Management Toolset for ii

17.1 IBM Rational Developer for Power Systems Software

IBM Rational Developer for Power Systems Software is an integrated development environment for developing applications to run on IBM i, AIX, and Linux. It supports the common programming languages that are used on Power Systems, including Java, C/C++ on Linux, C/C++ and COBOL on AIX, and RPG, COBOL, DDS, and CL on IBM i.

When used in combination with IBM Power Systems compilers and IBM Rational Team Concert, IBM Rational Developer for Power Systems Software provides a comprehensive application development environment, including compilers, development tools, and collaborative application lifecycle management.

The following sections focus on Rational Developer for Power Systems Software features for IBM i platform:

- ▶ RPG and COBOL Development Tools
- ▶ Rational Team Concert client integration for IBM i
- ▶ Enhancements in Version 8.0.3
- ▶ Version 8.x Fix Packs
- ▶ Migration to Rational Developer for Power Systems V8.5
- ▶ Enhancements in Version 8.5


17.1.1 RPG and COBOL Development Tools

Rational Developer for Power Systems offers the following RPG and COBOL Development Tools:

- ▶ Remote System Explorer
- ▶ IBM i Projects
- ▶ Application Diagram
- ▶ Screen Designer
- ▶ Report Designer
- ▶ Integrated i Debugger
- ▶ IBM i Web Services and Java Tools

The Remote System Explorer, which is shown in Figure 17-1, is an enhanced and more flexible workstation version of the Programming Development Manager. It is a workbench perspective that provides access to all development resources of your IBM i server.

- ▶ Remote Connection to IBM i server
- ▶ Manage IBM i objects
- ▶ Manage library lists
- ▶ Manage jobs
- ▶ Manage commands and shells
- ▶ Manage user actions
- ▶ Manage objects in “Object Table View”
- ▶ Editing, compiling, and debugging applications



IBM i Projects, which is shown in Figure 17-2, allow for disconnected development. A network connection is required only when code updates or a build are needed, or when you must view remote resources for a project.

In disconnected mode, you work on files locally, and upload them to the server after you finish. While you work in disconnected mode, you can still check source code for syntax and semantic error and connect only to submit a compile when you are ready to create the program object.




Figure 17-2 i Projects perspective

Application Diagram provides a graphical view of the different resources in an IBM i native application and their relationships to each other.

There are two different diagrams that you can look at in the Application Diagram view:

- ▶ **Source Call Diagram**
This diagram takes ILE RPG, ILE COBOL, and CL sources as input and displays a call graph that shows the subroutine and procedure calls.
- ▶ **Program Structure Diagram**
This diagram takes program and service program objects as input and displays the binding relationships between them and the modules that are bound to each program and service program.

Screen Designer

Screen Designer is now an official component. It provides the capability to graphically design and modify the content of DDS display files. Screen Designer, shown in Figure 17-3, provides an integrated palette for easy access to design items, and a preview page.




Figure 17-3 Screen Designer perspective

Report Designer

You can use the Report Designer to graphically design and modify the content of DDS printer files. The Report Design window provides an integrated palette for easy access to design items.

You can use the Report Designer to group individual records and see how this group of records would appear on the printed page. In addition, you can specify default data for each output field, and specify which indicators are on or off.

Integrated i Debugger

You can use Integrated i Debugger to debug an application that is running on an IBM i system. It provides an interactive graphical interface that makes it easy to debug and test your IBM i programs.

Integrated i Debugger provides many program debugging control features, including breakpoints setup, watching variables, step through program instructions, examining variables, and examining the call stack.

IBM i Web Services and Java Tools

You can use IBM i web development tools to create new e-business applications that use a web-based front end to communicate with the business logic in an ILE or non-ILE language program on an IBM i server.

The Web Services wizard works in the context of a web project and allows for creation, deployment, testing, generation of a proxy, and publication to a Universal Description, Discovery, and Integration (UDDI) registry of Web Services.

Requirement: The IBM i Web Services and Java Tools can be installed only if the appropriate prerequisite is installed.

17.1.2 Rational Team Concert client integration for IBM i

Rational Developer for Power Systems Software provides support for integration with the Rational Team Concert client for IBM i.

When Rational Developer for Power Systems Software is used in combination with IBM Power Systems compilers and Rational Team Concert, it provides a comprehensive application development environment, including compilers, development tools, and collaborative application lifecycle management.

For more information about Rational Team Concert, and integration with Rational Developer for Power Systems Software, see 17.2, “IBM Rational Team Concert” on page 611.

Requirement: The IBM Rational Team Concert V3.0 client product must be installed before you install the client integration.

17.1.3 Enhancements in Version 8.0.3

The Rational Developer for Power Systems Software solution offers integrated development environments for the major development workloads on IBM i, AIX, and Linux.

Here are enhancements for IBM i in Version 8.0 and Version 8.0.3:

- ▶ New Power Tools
- ▶ Secure connections to IBM i development platforms using a Secure Sockets Layer (SSL) connection
- ▶ Usability improvements to the Screen Designer, Report Designer, and Remote Systems LPEX Editor functions
- ▶ New fixed-term and token licensing options

New Power Tools

Rational Developer for Power Systems V8.0 adds new feature that combines IBM i development tools with IBM Rational Application Developer Standard Edition for WebSphere Software V8.

This new enhancement is ideal for organizations that are developing solutions that are integrating applications that are written in RPG and COBOL on IBM i with Web Services and web front ends that use Java, Java EE, Web 2.0, SOA, or Portal.

Power Tools allow developers to have all their development tools integrated into one work environment. Power Tools also reduce the operation costs by deploying a single development workbench image to all developers.

Version 8.0 enhancements for AIX and Linux

For a list of Rational Developer for Power Systems enhancements on AIX and Linux, see the following product announcement letter website:

http://www-01.ibm.com/common/ssi/ShowDoc.jsp?docURL=/common/ssi/rep_ca/4/897/ENUS211-344/index.html&lang=en

17.1.4 Version 8.x Fix Packs

Information about any required Power Systems server PTFs can be found by completing the following steps in IBM Rational Developer for Power Systems Software:

1. Open the Remote System Explorer perspective by clicking **Window** → **Open Perspective** → **Other** → **Remote System Explorer**.
2. Create a connection to your IBM i by expanding **New Connection** and double-clicking **IBM i**.
3. In Remote System Explorer, expand your connection, right-click **Object**, and click **Verify Connection**, as shown in Figure 17-4.




Figure 17-4 Verify connection

4. A window opens that shows which required PTFs are installed on the system and which ones are missing, as shown in Figure 17-5.




Figure 17-5 Verify PTFs

For information about the latest product Fix Packs, see the Rational Developer for Power Systems Software Support website (Downloads tab):

http://www-947.ibm.com/support/entry/portal/Overview/Software/Rational/Rational_Developer_for_Power_Systems_Software/

17.1.5 Enhancements in Version 8.5

Here are the Version 8.5 enhancements:

- For IBM i:
 - Live outline view for RPG applications
 - Enhancements to the i Project Perspective
- For all operating systems:
 - Expanded set of data development tools in the Data Perspective
 - Support for build specifications for multiple contexts (such as for multiple promotion levels or even multiple target platforms)
 - Service reliability enhancements
 - 64-bit desktop operating system support (the product can now use a larger memory space)

Version 8.5 enhancements for AIX and Linux

For a list of Rational Developer for Power Systems enhancements on AIX and Linux, and more IBM i details, see the following product announcement letter website:

<http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=an&subtype=ca&appname=g pateam&supplier=877&letternum=ENUSZP12-0219>

17.1.6 Migration to Rational Developer for Power Systems Software V8.5

You can migrate projects to Rational Developer for Power Systems V8.5 from previous versions. The migration process is seamless and transparent; you must use only your existing workspace and continue working with your projects with the new software version.

Migration to V8.5 can be done through Rational Developer for Power Systems Software. Migration is supported from Version 7.5.x, Version 7.6.x, or Version 8.0.x.

Migration of projects from earlier releases is not supported because of additional features that were added.

17.2 IBM Rational Team Concert

The Rational solution for Collaborative Lifecycle Management (CLM) provides integrations across the Change and Configuration Management, Requirements Management, and the IBM Quality Management Jazz™-based applications, to connect the work of analysts with development and test teams.

Rational Team Concert is the core of the Rational CLM solution. Rational Team Concert is a team collaboration tool that supports cross-platform development and features native hosting of the Jazz Team Server. Rational Team Concert includes an integrated set of collaborative software delivery lifecycle tools for development, including source control, change management, and build and process management.

Rational Team Concert has an open, extensible architecture that supports a broad range of desktop clients, IDEs, languages, and platforms, as shown in Figure 17-6.




Figure 17-6 Rational Team Concert architecture

17.2.1 Integration with Rational Developer for Power Systems Software

Rational Team Concert integrates with Rational Developer for Power Systems Software to provide IBM i developers an integrated collaborative application development environment. This integrated solution provides the value of the team collaboration capabilities of Rational Team Concert with the individual development environment of Rational Developer for Power Systems Software.

Using Rational Team Concert and Rational Developer for Power Systems Software together, software development teams can develop IBM i and AIX applications using the tools that are provided by Rational Developer for Power Systems Software and the planning, team collaboration, build, source control management, defect tracking, and deployment tools that are provided by Rational Team Concert.

Using the IBM i Projects perspective available with Rational Developer for Power Systems Software, Rational Team Concert and Rational Developer for Power Systems Software work together so that you can share and modify files that are managed by Jazz based source control, in addition to files on the remote IBM i system.

Installing and configuring the integration

Extra Installation and configuration steps are required to integrate Rational Developer for Power Systems Software and Rational Team Concert.

There are two ways that you can install the two products by using IBM Installation Manager:

- ▶ Installing the two products at the same time or in the following order:
 - a. Install Rational Team Concert.
 - b. Install Rational Developer for Power Systems Software with Rational Team Concert client integration for IBM i feature.
- ▶ Installing the products in the following order:
 - a. Install Rational Developer for Power Systems Software.

Remember: You cannot install the Rational Team Concert client integration for IBM i feature, because the feature is not available if Rational Team Concert is not present.

- b. Install Rational Team Concert.
- c. Start IBM Installation Manager and use the **Modify** option to add Rational Team Concert client integration for IBM i feature to Rational Developer for Power Systems Software product, as shown in Figure 17-7.




Figure 17-7 Rational Team Concert Client integration

Sharing i Projects in Rational Team Concert

Sharing IBM i Projects in Rational Team Concert is no different from sharing any other type of Eclipse project. You can manage Power Systems Software source code with the IBM i Projects feature.

To make your IBM i Project available to other team members, complete the following steps:

1. From the menu for your i Project, click **Team** → **Share Project**, as shown in Figure 17-8.




Figure 17-8 Share IBM i project

2. Select **Jazz Source Control**, as shown in Figure 17-9, and click **Next**.




Figure 17-9 Jazz Source Control Repository

3. Specify the repository workspace to use.
4. Click **Finish**.

Consideration: If your IBM i Project contains any file types or folders that are ignored by Eclipse, the wizard prompts you to review them.

The IBM i Project is now added to a component in a repository workspace. Changes that you make in the i Project are copied to the repository workspace when you check them in.

17.2.2 What is new in Rational Team Concert V3.0.1

Here are some highlights of what is new in Rational Team Concert V3.0.1:

- ▶ Ability to define work item and SCM read access per team area
- ▶ Rich text support for work items
- ▶ Ability to delete work items
- ▶ Filtering to quickly locate work items in plans and dashboards
- ▶ Automate delivery of changes on build success
- ▶ Work item-based promotion model¹
- ▶ Enhanced deployment support for IBM System z¹
- ▶ New Lotus Connections Integration
- ▶ Enhanced Rational Team Concert Client for Visual Studio IDE with CLM support

¹ Available in Rational Team Concert Developer for IBM Enterprise Platforms.

- ▶ Role-based licensing with built-in read access across Rational Team Concert, Rational Requirements Composer, and Rational Quality Manager
- ▶ Cross product reporting and lifecycle queries
- ▶ Flexible deployment options for CLM applications

17.2.3 What is new in Rational Team Concert V4.0

Here are some of the new features of Rational Team Concert V4.0:

- ▶ CLM 2012: Clustering for HA, incremental upgrades, and faster upgrades
- ▶ Planning: Planning across projects, support for Kanban, and an improved MS-Project Import/Export function.
- ▶ Work items: More flexible access control, user input on queries, and management of external HTTP data sources
- ▶ Source Control Management (SCM):
 - Improved variant management: You can locate change sets across release streams.
 - Flexible workspace layout and control: You can load rules from files.
 - Produce new version of anything outside the IDE through Windows Shell SCM integration.
 - Improved CLI: There are more CLI commands and enhanced scripting through the Enhanced ISPF client
 - There is better integration of SCM with MSCCI tools.
 - You can load with dependencies through sparse loading.²
- ▶ Build:
 - There is ready for use support for Jenkins and Hudson.
 - There is improved build, promote, and deploy functionality for IBM System z® and IBM i

For more information about these topics, see the following resources:

- ▶ Projects on jazz.net, found at:
<http://jazz.net/projects/rational-team-concert/>
- ▶ A wiki that includes tutorials and articles on various topics, found at:
<https://jazz.net/wiki/bin/view/Main/RTCHome>
- ▶ An end to end tutorial, found at:
<https://jazz.net/wiki/bin/view/Main/RTCpHome>
- ▶ ALM on IBM developerWorks, found at:
<https://www.ibm.com/developerworks/rational/community/alm.html>

² Available in Rational Team Concert Developer for IBM Enterprise Platforms.

17.3 IBM Rational Developer for i for SOA Construction

Rational Developer for i for SOA Construction (RDi SOA) is a comprehensive application development and modernization solution. It is based on Eclipse. It provides a powerful, flexible, and extensible workbench environment for IBM i development with support for RPG, COBOL, CL, DDS, SQL, C++, Java, and EGL (an open source business application language that was developed on Eclipse).

RDi SOA Version 8.5 combines the following products:

- ▶ IBM Rational Developer for Power Systems Software: RPG and COBOL Development Tools for i V8.5
- ▶ IBM Rational Business Developer V8.5
- ▶ IBM Rational Host Access Transformation Services (HATS) Toolkit V8.5

It accelerates IBM i development teams who must perform the following actions:

- ▶ Create, maintain, and enhance RPG and COBOL applications on IBM i.
- ▶ Extend RPG and COBOL applications to the web and to web services to enable their integration with other applications, implement business-to-business integrations, and provide users with a richer and more productive interface.
- ▶ Reuse existing applications and integrate them into automated business processes and new applications.

RDi SOA V8.5 helps you with the following tasks:

- ▶ Reduce development costs.
- ▶ Simplify the development of service-oriented architecture (SOA) applications.
- ▶ Modernize and integrate existing systems.
- ▶ Reuse, integrate, and extend valuable existing IT assets (such as RPG and COBOL programs and service programs) using wizards that generate web services interfaces to the existing programs.

Availability of features: Although Host Access Transformation Services (HATS) Toolkit V8.5 is available as a download at no additional cost, it is included in Rational Developer for i for SOA Construction for your convenience. Rational HATS for 5250 Applications (which requires a separate runtime license) offers richer user interfaces and flexible web services options.

For more information, see the Rational Developer for i for SOA Construction website at:

<http://www-01.ibm.com/software/awdtools/developer/rdisoa/>

17.4 IBM Rational Development Studio for i

IBM Rational continues to have a tight collaboration with IBM Systems Technology Group (STG), IBM Software Group (SWG), and IBM Research Group to provide compilers that use the underlying hardware and operating systems for the Power architecture. The RPG, COBOL, C, and C++ ILE compilers for IBM i are enhanced for IBM i 7.1 with new features and functions. On April 13, 2010, the WebSphere Development Studio for IBM i was rebranded into Rational Development Studio for i (5770-WDS), as illustrated in Figure 17-10.




Figure 17-10 Rebranding of WebSphere Development Studio for IBM i to Rational Development Studio for i

Application Development ToolSet (ADTS) and Original Program Model (OPM) Compilers were previously stabilized. However, Rational Development Studio for i V7.1 does include enhancements for the ILE RPG and ILE COBOL compilers.

17.4.1 Source code protection option

Protecting your code is essential. Unfortunately, there are times when one of your customers encounters a problem. The easiest way to figure out the problem is to put a debuggable version of your code onto the system, but leaves your code unprotected. So you either expose your code or figure out another way to diagnose the problem.

In Version 7.1, the ILE compilers (RPG, COBOL, CL, C, and C++) and precompilers have a new parameter that you can use to encrypt your debug views. You can ship debuggable code and know that your code is not exposed.

With the **DBGENCKEY** compiler option, you can specify the encryption key that is used to encrypt the program source that is embedded in debug views. The debugger requires the user to enter the debug encryption key before the views are decrypted. The length of the key can be 1 - 16 bytes. A key of length of 1 - 15 bytes is padded to 16 bytes with blanks for the encryption. Specifying a key of length zero is the same as specifying *NONE.

If the key contains any characters that are not invariant over all code pages, the user must ensure that the target system uses the same code page as the source system. Otherwise, the key might not match and decryption might fail. If the encryption key must be entered on systems with different code pages, the key should be made of characters from the (EBCDIC) invariant character.

To use this capability to protect data source code while allowing the usage of the debug view, complete the following steps:

1. Encrypt the debug view so that the debug view is only visible if the person knows the encryption key by running the following command:
`CRTBNDRPG MYPGM DBGENCKEY('my secret code')`
2. Then, either run **STRDBG MYPGM DBGENCKEY('my secret code')** or **STRDBG MYPGM** and wait to be prompted for the encryption key.

17.4.2 ILE RPG compiler

The following list details the ILE RPG compiler enhancements available with Rational Development Studio for i V7.1:

- ▶ XML processing support with new ILE RPG options for XML-INTO for improved flexibility
- ▶ Increased developer productivity with ILE RPG built-in support for searching and sorting data structure arrays
- ▶ New built-in functions
- ▶ Support for ALIAS names
- ▶ Optional prototypes
- ▶ Implicit parameter conversion between alphanumeric, USC-2, and graphic data types

Using Source Entry Utility (SEU): For customers using SEU to edit the ILE RPG source, the syntax checkers do not recognize any features that were added after V6R1.

New options for XML-INTO

There are two new options for XML-INTO (also available through a PTF for IBM i 6.1):

datasubf	You can use this option to name a subfield that receives the text data for an XML element that also has attributes. Here is an example: <code><emp type="regular" ID="13573">John Smith</emp></code>
countprefix	This option reduces the need for you to specify the allowmissing=yes option. It specifies the prefix for the names of the additional subfields that receive the number of RPG array elements or non-array subfields that are set by the XML-INTO operation.

New options for XML-INTO are shown in Example 17-1.

Example 17-1 New options for XML-INTO

D* RPG coding to get the information with one XML-INTO operation:
D* xml file content: `<emp type="regular" id="13573">John Smith</emp>`

```

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

```

```

/free
XML-INTO emp %xml('emp.xml' :
                  : 'datasubf=value doc=file');
// emp.id = 13573
// emp.type = 'regular'
// emp.value = 'John Smith'

```

17.4.3 Sorting and searching data structures

You can sort and search data structure arrays using one of the subfields as a key, as shown in Example 17-2.

Example 17-2 Sort and search data structures

```

#Sort a data structure array using one subfield as a key
// sort by name
SORTA info(*).name;

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

#Search a data structure array using one subfield as a key
// search for a name
pos = %LOOKUP('Jack' : info(*).name);

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

```

An array can be sorted in ascending order by using SORTA(A) and descending order by using SORTA(D). The array cannot be a sequenced array (ASCEND or DESCEND keyword), as shown in Example 17-3.

Example 17-3 Sort an array in descending order

```

//Sort the salary in descending order
SORTA(D) salary;

```

New and updated built-in functions

There are several updated built-in functions that simplify string management tasks:

%SCANRPL	Scans and replaces values.
%LEN	Gets and sets variable lengths.
%PARMNUM	Gets the parameter position.

The %SCANRPL built-in function scans for all occurrences of a value within a string and replaces them with another value, as shown in Example 17-4.

Example 17-4 %SCANRPL

```

// Replace NAME with 'Tom'
string1 = 'See NAME. See NAME run. Run NAME run.';
string2 = %ScanRpl('NAME' : 'Tom' : string1);
// string2 = 'See Tom. See Tom run. Run Tom run.'

```

The %LEN function can be used with a new optional second parameter *MAX to obtain the maximum number of characters for a varying-length character, UCS-2, or Graphic field.

The %PARMNUM built-in function returns a parameter's position in the parameter list, as shown in Example 17-5.

Example 17-5 %PARMNUM

```
CEEDOD (2 : more parms);           // hard to understand
CEEDOD (%PARMNUM(city) : more parms); // better
```

Soft-coding the parameter's number makes the code easier to read and maintain.

17.4.4 ALIAS support

The prototypes are supported on D specs for any externally described data structure. There is the following ALIAS support:

- ▶ As shown in Example 17-6, you can use the ALIAS keyword on a Definition specification to indicate that you want to use the alternative names for the subfields of externally described data structures.
- ▶ You can use the ALIAS keyword on a File specification to indicate that you want to use the alternative names for LIKERECD data structures that are defined from the records of the file.

Example 17-6 ALIAS keyword

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

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

Optional prototypes

If a program or procedure is not called by another RPG module, it is optional to specify the prototype. The prototype can be omitted for the following types of programs and procedures:

- ▶ A program that is intended to be used only as an exit program or as the command-processing program for a command.
- ▶ A program that is intended to be called only from another programming language.
- ▶ A procedure that is not exported from the module.
- ▶ A procedure that is exported from the module but intended to be called only from another programming language.

Here are some programs and procedures that do not require a prototype:

- ▶ An exit program, or the command-processing program for a command
- ▶ A program or procedure that is never intended to be called from RPG
- ▶ A procedure that is not exported from the module

Example 17-7 shows an example of an optional prototype.

Example 17-7 Optional prototypes

```

H main(hello)
P hello          b
D                pi          extpgm('HELLO')
D  name          10a      const
  /free
    sayHello();
...

P sayHello      b
  /free
    dsply ('Hello ' + name);

```

Implicit unicode conversion for parameters

Implicit CCSID conversion is now supported for prototyped parameters that are passed by VALUE and by a read-only reference (CONST). This conversion reduces the number of code changes that must be made when a database field is changed from alphanumeric or DBCS to unicode (UCS-2 or UTF-16).

In Example 17-8, there is only a “makeTitle” procedure with a UCS-2 parameter and a return value. If the passed parameter is alpha or DBCS, it is converted to UCS-2 on the call. The procedure works with the UCS-2 parameter and returns a UCS-2 value. This returned value can then be converted on assignment to alpha or DBCS, if necessary.

Example 17-8 Implicit unicode conversion for parameters

```

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

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

```

17.4.5 Performance improvement when returning large values

To obtain faster return values, you can now define a procedure with the RTNPARM keyword (see Example 17-9) to handle the return value as a hidden parameter. The RTNPARM keyword applies both to a prototype definition and to a procedure-interface definition.

Example 17-9 RTNPARM keyword

```

D getFileData    pr          a  varying len(1000000)
D                                     rtnparm
D  file          a  const varying len(500)
D data          S  a  varying len(1000)
D  file          a  const varying len(500)
D data          S  a  varying len(1000)
  /free
    data = getFileData ('/home/mydir/myfile.txt');

```

When a procedure is prototyped to return a large value, especially a large varying value, the performance for calling the procedure can be improved by defining the procedure with this keyword.

The impact on performance because of the RTNPARM keyword varies from having a small negative impact to having a large positive impact. There can be a small negative impact when the prototyped return value is relatively small, such as an integer, or a small data structure. There is improvement when the prototyped return value is a larger value, such as a 32767-byte data structure. The performance improvement is most apparent when the prototyped return value is a large varying length string, and the actual returned value is relatively small. For example, the prototype defines the return value as a 1 million byte varying length character string, and the value 'abc' is returned.

Using RTNPARM for a procedure prototype can also reduce the amount of automatic storage that is required for other procedures that contain calls to that procedure. For example, if procedure MYCALLER contains a call to procedure MYPROC that returns a large value, procedure MYCALLER requires additional automatic storage (even if MYCALLER does not actually call procedure MYPROC at run time). In certain cases, procedure MYCALLER cannot compile because of excessive automatic storage requirements; in other cases, MYCALLER is not able to be called because the total automatic storage on the call stack exceeds the maximum. Using RTNPARM avoids this problem with additional automatic storage.

RTNPARM considerations:

- ▶ The additional parameter is passed as the first parameter.
- ▶ The %PARMS and %PARMNUM built-in functions include the additional parameter in the parameter count. When the RTNPARM keyword is specified, the value that is returned by %PARMNUM is one higher than the apparent parameter number.
- ▶ When you call APIs that require a parameter number, such as CEEDOD or CEETSTA, you must account for the extra first parameter. For example, if your procedure has three parameters, and you want to find the length of the third parameter as it appears in your parameter list, you must ask for information about the fourth parameter. If you use the %PARMNUM built-in function to return the correct parameter number for calling these APIs, you do not need to worry about manually determining the correct parameter number.
- ▶ When the calling procedure is written in a language other than RPG, the caller must code the call as though the procedure has no return value and there is an additional first parameter that is passed by reference with the same type as the RPG return value.
- ▶ Similarly, when the called procedure is written in a language other than RPG, the procedure must be coded without a return value, and having an additional first parameter that is passed by reference with the same type as the RPG return value.
- ▶ When RTNPARM is specified for the procedure, the maximum number of prototyped parameters is 398.
- ▶ The RTNPARM keyword is not allowed for a Java method call.

Teraspace storage model

Definition: *Teraspace* is a large temporary space that is local to a job. A teraspace provides a contiguous address space but might consist of many individually allocated areas, with deallocated areas in between. Teraspace exists no longer than the time between job start and job end. A teraspace is not a space object, which means that it is not a system object, and that you cannot refer to it by using a system pointer. However, teraspace is addressable with space pointers within the same job.

The teraspace storage model has the following features:

- ▶ Much higher limits for automatic storage.
- ▶ Can compile *CALLER programs with **STGMDL(*INHERIT)** so that they can be called from either single-level or teraspace programs.
- ▶ RPG's %ALLOC and %REALLOC can allocate teraspace with a much higher limit.
- ▶ Teraspace allocations are the default in the teraspace storage model.
- ▶ You can specify H-spec ALLOC(*TERASPACE) to have teraspace allocations in any storage model.

17.4.6 ILE COBOL compiler

Here are the ILE COBOL compiler enhancements available with Rational Development Studio for i V7.1:

- ▶ The new COMP-5 data type is supported.
- ▶ You can use a non-numeric literal in the VALUE clause for a national data item.
- ▶ There are new **PROCESS** statement options:
 - **ACTGRP**
This option is now available as a **PROCESS** statement parameter with the option values of STGMDL, NEW, or CALLER.
 - **NEVEROPTIMIZE**
This option is now available as a **PROCESS** statement option.
 - **STGMDL**
This option is now available as a **PROCESS** statement parameter with the option values of INHERIT, SNGLVL, or TERASPACE.
 - **XMLGEN**
This option is now available as a **PROCESS** statement parameter with the option values of NOKEEPFILEOPEN / KEEPFILEOPEN or NOASSUMEVALIDCHARS / ASSUMEVALIDCHARS.

Using SEU: For customers using SEU to edit ILE COBOL source, the syntax checkers do not recognize any features that were added after V6R1.

COMPUTATIONAL-5 or COMP-5 now supported

COBOL supports the COMP-4 binary type, which does not handle the full range of the binary value. For example, a 2 byte COMP-4 supports values -9999 - 9999, when a 2-byte integer supports values -32768 - 32767. The **NOSTDTRUNC** process option can be specified to force all COMP-4 values to be treated as true integers.

The COMP-5 type is a true integer. This type is a native binary data type that is supported by the **USAGE** clause. COMP-5 data items are represented in storage as binary data, and can contain values up to the capacity of the native binary representation (2, 4, or 8 bytes).

When numeric data is moved or stored into a COMP-5 item, truncation occurs at the binary field size rather than at the COBOL picture size limit. When a COMP-5 item is referenced, the full binary field size is used in the operation.

COMP-5 is supported by COBOL on IBM System z. This support enhances portability to or from COBOL on other IBM platforms and operating systems.

Table 17-1 shows the equivalent SQL data types for the COBOL COMP-5 data type:

Table 17-1 SQL equivalent data types

COBOL data type	SQL data type	SQL description
01 name PIC S9(4) COMP-5.	SMALLINT	16-bit signed integer
01 name PIC S9(9) COMP-5.	INTEGER	32-bit signed integer
01 name PIC S9(18) COMP-5.	BIGINT	64-bit signed integer

- Encrypted debug view (see 17.4.1, “Source code protection option” on page 618).

Allows programmers to include a debug view with their application that is only visible with an encryption key.

OPTIMIZE(*NEVER) support to compile large programs

The new COBOL compiler command parameter **OPTIMIZE(*NEVER)** allows large COBOL programs to be compiled. This parameter is present in both **CRTBNDCBL** and **CRTCBLMOD** CL commands.

The **PROCESS** option **NEVEROPTIMIZE** is also added, and can be specified within the COBOL source file.

OPTIMIZE(*NEVER) and **NEVEROPTIMIZE** reduces the size of the generated code by preventing the COBOL compiler from generating the information necessary to optimize the program.

Teraspace storage model

The storage model for a program / module can now be specified using the new **CRTBNDCBL** / **CRTCBLMOD** parameter **STGMDL** with the option values:

- ***SNGLVL**

This value specifies that the program / module is to be created with a single-level storage model.

- ***TERASPACE**

This value specifies that the program / module is to be created with the teraspace storage model.

- ***INHERIT**

This value specifies that the program / module is to inherit the storage model of its caller so that they can be called from either single-level or teraspace programs.

Additionally, the activation group parameter **ACTGRP** on the **CRTBNDCBL** command now has a new default option value. When you specify **STGMDL(*TERASPACE)**, the program is activated in the **QILETS** activation group. For all other storage models, the program is activated in the **QILE** activation group when it is called.

XML GENERATE performance improvements

Performance improvements were made for **XML GENERATE** when the **APPEND** option is specified. Users who have many data records to be appended to a data structure or to a stream file benefit from these changes. The improvements include the addition of the new **PROCESS** statement parameter **XMLGEN** with the following option values:

- ▶ **NOKEEPFILEOPEN / KEEPFILEOPEN**

Specify **KEEPFILEOPEN** to indicate that the XML stream file is to be left open and not closed when the **XML GENERATE** statement is complete, so that subsequent **XML GENERATE FILE-STREAM APPEND** statements can quickly append data to the stream file.

- ▶ **NOASSUMEVALIDCHARS / ASSUMEVALIDCHARS**

Specify **ASSUMEVALIDCHARS** to have **XML GENERATE** bypass the checking for special characters (less than "<", greater than ">", ampersand "&", and the single and double quotation mark symbols), and for characters that are not supported by XML that requires being generated as hexadecimal. Otherwise, normal checking is done with the default **NOASSUMEVALIDCHARS**.

Example 17-10 shows the usage of a **PROCESS** statement with **XML GENERATE**.

Example 17-10 Improved performance for XML GENERATE

```
PROCESS XMLGEN(KEEPFILEOPEN).
```

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

17.4.7 ILE C and ILE C++

For ILE C and ILE C++ compilers, the following additions and modifications were made in IBM i 7.1:

- ▶ ILE C/C++ predefined macros
- ▶ **do_not_instantiate** and **namemanglingrule** pragmas
- ▶ ILE C++ decimal floating-point support
- ▶ ILE C++ compiler other enhancements
 - **__func__** predefined identifier.
 - Hexadecimal floating-point literals
 - Variable length arrays
 - Empty arguments for function-like macros
 - Variable number of arguments for function-like macros

For more information about Rational Development Studio for i, go to:

<http://www-01.ibm.com/software/rational/products/devstudio/i/>

Predefined macros

These macros can be grouped either for C or for C++. Most of them are new, and others were modified.

C macros

Here are the predefined macros for C:

- ▶ `__C99_CPLUSCMT` indicates support for C++ style comments. You can define it when the **LANGLVL(*EXTENDED)** compiler option is in effect.
- ▶ `__IBMC__` indicates the version of the C compiler. It returns an integer of the form *VRM*, where V represents the version, R the release, and M the modification level. For example, using the IBM i 7.1 compiler with the **TGTRLS(*CURRENT)** compiler option, `__IBMC__` returns the integer value 710.
- ▶ `__ILEC400__` indicates that the ILE C compiler is being used.
- ▶ `__ILEC400_TGTVRM__` is functionally equivalent to the `__OS400_TGTVRM__` macro.
- ▶ `__SIZE_TYPE__` indicates the underlying type of `size_t` on the current platform. For IBM i, it is unsigned int.

C++ macros

Here are the predefined macros for C++:

- ▶ `__BOOL__` indicates that the `bool` keyword is accepted.
- ▶ `__cplusplus98__interface__` can be defined when the **LANGLVL(*ANSI)** compiler option is specified.
- ▶ `__C99_COMPOUND_LITERAL` indicates support for compound literals and can be defined when the **LANGLVL(*EXTENDED)** compiler option is in effect.
- ▶ `__C99_FUNC__` indicates support for the `__func__` predefined identifier and can be defined when the **LANGLVL(*EXTENDED)** compiler option is in effect.
- ▶ `__C99_HEX_FLOAT_CONST` indicates support for hexadecimal floating constants and can be defined when the **LANGLVL(*EXTENDED)** compiler option is in effect.
- ▶ `__C99_PRAGMA_OPERATOR` indicates support for the `_Pragma` operator and can be defined when the **LANGLVL(*EXTENDED)** compiler option is in effect.
- ▶ `__C99_RESTRICT` indicates support for the C99 restrict qualifier and can be defined when the **LANGLVL(*EXTENDED)** compiler option is in effect.
- ▶ `__C99_VARIABLE_LENGTH_ARRAY` indicates support for variable length arrays and can be defined when the **LANGLVL(*EXTENDED)** compiler option is in effect.
- ▶ `__IBMCPP__` indicates the version of the AIX Xpertise Library C++ compiler upon which the ILE C++ compiler is based. It returns an integer that represents the compiler version. For example, using the IBM i 7.1 compiler with the **TGTRLS(*CURRENT)** compiler option, `__IBMCPP__` returns the integer value 900, which means the ILE C++ compiler is based on the Xpertise Library C++ V9.0 compiler.
- ▶ `__IBM__ALIGN` indicates support for the `__align` specifier.
- ▶ `__IBM_ATTRIBUTES` indicates support for type, variable, and function attributes and can be defined when the **LANGLVL(*EXTENDED)** compiler option is in effect.
- ▶ `__IBM_COMPUTED_GOTO` indicates support for computed GOTO statements and can be defined when the **LANGLVL(*EXTENDED)** compiler option is in effect.

- ▶ `__IBM_EXTENSION_KEYWORD` indicates support for the `__extension__` keyword and can be defined when the **LANGLVL(*EXTENDED)** compiler option is in effect.
- ▶ `__IBM_LABEL_VALUE` indicates support for labels as values and can be defined when the **LANGLVL(*EXTENDED)** compiler option is in effect.
- ▶ `__IBM_LOCAL_LABEL` indicates support for local labels and can be defined when the **LANGLVL(*EXTENDED)** compiler option is in effect.
- ▶ `__IBM_MACRO_WITH_VA_ARGS` indicates support for variadic macro extensions and can be defined when the **LANGLVL(*EXTENDED)** compiler option is in effect.
- ▶ `__NO_RTTI__` can be defined when the **OPTION(*NORTTI)** compiler option is in effect.
- ▶ `__OPTIMIZE__` indicates the level of optimization in effect. The macro is undefined for `OPTIMIZE(10)`. For other `OPTIMIZE` settings, the macro is defined with 2 for `OPTIMIZE(20)`, 3 for `OPTIMIZE(30)` and 4 for `OPTIMIZE(40)`.
- ▶ `__RTTI_DYNAMIC_CAST__` can be defined when the **OPTION(*RTTIALI)** or **OPTION(*RTTICAST)** compiler option is specified.
- ▶ `__RTTI_TYPE_INFO__` can be defined when the **OPTION(*RTTIALI)** or **OPTION(*RTTITYPE)** compiler option is specified.

C and C++

Here are the predefined macros for C and C++:

- ▶ `__BASE_FILE__` indicates the fully qualified name of the primary source file.
- ▶ `__IBM_DFP__` indicates support for decimal floating-point types and can be defined when the **LANGLVL(*EXTENDED)** compiler option is in effect.
- ▶ `__IBM_INCLUDE_NEXT` indicates support for the `#include_next` preprocessing directive.
- ▶ `__IBM_TYPEOF__` indicates support for the `__typeof__` or `typeof` keyword. This macro is always defined for C. For C++, it is defined when the **LANGLVL(*EXTENDED)** compiler option is in effect.
- ▶ `__IFS_IO__` can be defined when the **SYSIFCOPT(*IFSIO)** or **SYSIFCOPT(*IFS64IO)** compiler option is specified.
- ▶ `__IFS64_IO__` can be defined when the **SYSIFCOPT(*IFS64IO)** compiler option is specified. When this macro is defined, `_LARGE_FILES` and `_LARGE_FILE_API` are also defined in the relevant IBM-supplied header files.
- ▶ `__LONGDOUBLE64` indicates that the size of a long double type is 64 bits. This macro is always defined.
- ▶ `LONG_LONG` indicates support for IBM long long data types and can be defined when the **LANGLVL(*EXTENDED)** compiler option is in effect.
- ▶ `__POSIX_LOCALE__` can be defined when the **LOCALETYPE(*LOCALE)**, **LOCALETYPE(*LOCALEUCS2)**, or **LOCALETYPE(*LOCALEUTF)** compiler option is specified.
- ▶ `__UCS2__` can be defined when the **LOCALETYPE(*LOCALEUCS2)** compiler option is specified.
- ▶ `__UTF32__` can be defined when the **LOCALETYPE(*LOCALEUTF)** compiler option is specified.
- ▶ `__C99_MACRO_WITH_VA_ARGS` indicates support for function-like macros with variable arguments. It is defined when the **LANGLVL(*EXTENDED)** compiler option is in effect.
- ▶ `__DIGRAPHS__` indicates support for digraphs.

Pragmas

The `#pragma` preprocessor directive allows each compiler to implement compiler-specific features that can be turned on and off with the `#pragma` statement. The `do_not_instantiate` and `namemanglingrule` pragmas are included in IBM i 7.1.

do_not_instantiate

The `#pragma do_not_instantiate` directive suppresses the instantiation of a specified entity. It is typically used to suppress the instantiation of an entity for which a specific definition is supplied. If you are handling template instantiations manually (that is, compiler options **TEMPLATE(*NONE)** and **TMPLREG(*NONE)** are in effect), and the specified template instantiation exists in another compilation unit, using `#pragma do_not_instantiate` ensures that you do not get multiple symbol definitions during the link step.

namemanglingrule

Name mangling or name decoration is a technique that is used to solve various problems that are caused by the need to resolve unique names for programming entities. You can use it to encode extra metadata information in the name of a function, structure, class, or another data type to pass more semantic information from the compilers to linkers. Most of the time, you need it when the language allows entities to be named with the same identifier if they occupy another namespace, which is typically defined by a module, class, or explicit namespace directive.

The `#pragma namemanglingrule` directive provides fine-grained control over the name mangling scheme in effect for selected portions of source code, specifically regarding the mangling of cv-qualifiers in function parameters. You can use it to control whether top-level cv-qualifiers are mangled in function parameters or whether intermediate-level cv-qualifiers are considered when the compiler compares repeated function parameters for equivalence.

Definition: cv stands for *constant* and *volatile*.

ILE C++ decimal floating-point support

Decimal (the classic day-to-day base 10) data is widely used in commercial and financial applications. However, most computer systems have only binary (base two) arithmetic, using 0 and 1 to represent numbers. There are two binary number systems in computers: integer (fixed-point) and floating-point. Unfortunately, decimal calculations cannot be directly implemented with binary floating-point. For example, the value 0.1 needs an infinitely recurring binary fraction, but a decimal number system can represent it exactly as one tenth. So, using binary floating-point cannot ensure that results are the same as those using decimal arithmetic.

In general, decimal floating-point operations are emulated with binary fixed-point integers. Decimal numbers are traditionally held in a binary-coded decimal (BCD) format. Although BCD provides sufficient accuracy for decimal calculation, it imposes a heavy cost in performance because it is usually implemented in software.

IBM POWER6 and POWER7 processor-based systems provide hardware support for decimal floating-point arithmetic. POWER microprocessor core includes the decimal floating-point unit that provides acceleration for the decimal floating-point arithmetic.

ILE C++ compiler enhancements that are available with Rational Development Studio for i V7.1 include decimal floating-point support. The support has the following features:

- ▶ Allows floating-point computations to be performed by using decimal arithmetic (base 10).
- ▶ Avoids potential rounding errors when you convert binary floating-point data to / from human readable formats.
- ▶ Conforms to the decimal formats and arithmetic that is described in the IEEE 754-2008 Standard for Floating-Point Arithmetic.
- ▶ Adds support to the ILE C++ compiler, which is based on Draft Technical Report 24732 submitted to the ISO/IEC JTC1/SC22/WG14 Programming Language C committee.
- ▶ New data types:
 - `_Decimal32`, 4 bytes, 7 digits precision, and -95/+96 exponent
 - `_Decimal64`, 8 bytes, 16 digits precision, and -383/+384 exponent
 - `_Decimal128`, 16 bytes, 34 digits precision, and -6143/+6144 exponent
- ▶ Provides conversions to / from C++ built-in data types, such as integers and binary floating-point types
- ▶ Includes the **DECFLTRND** option for the C++ compiler commands (**CRTCPPMOD** and **CRTBNDCPP**) to control compile-time decimal floating-point rounding mode.

ILE C/C++ compile-time round options

With the **DECFLTRND** compiler option, you can specify the compile-time rounding mode for the evaluation of constant decimal floating-point expressions. This option does not affect the runtime decimal floating-point rounding mode, which is set using the `setca` built-in function.

Here are the possible values for this option:

- ▶ ***HALFEVEN**

This setting is the default setting. It rounds to the nearest value, but, when in a tie, it chooses an even number. For example, 5.22 rounds to 5.2, 5.67 rounds to 5.7, and 5.55 and 5.65 round to 5.6.

- ▶ ***DOWN**

This value rounds toward zero, or truncates the result. For example, 5.22 rounds to 5.2, 5.67 rounds to 5.6, 5.55 rounds to 5.5, and 5.65 rounds to 5.6.

- ▶ ***UP**

This value rounds toward away from zero. For example, 5.22 rounds to 5.3, 5.67 rounds to 5.7, 5.55 rounds to 5.6, and 5.65 rounds to 5.7.

- ▶ ***HALFUP**

This value rounds toward to the nearest value and, in a tie, rounds away from zero. For example, 5.22 rounds to 5.2, 5.67 rounds to 5.7, 5.55 rounds to 5.6, and 5.65 rounds to 5.7.

- ▶ ***HALFDOWN**

This value rounds toward to the nearest value and, in a tie, rounds toward zero. For example, 5.22 rounds to 5.2, 5.67 rounds to 5.7, 5.55 rounds to 5.5, and 5.65 rounds to 5.6.

- ▶ ***FLOOR**

This value rounds toward negative affinity. For example, 5.22 rounds to 5.2, 5.67 rounds to 5.6, 5.55 rounds to 5.5, and 5.65 rounds to 5.6.

- ▶ ***CEILING**

This value rounds toward positive infinity. For example, 5.22 rounds to 5.3, 5.67 rounds to 5.7, 5.55 rounds to 5.6, and 5.65 rounds to 5.7.

17.5 IBM Rational Open Access: RPG Edition

Rational Open Access: RPG Edition provides a way for RPG programmers to use the simple and well-understood RPG I/O model to access resources and devices that are not directly supported by RPG.

Open Access opens RPG's file I/O capabilities, allowing anyone to write innovative I/O handlers to access other devices and resources, such as:

- ▶ Browsers
- ▶ Mobile devices
- ▶ Cloud-computing resources
- ▶ Web services
- ▶ External databases
- ▶ XML files
- ▶ Spreadsheets

An Open Access application has three parts:

1. An RPG program that uses normal RPG coding to define an Open Access file and use I/O operations against the file.
2. A handler procedure or program that is called by Open Access to handle the I/O operations for the file.
3. The resource or device that the handler is using or communicating with.

Open Access is the linkage between parts 1 and 2. Licensed program 5733-OAR is required to use Open Access at run time.

Figure 17-11 illustrates the three parts of an Open Access for RPG solution.




Figure 17-11 Rational Open Access: RPG Edition

Open Access does not provide handlers. A handler can be customer-developed or it can be provided by another provider, such as an ISV. The following list details the characteristics of a handler:

- ▶ A handler is a program or a procedure in a service program.
- ▶ A handler can be a generic handler that can handle any file of that device type, or it can be a handler that is specifically written to handle a particular “file”.
- ▶ A handler is not required to support every operation that RPG allows for that type of file. If the handler does not support a specific operation, then the RPG programmer must not code that operation. For example, for a PRINTER file, if the handler does not support the Force-End-Of-Data operation, then the RPG programmer does not code an FEOD operation for the file.

17.5.1 How to use Rational Open Access

There are two ways to use Rational Open Access: RPG Edition:

1. Handler After
2. Handler First

The following two sections describe how each approach is done.

Handler After

The handler is written after the application is written. For example, an existing application that uses 5250 display files is modified to use Open Access for the WORKSTN files, where:

- ▶ The RPG program is modified by adding the **HANDLER** keyword to the WORKSTN files.
- ▶ The handler must handle all the operations and requirements of the existing RPG program.
- ▶ This type of handler is often provided by an outside expert, such as a software tool vendor or IBM Business Partner.

Handler First

The handler is written before the application is written. For example, the RPG programmer wants to use a web service that returns information for a specific set of criteria, where:

- ▶ The handler provider creates a keyed database file matching the web service.
- ▶ The handler provider can tell the RPG programmer what I/O operations that the handler supports.
- ▶ The RPG programmer codes the RPG program, using the file as an externally described keyed DISK file, with the **HANDLER** keyword to identify the Open-Access handler.
- ▶ The handler uses externally described data structures that are defined from the same file.
- ▶ This type of handler may be written by the same RPG programmer who uses the Open Access file, or it may be provided by an outside expert.

Open Access Example

Example 17-11 illustrates a web service that provides the weather based on City Name.

Example 17-11 Web service example

```
* Define the Open I/O file to use handler cityWeather
* in service program WEBSVCS/CITYWTHR
FweatherF IF E          K DISK   HANDLER('WEBSVCS/CITYWTHR(cityWeather)' (2) and (3)
F                                     : commArea)
F                                     EXTDESC('WEBSVCS/CITYWTHR')      (4)
* Data structure to receive the weather information
D weatherDs          ds          likerec(wthrRec)
* Definitions for communication-area parameter that allows the RPG
* programmer to communicate directly with the handler
/copy WEBSVCS/QRPGLESRC,CITYWTHR
D commArea          ds          likeds(cityWeatherComm_t)
/free
// set up communication area parameter; Programmer is telling
// handler to deliver temperature in Celcius
commArea.temperatureUnit = CITYWTHR_CELCIUS;
// the CHAIN operation will call the handler
chain ('Chicago') wthrRec weatherDs;          (1)
if %found;
// process the data in weatherDs
```

Where:

1. A city name is input as a key.
2. An RPG keyed DISK “file” is used as the interface to the web service.
3. The handler provider chooses to handle OPEN, CLOSE, and CHAIN.
4. The handler provider puts all the required objects into library WEBSVCS, including:
 - A service program that contains the handler procedure
 - A binding directory for the RPG program to locate the service program with the handler
 - A /copy file to define the handler-specific information

An RPG program using **PF CITYWTHR** defines records and key information.

The following actions then occur:

1. RPG implicitly opens the “file” and calls handler cityWeather to handle the CHAIN operation.
2. The handler sets up information that is needed by the handler-specific data structure so that it can handle the CHAIN operation correctly. In this case, perhaps a socket to the web service URL is used.
3. When the RPG program does a CHAIN operation, RPG calls the handler again to handle the CHAIN operation.
4. After the handler returns, RPG determines the value of %FOUND according to the feedback from the handler.
5. When the RPG program ends with LR on, RPG implicitly closes the file and calls the handler again to handle the CLOSE operation. The handler closes the socket to the web service URL.

The handler service program uses **PF CITYWTHR** to define the records and key information, as shown in Example 17-12.

Example 17-12 Handler service program

```
H NOMAIN
  /copy WEBSVCS/QRPGLESRC,CITYWEATHER
  /copy QOAR/QRPGLESRC,QRNOOPENACC
D wthrKey      e ds          extname('WEBSVCS/CITYWTHR':*KEY)
D wthrData     e ds          extname('WEBSVCS/CITYWTHR':*INPUT)
  * The Handler:
P cityWeather...
P              b            export
D cityWeatherChain...
D              pi
D  info        likeds(QrnOpenAccess_T)
D wthrInfo     ds           likeds(cityWeatherComm_t)
D              based(info.userArea)
D key          ds           likeds(wthrKey)
D              based(info.key)
D data         ds           likeds(wthrData)
D              based(info.inputBuffer)
  /free

...
if info.rpgOperation == QrnOperation_CHAIN;
  . . . call the web service using the key ... (not shown here) . . .
  if an error occurred . . .
    info.rpgStatus = 1299;
```

```

else;
  // set the data values from the info returned by the web service
  data.precip = . . .
  data.temp = . . .
  // use the wthrInfo communication area to find out how the RPG
  // programmer wants to get the temperature
  if wthrInfo.temperatureUnits = CITYWTHR_CELCIUS;
    data.temp = (dataParm.temp * (9/5)) + 32;
  endif;

```

17.5.2 IBM Rational Open Access: RPG Edition withdrawn

Effective May 8, 2012, IBM withdrew from marketing the following programs and features, which were licensed under the IBM International Program License Agreement:

- ▶ 5733-OAR 1.1.0 IBM Rational Open Access: RPG Edition
- ▶ 5733-OAR 1.1.1 IBM Rational Open Access: RPG Edition

To encourage the adoption of the technology of Rational Open Access: RPG Edition and enhance the value of the IBM i platform and the RPG language, IBM removed the additional license entitlement requirement for Rational Open Access: RPG Edition for customers of IBM i and IBM Rational Development Studio for i. IBM i 6.1 and 7.1 Program Temporary Fixes (PTFs) are available to enable Open Access: RPG Edition applications to run in the absence of a Rational Open Access runtime license, and to move the technology of Open Access: RPG Edition into Rational Development Studio for i. Here are descriptions of the relevant PTFs:

- ▶ A PTF for the RPG runtime function that checks for the 5733-OAR license: With this PTF, the function does not check for the 5733-OAR license, but instead returns the value that indicates that the license is valid.
- ▶ A PTF for the RPG compiler: With this PTF, the compiler does not generate a call to the RPG runtime license-checking function as part of the generated code for opening an Open Access file.
- ▶ A PTF to make the QOAR library part of the Rational Development Studio for i product: If you do not have 5733-OAR, you must apply this PTF to obtain the include files that are used for developing handlers. If you have 5733-OAR, it is not necessary to apply this PTF. However, any further updates or fixes to the include files are available only through additional PTFs for the compiler product 57xx-WDS. 5733-OAR must be removed from the system before the PTF is applied.

As of February 1, 2012, running an ILE RPG program using an Open Access file is available, at no additional charge, to everyone who is licensed to use IBM i 6.1 and 7.1. Also, compiling an ILE RPG program using an Open Access file is available, at no additional charge, to everyone who is licensed to use the ILE compilers feature of either IBM Rational Development Studio for i 7.1 or IBM WebSphere Development Studio for i 6.1

For the announcement letter and more PTF information, visit the RPG Cafe at:

https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/We13116a562db_467e_bcd4_882013aec57a/page/Open%20Access%20announcement

17.5.3 Open Access requirements

Here are the Open Access requirements for IBM i 6.1 and IBM 7.1:

- ▶ IBM i 6.1: IBM WebSphere Development Studio for i V6.1 (5761-WDS) or RPG run time V6.1 (5761-SS1) with applicable PTFs
- ▶ IBM i 7.1: IBM Rational Development Studio for i V7.1 (5770-WDS) or RPG run time V7.1 (5770-SS1) with applicable PTFs
- ▶ For IBM i 6.1
 - POWER5/5+: 9407-515, 9406-520, 9406-525, 9406-550, 9406-570, 9406-MMA, 9406-595, 9406-800, 9406-810, 9406-825, 9406-870, and 9406-890
 - BladeCenter: 7998-61X
- ▶ For IBM i 7.1
 - POWER7: 8233-E8B, 9117-MMB, 9179-MHB, 8406-70Y, or 8406-71Y
 - POWER6/6+: 9407-M15, 9408-M25, 9409-M50, 8203-E4A, 8204-E8A, 8234-EMA, 8261-E4S, 9406-MMA, 9117-MMA, and 9119-FHA
 - BladeCenter: 7998-60X, 7998-61X, or 7778-23X
 - POWER5/5+: 9405-520, 9407-515, 9406-520, 9406-525, 9406-550, 9406-570, or 9406-595

For more information about handlers, visit the RPG Café at:

- ▶ <http://www-949.ibm.com/software/rational/cafe/community/rpg>
- ▶ <http://www-949.ibm.com/software/rational/cafe/docs/DOC-3414>

For more information about Open Access RPG Edition, go to

<http://www-01.ibm.com/software/rational/products/openaccess/>

17.6 IBM Rational Application Management Toolset for i

Rational Application Management Toolset for i provides IBM i system administrators and other advanced users with a lightweight set of tools for the common tasks of working with library file system objects and for creating and maintaining Command Language (CL) programs.

Application Management Toolset for i is a subset of the Application Development Toolset (ADTS), which is sold as part of WebSphere Development Studio for i V6.1 or as part of Rational Development Studio for i V7.1. Two of the key components of ADTS are Programming Development Manager (commonly known as PDM) and Source Entry Utility (commonly known as SEU).

The new Application Management Toolset for i includes these two components, in modified form, as illustrated in Figure 17-12.




Figure 17-12 Application Management Toolset for i includes modified components of ADTS

The version of SEU that is included in Application Management Toolset for i supports only editing of the CL source. It does not support editing of source members that are written in other languages, such as RPG, COBOL, C, C++ or DDS. Like SEU, this editor provides language-sensitive features, such as syntax checking and prompting for CL source members.

Application Management Toolset for i supports the operating system member types CL, CLLE, CLP, TXT, and CMD in the **EDTCLU** (same as **STRSEU**) command.

The version of PDM that is included in Application Management Toolset for i can be used to browse, move, filter, and manipulate objects of any type, but it enables only software development options (such as Edit and Compile) for CL objects.

Application Management Toolset for i supports the following functions from PDM:

- ▶ All the menu functions of **STRPDM** (new command **STRAMT**)
- ▶ All the functions of **WRKLIBPDM** (new command **WRKLIBAMT**)
- ▶ All the functions of **WRKOBJPDM** (new command **WRKOBJAMT**), including **FNDSTRPDM** (new command **FNDSTRAMT**), except for:
 - No option 18 to call DFU.
 - No option 34 to call ISDB.
 - No option 54 to call CMPPFM.
- ▶ All the functions of **WRKMBRPDM** (new command **WRKMBRAMT**), including **FNDSTRPDM** (new command **FNDSTRAMT**), with the following exceptions:
 - Option 2 (Edit) uses the new command **EDTCLU**, which supports only the CL, CLLE, CLP, TXT, and CMD member types.
 - No option 17 to call SDA.

- No option 19 to call RLU.
- No option 54 to call CMPPFM.
- No option 55 to call MRGSRC.

None of the other components from ADTS are included with Rational Application Management Toolset for i.

17.6.1 Application Management Toolset for i Licensing

IBM Rational Application Management Toolset for i V6.1.0 (5761-AMT) is licensed per processor, for unlimited usage on that processor by any number of persons. The license is priced according to the software tier of the machine on which Application Management Toolset for i is used.

Like WebSphere Development Studio for i and Rational Development Studio for i, ongoing maintenance and support costs for Application Management Toolset for i are included in the IBM i system Software Maintenance agreement (SWMA).

17.6.2 Application Management Toolset for i Requirements

- ▶ Hardware requirements:
IBM Rational Application Management Toolset for i V6.1 supports all the hardware models that support IBM i 6.1 and 7.1.
- ▶ Software requirements:
IBM Rational Application Management Toolset for i V6.1 supports IBM i 6.1 and 7.1.

17.6.3 Accessing Application Management Toolset for i

The main menu for Application Management Toolset for i (AMT) (Figure 17-13) can be accessed by running **STRAMT**.

Application Management Toolkit (AMT)

Select one of the following:

1. Work with libraries
2. Work with objects
3. Work with members
9. Work with user-defined options

Selection or command
===>

F3=Exit	F4=Prompt	F9=Retrieve	F10=Command entry
F12=Cancel	F18=Change defaults		

Figure 17-13 Main menu for Application Management Toolkit for i



IBM Navigator for i V7.1

This chapter introduces the following enhancements for the brand new IBM Navigator for i:

- ▶ New name: IBM Navigator for i
- ▶ Signing on to IBM Navigator for i
- ▶ Figure 18-3 on page 642
- ▶ Dynamic navigation
- ▶ Modern tables
- ▶ 5250 emulation-related enhancements
- ▶ Set Target System enhancements
- ▶ Database enhancements
- ▶ New journal management enhancements
- ▶ Integrated Server Administration enhancements
- ▶ Printer output enhancements
- ▶ File system enhancements
- ▶ Networking enhancements
- ▶ Disk management enhancements
- ▶ Tape support enhancements
- ▶ Performance enhancements
- ▶ JS1- Advanced Job Scheduler for i enhancements
- ▶ Backup Recovery Media Services: Plug-in enhancements
- ▶ More information

18.1 New name: IBM Navigator for i

The interface that was previously known as IBM Systems Director Navigator for IBM i underwent a significant transformation and now has tremendous usability and performance enhancements. Its new name is *IBM Navigator for i*.

IBM Navigator for i is the strategic application for IBM i administration tasks, and new administration tasks are supported only by the web application, and will not be added to the System i Navigator Windows client application.

Unavailable tasks: The following tasks are *not* available within IBM Navigator for i:

- ▶ Run SQL Scripts
- ▶ Visual Explain
- ▶ Database Navigator maps
- ▶ SQL Assist
- ▶ AFP Manager

IBM Navigator for i is included as part of the base operating system 5770SS1 - option 3 (Extended Base Directory Support). There is no additional charge for obtaining it.

The following items must be installed on your system or partition:

- ▶ IBM HTTP Server for i - 5770DG1
- ▶ Qshell 5770SS1 - option 30
- ▶ Portable App Solutions Environment (PASE) 5770SS1 - option 33
- ▶ IBM Developer Kit for Java - 5761JV1
- ▶ J2SE 5.0 32 bit 5761JV1 - option 8
- ▶ Host Servers 5770SS1 - option 12

18.2 Signing on to IBM Navigator for i

You can access the IBM Navigator for i interface by typing `systemname:2001` within the URL address field of your browser. A sign-on window opens, as shown in Figure 18-1.



Figure 18-1 IBM Navigator for i - Sign on window

If your password is expired, you can update the password on the IBM Navigator for i sign-on window, as shown in Figure 18-2.



The image shows the IBM Navigator for i sign-on window. At the top right is the IBM logo. The main heading is "IBM® Navigator for i". Below this, there are two red error messages: "Login failed. Check the user ID and password and try again." and "Password is expired." To the left of the input fields is a circular icon with an 'i' and the text "for Business". The input fields are labeled "User ID:", "Password:", "New password", and "New password (to verify)". A "Log in" button is at the bottom left.

Figure 18-2 IBM Navigator for i - expired password

After you sign on, the Welcome window opens, as shown in Figure 18-3.




Figure 18-3 IBM Navigator for i - Welcome page - Independent tabs

18.3 Independent tabs

Because there are independent tabs that have their page contents cached, you can quickly switch back and forth between tabs without a refresh on the server side. You can also open a tab in a new window, as shown in Figure 18-4. Right-click the **Active Jobs** tab and click **Open in a new window**. A new browser window opens and displays the Active Jobs.

Figure 18-4 shows the IBM Navigator for i interface with the 'Active Jobs' tab selected. A red box highlights the 'Active Jobs' tab in the top navigation bar. A red arrow points to the 'Open in new window' option in the context menu that appears when right-clicking the 'Active Jobs' tab. The main content area displays a table of active jobs with columns: Job Name, Detailed Status, Current User, Type, CPU %, Run Priority, and Thread Count. The table lists various system jobs like Qbatch, Qcmn, Qacscot, Qlzpssrv, Qnmapingd, Qnmarexecd, Qnpservr, Qzrcsrvr, Qzscsrvr, Qctf, Qsysscd, Qhttpsrv, Admin, and Admin1-3. The bottom of the screen shows pagination information: '1 - 100 of 375 items' and a list of page numbers: '5 | 10 | 25 | 50 | 100 | All'.

Figure 18-4 Active Jobs tab - Open in new window

18.4 Dynamic navigation

With IBM Navigator for i, you have a quick and dynamic navigation vehicle with performance and usability that is similar to the IBM i Navigator client environment, but without the need to have an installation and do any management on the client side.

The containers open quickly, and you can reach endpoints with fewer clicks. It is also possible to dynamically update the content by clicking the **Reload** option to the right of the group within the container, as shown in Figure 18-5.




Figure 18-5 Dynamic navigation and update

Before this update, the left navigation was a static list with only core functionality enabled. You now have dynamic navigation, making the left frame a true navigation area. This functionality changes how you can manage your systems and partitions.

18.5 Modern tables

IBM Navigator for i has the following new features that are related to tables and filters, as shown in Figure 18-6:

- ▶ Click column to sort.
- ▶ Right-click any row to display a list of functions.
- ▶ Fast filtering (by selecting the quick filter option).
- ▶ Hover on any row to see details.
- ▶ Select how many entries to show.
- ▶ Fast path to the page you want to see.




Figure 18-6 Enhancements to tables and filters

Note: Do *not* press Enter when you use the Quick filter option, as the window updates automatically when you enter the filtering argument. If you press Enter, you discard the search argument and do not get the correct output.

18.6 5250 emulation-related enhancements

You can now access the same 5250 emulation portlet that is included with IBM i Access for Web by using the System option within IBM i management in IBM Navigator for i, as shown in Figure 18-7. This option makes it easier to use one single interface for managing IBM i systems and partitions.



Figure 18-7 5250 emulation using the System option in IBM i management

However, as you can see in Figure 18-8, you cannot customize the 5250 emulation portlet from within the IBM Navigator for i console.




Figure 18-8 5250 emulation session

18.7 Set Target System enhancements

IBM Navigator for i is enhanced so that you can work with multiple systems from a single interface using the Set Target System feature. You can manage an IBM i system or partition by selecting the target system. After this system or partition is selected, all data returned to the IBM Navigator for i interface is for this system or partition.

IBM Navigator for i can manage a target IBM i 5.4, 6.1, or 7.1 system or partition. The options that are available on that target partition or system can vary depending on the IBM i release that is on that target system or partition.

As shown in Figure 18-9, click the **Set Target System** task in the task navigation area to specify the system or partition to be managed. Select the system or partition name, the corresponding user ID and password, and click **OK**.




Figure 18-9 Setting the target system to manage another system or partition

This example shows IBM Navigator for i on a IBM i 7.1 system. Figure 18-9 shows that the select target system selection changed to another system, which in this case is running IBM i 6.1. You can then access the IBM i 6.1 functions available within IBM Navigator for i.

With the Set Target System feature, the IBM Navigator for i management server runs in one place. One single browser can be used to manage multiple environments and management is extended to IBM i 5.4 environments.

To go back to the original IBM i 7.1 system, click **Manage local system** and click **OK**, as shown in Figure 18-10.




Figure 18-10 Manage a local system

18.8 Database enhancements

This section describes the following database enhancements:

- ▶ On Demand Performance Center enhancements
- ▶ Health center enhancements
- ▶ Database management enhancements
- ▶ Database Maintenance category enhancements
- ▶ Long SQL schema names enhancements
- ▶ OmniFind text search functions enhancements

18.8.1 On Demand Performance Center enhancements

This section describes the following enhancements:

- ▶ New centralized authority model
- ▶ Enhanced filtering support
- ▶ SQL performance monitor enhancement
- ▶ Enhancements to Index Advisor and the indexes list

New centralized authority model

Support is added in IBM i 7.1 for a centralized authority model that is used when you work with all the SQL performance tools.

Previously, a system security officer needed to grant the *JOBCTL user special authority so that database analysts and database administrators could use the database tools. Because the *JOBCTL authority allows a user to change many system critical settings that are unrelated to database activity, it was not an easy decision for security officers to grant this authority. In certain cases, *JOBCTL was not granted to database analysts, thus prohibiting the usage of the full set of database tools.

In IBM i 7.1, the security officer can authorize access to the database analysis tools, and the SQL Plan Cache. DB2 for i takes advantage of the function usage capability available in the operating system.

A function usage group that is called *QIBM_DB* was created. In IBM i 7.1, there are four function IDs in the QIBM_DB group:

- ▶ QIBM_DB_SQLADM (IBM i Database Administrator tasks)
- ▶ QIBM_DB_SYSMON (IBM i Database Information tasks)
- ▶ QIBM_DB_DDMDRDA (DDM ad DRDA Application Server Access)
- ▶ QIBM_DB_ZDA (Toolbox Application Server Access)

The security officer can grant authorities by using either of the following methods:

- ▶ Granting *JOBCTL special authority
- ▶ Authorizing a user or group to the IBM i Database Administrator Function through Application Administration in IBM Navigator for i

To authorize a user or group within IBM Navigator for i, complete the following steps, as shown in Figure 18-11:

1. Expand **IBM i Management**.
2. Click the **System** category.
3. Click **Application Administration**.




Figure 18-11 Application Administration

In the Application Administration window, complete the following steps, as shown in Figure 18-12:

1. Click **Host Applications** on the left side menu.
2. Expand **IBM i**.
3. Expand **Database**.
4. Click **Customize** from the menu to the right of Database Administrator.




Figure 18-12 Database Administrator - Customize

You can customize the access to the Database Administrator functions, as shown in Figure 18-13.




Figure 18-13 Customizing access

Tip: You can use the Change Function Usage (**CHGFCNUSG**) command, with a function ID of QIBM_DB_SQLADM, to change the list of users that are allowed to perform database administration operations. The function usage controls which groups or specific users are allowed or denied authority. The **CHGFCNUSG** command also has a parameter (**ALLOBJAUT(*USED)**) that can be used to grant function usage authority to any user that has *ALLOBJ user special authority.

The access customization configuration for the database administration operations can also be performed in a similar way for database information-related operations. The same is true for the DDM and DRDA Application Server Access and the Toolbox Application Server Access.

Enhanced filtering support

There are more filters for the SQL performance monitors. The new filters include the client registers filters that are shown in Figure 18-14. They allow even more granularity to help reduce the amount of data collected.

SQL Performance Monitor Wizard - 9-5-90-243(Zh22dp1)

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

If you would like to limit the amount of data collected specify which filters to use:

- ☐ Initial number of records: 0
- ☐ Minimum estimated query runtime: 0
- ☐ Minimum estimated temporary storage: 0
- ☐ Job name: QZDASOINIT
- ☐ Job user: QUSER
- ☐ Current user or group profile: TOMV
- ☐ Client location:
- ☐ Local port: Use entry from below
- ☐ Query Governor limits: Always collect information when exceeded
- ☐ Statement result or SQL code: Error or warning

Client registers

- ☐ Accounting string:
- ☐ Application name:
- ☐ Program name:
- ☐ Client user:
- ☐ Workstation:
- ☐ Statements that access these objects:

Figure 18-14 Filter the database monitor collection

SQL performance monitor enhancement

SQL Details for a Job can start a SQL Performance Monitor from the jobs list. You can select **Start SQL Performance Monitor** from the Select Action menu on the SQL Details for Jobs window, as shown in Figure 18-15.

SQL Details for Jobs - 9-5-90-243

--- Select Action ---

Select	Name	Job	Current User	Type	Detailed Status
○	QZDASOINIT	QPGMR	Batch	Waiting for event	

Status: Completed

SQL statement and details:
No statement available

Detail

Detail	Value
Job	156956/Qpgmr/Qzserver
Time when SQL information was retrieved	12:13:05 PM

Page 1 of 1 1 Go Rows 2 Total: 2

Figure 18-15 Start SQL Performance Monitor from SQL details for a job

Enhancements to Index Advisor and the indexes list

Within IBM Navigator for i, there are several enhancements to Index Advisor and the indexes list.

Index Advisor enhancements

IBM Navigator for i has the following maintained temporary index (MTI) information within the Index Advisor, as shown in Figure 18-16:

- ▶ **First Advised**
Date / time when a row is first added to the Index Advisor table for this advice.
- ▶ **MTI Created**
Number of times that this specific MTI was created by the optimizer. MTIs do not persist across system IPLs.
- ▶ **MTI Used**
Number of times that this specific MTI was used by the optimizer.
- ▶ **MTI Last used**
The time stamp that represents the last time this specific MTI was used by the optimizer to improve the performance of a query. This field can be blank, indicating that an MTI that exactly matches this advice was never used by the queries that generated this index advice.
- ▶ **MTI Used for Statistics**
The number of times that this specific MTI was used by the optimizer for statistics.
- ▶ **MTI Last Used for Statistics**
The time stamp that represents the last time that this specific MTI was used by the optimizer to obtain statistics for a query.


The screenshot shows the 'Index Advisor - 9:5-90-243' window. The title bar indicates the database is 'Zh22dp1' and the view is 'Advised Indexes for QUSRSYS'. The table has columns: 'Table for Which Index was Advised', 'Schema', 'Advised Index Type', 'Rows in Table when Advised', 'MTI Used', 'MTI Created', 'MTI Last Used', 'MTI Used for Statistics', and 'MTI Last Used for Statistics'. The last five columns are highlighted with a red border. The table lists several indexes, including QATOVD1, QATOVD1RGRP, QAPMCCCNAC, QASZRAIRD, QATOVD1R, QATOVDUCP, QASZRAIRC, QAPMCCCNAC, QATOVDN2, QAUAGDBPMD2, QAUAGDBPMD2, QAUAGDBPMD2, and QATOVDN3. The status bar at the bottom shows '1 - 68 of 68 items' and a 'Close' button.

Table for Which Index was Advised	Schema	Advised Index Type	Rows in Table when Advised	MTI Used	MTI Created	MTI Last Used	MTI Used for Statistics	MTI Last Used for Statistics
QATOVD1	QUSRSYS	Binary Radix	32	1,232	19	9/24/12 9:56:45 PM	9	9/24/12 9:55:57 PM
QATOVD1RGRP	QUSRSYS	Binary Radix	21	3,037	16	9/24/12 9:56:45 PM	20	9/24/12 9:55:58 PM
QAPMCCCNAC	QUSRSYS	Binary Radix	64	0	0		0	
QASZRAIRD	QUSRSYS	Binary Radix	1,244	0	0		0	
QATOVD1R	QUSRSYS	Binary Radix	24	0	0		0	
QATOVDUCP	QUSRSYS	Binary Radix	31	258	34	9/19/12 7:08:24 AM	2	9/19/12 7:08:24 AM
QASZRAIRC	QUSRSYS	Binary Radix	135	0	0		0	
QAPMCCCNAC	QUSRSYS	Binary Radix	64	0	0		0	
QATOVDN2	QUSRSYS	Binary Radix	7	0	0		0	
QAUAGDBPMD2	QUSRSYS	Binary Radix	32	1	0	9/23/12 10:33:15 PM	1	9/23/12 10:33:15 PM
QAUAGDBPMD2	QUSRSYS	Binary Radix	32	53	1	9/24/12 9:34:08 PM	0	
QAUAGDBPMD2	QUSRSYS	Binary Radix	33	0	0		0	
QATOVDN3	QUSRSYS	Binary Radix	10	0	0		0	

Figure 18-16 Index Advisor - MTI information

Indexes list enhancements

The indexes list includes information for the last time the index was built, as shown in Figure 18-17.



Indexes for QSYS2.DBMON_QUERIES - 9

Database: Zh22dp1 Indexes for QSYS2.DBMON_QUERIES


Name	Schema	Type	Last Build	Text
No filter applied				
Q_QSYS2_QSQPDTBL_QUERYID_00001	QSYS2	Primary Key Constraint	9/17/12 2:14:09 AM	

Tooltip for Q_QSYS2_QSQPDTBL_QUERYID_00001:

- Name: Q_QSYS2_QSQPDTBL_QUERYID_00001
- Creation Date: 9/17/12 2:14:09 AM
- Current Allocated Pages: 2
- Current Key Values: 377

Figure 18-17 Last built index information

The indexes list also includes columns for SSD (Media Preference) and Keep In Memory values, as shown in Figure 18-18.



Indexes for QSYS2.DBMON_QUERIES - 9

Database: Zh22dp1 Indexes for QSYS2.DBMON_QUERIES

Name	Schema	Type	Last Build	Media Preference	Keep in Memory	Text
No filter applied						
Q_QSYS2_QSQPDTBL_QUERYID_00001	QSYS2	Primary Key Constraint	9/17/12 2:14:09 AM	Any	No	

Figure 18-18 Index columns for SSD and Keep In Memory values

For more information related to the DB2 Media Preference, see 9.3.1, “DB2 media preference” on page 398.

Tip: You might need to add the columns so that they show up, as they are not displayed by default. Select the **Columns** option in the Actions menu to add those columns.

18.8.2 Health center enhancements

In Figure 18-19, an example is shown where an SQL0901 category was added to the Environmental Limits tab, which helps to log situations where the database engine has an unknown error and more data must be collected.

Health Center -

Environmental Limit - 3/22/10 3:46:24 PM						
Overview	--- Select Action ---					
Environmental Limits						
Activity						
Size Limits						
Design Limits						
	Select	Environmental Limit - 3/22/10 3:46:24 PM	Value	Percent of Limit	Status	
	☐	Maximum number of activation groups to use SQL per job				
	☐	Maximum number of active descriptors per job				
	☐	Maximum length of SQL statement per job (2 MB)				
	☒	SQL0901 UNEXPECTED ERROR:1022 MESSAGE: SQL0901 F/QSQROUTS-MODULE/QSQCLNUP-PROCEDURE/SQROUTE_CLEANUP-STMT/4009 T/QSQROUTS-MODULE/QSQCLNUP-PROCEDURE/SQROUTE_CLEANUP-STMT/4009MESSAGE: SQL0901 F/QSQROUTS-MODULE/QSQCLNUP-PROCEDURE/SQROUTE_CLEANUP-STMT/4009 T/QSQROUTS-MODULE/QSQCLNUP-PROCEDURE/SQROUTE_CLEANUP-STMT/4009MESSAGE: CPC2956 F/QDBCHGFI-MODULE/QDBCHGFI-PROCEDURE/RESIGNAL_COMP-STMT/36349 T/QDBCHGFI-MODULE/QDBCHGFI-PROCEDURE/CPY_DATA-STMT/28458MESSAGE: SQL0117 F/QSQRUN1-MODULE/QSQPREP-PROCEDURE/CLEANUP-STMT/9947 T/QSQRUN1-MODULE/QSQPREP-PROCEDURE/CLEANUP-STMT/9947MESSAGE: SQL6789 F/QSQIMAIN-INS/1668 T/QSQIMAIN-INS/1668MESSAGE: SQL0443 F/QSQRUN4-MODULE/QSQCALLSP-PROCEDURE/CLEANUP-STMT/35802 T/QSQRUN4-MODULE/QSQCALLSP-PROCEDURE/CLEANUP-STMT/35802MESSAGE: SQL0443 F/QSQRUN4-MODULE/QSQCALLSP-PROCEDURE/CLEANUP-STMT/35802 T/QSQRUN4-MODULE/QSQCALLSP-PROCEDURE/CLEANUP-STMT/35802MESSAGE: SQ20400 F/QSQXSRMCM-PROCEDURE/CLEANUP-STMT/12547 T/QSQXSRMCM-PROCEDURE/CLEANUP-STMT/12547MESSAGE: SQ				
	☐	057826/HOSPERS/QPADEV000F (HOSPERS)		0	No maximum	☒ Normal
	☐	SQL0901 UNEXPECTED ERROR:1022 MESSAGE: SQL0901 F/QSQROUTS-MODULE/QSQCLNUP-PROCEDURE/SQROUTE_CLEANUP-STMT/4009 T/QSQROUTS-MODULE/QSQCLNUP-PROCEDURE/SQROUTE_CLEANUP-STMT/4009MESSAGE: CPC2956 F/QDBCHGFI-MODULE/QDBCHGFI-PROCEDURE/RESIGNAL_COMP-STMT/36349 T/QDBCHGFI-MODULE/QDBCHGFI-PROCEDURE/CPY_DATA-STMT/28458MESSAGE: SQL0117 F/QSQRUN1-MODULE/QSQPREP-PROCEDURE/CLEANUP-STMT/9947 T/QSQRUN1-MODULE/QSQPREP-PROCEDURE/CLEANUP-STMT/9947MESSAGE: SQL6789 F/QSQIMAIN-INS/1668 T/QSQIMAIN-INS/1668MESSAGE: SQL0443 F/QSQRUN4-MODULE/QSQCALLSP-PROCEDURE/CLEANUP-STMT/35802 T/QSQRUN4-MODULE/QSQCALLSP-PROCEDURE/CLEANUP-STMT/35802MESSAGE: SQL0443 F/QSQRUN4-MODULE/QSQCALLSP-PROCEDURE/CLEANUP-STMT/35802MESSAGE: SQ20400 F/QSQXSRMCM-PROCEDURE/CLEANUP-STMT/12547 T/QSQXSRMCM-PROCEDURE/CLEANUP-STMT/12547MESSAGE: SQ20400 F/QSQRUN4-MODULE/QSQCALLSP-PROCEDURE/CLEANUP-STMT/35802 T/QSQRUN4-MODULE/QSQCALLSP-PROCEDURE/CLEANUP-STMT/35802MESSAGE: MCH3601 F/AiUpca				

Figure 18-19 SQL0901 category added to the Environmental Limits tab

Random and sequential I/O counts are collected and displayed in the Activity tab, as shown in Figure 18-20.




Figure 18-20 Activity tab - random and sequential reads

Sequential reads: The sequential reads are available on the next page of the Activity tab.

18.8.3 Database management enhancements

Support was added in IBM i 7.1 for the following new database features:

- ▶ XML Data type column support
- ▶ XML Schema Repository
- ▶ Global variables
- ▶ Array support
- ▶ FIELDPROC support
- ▶ Three-part names support
- ▶ Procedure creation and definition
- ▶ Function creation and definition
- ▶ Viewing and editing table contents

- ▶ SQL privilege statements
- ▶ Several usability enhancements

XML Data type column support

Support is now available for XML Data type column within the IBM Navigator for i interface, as shown in Figure 18-21.

The screenshot shows the 'New Column - 9(Zh22dp1)' dialog box. The 'Data type' dropdown menu is set to 'XML' and is highlighted with a red box. A red arrow points to the 'Data type' dropdown. Other fields include 'Column name: TOMV_TEST_XML', 'System name: System-generated', 'Allocate: 0', 'Encoding: Data type default', 'Text: XML Data Type', 'Heading line 1: ', 'Heading line 2: ', 'Heading line 3: ', 'Default value: No default', and checkboxes for 'Nullable' (checked), 'Implicitly hidden', and 'Implements field procedure'. 'Add' and 'Close' buttons are at the bottom.

Figure 18-21 XML Data type column support

For more information related to this support, see 6.2.1, “XML support” on page 155.

XML Schema Repository

An XML Schema Repository (XSR) was introduced to store information about an XML schema to allow for document validation or decomposition (shredding).

With IBM i 7.1, support was added for listing and working with XML Schema Repositories. However, there is no support to create an XML Schema Repository using this IBM Navigator for i interface.

Global variables

Support is available for global variables within the IBM Navigator for i interface, as shown in Figure 18-22.




Figure 18-22 Global variable support

For more information related to global variables, see 6.2.3, “Creating and using global variables” on page 162.

Array support

Figure 18-23 shows the support that is available for array types within the IBM Navigator for i interface.




Figure 18-23 Array type support

For more information related to array support, see 6.2.4, “Support for arrays in procedures” on page 163.

FIELDPROC support

IBM i 7.1 has FIELDPROC support at the column level. To use this support, select the Implements Field procedure option when you define the column when you create or alter a table. As shown in Figure 18-24, specify the schema and the corresponding procedure name.

The screenshot shows the 'New Column - 9(Zh22dp1)' dialog box. The 'Column name' is 'COLUMN_TVE_FIELDPROC'. The 'System name' is 'System-generated'. The 'Data type' is 'CHARACTER' with a 'Length' of '1' and 'Encoding' set to 'Data type default'. The 'Text' field contains 'FIELDPROC TESTING TVE'. There are three empty 'Heading line' fields. The 'Nullable' checkbox is checked, and 'Implicitly hidden' is unchecked. The 'Default value' is 'No default'. The 'Implements field procedure' checkbox is checked and highlighted with a red box and an arrow. Below this, the 'Field procedure specification' section has 'Schema' set to 'Use entry from below', 'Program' set to '[Empty]', and an empty 'Parameter list' field. 'Add' and 'Close' buttons are at the bottom.

Figure 18-24 FIELDPROC column support

For more information related to FIELDPROC support, see 6.2.6, “FIELDPROC support for encoding and encryption” on page 166.

Three-part names support

Figure 18-25 shows an example of how to define a three-part alias by referencing a table on a remote system. The corresponding SQL statement is also displayed.




Figure 18-25 Three-part name support for an alias

Support availability: This support is available for a table, a view, or an alias.

Procedure creation and definition

You can create and define a procedure from within the Actions menu, as shown in Figure 18-26.




Figure 18-26 New procedure

There are two options available:

- ▶ SQL procedure
- ▶ External procedure

New menu option: There is also a new menu option called “New based on...”, which you can use to build a similar procedure from an existing one.

Function creation and definition

It is possible to create a function from within the Actions menu, as shown in Figure 18-27.




Figure 18-27 Create Function

There are three options available:

- ▶ Sourced
- ▶ SQL
- ▶ External

Viewing and editing table contents

You can view and edit table contents from within the Actions menu, as shown in Figure 18-28.

The screenshot shows the IBM Navigator for i interface. The top panel displays the 'Tables - 9.5.166.245' window with the 'Database: Zh22dp1 Schema: QIWS' selected. The 'Actions' menu is open, and the 'View Contents' option is highlighted with a red box. A red arrow points from this option to the 'Contents of QIWS.QCUSTCDT - Localhost(Zh22dp1)' window below. This window shows a table with columns: CUSNUM, LSTNAM, INIT, STREET, CITY, and STATE. The table contains 10 rows of data, each with a checkbox in the first column.

	CUSNUM	LSTNAM	INIT	STREET	CITY	STATE
☐	938472	Henning	G K	4859 Elm Ave	Dallas	TX
☐	839283	Jones	B D	21B NW 135 St	Clay	NY
☐	392859	Vine	S S	PO Box 79	Broton	VT
☐	938485	Johnson	J A	3 Alpine Way	Helen	GA
☐	397267	Tyron	W E	13 Myrtle Dr	Hector	NY
☐	389572	Stevens	K L	208 Snow Pass	Denver	CO
☐	846283	Alison	J S	787 Lake Dr	Isle	MN
☐	475938	Doe	J W	59 Archer Rd	Sutter	CA
☐	693829	Thomas	A N	3 Dove Circle	Casper	WY
☐	593029	Williams	E D	485 SE 2 Ave	Dallas	TX
☐	192837	Lee	F L	5963 Oak St	Hector	NY
☐	583990	Abraham	M T	392 Mill St	Isle	MN

Figure 18-28 View and edit table contents

SQL privilege statements

Generate SQL supports the option to include SQL privilege statements. This output option is on by default, as shown in Figure 18-29.




Figure 18-29 Generate SQL - including SQL privileges

Several usability enhancements

In IBM i 7.1, it is possible to save list contents to various output formats. In Figure 18-30, the list contents of all schemas within the database is saved into a comma-separated value file within the IFS of the system.




Figure 18-30 Save list contents

Figure 18-31 shows an example where it is specified that, from the list of procedures within a library, only those containing a specific text string within their name are displayed.




Figure 18-31 Include list support for procedures

List support: This list support is also available for Tables, Indexes, Aliases, Views, Functions, Triggers, Index Advice, Condensed Index Advice, SQL Plan Cache Snapshots, SQL Plan Cache Event Monitors, Schemas, Database Transactions, and Global Transactions.

18.8.4 Database Maintenance category enhancements

Figure 18-32 shows the new Database Maintenance category that you can use to organize database maintenance lists for table reorganizations, index builds, text search index builds, table alters, and check pending constraints.




Figure 18-32 New maintenance category for database maintenance

Progress pages were added for index builds, text search index builds, and table alters. The window that is shown in Figure 18-33 shows the progress that occurs as the index is built.

The screenshot shows the 'Index builds' window for database D1040f40. It contains a table with one entry: USED_IX1, TOMV, Index, In progress, 3/23/10 2:24:38 PM. Below the table is a 'Close' button. A red arrow points from the 'USED_IX1' cell to the 'Work with Job' button in the details pane of the 'TOMV.USED_IX1 - Localhost(D1040f40)' window.

Index builds

Refresh Database: D1040f40

--- Select Action --- Go

Select	Name	Schema	System Name	System Schema	Type	Status	Date Started
☐	USED_IX1	TOMV	USED_IX1	TOMV	Index	In progress	3/23/10 2:24:38 PM

Page 1 of 1 Rows: 1 Total: 1 Filtered: 1 Selected: 0

Close

TOMV.USED_IX1 - Localhost(D1040f40)

Status: In progress

Select	null	Number of Rows	Percent	Elapsed Time
☐	Key build	4000000	81.49	40.374552
☐	Merge phase			
☐	Process changed rows			

Details:

Detail	Value
Index build operation:	CREATE INDEX USED_IX1
Index build job:	005446/Quser/Qzdasoinit
Parallel degree requested:	No parallel processing
Parallel degree used:	0
Estimated index build time:	00:01:23
Index build start time:	3/23/10 2:24:38 PM
Elapsed time:	40.542804

Job Log Refresh **Work with Job** Close

Figure 18-33 Progress page for index builds

Work with Job: From within this same interface, you can get to the corresponding job by clicking the **Work with Job** button.

In addition, history information is shown in the table reorganization progress window, as shown in Figure 18-34.




Figure 18-34 Table reorganization - history information

The Total Elapsed Time field shows the time of all the history entries plus the current entry. For example, if you start reorganizing a table, you see an entry in the history section for the current run. It is updated and shows in green.

If you choose to suspend that instance of the reorganization and resume it later, you see a row in the history section for the previous instance, and a new row for this instance.

The Total elapsed time value then includes both the previous instance of the reorganization, plus this current instance of the reorganization. The history applies to only the history of the reorganization for one instance of the reorganization of this table. It does not show the prior history of completed reorganizations of this table.


From within this same interface, you can now further drill down into the corresponding job by clicking the **Work with Job** button.

Improved functionality: The reorganize progress in releases before IBM i 7.1 required that you find the table and select to reorganize it to see if it is being reorganized. This process is an easier way to accomplish the same task, because the Table Reorganizations option in the Database Maintenance category is now available.

18.8.5 Long SQL schema names enhancements

Support for long schema names was added throughout the database pages to support the porting of database applications from other databases, and to support more descriptive schema names.

Figure 18-35 defines a new Long Schema Name within an IBM i 7.1 database.



The image shows a 'New Schema' dialog box for database '9(Zh22dp1)'. It contains the following fields and controls:

- Name:** A text input field containing 'LONG_NAME_TVE_TEST'.
- Options:** Two checkboxes. The first, 'Add to displayed list of schemas', is checked. The second, 'Create as a standard library', is unchecked.
- Create in:** A dropdown menu showing 'System disk pool'.
- Text:** A text input field containing 'Long Name Support for schemas'.
- Buttons:** 'Show SQL', 'OK', and 'Cancel' at the bottom.

Figure 18-35 Long name support for schemas

18.8.6 OmniFind text search functions enhancements

OmniFind is a text search product that IBM i users can use to search through unstructured text that is stored in a column of a DB2 for i table. The text that is stored in the column can be either simple character text, an XML document, or any of several types of rich text documents, such as a PDF or a .doc file. You can use this product to index unstructured data without parsing it into a structured form, such as an SQL table.

Figure 18-36 shows the support that enables the management of OmniFind text search servers and indexes.




Figure 18-36 OmniFind text search support

The support for OmniFind text search in DB2 adds simpler access to non-structured data that is often stored in XML format.

For more information related to OmniFind text search, see the IBM i 7.1 Information Center at:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=%2Frzash%2Frzashkickoff.htm>

18.9 New journal management enhancements

You can now manage journal environments through IBM Navigator for i. Here is a list of the additional functions added to journal management in i 7.1:

- ▶ Show Journalized Objects function
- ▶ Change Receivers function
- ▶ Remote Journals function
- ▶ Add Remote Journal function
- ▶ Properties function
- ▶ Activate function
- ▶ Deactivate function
- ▶ Remove function
- ▶ Properties function

For more information related to the Journal Management capabilities and enhancements, see 6.3, “Performance and query optimization” on page 182.

18.9.1 Show Journalized Objects function

The Show Journalized Objects function is added for a journal, as shown in Figure 18-37.




Figure 18-37 Showing journalized objects

This function displays the objects that are journaled to the journal, including files, data areas, data queues, libraries, and integrated file system objects.

18.9.2 Change Receivers function

The Change Receivers function is added for a journal, as shown in Figure 18-38.




Figure 18-38 Changing receiver

You can use this function to attach a new journal receiver to the journal and change the attributes of a journal, including the fixed-length data values, receiver maximum option, minimize options, and other options.

18.9.3 Remote Journals function

The Remote Journals function is added for a journal, as shown in Figure 18-39.




Figure 18-39 Remote journals related information

You can use this function to see a list of all the remote journals that are associated with a journal. Included in the list are the remote journal state and the delivery mode.

18.9.4 Add Remote Journal function

The Add Remote Journal function was added for a journal, as shown in Figure 18-40.




Figure 18-40 Adding remote journal

You can use this function to define a remote journal environment.

18.9.5 Properties function

The Properties function was added for a journal receiver, as shown in Figure 18-41.




Figure 18-41 Properties action for a journal receiver

You can use this function to view the attributes of a journal receiver. Included in the attributes are the sequence number of the first entry in the journal receiver, the minimize options, the fixed-length data options, and other options.

18.9.6 Activate function

The Activate function was added for remote journals, as shown in Figure 18-42.




Figure 18-42 Activating a remote journal

You can use this function to activate a remote journal and specify options.

18.9.7 Deactivate function

The Deactivate function was added for remote journals, as shown in Figure 18-43.




Figure 18-43 Deactivating a remote journal

You can use this function to deactivate a remote journal.

18.9.8 Remove function

The Remove action was added for remote journals, as shown in Figure 18-44.




Figure 18-44 Removal of a remote journal

You can use this function to remove a remote journal environment.

18.9.9 Properties function

The Properties function was added for remote journals, as shown in Figure 18-45.




Figure 18-45 Properties for a remote journal

You can use this function to view the properties of a remote journal connection, which includes the delivery mode, remote journal type, connection statistics, and so on.

18.10 Integrated Server Administration enhancements

The following Integrated Server Administration topics were added to IBM Navigator for i in IBM i 7.1:

- ▶ Create Server function
- ▶ New Based on ... (cloning) a Server function
- ▶ Delete Server function
- ▶ Launch Web Console function

For more information related to Integrated Server Administration, see Chapter 12, "Integration with IBM BladeCenter and IBM System x" on page 467.

18.10.1 Create Server function

The Create Server function was added so that you can create the IBM i hosting environment for an iSCSI attached BladeCenter or System x server, as shown in Figure 18-46.




Figure 18-46 Creating a server task

The Create Server wizard starts. This function is described in more detail in 12.6.1, “Create Server task” on page 475.

18.10.2 New Based on ... (cloning) a Server function

It is also possible to create a server that is based on an existing one. The New Based on ... (cloning) a Server function is shown in Figure 18-47. It starts the Create Server wizard and creates an iSCSI-attached integrated Windows server that is based on one that was previously installed.




Figure 18-47 New Based On ... (cloning) a Server task

18.10.3 Delete Server function

The Delete Server function was added for deleting the IBM i hosting environment for an integrated server, as shown in Figure 18-48.




Figure 18-48 Deleting a server

18.10.4 Launch Web Console function

The Launch Web Console function was added for launching the service processor web console for an iSCSI attached integrated server, as shown in Figure 18-49.




Figure 18-49 Launching the web console

Starting with IBM i 7.1, several iSCSI configuration functions were simplified. For more information about these functions, see 12.6, “IBM Navigator for i” on page 474.

18.11 Printer output enhancements

In IBM i 7.1, new functions were added to the View as PDF and Export as PDF functions.

18.11.1 View as PDF function

The View as PDF function was added so that you can view an AFPDS or SCS printer output file, as shown in Figure 18-50.




Figure 18-50 Printer output - View as PDF

You can use this function to open and view the contents of a printer output file using Adobe Acrobat Reader.

18.11.2 Export as PDF function

The Export as PDF function was added to export an AFPDS or SCS printer output file, as shown in Figure 18-51.




Figure 18-51 Printer output - Export as PDF

You can use this function to save the contents of a printer output file to the following sources:

- ▶ Your client desktop or file system
- ▶ An output queue
- ▶ The Integrated File System (IFS)
- ▶ An attachment to an email

Requirement: For the latter three options, the Infoprint Server licensed program (5722-IP1) is required. Users can use the native IBM Transform Services for i (5770-TS1) licensed program to export to IFS, but they must map a network drive to the IFS and then select the first option, as shown in Figure 18-51 on page 684.

18.12 File system enhancements

A temporary file system is supported. The temporary file system can be created by creating a user-defined file system (UDFS) by specifying “.TMPUDFS” as the extension instead of the default “.UDFS” for the “UDFS name” field. See Figure 18-52.

New User-Defined File System - 192.168.144.10

Auxiliary storage pool: /dev/QASP01

UDFS name: QASP01.TMPUDFS

Description: Temp File System test

Options

☐ Restrict rename and unlink

☐ Case sensitive object names

Default disk space allocation: Normal

Default memory allocation: Normal

Audit objects created in UDFS: Use system value

Scan objects created in UDFS: Inherit from parent

Default file format: Type 2 stream file

OK Cancel

Figure 18-52 Temporary UDFS creation

For more information related to temporary user-defined file system support, see 20.3, “Temporary user-defined file systems” on page 783.

18.13 Networking enhancements

Internet Protocol version 6 (IPv6) is enabled for Remote Access Services, as shown in Figure 18-53.

The screenshot shows the 'RMTACCESS - localhost' configuration window. On the left is a sidebar with tabs: General, Connection, Authentication, TCP/IP IPv4 Settings, *TCP/IP IPv6 Settings (selected), and Other. The main area is divided into sections: 'Enable IPv6' (checked), 'Local IP address' (with a text box containing '1::ff' and a red arrow pointing to it), 'Local interface' (with a dropdown menu set to 'Generate'), and 'Routing'. Under 'Routing', there is a section 'Allow remote system to access other networks (IP forwarding)' with radio buttons for 'Yes' and 'No'. The 'No' option is selected, and it has a checked sub-option 'Accept a default route'. Below this is a checkbox 'Define additional IP addresses based on the user ID of the remote system' which is unchecked. At the bottom right is an 'Advanced' button.

Figure 18-53 IPv6 support for remote access services

IKE version 2 is enabled for VPN, as shown in Figure 18-54.

The screenshot shows the 'Connection Scenario' wizard. It has a title bar and a main area with a diagram on the left showing a server and a client connected by a line with question marks. The text on the right says: 'This wizard allows you to create connection groups so that your system can communicate securely with other hosts or gateways. What type of connection group would you like to create?'. There are four radio button options: 'Connect your gateway to another gateway', 'Connect your gateway to another host', 'Connect your host to another gateway' (selected), and 'Connect your host to another host'. Below this is another question: 'What is the Internet Key Exchange version that this connection will use?'. A dropdown menu shows 'IKE version 2' selected, with a red arrow pointing to it. At the bottom are buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

Figure 18-54 Internet key exchange V2 enablement for VPN

For more information related to networking enhancements in IBM i 7.1, see Chapter 10, “Networking enhancements” on page 413.

18.14 Disk management enhancements

There are several enhancements within IBM Navigator for i within IBM i 7.1 that are related to disk management:

- ▶ Graphical view enhancements
- ▶ Start / Stop encryption on disk pools enhancements
- ▶ Asynchronous delivery for Geographic Mirroring enhancements

18.14.1 Graphical view enhancements

To access the graphical view, complete the following steps:

1. From the Welcome window, click **Configuration and Service** → **Disk Units**, as shown in Figure 18-55.




Figure 18-55 Configuration and service - Disk Units

2. From the Actions menu, click **Graphical View**, as shown in Figure 18-56.




Figure 18-56 Disk Units - Graphical View menu

The graphical view opens, as shown in Figure 18-57.




Figure 18-57 Disk Units - Graphical View window

18.14.2 Start / Stop encryption on disk pools enhancements

The Start / Stop encryption on disk pools feature is now enabled within IBM Navigator for i. You can use a change encryption key when the encryption starts, as shown in Figure 18-58.




Figure 18-58 Disk management - Start / Stop encryption

18.14.3 Asynchronous delivery for Geographic Mirroring enhancements

IBM Navigator for i also provides asynchronous communications between the source and target servers for geographic mirroring, as shown in Figure 18-59.




Figure 18-59 Asynchronous delivery mode

For more information related to asynchronous transmission delivery for geographic mirroring, see 5.1.6, “Asynchronous geographic mirroring” on page 133.

18.15 Tape support enhancements

To access the new Tape Devices option, in the Welcome window, click **Configuration and Service** → **Show All Configuration and Service tasks**, as shown in Figure 18-60.




Figure 18-60 Showing configuration and service tasks

In the Tape Devices menu, you have the following options available:

- ▶ Stand-Alone Devices offers the following possibilities:
 - Make (un)available
 - Look into the properties
 - Upgrade the firmware
- ▶ Tape Image Catalogs offers the following possibilities:
 - Add or list volumes
 - Look at the properties

- ▶ Tape Libraries offer the following possibilities:
 - Make (un)available
 - Look into the properties
- ▶ Create Image Catalog
- ▶ Create Virtual Device

18.16 Performance enhancements

When we describe the performance enhancements within IBM Navigator for i, we assume that the IBM HTTP Server for i 7.1 Group PTF (SF99368) is installed on your system. For more information, see the following website:

http://www-912.ibm.com/s_dir/sline003.nsf/554c38c4848b77f2862567bd0046e003/56ff79c9506270c98625771400478c7f?OpenDocument

Figure 18-61 shows the IBM Navigator for i interface that is used to work with Performance-related tasks within IBM i 7.1.




Figure 18-61 Performance within IBM Navigator for i

Several enhancements were made to the Performance Data Investigator (PDI), which can be accessed by selecting the **Investigate Data** task, as shown in Figure 18-62.




Figure 18-62 Investigate Data option

For more in-depth information about working with Performance-related tasks within IBM i 7.1, see Chapter 7, “Performance tools” on page 247.

18.16.1 General health indicators

This content package helps a system administrator, an IBM i performance specialist, or a general user learn whether the partition is performing properly from a general performance management perspective or if there are areas of the system that must be investigated.

Main system resources and components (such as processor, DASD, and memory) and communications are analyzed. The results are displayed graphically. The main source of data for analysis is the Collection Services performance data files.

The new content package, which deals with the general health of your partition, is shown in Figure 18-63.




Figure 18-63 Health indicators

The following subsets of information are now available:

- ▶ System Resources Health Indicators
- ▶ CPU Health Indicators
- ▶ Disk Health Indicators
- ▶ Memory Pools Health Indicators
- ▶ Response Time Health Indicators

System Resources Health Indicators

You can use this perspective to determine the general health of the main system resources according to the defined thresholds. Figure 18-64 shows health indicators for CPU, Disk, Memory, and communications for the entire selected collection.




Figure 18-64 System Resources Health Indicators

You can use this perspective to quickly determine the percentage of intervals that exceeded the various defined thresholds for CPU, Disk, Memory Pools, and Response Time.

From the System Resources Health Indicators perspective, you can navigate to the following new perspectives:

- ▶ CPU Health Indicators
- ▶ Disk Health Indicators
- ▶ Memory Pools Health Indicators
- ▶ Response Time Health Indicators

Each of these perspectives consists of key metrics that are associated with the system resource that is being investigated. The perspectives are available through the Actions menu, as shown in Figure 18-65, and are described in more detail in the following sections.




Figure 18-65 System Resources Health Indicators drill-down menu

CPU Health Indicators

You can use this perspective to determine the proportion of intervals where processor health indicators exceeded the defined thresholds. The chart in Figure 18-66 shows you processor health indicators by analyzing all collection time intervals according to the defined thresholds for processors.




Figure 18-66 CPU Health Indicators

From the CPU Health Indicators perspective, you can navigate to the following perspectives:

- ▶ CPU Utilization and Waits Overview
- ▶ CPU Utilization Overview
- ▶ Interactive Capacity CPU Utilization

At the top of the CPU Utilization and Waits Overview window, you can click **View** → **Show System Information** to get the system-related information displayed in the heading of the window, as shown in Figure 18-67.




Figure 18-67 Show System Information

Disk Health Indicators

The perspective in Figure 18-68 shows disk health indicators by analyzing all collection time intervals according to the defined thresholds for the disk. This chart can be used to determine the proportion of intervals where the disk health indicators exceeded the defined thresholds.



Figure 18-68 Disk Health Indicators

From the Disk Health Indicators perspective, you can navigate to the following perspectives:

- ▶ Resource Utilization Overview
- ▶ Disk Overview by Disk Pools
- ▶ Disk Details by Disk Pools

Memory Pools Health Indicators

The perspective in Figure 18-69 shows memory pools health indicators by analyzing all collection time intervals according to the defined thresholds for memory pools. This chart can be used to determine the proportion of intervals where the memory pools health indicators exceeded the defined thresholds.




Figure 18-69 Memory Pools Health Indicators

From the Memory Pools Health Indicators perspective, you can navigate to the following perspectives:

- ▶ Resource Utilization Overview
- ▶ Page Faults Overview

Response Time Health Indicators

Figure 18-70 shows response time health indicators by analyzing all collection time intervals according to the defined thresholds for response time. This chart can be used to determine the proportion of intervals where the response time exceeded the defined thresholds.




Figure 18-70 Response Time Health Indicators

From this perspective, you can navigate to the 5250 Display Transactions Overview perspective.

Defining the health indicators thresholds

With IBM i 7.1, the health indicators can be user-defined if the IBM shipped defaults do not meet your business requirements. They can be configured by selecting the **Define Health Indicators** action from the Select Action list, as shown in Figure 18-71.




Figure 18-71 Defining the disk health indicators threshold

This new action is available from the following perspectives:

- ▶ System Resources Health Indicators
- ▶ CPU Health Indicators
- ▶ Disk Health Indicators
- ▶ Memory Pools Health Indicators
- ▶ Response Time Health Indicators

Figure 18-71 shows how to modify the disk health indicators thresholds by specifying the current threshold values as 10 for the Warning field and 20 for the Action field for the Average Disk Response Time.

By correctly defining those threshold values, you have a point of reference to visually determine if the system is between its defined control limits. Figure 18-72 has an overview with the following Average Disk Response Time information:

- ▶ The green area shows the percentage of intervals where the Average Disk Response Time is below the defined Warning Threshold value.
- ▶ The yellow area shows the percentage of intervals where the Average Disk Response Time is above the defined Warning Threshold value and below the defined Action Threshold value.
- ▶ The red area shows the percentage of intervals where the Average Disk Response Time is above the defined Action Threshold value.




Figure 18-72 Threshold that is reached for average disk response time

It is possible to define various thresholds for a specific Graphical Health Indicator graph. Figure 18-73 shows an example of the CPU Health Indicators, where thresholds for each of the following are defined:

- ▶ Partition CPU Utilization
- ▶ Jobs CPU Utilization
- ▶ Interactive CPU Utilization

Define Health Indicators					
CPU Health Indicators	Available Indicators [Empty] Add >> Remove << Selected Indicators Partition CPU Utilization Jobs CPU Queuing Percent Interactive CPU Utilization Current Threshold Values <table> <tr> <td>Warning</td> <td>85</td> </tr> <tr> <td>Action</td> <td>90</td> </tr> </table>	Warning	85	Action	90
Warning	85				
Action	90				

Figure 18-73 Multiple thresholds can be defined

All the individual thresholds can be added, removed, and tailored to your own specifications.

18.16.2 Viewing cross-partition processing

Collection Services can collect high-level cross-partition processor performance metrics for all logical partitions on the same single physical server regardless of the operating system that is running.

Collections Services is available on POWER6 and later servers with a minimum firmware level of xx340_061.

When this data is available, it can be viewed through several perspectives that are found under the Physical System menu, as shown in Figure 18-74.



Figure 18-74 Perspectives - Physical System

There are several perspectives available under Physical System:

- ▶ Logical Partitions Overview
- ▶ Donated Processor Time by Logical Partition
- ▶ Uncapped Processor Time Used by Logical Partition
- ▶ Virtual Processor Pool Utilization
- ▶ Physical Processors Utilization by Physical Processor
- ▶ Dedicated Processors Utilization by Logical Partition
- ▶ Physical Processors Utilization by Processor Status Overview
- ▶ Physical Processors Utilization by Processor Status Detail
- ▶ Shared Memory Overview

The performance collection must be turned on for the IBM i partition to collect the data. This option must be enabled on the HMC, as shown in Figure 18-75.




Figure 18-75 Allowing the performance information collection on the HMC

To turn on this option, from the HMC, complete the following steps:

1. Select **Systems Management** → **Servers**.
2. Click your IBM i system.
3. Select the partition profile.
4. Click **Properties**.
5. Click the **Hardware** tab.
6. Click **Processors**.
7. Click **Allow performance information collection**.

Logical Partitions Overview

Figure 18-76 shows configuration data and processor usage for all logical partitions on the system, including operating system, number of virtual processors, partition memory, donated processor time, uncapped processor time that is used, and so on.




Figure 18-76 Logical Partitions Overview

Donated Processor Time by Logical Partition

Figure 18-77 shows the processor time that is donated by dedicated processor logical partitions that are configured to donate unused processor cycles.




Figure 18-77 Donated Processor Time by Logical Partition

Uncapped Processor Time Used by Logical Partition

Figure 18-78 shows the uncapped processor time that was used by logical partitions in excess of their entitled processing capacity.




Figure 18-78 Uncapped Processor Time Used by Logical Partition

Virtual Shared Processor Pool Utilization

Figure 18-79 shows the processing capacity available and processing capacity usage for virtual shared processor pools.




Figure 18-79 Virtual Shared Processor Pool Utilization

Physical Processors Utilization by Physical Processor

Figure 18-80 shows the usage percentage for each physical processing unit.




Figure 18-80 Physical Processor Utilization by Physical Processor

Dedicated Processors Utilization by Logical Partition

Figure 18-81 shows the dedicated processor usage by logical partitions.




Figure 18-81 Dedicated Processors Utilization by Logical Partition

Physical Processors Utilization by Processor Status Overview

Figure 18-83 on page 713 shows a summary of processors usage for the entire collection, identifying usage for dedicated processors and for shared processors.




Figure 18-82 Physical Processors Utilization by Processor Status Overview

Physical Processors Utilization by Processor Status Detail

Figure 18-83 shows processor usage over time. It shows the dedicated processors and the shared processors.

Physical Processors Utilization by Processor Status Detail					
Perspective Edit View History					
☐ ☐ ☐ ☐ ☐ ☐ --- Select Action ---					
Select	Date - Time	Interval Number	Processor Status	Processor Identifier	Average Processor Utilization (Percent)
☐	Dec 16, 2008 11:06:15 PM	1		0000	14.11
☐	Dec 16, 2008 11:06:15 PM	1		0001	13.94
☐	Dec 16, 2008 11:06:15 PM	1		0002	7.94
☐	Dec 16, 2008 11:06:15 PM	1		0003	3.6
☐	Dec 16, 2008 11:06:30 PM	2		0000	3.29
☐	Dec 16, 2008 11:06:30 PM	2		0001	1.23
☐	Dec 16, 2008 11:06:30 PM	2		0002	0.24
☐	Dec 16, 2008 11:06:30 PM	2		0003	0.13
☐	Dec 16, 2008 11:06:45 PM	3		0000	4.86
☐	Dec 16, 2008 11:06:45 PM	3		0001	2.15
☐	Dec 16, 2008 11:06:45 PM	3	Processor Status	0002	0.56
☐	Dec 16, 2008 11:06:45 PM	3		0003	0.09
Total: 300			Filtered: 300		

Figure 18-83 Physical Processors Utilization by Processor Status Detail

Shared Memory Overview

Figure 18-84 shows shared memory pool use by the partition on an interval-by-interval basis. This information can be used to find the partition's usage of the memory shared pool and metrics that are the sum of the activity that is caused by all partitions using the memory shared pool.




Figure 18-84 Shared Memory Overview

Performance collection must be turned on for the IBM i partition to collect the data. This option must be enabled on the HMC, as shown in Figure 18-85.




Figure 18-85 Allowing performance information collection on the HMC

To activate performance collection, from the HMC, complete the following steps:

1. Click **Systems Management** → **Servers**.
2. Click your IBM i system.
3. Select the partition profile.
4. Click **Properties**.
5. Click the **Hardware** tab.
6. Click **Processors**.
7. Click **Allow performance information collection**.

18.16.3 New Database package

There is now a robust graphical interface for the Database perspective that is enabled by Collection Services support, as shown in Figure 18-86:

- ▶ New perspectives are available within the Database perspective
- ▶ Support for SQL Plan Cache & SQL Performance Monitor database performance files

Required program: IBM i 7.1 5770-PT1 licensed program Performance Tools - Manager Feature - Option 1 must be installed to use this interface.




Figure 18-86 PDI - Investigate Data - Database

Database I/O

For the Physical I/O perspective, the following detailed views are available:

- ▶ Physical Database I/O Overview
- ▶ Physical Database I/O by Job or Task
- ▶ Physical Database I/O by Thread or Task
- ▶ Physical Database by Generic Job or Task
- ▶ Physical Database by Job User Profile
- ▶ Physical Database by Job Current User Profile
- ▶ Physical Database by Subsystem
- ▶ Physical Database by Server Type

Physical Database I/O Overview

This chart shows physical database I/O by SQL-related and non-SQL related synchronous and asynchronous read and write rates per second.

Physical Database I/O by Job or Task

This chart shows the breakdown of SQL-related and non-SQL related, synchronous and asynchronous, and read and write physical database I/O requests as an I/O rate per second grouped by job or task and ranked by the largest contributors. Use this chart to select contributors for further detailed investigation.

Physical Database I/O by Thread or Task

This chart shows the breakdown of SQL-related and non-SQL related, synchronous and asynchronous, and read and write physical database I/O requests as an I/O rate per second grouped by thread or task and ranked by the largest contributors. Use this chart to select contributors for further investigation.

Physical Database I/O by Generic Job or Task

This chart shows the breakdown of SQL-related and non-SQL related, synchronous and asynchronous, and read and write physical database I/O requests as an I/O rate per second grouped by the first six characters of the job or task name and ranked by the largest contributors. Use this chart to select contributing groups for further investigation.

Physical Database by Job Current User Profile

This chart shows the breakdown of SQL-related and non-SQL related, synchronous and asynchronous, and read and write physical database I/O requests as an I/O rate per second grouped by the job user profile and ranked by the largest contributors. Use this chart to select contributing groups for further investigation.

Physical Database by Subsystem

This chart shows the breakdown of SQL-related and non-SQL related, synchronous and asynchronous, and read and write physical database I/O requests as an I/O rate per second grouped by subsystem and ranked by the largest contributors. Use this chart to select contributing groups for further investigation.

Physical Database by Server Type

This chart shows the breakdown of SQL-related and non-SQL related, synchronous and asynchronous, and read and write physical database I/O requests as an I/O rate per second grouped by server type and ranked by the largest contributors. Use this chart to select contributing groups for further investigation.

This new breakdown makes it definitely much easier to notice *what changed*.

The graph in Figure 18-87 shows an example of the Physical Database I/O Overview chart.




Figure 18-87 Physical Database I/O Overview

The graph in Figure 18-88 shows an example of the Physical Database I/O by Job or Task chart.




Figure 18-88 Physical Database I/O by Job or Task

SQL Performance Data

You can now find SQL Plan Cache Snapshot, SQL Plan Cache Event Monitor, and SQL Performance Monitor database files under the PDI Collection Name List, as shown in Figure 18-89.




Figure 18-89 Collection Name list

Select these collections with the SQL Performance Data perspectives to view several high-level charts for a specific SQL Performance Monitor, SQL Plan Cache Snapshot, or SQL Plan Cache Event Monitor.

The SQL Plan Cache & SQL Performance Monitor folders each have two perspectives:

- ▶ SQL Overview
- ▶ SQL Attribute Mix

SQL Overview perspective

This perspective gives a comprehensive picture of how queries are running overall. The SQL Overview perspective for both SQL Plan Cache and SQL Performance Monitor database performance files has 11 views:

- ▶ Query Time Summary
- ▶ Open Summary
- ▶ Open Type Summary
- ▶ Statement Usage Summary
- ▶ Index Used Summary
- ▶ Index Create Summary
- ▶ Index Advised
- ▶ Statistics Advised
- ▶ MQT Use
- ▶ Access Plan Use
- ▶ Parallel Degree Usage

In Figure 18-90, you can see an example of the Query Time Summary view with a corresponding Open Type summary view.




Figure 18-90 SQL Overview Perspective - views

SQL Attribute Mix Perspective

This perspective provides an indication of the nature of the users applications and queries by displaying the attribute mix. The Attribute Mix perspective has the following views available:

- ▶ Statement Summary
- ▶ Statement Type Summary
- ▶ Isolation Level Summary
- ▶ Allow Copy Data Summary
- ▶ Sort Sequence Summary
- ▶ Close Cursor Summary
- ▶ Naming Summary
- ▶ Optimization Goal
- ▶ Blocking Summary

In Figure 18-92 on page 723, you can see an example of the Statement Summary view with a corresponding Statement Type Summary view.




Figure 18-91 SQL Attribute Mix perspective - views

18.16.4 QAPMCONF perspective

A new perspective view of the QAPMCONF database file was added. It can be found within the Collections Services content package under the Collection Services Database Files folder, as shown in Figure 18-92.




Figure 18-92 QAPMCONF perspective view

The window that is shown in Figure 18-93 contains general information about the collections options and information about the system on which the data was collected.

QAPMCONF

Perspective » Edit » View » History »

Collection	Time	System
Name(s): Q257000012	Start: Sep 14, 2011 12:00:12 AM	Name: MCV7R1
Library: QPFRDATA	End: Sep 15, 2011 12:00:15 AM	Release: V7R1M0
Type: Collection Services File Based Collection		

QAPMCONF Panel View

Library Name:	QPFRDATA	Processor Firmware Time:	No
Member Name:	Q257000012	Task Threshold Value (ms):	1,000
Start Time:	Sep 14, 2011 12:00:12 AM	Secondary Thread Thresh (ms):	1,000
Model Number:	MMA	Disk Response Time Boundary 1 (us):	15
System Type:	9408	Disk Response Time Boundary 2 (us):	250
Partition Memory (KB):	25165824	Disk Response Time Boundary 3 (us):	1,000
Comm Data Collected:	N	Disk Response Time Boundary 4 (us):	4,000
Machine Serial Number:	10-40F40	Disk Response Time Boundary 5 (us):	8,000
Response Time Boundary 1 (ms):	1000	Disk Response Time Boundary 6 (us):	16,000
Response Time Boundary 2 (ms):	2000	Disk Response Time Boundary 7 (us):	64,000
Response Time Boundary 3 (ms):	4000	Disk Response Time Boundary 8 (us):	256,000
Response Time Boundary 4 (ms):	8000	Disk Response Time Boundary 9 (us):	500,000
System ASP Capacity (KB):	895,844,352	Disk Response Time Boundary 10 (us):	1,024,000
Checksum Protection On:	N	Hypervisor Memory (MB):	7,424
Virtual Processors:	2	SMT Hardware Threads:	0
Installed Processors:	16	Time Interval (minutes):	1
Remote Response Boundary 1 (ms):	-	Interactive Limit (%):	100.00
Remote Response Boundary 2 (ms):	-	Time Interval (seconds):	60
		Interactive Threshold (%):	100.00

Figure 18-93 QAPMCONF perspective

18.16.5 QAPMJOBMI perspective

New fields are added to the QAPMJOBMI database file. These fields can be found within the Collections Services content package under the Collection Services Database Files folder, as shown in Figure 18-94.




Figure 18-94 QAPMJOBMI perspective

The window that is shown in Figure 18-95 contains more information that enables improved Native versus SQL database performance isolation:

- ▶ SQL clock time (total time in SQ and below) per thread (microseconds)
- ▶ SQL unscaled CPU per thread (microseconds)
- ▶ SQL scaled CPU per thread (microseconds)
- ▶ SQL synchronous database reads per thread
- ▶ SQL synchronous non-database reads per thread
- ▶ SQL synchronous database writes per thread
- ▶ SQL synchronous non-database writes per thread
- ▶ SQL asynchronous database reads per thread
- ▶ SQL asynchronous non-database reads per thread
- ▶ SQL asynchronous database writes per thread
- ▶ SQL asynchronous non-database writes per thread

QAPMJOBMI

Perspective Edit View History

--- Select Action ---

Select	SQL Asynchronous Database Writes	SQL Asynchronous Non Database Reads	SQL Asynchronous Non Database Writes	SQL Clock Time	SQL Synchronous Database Reads	SQL Synchronous Database Writes	SQL Synchronous Non Database Reads	SQL Synchronous Non Database Writes
☐	0	0	0	0	0	0	0	0
☐	2	0	5	6793	3	0	5	5
☐	0	0	0	0	0	0	0	0
☐	0	0	0	0	0	0	0	0
☐	0	0	0	0	0	0	0	0
☐	0	0	0	0	0	0	0	0
☐	0	0	0	0	0	0	0	0
☐	2	0	4	33053	2	0	0	4
☐	0	0	0	0	0	0	0	0
☐	0	0	0	0	0	0	0	0
☐	0	0	0	0	0	0	0	0
☐	0	0	0	0	0	0	0	0
☐	0	0	0	0	0	0	0	0
☐	0	0	0	0	0	0	0	0
☐	0	0	0	0	0	0	0	0

Total: 76460 Filtered: 76460

Done Options Save As...

Figure 18-95 QAPMJOBMI - native versus SQL Database performance isolation

18.16.6 Create Performance Data Report function

Within the Performance category, there is a new Create Performance Data Report function. To use this function, from the left navigation pane, click **Performance** → **All Tasks** → **Reports** → **Create Performance Data Report**, as shown in Figure 18-96.




Figure 18-96 Reports - Create Performance Data Report

You have definitions that provide predefined sets of perspectives to create for a specific collection. A report definition can be selected along with the collection you are interested in to generate a .pdf or a .zip file of the perspectives, as shown in Figure 18-97. This definition makes it easy to export multiple charts or tables for collection at one time.




Figure 18-97 Create performance data report - definition

18.16.7 Image and file export

Within Performance Data Investigator, you can now export an image (charts only), a comma delimited file, or a tab delimited file. Figure 18-98 shows an example of exporting CPU Utilization and Waits Overview information to a comma delimited file. You can insert this data into a spreadsheet for later manipulation.




Figure 18-98 Exporting an image, comma delimited file, or a tab delimited file

This feature makes it possible to print or export your data to store it for later reference.

The export window has two main choices:

- ▶ **Format**, which has three options that can be specified by the user:
 - An image (*.png or .jpeg) file.
 - Comma-separated values (CSV), which is the default format.
 - A tab delimited (*.txt) file.
- ▶ **Data Range**, which has three options that can be specified by the user:
 - All data
This option exports the entire data set to the chosen format.
 - Displayed data
This option is the default and produces smaller results, but can omit off-window information.
 - User-defined range
You can use this option to specify exactly which records and which series you want to export.

18.16.8 Sizing the next upgrade

The Investigate Data feature can now send data from your session to the Workload Estimator for use in sizing a future system using current performance characteristics.

The IBM Systems Workload Estimator is a web-based sizing tool for IBM Power Systems, System i, System p®, and System x, and is available at the following website:

<http://www.ibm.com/systems/support/tools/estimator>

Use this tool to size a new system, to size an upgrade to an existing system, or to size a consolidation of several systems.

Figure 18-99 shows an example where measured data is taken from the collection services (in this case, from CPU Utilization and Waits Overview) and **Size next upgrade** is selected from the Actions drop-down menu. This action opens the IBM Systems Workload Estimator window.




Figure 18-99 Size next upgrade

It is possible to use filters to have only a subset of collected performance data sent to the Workload Estimator. The filtering can be done by time range. Complete the following steps to accomplish this task:

1. From the CPU Utilization Overview perspective that is shown in Figure 18-100, click the first icon at the upper left of the chart to enable the select tool.
2. Select minimum and maximum points in the chart. These points are shown as two red signs that are selected, in our case, at 6:15 a.m. and 12:15 a.m.

Important: You must select bars of the same metric (the color / pattern must match) or each selection clears the previous selections.




Figure 18-100 Time range

3. Select the **Size next upgrade** from the Action menu, as shown in Figure 18-101.
A window where the time range is defined as filtered before opens.




Figure 18-101 Subset of performance-related data sent to WLE

4. Click **OK** to send this subset of the data to Workload Estimator.

Note: The Size Next Upgrade action to start Workload Estimator (WLE) is now updated so that the metrics for Disk Read IO (bytes) and Write IO (bytes) calculations give more accurate information.

18.16.9 Ethernet Link Aggregation Modified perspectives

The Ethernet Protocol Overview perspective is available in the Collection Services package under the Communications folder.

The perspective charts now use a breakdown dimension to distinguish different lines, as shown in Figure 18-102.




Figure 18-102 Ethernet Protocol - aggregated link

In the Ethernet Lines Overview chart, you can click **Communications Overview for One Ethernet Line** from the Actions menu, as shown in Figure 18-103.




Figure 18-103 Communications Overview for One Ethernet Line

Port Resource Names: Within the QAPMETH file, the Port Resource Name is added to the end of the file after the ETMUPF field. There is one record per line per port per interval. With the introduction of link aggregation, line names do not need to be unique within an interval. A Port Resource Name should be used to uniquely associate records across intervals.

For more information related to Ethernet Link Aggregation, see 10.8, “Ethernet link aggregation” on page 430.

18.16.10 Workload Group perspectives

A new perspective named Dispatch Latency Totals by Workload Group is available in the Collection Services package under the Workload Group folder.

This chart shows an overview of workload group dispatch latency. It shows the total delay time for each workload group, which is the amount of time that threads that were ready to run could not be dispatched because of the group's maximum concurrent processor limit.

It has two available views (graphs):

- ▶ Dispatch Latency by Workload Group
- ▶ Dispatch Latency by Workload Group and Processors Assigned.

The Display Latency Totals by Workload Group view is shown in Figure 18-104.




Figure 18-104 Dispatch Latency Totals by Workload Group

The Dispatch Latency Totals by Workload Group and Processors Assigned view is shown in Figure 18-105.




Figure 18-105 Dispatch Latency by Workload Group and Processors Assigned

From these views, you can select a Workload Group and drill down to either of the following views:

- ▶ Dispatch Latency Totals by Thread for one Workload Group.
- ▶ Dispatch Latency for One Workload Group.

The workload capping functionality prevents a subsystem from getting more than the capacity of the number of cores that are allocated to that workload capping group. There is nothing that causes or ensures that the threads of a capped subsystem receive the full throughput allowed. The data can be analyzed to determine whether the workload capping group limit on the number of processors that are assigned are causing wait time for the threads to complete. This function does not ensure that with more processor cores that the workload group does not encounter other bottlenecks because the system is busy with other uncapped subsystem threads.

The purpose of the PDI perspectives is to show the statistics of the Workload Groups as collected by Collection Services.

For more information about Workload Groups, see 20.7, “IBM i workload groups” on page 791 and the IBM i developerWorks Wiki - IBM i Technology Updates - IBM i workload groups at:

<https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/IBM%20i%20Technology%20Updates/page/IBM%20i%20workload%20groups>

18.16.11 Java perspectives

You now have the following available Java perspectives:

- ▶ IBM Technology for Java Memory Overview
- ▶ IBM Technology for Java Memory by Job

The charts contain general information for active Java virtual machines (JVM).

Figure 18-106 shows the IBM Technology for Java Memory Overview perspective that shows average heap sizes and memory over the entire time of the collection broken down by job. You can use this perspective to discover what jobs are using the memory so that you can look at them in more depth afterward.




Figure 18-106 IBM Technology for Java - Memory Overview

You can now select the perspective IBM Technology for Java Average Memory by Job from the Java folder and also as a drill-down menu under the existing IBM Technology for Java Memory Overview, as shown in Figure 18-107.




Figure 18-107 IBM Technology for Java - Memory Overview by Job

This chart shows the average heap sizes and memory over the entire time of the collection broken down by job. You can use this chart to discover what jobs are using the memory so you can look at them later.

A new perspective named IBM Technology for Java Memory for One Job is available as a drill-down menu under the IBM Technology for Java Average Memory by Job perspective. You can use this perspective to look at the heap and memory usage over time for one selected job, as shown in Figure 18-108.




Figure 18-108 IBM technology for Java - Memory for One job

18.16.12 Disk Response Time charts

The Disk Response Time charts are a new addition to the IBM i 7.1 Collection Services Perspectives. These charts can be accessed by clicking **Disk** → **Disk Response Time** → **Detailed**, as shown in Figure 18-109. The charts show disk I/O metrics segmented by response time and are only available for data that is collected on IBM i 7.1 or later.




Figure 18-109 Disk Response Time charts

Within this perspective, the following detailed options are available:

- ▶ Disk I/O Rates Overview
- ▶ Disk I/O Rates Overview with Cache Statistics
- ▶ Disk I/O Average Response Time Overview
- ▶ Disk I/O Total Response Time Overview
- ▶ Disk I/O Total Service Time Overview

Disk I/O Rates Overview

Figure 18-110 shows the disk I/O that is segmented by the number of I/Os that occurred when the response time was in specific ranges, and the average response time and average service time.




Figure 18-110 Disk I/O Rates Overview - detailed

Disk I/O Rates Overview With Cache Statistics

Figure 18-111 shows the disk I/O that is segmented by the number of I/Os that occurred when the response time was in specific ranges, and cache statistics.




Figure 18-111 Disk I/O Rates Overview with Cache Statistics - detailed

Disk I/O Average Response Time Overview

Figure 18-112 shows the disk average response time that is segmented by the number of I/Os that occurred when the response time was in specific ranges, and average response time and average service time.




Figure 18-112 Disk I/O Average Response Time Overview - detailed

Disk I/O Total Response Time Overview

Figure 18-113 shows the disk total response time that is segmented by the number of I/Os that occurred when the response time was in specific ranges, and average response time and average service time.




Figure 18-113 Disk I/O Total Response Time Overview - detailed

Disk I/O Total Service Time Overview

Figure 18-114 shows the disk total service time that is segmented by the number of I/Os that occurred when the response time was in specific ranges, and average response time and average service time.




Figure 18-114 Disk I/O Total Service Time Overview - detailed

For more information about these enhancements, see "QAPMDISKRB" on page 253.

18.16.13 Interactive perspective development

You can develop new charts and tables in Investigate Data. You can add views, modify data series information, and modify SQL statements without leaving the Investigate Data window.

Adding views

To add views, complete the following steps:

1. At the bottom of the Investigate Data Perspectives window, click **Options**. The window that is shown in Figure 18-115 opens.




Figure 18-115 Investigate Data perspective - Options

2. Select the **Enable Design Mode** option and click **OK**. The options that are chosen here are persistent across connections for the user.

3. Select the **New Package** icon at the top of the window that is shown in Figure 18-116.




Figure 18-116 Creating a custom content package

4. On the New Package window, specify a name, and click **OK**. The package is created and displayed.

5. Click **New Perspective**. The New Perspective window (Figure 18-117) opens.




Figure 18-117 Add data series - Thresholds

6. Specify a name and a description for the perspective and click **Add** to add a view (Table or Chart) with a Data series and a corresponding Threshold.

7. Click **OK** in the Add View window and click **OK** in the New Perspective window. The customized perspective is saved and shows up in the main perspective tree list for further reference, as shown in Figure 18-118.




Figure 18-118 Customized package and perspectives added to the main perspective tree list

Modifying Data Series information

It is also possible to modify the Data Series information, as shown in Figure 18-119, by clicking **Edit** in the View pane in the Edit Perspective window.




Figure 18-119 Modifying data series

After you modify the view, click **OK** to save the view with the changed information.

Modifying an SQL statement

You can use this feature to view and modify the current SQL and rerender its output to the window. Figure 18-120 shows an example where **Modify SQL** is selected from the Actions drop-down menu on the CPU Utilization and Waits Overview chart. A window opens, where the user can modify the SQL statements. Click **OK** to save the modifications.




Figure 18-120 Modifying the SQL

Formatting enhancements were made to the Modify SQL window, making the text easier to read and understand. The Modify SQL action is available for charts and tables.

There is a new option that available on the Options panel that you can use to see SQL error messages when you use the Modify SQL function, as shown in Figure 18-121.

Investigate Data - Performance Data Investigator

Options

- ☒ Use patterns Use patterns where applicable in charts.
- ☒ Show charts Whenever possible, show charts instead of tables.
- ☒ Enable design mode Enable advanced features allowing design and development of new content.
- ☒ Show help Show help messages for many tasks.
- ☒ Show SQL error messages Show SQL error messages to user.
- ☐ Set table size Rows: Columns: Specify the number of visible rows and columns shown for tables.

Default library Specify the default library that will be used when a collection is selected.

- ☐ Use Collection Services configured library
- ☒ Use last visited library
- ☐ Use library:

OK Cancel

Figure 18-121 Modify SQL - new option

18.16.14 Metric finder

With the significant number of metrics that are provided by the Investigate Data feature, sometimes knowing which perspective contains the metric you want to see is not a simple task.

You can use the Metric Finder function to display perspectives that are based on a specific metric. This function is useful when you know what type of information you are looking for but do not know where to start, or where certain metrics are included.

To start this feature, click the **Search** button when you start Investigate Data, as shown in Figure 18-122. Then, choose your perspective by searching for a metric name.

Tip: You can specify a filter to limit the metric names included in the drop-down menu. The filter facilitates your search for one specific metric, without knowing its exact name. After you enter a filter, click **Apply Filter** to update the metric name list. After a perspective is selected, you can display it by clicking the **Display** button.




Figure 18-122 Metric finder

The tree-format of perspectives is replaced with the new search functionality. The Search button is replaced by a List button and is used to revert to the normal panel afterward.

18.16.15 Performance Explorer content package

A simple Performance Explorer content package is included to start the analysis of Performance Explorer data, as shown in Figure 18-123. You can use this feature to analyze the performance of an application program. You can use the Performance Explorer content package to view the PEX database files.




Figure 18-123 Performance Explorer content package

The following PEX Profile perspectives provide functions that are similar to what Profile Data Trace Visualizer (PDTV) offers:

- ▶ Profile by Procedure
- ▶ Profile by Component
- ▶ Hierarchical Trace Profile
- ▶ Job/Thread List

For more information related to PDTV, see the following website:

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

18.16.16 New metrics

Many perspectives were added to Investigate Data, and many new metrics were added to existing perspectives:

- ▶ SAN
- ▶ Virtual I/O
- ▶ Energy Management
- ▶ Communication data

SAN metrics

The Disk Response Time Charts are new in IBM i 7.1 Collection Services Perspectives, as shown in Figure 18-124. This feature is accessed by clicking **Disk** → **Disk Response Time** → **Detailed**. The charts show disk I/O metrics (including SAN-attached disks) that are segmented by response time and are only available for data that is collected on IBM i 7.1 or later.




Figure 18-124 SAN metrics

For more information about these buckets, see “QAPMDISKRB” on page 253.

Virtual I/O metrics

These charts show I/O operations rates and throughput for virtual I/O devices. Figure 18-125 shows a view of virtual I/O devices categories (Adapter, Disk, Optical, and Type).




Figure 18-125 Virtual I/O metrics

Figure 18-126 shows a chart with the Virtual I/O Overview displayed by device type.




Figure 18-126 Virtual I/O overview by device type

Energy management metrics

POWER6 systems had IBM EnergyScale™ features available (such as Power Saver Mode), which made the processor frequency and voltage vary over time.

For more information related to EnergyScale functionality on POWER6 and POWER7 processor-base systems, see the following website:

<http://www-03.ibm.com/systems/power/hardware/whitepapers/energyscale.html>

Use Active Energy Manager V4.2 (AEM) with IBM Systems Director 6.1 to set the power savings value for IBM Power Systems running IBM i 7.1. This power savings value is then used to achieve a balance between the power consumption and the performance of the Power Systems system.

For more information related to the AEM, see the following website:

http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/aem_420/rb0_main.html

Charts now have the Scaled CPU time displayed along with the CPU time. This metric is shown in the CPU Utilization Overview, as shown in Figure 18-127.




Figure 18-127 CPU use overview

Figure 18-128 shows the CPU rate (Scaled CPU: Nominal CPU Ratio) for a specific period.




Figure 18-128 CPU rate (Scaled CPU: Nominal CPU Ratio)

Communication metrics

These charts show communication lines traffic and errors for the active protocols on the partition. You can use them for a view of communication protocols and specific line information about the partition. Figure 18-129 has an example of the Ethernet Protocol Overview, where the kilobytes transmitted and received per second are shown for all Ethernet lines on the system, and the number of retries of frame and MAC errors.




Figure 18-129 Ethernet protocol overview

18.16.17 Miscellaneous enhancements

Miscellaneous changes were made to improve the overall capabilities of the performance tasks:

- ▶ Added date and time information to the Collection Name drop-down list.
- ▶ Improved integration with Active Jobs.
- ▶ Improved integration with System and Disk Status.
- ▶ New collection information at the top of each perspective.
- ▶ A new menu bar that allows for quicker navigation.
- ▶ More complete history data.
- ▶ Drill down information.

Added date and time information to the Collection Name list

The date and time information is added to the Collection Name drop-down list on the main Investigate Data window when you select a collection, as shown in Figure 18-130.




Figure 18-130 Collection name - date and time information

Improved integration with active jobs

To use this enhancement, from within the Performance category, select **All Tasks** and select **Active Jobs** from the list, as shown in Figure 18-131.




Figure 18-131 Improved integration with active jobs

From within the Active Jobs window, it is possible to go directly to performance-related information for a selected job, as shown in Figure 18-132.




Figure 18-132 Active Jobs Integration

Remember: When you click **Investigate Job Wait Data**, the chart that is rendered includes only data that pertain to that specific job. If you click **Investigate Jobs Wait Data**, the chart that is rendered includes all jobs that were active during that collection.

The following performance-related selections are available either from the menu at the right side of the selected job or from the Actions list at the top:

- ▶ Elapsed Performance Statistics
- ▶ Investigate Job Wait Data
- ▶ Start Job Watcher

Improved integration with system status and disk status

To use this enhancement, from the Performance category, click **All Tasks** and then either click **Disk Status** or **System Status** from the list, as shown in Figure 18-133.




Figure 18-133 Improved integration with disk status and system status

From within System Status and Disk Status, you can go directly to the performance-related information, as shown in Figure 18-134.




Figure 18-134 System Status and Disk Status integration

The following performance-related options are available:

- ▶ System Resources Health Indicators (from the General Tab on System Status)
- ▶ CPU Health Indicators (from the Processors tab on System Status)
- ▶ Memory Pools Health Indicators (from the Memory Tab on System Status)
- ▶ CPU Health Indicators (from the Processors Tab on System Status)
- ▶ Disk Health Indicators (from the Disk Space Tab on System Status)
- ▶ Investigate Disk Data (from the Actions list on Disk Status)
- ▶ Start Disk Watcher (from the Actions list on Disk Status)

New collection context at the top of each perspective

A new feature in Performance Data Investigator is that display collection information or collection details information can be displayed when you view a specific perspective.

By default, the name of the viewed perspective is shown as the title of the perspective window.

You might not want this information to be displayed. If so, see the View menu that has the **Show Context** check box, as shown in Figure 18-135. If you clear this check box, this information is hidden. This choice is preserved across sessions.




Figure 18-135 Collection information and collection details information

New menu bar

In IBM i 7.1, a menu system was added at the top of every perspective, which allows for quicker navigation, as shown in Figure 18-136.




Figure 18-136 New menu bar

The new menu at the top of the perspective has the same actions available as the one at the bottom, but it is available without needing to scroll down because the actions sometimes might be rendered off-window. As such, it improves the availability of the options.

More complete history data

In IBM i 7.1, the complete History data is available, as shown in Figure 18-137, when you click **History** in the menu bar at the top. Start from the Memory Pools Health Indicators perspective and drill down to the Page Faults Overview perspective and then to the Page Faults by Job User Profile perspective.




Figure 18-137 History data

From within that menu bar, it is possible to go back to the previous perspectives by clicking the corresponding item in the History Data.

Displaying holder information

From within the IBM Navigator for i, the Display holder information is available after you identify a thread or task that is holding others.

Within Investigate Data, you can click **Job Watcher** → **Waits** → **All Waits by thread or task Sorted by Non-idle waits**, as shown in Figure 18-138.




Figure 18-138 All Waits by thread or task Sorted by Non-idle waits

You can select one thread or task and click **All Waits for One Thread or Task** from the Actions menu, as shown in Figure 18-139.




Figure 18-139 All Waits for One Thread or Task

You now select one interval and perform the following actions, as shown in Figure 18-140:

- ▶ Click **Interval Details for One Thread or Task**.
- ▶ Click the **Show Holder** button for the thread or task.




Figure 18-140 Show holder information

When you click the **Show Holder** button, the holding job or task information is displayed.

Selected items on charts: On certain charts (for example, Waits for one job or task), you see which job, task, or thread was selected for drill-down.

18.17 JS1- Advanced Job Scheduler for i enhancements

In IBM i 7.1, IBM Navigator for i has an Advanced Job Scheduler function that is equivalent to the System i Navigator client.

For more information that is related to Advanced Job Scheduler enhancements, see Chapter 13, "IBM Advanced Job Scheduler for i enhancements" on page 487.

18.18 Backup Recovery Media Services: Plug-in enhancements

IBM Navigator for i had limited Backup Recovery Media Services (BRMS) functionality in IBM i 6.1. Its capabilities were expanded into a full-featured BRMS interface, effectively bringing this web interface into parity with the client-based System i Navigator product.

For more information related to Backup Recovery Media Services enhancements, see Chapter 4, “Backup and recovery” on page 51.

18.19 More information

Here are some additional resources about the topics that are described in this chapter:

- ▶ Connecting to IBM i - IBM Navigator for i PDF, found at the following website:
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzatg/rzatgdirect.or.pdf>
- ▶ *IBM Systems Director Navigator for i*, SG24-7789
- ▶ *End to End Performance Management on IBM i*, SG24-7808



IBM i Access

This chapter describes the following IBM i Access topics for IBM i 7.1:

- ▶ IBM i Access for Windows
- ▶ IBM i Access Client Solutions
- ▶ IBM Navigator for i
- ▶ System i Navigator
- ▶ IBM i Access for Web enhancements
- ▶ IBM System i Access for Wireless

19.1 IBM i Access for Windows

IBM i Access for Windows is a key offering of the IBM i Access Family. It provides Microsoft Windows connectivity to IBM i platforms. It extends IBM i applications and resources to a Windows desktop through a GUI. This section describes the IBM i Access for Windows changes for IBM i 7.1.

For more information about IBM i Access for Windows, go to the following website:

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

19.1.1 Installation enhancements

The following list details installation enhancements for IBM i Access for Windows for IBM i 7.1:

- ▶ Help files were converted to HTML help format. A separate download is no longer required on Windows Vista and later operating systems to display the help.
- ▶ New SSL signer certificates are merged automatically into the certificate management key database file during an upgrade installation.
- ▶ Install Language support was expanded. Secondary languages can be selected during the custom installation, and Arabic language support was added.

Support changes: Here are some additional changes:

- ▶ Support for Windows Server 2008 R2 and Windows 7 was added.
- ▶ Support for Windows 2000 was removed.
- ▶ Support for 64-bit Itanium processors was removed.
- ▶ The \QIBM\ProdData\Access\Windows\Image64i install directory was removed.

19.1.2 .NET Data Provider enhancements

The IBM i Access for Windows .NET Data Provider allows .NET managed programs to access the IBM i database files using SQL.

Here are the .NET Data Provider enhancements for IBM i Access for Windows for IBM i 7.1:

- ▶ 128-byte schema names
- ▶ Support for the IBM i XML Data type
- ▶ Connection property to configure Concurrent Access Resolution
- ▶ Support for multi-row **UPDATE**, **DELETE**, and **MERGE** statements
- ▶ Support Visual Studio 2008
- ▶ Online help through Visual Studio

19.1.3 OLE Data Provider enhancements

The IBM i Access for Windows OLE Data Provider supports record-level access and SQL access to IBM i database files. The following list details the OLE Data Provider enhancements for IBM i Access for Windows for IBM i 7.1:

- ▶ 128-byte schema names
- ▶ Support for the IBM i XML Data type
- ▶ The connection property that is used to configure Concurrent Access Resolution

19.1.4 Windows ODBC Driver enhancements

The IBM i Access for Windows ODBC Driver provides the application programming interfaces to connect to a database management system, run SQL statements, and to retrieve data. Here are the ODBC Driver enhancements for IBM i Access for Windows for IBM i 7.1:

- ▶ 128-byte schema names
- ▶ Support for the IBM i XML Data type
- ▶ The connection property that is used to configure Concurrent Access Resolution
- ▶ Support for multi-row **UPDATE**, **DELETE**, and **MERGE** statements

19.1.5 Data Transfer enhancements

The IBM i Access for Windows Data Transfer function is enhanced to support 128-byte schema names on IBM i 7.1.

19.1.6 Personal Communications Emulator enhancements

IBM i Access for Windows includes a 5250 emulator to access and run host applications on a Windows desktop. The 5250 emulator included with IBM i Access for Windows for IBM i 7.1 provides updated Display and Printer Emulation that are based on Personal Communications V6, with the following additional functions:

Autocopy	Automatically copies the selected text to the clipboard.
Copy As Image	Copies the marked area in the presentation space as a bitmap into the clipboard.
Find text in Presentation Space	Finds text in the Presentation Space of a display session. The text that is found is highlighted on the window. The search can be case-sensitive.
Quick connect for Telnet connections	You can connect a Telnet session quickly through the Quick Connect Bar by configuring only the Host, Port, and LU Name/Workstation ID. For more information, see <i>Administrator's Guide and Reference</i> , found at http://pic.dhe.ibm.com/infocenter/pcomhelp/v6r0/index.jsp?topic=%2Fcom.ibm.pcomm.doc%2Fbooks%2Fhtml%2Fadmin_guide.htm .

19.1.7 Direct Attach Operations Console withdrawal

Support for local console direct attached was removed from the Operations Console for IBM i Access for Windows for IBM i 7.1.

With IBM i 7.1, Operations Console supports only the local console on a network (LAN) configuration.

For more information, go to the IBM i Planning website at:

<http://www-947.ibm.com/systems/support/i/planning/upgrade/v6r1/planstmts.html>

19.2 IBM i Access Client Solutions

IBM i Access Client Solutions is the newest member of the IBM i Access Family (5733-XJ1). It provides a Java -based, platform-independent interface that runs on most operating systems that support Java 6.0 or higher, including Linux, Mac, and Windows. IBM i Access Client Solutions consolidates the most commonly used tasks for managing your IBM i into one simplified location.

It has the following features:

- ▶ 5250 display and printer emulation that is based on IBM Host-on-Demand.
- ▶ 5250 Session Manager support that is modeled after IBM Personal Communications Session Manager.
- ▶ Multiple language support for concurrent 5250 emulation sessions on the same client.
- ▶ Data Transfer similar to IBM i Access for Windows Data Transfer plus support for OpenDocument spreadsheet (*.ods), Excel Workbook (*.xlsx), and other file formats.
- ▶ Virtual Control Panel with a graphical interface to the IBM i control panel.
- ▶ 5250 emulation for LAN Console.
- ▶ Consolidation of hardware management interface configurations, including ASMI, IVM, and HMC.

IBM i Access Client Solutions uses the same IBM i host servers as the other IBM i Access Family products and requires the same IBM i Access Family license (XW1) to use the 5250 emulation and Data Transfer features.

At the time of the writing of this book, IBM i Access Client Solutions is a technology preview. You can download it, try it, and give your feedback for improvements or additions into the final product. However, most of the core functionality is there. You can test and evaluate the deployment flexibility of this member of IBM i Access Family using your desktop operating system of choice, such as Windows XP, Windows 7, Mac OS, or Linux distributions, such as Red Hat, SUSE, or Ubuntu, to name a few.

For more information about IBM i Client Access Solutions, go to the following website:

<http://www-03.ibm.com/systems/power/software/i/access/solutions.html>

19.2.1 Prerequisites

Before you start using IBM i Client Access Solutions, you must meet a few prerequisites.

Workstation prerequisites

IBM i Access Client Solutions runs on most operating systems that support Java V6.0 or later, including various versions of Linux, Mac OS, and Windows.

One way to check the version of Java that is installed on your system is to open a command prompt and then run the following command:

```
Java -version
```

Prerequisites (connectivity to IBM i)

IBM i Access Client Solutions can connect to IBM i 5.4 and later.

The following features require IBM i 6.1 or later:

- ▶ Navigator for i
- ▶ 5250 Console
- ▶ Virtual Control Panel

If you use the 5250 Console or the Virtual Control Panel feature for IBM i 6.1, load and apply the following PTFs:

- ▶ MF55543
- ▶ MF55549

If you use the 5250 Console or the Virtual Control Panel feature for IBM i 6.1 with machine code 6.1.1, load and apply the following PTFs:

- ▶ MF55540
- ▶ MF55547

If you use the 5250 Console or the Virtual Control Panel feature for IBM i 7.1, load and apply the following PTFs:

- ▶ MF55485
- ▶ MF55538

19.3 IBM Navigator for i

The IBM i long-term strategy for Systems Management is provided by IBM Navigator for i (IBM Systems Director Navigator for i).

Its latest enhancements include better performance and a renewed leading web interface, which provides usability that is comparable to Systems i Navigator (Windows client).

At the time of the writing of this book, most of the functions, including plug-ins, are available through the web interface (IBM Navigator for i) and Windows client interface (System i Navigator), with the following exceptions:

- ▶ Here are the functions that are available in the web interface only:
 - Performance Data Investigator (PDI)
 - PowerHA
 - Clustering Resource Services

Cluster Resource Services plug-in: Cluster Resource Services plug-in for System i Navigator from the High Availability Switchable Resources licensed program (IBM i option 41) was removed in IBM i 7.1.

- ▶ Here are the functions that are available in the Windows Client interface only:
 - Management Central functions
 - Database functions that involve graphics and charts:
 - SQL Scripts
 - Visual Explain
 - Database Navigator
 - SQL Assist

- AFP Manager:
 - Resources
 - PSF configurations
 - Font Mapping tables

IBM Systems Director: IBM Systems Director is the alternative to Management Central for multisystem tasks on the web.

For more information about Navigator for i, see Chapter 18, “IBM Navigator for i V7.1” on page 639.

19.4 System i Navigator

System i Navigator is the Windows client for IBM i Systems Management. It is packaged as part of the IBM i Access for Windows product. It is the more traditional Systems Management console, but continues to be supported and enhanced.

For more information about Navigator for i and System i Navigator, go to:

<http://www-03.ibm.com/systems/i/software/navigator/index.html>

19.5 IBM i Access for Web enhancements

IBM i Access for Web is an IBM i Access Family offering that provides browser-based access to IBM i resources. The resources include printing, messages, jobs, databases, and a 5250 emulator. The following list details the IBM i Access for Web changes for IBM i 7.1:

- ▶ IBM i Access for Web in a web application server environment:
 - IBM i 7.1 or later support for the AFP to PDF Transform.
 - IBM i 7.1 or later support for WebSphere Application Server V8.0 for IBM i (Base, Express, and Network Deployment editions).
 - IBM i 6.1 or later support for WebSphere Application Server V7.0 for IBM i (Base, Express, and Network Deployment editions).
 - IBM i 6.1 or later support for IBM integrated WebSphere Application Server for i instances.
 - Support for the WebSphere Application Server V6.0 for i5/OS (Base, Express, and Network Deployment editions) was removed.
- ▶ IBM i Access for Web in a portal environment:
 - IBM i 7.1 or later support for the AFP to PDF Transform.
 - IBM i 6.1 or later support for IBM WebSphere Portal V7.0, which requires PTF SI40905 or later for 5770-XH2.
 - Support for the IBM WebSphere Portal V6.0 and IBM WebSphere Portal V5.1 environments was removed.

For more information about IBM i Access for Web, go to the following website:

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

19.5.1 AFP to PDF Transform

IBM i Access for Web offers an alternative for generating a PDF document from a spooled file through the AFP to PDF Transform support on target systems with IBM i 7.1 or later. The following list details requirements for the AFP to PDF Transform support:

- ▶ The target system has IBM i 7.1 or later.
- ▶ 5770-TS1 *BASE (IBM Transform Services for i) and option 1 (Transforms – AFP to PDF Transform) are installed on the target system.
- ▶ 5770XH2 IBM i Access for Web or later is installed on the target system.

For IBM i Access for Web in a web application server environment, a new preference that is called *Use AFP to PDF Transform* is available to control whether the administrator or user can use the AFP to PDF Transform. The possible values are Yes or No (the default is Yes). To see this new preference, click **Customize** → **Preferences** → **View all preferences** and look under the Print category. The usage of this preference is similar to the usage of the “Use Infoprint Server if installed” preference in versions earlier than IBM i 7.1.

For IBM i Access for Web in a Portal environment, the JSR168 Print portlets are enhanced to use the AFP to PDF Transform. The PDF Output Settings (available in edit mode) have two new options:

- ▶ Use Infoprint Server if installed
The possible values are “Yes” and “No” (the default value is “Yes”).
- ▶ Use AFP to PDF Transform if installed
The possible values are “Yes” and “No” (the default value is “Yes”).

These options control the usage of Infoprint Server and the AFP to PDF Transform. They are similar to the policies and preferences that are used with the servlets. The IBM specification portlets are not enhanced.

The various methods of generating a PDF document from a spooled file are used in the following order:

- ▶ Infoprint Server is used if allowed (policy / preference) and if 5722IP1 *BASE is installed on the target system.
- ▶ AFP to PDF Transform is used if allowed (policy / preference) and if the target system has the following releases:
 - V7R1 or later IBM i
 - V7R1 or later 5770TS1 *BASE and option 1
 - V7R1 or later 5770XH2
- ▶ A built-in transformation that generates a PDF document with each page of the spooled file as an image.

To use the AFP to PDF Transform, click **Print** → **Printer output** and select the View PDF icon, as shown in Figure 19-1.




Figure 19-1 AFP to PDF Transform View PDF icon

The **PDF Output Settings** window opens (Figure 19-2).




Figure 19-2 PDF Output Settings window

The Destination drop-down menu offers the following options:

- **Browser**

The PDF document is sent to the browser. Based on the browser's settings, the document is either displayed by a plug-in or a File Save window opens.

- **Personal folder**

The PDF document is sent to one or more IBM i Access for Web Personal folders.

- **Output Queue**

The PDF document is sent to an output queue. Use the Settings button to specify the target output queue.

For more information about IBM i Access for Web AFP to PDF Transform, go to the IBM i Information Center at the following website:

<http://publib.boulder.ibm.com/series/>

19.6 IBM System i Access for Wireless

System i Access for Wireless (5722-XP1) is not supported on IBM i 7.1. For more information about this topic, go to the IBM i Planning website at the following address:

<http://www-947.ibm.com/systems/support/i/planning/upgrade/v6r1/planstmts.html>



Miscellaneous enhancements

This chapter describes the following IBM i 7.1 changes or enhancements:

- ▶ Licensed product structures and sizes
- ▶ Changed or new CL commands and APIs
- ▶ Temporary user-defined file systems
- ▶ Watch for Event function (message enhancements)
- ▶ IBM Tivoli Directory Server for IBM i enhancements
- ▶ Automating extra IPLs for PTF installation
- ▶ IBM i workload groups
- ▶ IBM i on a Flex Compute Node

20.1 Licensed product structures and sizes

The IBM i 7.1 operating system and licensed product (LPP) sizes are documented in the IBM i 7.1 Information Center in the section “Licensed program releases and sizes” at the following website:

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzahc/rzahcswslpsze.htm>

The DVD installation media is consolidated for IBM i 7.1 into three sets of multiple language version media that support a total of 51 globalizations.

With IBM i 7.1, there is no offline installation version of the IBM i Information Center that is shipped on physical media. The IBM i Information Center is available online at the following website:

<http://public.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp>

The following changes to the licensed product (LPP) structure are implemented in IBM i 7.1:

- ▶ DHCP moved from the base OS to option 31 “Domain Name System”.
- ▶ The Clusters GUI was withdrawn from option 41 “HA Switchable Resources” and is available with PowerHA for i (5770-HAS).
- ▶ IBM HTTP Server i (DG1) option 1 “Triggered Cache Manager” was removed.
- ▶ IBM Toolbox for Java (JC1) moved to 5770-SS1 option 3 “Extended Base Directory Support”.
- ▶ IBM Developer Kit for Java (JV1) options 6 (JDK 1.4) and 7 (JDK 5.0) are no longer supported. J2SE 6.0 32 bit is the default JVM in IBM i 7.1.
- ▶ Extended Integrated Server Support for IBM i (5761-LSV) is no longer supported. Option 29 “Integrated Server Support” is available as a replacement.
- ▶ IBM System i Access for Wireless (5722-XP1) was withdrawn. The IBM Systems Director family provides similar systems management functionality.
- ▶ IBM Secure Perspective for System i (5733-PS1 and 5724-PS1) was withdrawn, although it continues to be available as a custom service offering only.
- ▶ The IBM WebSphere Application Server (5733-W61 and 5733-W70) minimum required levels for IBM i 7.1 are 6.1.0.29 and 7.0.0.7.

For more information about LPP changes, see the *IBM i Memo to Users* at the following website:

<http://public.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzaq9/rzaq9.pdf>

Before you plan an IBM i release upgrade, see the IBM i upgrade planning website, which provides planning statements about IBM i product changes or replacements:

<http://www-947.ibm.com/systems/support/i/planning/upgrade/v6r1/planstmts.html>

20.2 Changed or new CL commands and APIs

For a list with detailed information about changed or new CL commands and APIs in IBM i 7.1, see the CL command finder and API finder section in the IBM i 7.1 Information Center at the following websites:

- ▶ <http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/apiref/new.htm>
- ▶ <http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rbam6/rbam6whatsnew.htm>

20.3 Temporary user-defined file systems

New support for temporary user-defined file systems (UDFSs) is included in the IBM i 7.1 base code and available for IBM i 6.1 through PTF SI34983.

Temporary UDFSs can increase performance by reducing auxiliary storage operations. Applications that create and delete many temporarily used stream files can most benefit from using temporary UDFSs.

For temporary UDFSs, the system allocates only temporary storage. These temporary files and directories are automatically deleted after an IPL, unmount, or reclaim storage operation. Although regular (that is, permanent) UDFSs can be created in any ASP or IASP, the temporary UDFSs are supported in the system ASP only.

Normally, the /tmp IFS directory contains permanent objects that are not cleared when the system is restarted. To have /tmp on IBM i behave more like other platforms, a temporary UDFS can be mounted over /tmp so that it is cleared at system restarts. The files in a temporary UDFS should not contain critical data because it is not persistent.

The **CRTUDFS** command and IBM Navigator for i are enhanced to support the creation of temporary UDFSs through a new naming convention. Although names for permanent UDFSs must end with **.udfs**, the names for the new temporary UDFSs adhere to the naming convention of **/dev/QASP01/newname.tmpudfs**, as shown in Figure 20-1.

Create User-Defined FS (CRTUDFS)

Type choices, press Enter.

User-defined file system > '/dev/QASP01/mytmpfs.tmpudfs'

Public authority for data . . .	*INDIR	Name, *INDIR, *RWX, *RW...
Public authority for object . .	*INDIR	*INDIR, *NONE, *ALL...
+ for more values		
Auditing value for objects . . .	*SYSVAL	*SYSVAL, *NONE, *USRPRF...
Scanning option for objects . .	*PARENT	*PARENT, *YES, *NO, *CHGONLY
Restricted rename and unlink . .	*NO	*NO, *YES
Default disk storage option . .	*NORMAL	*NORMAL, *MINIMIZE, *DYNAMIC
Default main storage option . .	*NORMAL	*NORMAL, *MINIMIZE, *DYNAMIC

Additional Parameters

Case sensitivity	*MONO	*MIXED, *MONO
Default file format	*TYPE2	*TYPE1, *TYPE2

More...

F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

Figure 20-1 IBM i CRTUDFS command for creating temporary file systems

The following additional considerations apply for using temporary file systems:

- ▶ Temporary objects cannot be secured by authorization lists.
- ▶ User journaling of temporary objects is not allowed.
- ▶ Objects cannot be saved from, or restored into, a temporary file system.
- ▶ Extended attributes are not supported for temporary objects.
- ▶ Object signing of temporary objects is not allowed.
- ▶ Read-only mount of a temporary file system is not supported.
- ▶ The storage that is used for temporary objects is not accounted to the owning user profile or to any process.

For more information about temporary user-defined file systems, see the IBM i 7.1 Information Center at the following website:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/ifs/rzaaxudftempudfs.htm>

20.4 Watch for Event function (message enhancements)

The Watch for Event function allows a user exit program to be started in a server job when certain events occur. The watch exit program does not run in the source job where the event occurs; it runs in a QSCWCHPS job in subsystem QUSRWRK. The following events can be watched:

- ▶ Messages sent to a program message queue (job log) of a job
- ▶ Messages sent to a nonprogram message queue (for example, a standard message queue such as QSYSOPR or a user profile message queue)
- ▶ LIC log entry (VLOG)
- ▶ Product Activity Log (PAL) entry

The function can also be used to end a trace when a watched event occurs. Watch parameters exist for the following trace CL commands:

- ▶ Start Trace (**STRTRC**)
- ▶ Start Communications Trace (**STRCMNTRC**)
- ▶ Trace Internal (**TRCINT**)
- ▶ Trace TCP/IP Application (**TRCTCPAPP**)
- ▶ Trace Connection (**TRCCNN**)

The Watch for Event function was initially available for the **trace** commands in V5R3M0 of IBM i. Watches were generalized in the following release, in V5R4M0, so that the watches were no longer tied only to **trace** commands. The Start Watch (**STRWCH**) and Start Watch API (**QSCSWCH**) commands were created for the generalized support. Additionally, the Work with Watch (**WRKWCH**) command was created to view watches, and the End Watch (**ENDWCH**) and End Watch API (**QSCEWCH**) commands were created to end watches. Support to watch for messages and LIC log entries was added in V5R4M0. Support to watch for PAL entries was added in V6R1M0.

The IBM i Information Center contains exit program information to describe all the parameters that are passed to a watch or trace exit program.

20.4.1 Advantages of using watches

Watches can be used to accomplish the following tasks:

- ▶ Capture data for program debugging.
- ▶ Automate system management by performing a corrective action or start / end functions without human intervention.
- ▶ Stop a trace when an event occurs to minimize the amount of data that is collected and minimize the amount of time trace active and slowing performance.
- ▶ Real-time notification of events can replace functions that used a periodic polling technique.

20.4.2 New Watch for Event function for IBM i 7.1

The message portion of the Watch for Event function has several enhancements that are described in the following sections. In addition, two new APIs to obtain watch information programmatically were created. The End Watch command and API were expanded to allow generic names, including *ALL.

Watch for Message enhancements

Functions were added to expand the message attributes that are allowed when you start a watch. There are new options for the message ID, message type, and severity in Version 7.1. These new attributes pertain to watching for messages, not LIC log or PAL entries, and allow specific watches to be created:

- ▶ Previously, a specific message ID needed to be specified when you watch for a message. Now, immediate or impromptu messages can be watched. The text of an immediate message does not exist in a message file. An example of an immediate message is The order application is down for 5 minutes starting at 11PM to load a fix. or Are you ready to go to lunch?
- ▶ Because immediate messages are now supported, this situation enabled support to watch for *ALL messages sent to a nonprogram message queue (such as QSYSOPR) or all messages sent to a program message queue for a job (job log).
- ▶ For predefined messages (that exist in a message file), a generic message ID can be specified, such as CPF18*, which allows messages CPF1806, CPF1808, and CPF1809 to be handled by starting one watch session.

- ▶ Message Type is a new attribute that is used for watching for messages. It allows certain message types to be watched with or without regard to other watch attributes.

For example, a message ID can be sent as a diagnostic and escape message, but if you want to act only when the message is sent as an escape, a watch can be tailored to that condition.

Another example is the need to be notified when any message is sent to a nonprogram (standard) message queue that was created with the Create Message Queue (**CRTMSGQ**) command. A watch can be started to watch for *ALL messages sent to the standard message queue. In the past, a program needed to use the receive message function with a wait time to obtain the next message sent to a standard message queue.

- ▶ Messages can be watched based on message severity. Valid message severities values are 0 – 99. Five relational operators can be specified with a severity value:
 - Equal to (*EQ)
 - Greater than (*GT)
 - Less than (*LT)
 - Greater than or equal to (*GE)
 - Less than or equal to (*LE)

For example, if you care only about messages of severity 99, you can now watch for only those messages.

Figure 19-2 shows the Watch for Message keyword with the **STRWCH** command for IBM i 7.1, with the new message type, relational operator, and severity code fields near the bottom portion of the window.

Start Watch (STRWCH)		
Type choices, press Enter.		
Session ID		Name, *GEN
Watch program		Name
Library	*LIBL	Name, *LIBL, *CURLIB
Call watch program	*WCHEVT	*WCHEVT *STRWCH *ENDWCH
Watch for message:		
Message to watch	*NONE	Name, generic*, *NONE...
Comparison data		
Compare against		*MSGDTA, *FROMPGM, *TOPGM
Message type		*ALL, *COMP, *DIAG...
Relational operator		*GE, *EQ, *GT, *LT, *LE
Severity code		0-99
+ for more values		
Bottom		
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display		
F24=More keys		

Figure 20-2 Start Watch command

Two new APIs

The new watch APIs are Retrieve Watch List (QSCRWCHL) and Retrieve Watch Information (QSCRWCHI). QSCRWCHL obtains a list of watches on the system. Previously, the list of watches were viewed by using the Work with Watches (**WRKWCH**) command. QSCRWCHI returns information about a specific watch session. The information that is returned is similar to the data seen by running the **WRKWCH** command and using option 5 to display the details of a watch session.

End Watch additions

The End Watch function requires a session ID to be specified. The End Watch (**ENDWCH**) command and End Watch (QSCEWCH) API are enhanced to accept a generic session ID, including a value of *ALL. Previously, the valid values included a specific session ID name and the special value *PRV, which represented the watch session that was started most recently by the same user who is running the End Watch function. In IBM i 7.1, a generic name (such as TSTMSG*, WCH*, or *ALL) can be specified. The new values make it easier to end a group of watches. For example, a generic name of TSTMSG* specifies that all watch sessions with identifiers that begin with the prefix TSTMSG will be ended.

Recursive watches

Do not watch for common or high use events. For example, if the job started message (CPF1124) or job ended message (CPF1164) are watched, the system is significantly affected. This watch never ends, because when a watch exit program is called in the server job, that job ends and a new job is started to handle subsequent processing for watches. The ending of the job and starting of the new job causes the CPF1164 and CPF1124 to be generated, so a never-ending loop begins. A similar thing happens if *ALL messages in QHST were watched because those job messages for watch processing go to QHST. To prevent this recursive problem, start watches for specific message IDs or specify compare data or addition selection criteria to restrict the occurrence of the event.

20.5 IBM Tivoli Directory Server for IBM i enhancements

This section covers the enhancements to the IBM Tivoli Directory Server for IBM i (LDAP).

20.5.1 Creating suffix entries automatically when they are needed

The directory administrator can configure a new suffix dynamically and start adding entries beneath it. If the suffix entry does not exist, it is created when the first child entry is added.

20.5.2 Administrative roles

IBM Tivoli Directory Server for IBM i now implements a scheme where the root administrator can delegate tasks at a more granular level, which is based on the administrative roles of the users that are defined in the configuration file. These roles are applicable only to the admin group members. Six roles are supported by IBM i:

- ▶ Audit Administrator (AuditAdmin)
- ▶ Directory Data Administrator (DirDataAdmin)
- ▶ No Administrator (NoAdmin)
- ▶ Replication Administrator (ReplicationAdmin)
- ▶ Schema Administrator (SchemaAdmin)
- ▶ Password Administrator (PasswordAdmin)

20.5.3 User interface enhancements

Tivoli ships the Web Administration Tool interface with Version 6.2. The Web-enablement for LDAP interface on IBM i Navigator enhancement enables the usage of an LDAP management tool on IBM Navigator for i and IBM i Navigator Tasks for the Web.

20.5.4 Security enhancements

Attribute encryption can encrypt arbitrary attributes when they are stored in the underlying directory database.

20.5.5 New password encryption options

There are two new password encryption options that are supported:

- ▶ Salted SHA
- ▶ MD5

20.5.6 Pass-through authentication

If an LDAP client tries to bind to the Tivoli Directory Server and the credential is not available locally, the server attempts to verify the credential from an external directory server on behalf of the client.

20.5.7 Enhanced password policy to use global date and time for initialization

The proposed design change for the initialization of password policy attributes when the Password Policy function is first turned on is to introduce a new password policy entry attribute, `ibm-pwdPolicyStartTime`, which is added to the `cn=pwdPolicy` entry. This attribute is generated by the server when the administrator sends a request to turn on the Password Policy function. The current time is put into this attribute. This attribute is an optional attribute that cannot be deleted by a client request. It cannot be modified by a client request, except by administrators with administrative control. It can be replaced by a master server-generated request. The value of this attribute is changed when the Password Policy function is turned off and on by an administrator.

20.5.8 Multiple password policies

In this release, more options are available. In addition to the global password policy, each user in the directory can have their own individual password policy. Furthermore, to assist administrators, a group password policy is supported to enable effective password management.

20.5.9 Policy that is enforced for Digest-MD5 binds

The implementation of this feature ensures password policy rules such as account lockout, usage of grace logins, and a password expiration warning message is sent to a user when it uses DIGEST-MD5 bind as authentication mechanism.

In addition, the `ibm-slapdDigestEnabled` configuration option is added to enable and disable the DIGEST-MD5 bind mechanism.

20.5.10 Persistent search

Persistent search provides a function for clients to receive notification of changes that occur in the directory server by altering the standard LDAP search operation so that it does not end after the initial set of entries that match the search criteria are returned. Instead, LDAP clients can keep an active channel through which information about entries that change is communicated.

20.5.11 Replication configuration enhancements

The server configuration attributes master DN and password in the consumer server's configuration is now dynamic. For the extended operation readconfig, addition, deletion, or modification of entries that have an objectclass of ibm-slapdReplication/ibm-slapdSupplier is supported for the scopeValues of entire/entry/subtree.

20.5.12 Filtered replication

This enhancement allows the directory administrator to control what data is replicated to consumer servers by specifying which entries and attributes are to be replicated, based on the filters that are defined by the directory administrator.

20.5.13 Limiting the number of values that are returned by a search

The LDAP server provides a control that can be used on a search operation to limit the total number of attribute values that are returned for an entry and to limit the number of attribute values that are returned for each attribute in the entry.

20.5.14 Enhanced syntaxes and matching rules

Additional matching rule and syntax support (24 syntaxes and 17 matching rules) were added for new syntaxes and matching rules from RFC 2252, RFC 2256, and RFC 3698. Matching rules are not defined in any RFC, but are referenced in RFC 2798.

20.5.15 IASP enablement for Directory Server on IBM i

From IBM i 7.1 onward, the Directory Server on IBM i supports private IASPs, where:

- ▶ The support database library is in the IASP.
- ▶ The support change log library is in the IASP.

20.5.16 The ldapcompare utility

The `ldapcompare` utility compares the attribute value of an entry with a user provided value.

20.5.17 Providing a re-entrant LDAP C client library

The LDAP C client library is now re-entrant.

20.6 Automating extra IPLs for PTF installation

A new function in IBM i 6.1 and IBM i 7.1 automates any extra IPLs that are required for a technology refresh PTF or special handling of pre-apply PTFs during the PTF installation process. If an extra IPL is required for a technology refresh PTF, your PTF installation parameters are saved and used during the next IPL. Instead of seeing the “Confirm IPL for Technology Refresh or Special Handling PTFs” panel, you see a new English message CPF362E: “IPL required to complete PTF install processing”. If you select Automatic IPL=Y on the “Install Options for PTFs” panel, you do not see any messages or panels; the server shuts down. On the next normal IPL, your second “GO PTF” occurs during the “PTF Processing” IPL step in the SCPF job, and then a second IPL of the partition occurs automatically. So when the system undergoes an IPL the second time, all the way up to sign on, your PTFs are all activated and ready to go.

Your total PTF installation time is shorter because none of the system jobs start during the first IPL when the partition restarts. However, you have a longer IPL time because the system is doing the work you previously did interactively, that is, the second GO PTF to set all PTFs for delayed applies.

If you wonder why the IPL is taking so long and what it is doing, you can always open the console. When you open the console, you can open the “Operating System IPL in Progress” panel, which shows the “PTF Processing” IPL step as active. Previously, the “Applying PTFs” for the Activity were shown in this step, but now you also see “Loading PTFs” or “Setting IPL Action for PTFs” if the previous PTF installation was incomplete.

When all PTFs are set for delayed apply, you see the IPL requested by PTF processing status message at the bottom of the panel and then the partition restarts to apply the delayed LIC PTFs. The next time that you reach the “PTF Processing” IPL step, you see the usual “Applying PTFs” step and the IPL continues.

To take advantage of this new function, you must have the following PTF (PTF management code) temporarily applied *before* you perform your PTF installation:

- ▶ Version 7.1: SI43585 in HIPER PTF group SF99709 level 30 or higher
- ▶ Version 6.1: SI43939 in HIPER PTF group SF99609 level 94 or higher

For Version 7.1, if an IPL is required for a technology refresh PTF, the new function supports installing only from a virtual optical device or *SERVICE (PTFs downloaded electronically to save files). If you are installing from a physical optical device, you still must perform the extra IPL and second GO PTF manually. So, if you received your PTFs on physical DVDs, create an image catalog from the DVDs and use the new support.

20.7 IBM i workload groups

IBM i 7.1 now provides workload groups, which were formerly called workload capping. Workload groups can restrict a workload to a specified maximum number of processor cores within the partition it is running in.

A workload is defined as a job, subsystem, or product that is running on the IBM i system. The user or system administrator can define a workload group, assigning a specified number of processing cores to that group. The workload group is then assigned to a job or subsystem. After the assignment is done, the workload is limited to the defined number of processing cores. The system enforces this processing core assignment, ensuring that a job or all the jobs that are running (and threads) under the subsystem are not allowed to run on more processing cores than are designated. An example is shown in Figure 20-3. In Figure 20-3, the number of processing cores of the subsystem for Application #1 is set to three and for Application #2 is set to six. No matter how much Application #1 needs processing capabilities that exceed three, it is not allowed to run on more processing cores than three cores, even though IBM i has eight cores. The general concept is if a workload is designated to use a single core, the workload behaves as through it is actually running on a single processor core system.




Figure 20-3 Concept of Workload Groups

20.7.1 Example of how Workload Group works

A user has a batch job that is processor intensive. This user must run this job during the day but cannot afford to impact the performance of the production system. By assigning this batch job to a workload group, this job can be put into a “processing container” to help ensure that this job is kept to a limited amount of system capacity. If the workload group has a processor core limit of one, then the batch job and any threads that are running under that job are allowed to run only on a single processor core. If this job is running on a multiple threaded core, multiple threads can be running for that designated batch job, but only a single core is used at a time. This same concept also applies to jobs that are running under a subsystem that is assigned to a workload group. All jobs and their associated threads are limited to the number of processor cores that are specified in the workload group.

This new capability can help users get better control of the workloads on their systems along with ensuring products are using only a designated number of processor cores. Software vendors can take advantage of the workload group support as a new virtualization technology. A workload can be virtualized and licensed within a virtualized system. Product entitlements can be specified based on the usage of the product instead of the total processor cores of the LPAR.

Customers who want to take advantage of the enhanced licensing controls must register the specified products with the native IBM i License Management tool that facilitates both the registering and management of the enforcement of the workload groups. To help users manage and understand the performance of jobs that are running in a workload group, the performance metrics are updated to include metrics on workload groups.

To learn more about the workload groups support, visit the IBM i Information Center at the following website:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzaks/rzaksworkloadcapping.htm>

20.7.2 IBM Passport Advantage Sub-capacity Licensing support

IBM i Workload Groups is now an Eligible Virtualization Technology for IBM Passport Advantage® Sub-capacity Licensing. You can use Passport Advantage Sub-capacity Licensing to license an eligible software product for less than the full capacity of a server. It provides the licensing granularity that is needed to use various multi-core chip and virtualization technologies. As more customers take advantage of virtualization technologies in their server environments, IBM continues to introduce enhancements to the Subcapacity Licensing offering.

For details and current information about Sub-capacity Licensing terms, go to the following websites:

<http://www.ibm.com/software/lotus/passportadvantage/subcaplicensing.html>
<http://www.ibm.com/software/tivoli/products/license-metric-tool>
http://www.ibm.com/software/lotus/passportadvantage/pvu_licensing_for_customers.html

Configuring WebSphere MQ to use workload groups

You can learn how to configure your systems to use workload groups for the WebSphere MQ product. You can limit the amount of processing capacity available to this product to better fit your licensing needs based on the capacity that is being used for this product. To see the details, see the document *WebSphere MQ and Workload Groups Final*, found at:

<https://www.ibm.com/developerworks/mydeveloperworks/wikis/form/anonymous/api/library/beb2d3aa-565f-41f2-b8ed-55a791b93f4f/document/d0e23be6-8d9d-4739-9f8c-fbfcd73030f/attachment/b141097b-30b0-4c2a-8603-6f069736b9d0/media/WebSphere%20MQ%20and%20Workload%20Groups%20Final.pdf>

20.8 IBM i on a Flex Compute Node

IBM i can run on a Flex Compute Node since April 2012.

The following two models of the POWER processor-based Flex Compute Node support IBM i.

- ▶ IBM Flex System™ p260 Compute Node Model 7895-22X
- ▶ IBM Flex System p460 Compute Node Model 7895-42X

For the hardware details, see the following website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.7895.doc/product_page.html

The following releases of IBM i are supported on the POWER processor-based Flex Compute Node.

- ▶ IBM i 6.1 Resave RS610-10 with License Internal Code 6.1.1 Resave RS-611-H or later
- ▶ IBM i 7.1 with Technology Refresh 4 or later

To configure an IBM i environment on those Compute Nodes, VIOS is required. Any physical adapters that are installed on the node are owned by VIOS, which is also the case with POWER processor-based blade servers. To see the details of IBM i on a Flex Compute Node, including configuring IBM i, go to the developerWorks website at:

<https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/IBM%20i%20Technology%20Updates/page/IBM%20i%20on%20a%20Flex%20Compute%20Node>

You can also run IBM i on the POWER processor-based Compute Node in the IBM PureFlex™ System. When you order IBM PureFlex System, a dedicated management appliance, which is called Flex System Manager (FSM), is bundled with the IBM PureFlex System. FSM is used for managing and operating an IBM PureFlex System, including hardware, firmware, virtualization environment, and operating systems environment on a POWER processor-based Compute Node.



Upgrading IBM i

Consistent with OS/400 and i5/OS in the past, IBM i 7.1 supports n-2 for previous release compatibility. More importantly, this support is also what allows IBM i 5.4 and IBM i 6.1 to have an upgrade path to IBM i 7.1. However, when you upgrade from IBM i 5.4 to IBM i 7.1, there are special considerations that pertain to object conversions, which are much the same as when you upgrade to IBM i 6.1 from a previous release. This chapter highlights several of these planning details, although you should review the *i5/OS V6R1 Memo to Users* for more details and *IBM i Program Conversion: Getting Ready for 6.1 and Beyond*, REDP-4293, which provides a complete set of planning considerations.

This chapter also contains information that highlights new features and enhancements available in IBM i 7.1, which relate to installing, upgrading, distributing software, and maintenance options for IBM i, and changes in licensed programs (LPP) support when you upgrade to IBM i 7.1:

- ▶ Installing or upgrading IBM i
- ▶ Upgrading from i5/OS 5.4 to IBM i 7.1
- ▶ Media delivery changes
- ▶ IBM i network upgrade
- ▶ More considerations for upgrading to IBM i 7.1
- ▶ IBM i network installation

21.1 Installing or upgrading IBM i

The following sections describe planning considerations when you upgrade to IBM i 7.1. Starting in IBM i 7.1, the IBM i Information Center is no longer available on physical media that allow it to be installed on a System i system, a Power Systems system, or personal computer. You can access the IBM i Information Center at the following website:

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

Be familiar with the following essential websites and documents:

- ▶ *Memo to Users*. For IBM i 7.1, this memo is available at the following website:
<https://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzaq9/rzaq9.pdf>
- ▶ IBM i planning website:
<http://www-947.ibm.com/systems/support/i/planning/upgrade/index.html>
- ▶ *Installing, upgrading, or deleting IBM i and related software*, found at:
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp?topic=/rzahc/rzahc1.htm>

21.1.1 IBM i 7.1 support for POWER5 and later hardware

To determine whether your existing hardware is supported by IBM i 7.1, review Figure 21-1. IBM i 7.1 supports POWER5 and later hardware. For more information, see the Upgrade planning website at:

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

Servers	IBM i 5.4	IBM i 6.1	IBM i 7.1
POWER7 BladeCenter PS700, PS701, PS702, PS703, PS704		✓	✓
POWER7 Power 710, 720, 730, 740, 750, 770, 780, 795		✓	✓
POWER6 BladeCenter JS12, JS22, JS23, JS43		✓	✓
POWER6/6+ Power 520, 550*, 560, 570, 595	✓	✓	✓
POWER5/5+ 515, 520, 525, 550, 570, 595	✓	✓	✓
800, 810, 825, 870, 890	✓	✓	
270, 820, 830, 840	✓		

Figure 21-1 IBM i 7.1 hardware model support

Requirement: IBM i 6.1 or later is required to migrate to IBM i 7.1 for a POWER6+™ 550 system.

For enterprise clients, IBM i 7.1 is now supported on the 16-way through 256-way POWER7 795. IBM i supports up to 32 cores in a single partition. You can contact IBM Lab Services about an offering to grow beyond 32 cores in a single partition.

21.1.2 Planning statements

Planning statements provide insight into IBM current plans, directions, and intent, and are subject to change or withdrawal without notice. Statements pertaining to IBM i 7.1 can be found on the *Upgrade Planning statements* website at:

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

21.1.3 Supported upgrade paths

IBM i 7.1 supports n-2, which means a direct upgrade to IBM i 7.1 from IBM i 5.4 or IBM i 6.1 is supported. See Figure 21-2.




Figure 21-2 Supported upgrade paths to IBM i 7.1

21.1.4 Program temporary fixes for upgrading to IBM i 7.1

As you plan your IBM i installation or upgrade, ensure that you review Information APAR II14482 for IBM i 7.1, which is available in the preventive service planning (PSP) section.

To access the PSP, click **Technical Databases** → **Authorized Program Analysis Reports (APARs)** at the following website:

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

21.1.5 Pre-upgrade verification tool

This tool was initially introduced with IBM i 6.1 and was updated for IBM i 7.1. The tool runs on a Windows client and checks the IBM i environment to ensure that all requirements are completed to help ensure a successful upgrade. The tool is available at the following website:

http://www-912.ibm.com/s_dir/s1kbase.NSF/DocNumber/465353483

21.1.6 License program releases and sizes

Operating system and license program sizes are documented in the “License program releases and sizes” section at the following website:

<https://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzahc/rzahcswslpsze.htm>

21.1.7 Server firmware requirements

Before you upgrade, verify the server firmware level that is supported on your POWER5, POWER6, or POWER7 processor-based server. This firmware must be at a certain minimal level to support IBM i 7.1.

To find the firmware level and IBM i OS levels that are needed as prerequisite requirements for features that you have or plan to add to your system, go to the IBM Prerequisite website at:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzahc/rzahcverifyingfirmwarerequirements.htm>

21.1.8 IBM i 7.1 Resaves

IBM periodically releases new Resave media. Resaves are often Machine Code (57xx-999), but occasionally a Resave is made available for IBM i (formerly known as i5/OS) (57xx-SS1).

For the mapping between the SLIC Resave level and the IBM i Resave level, see Table 21-1.

Table 21-1 IBM i 7.1 Resave history (with 7.1.0 Machine Code)

Resave release date	Description	5770-999 Resave level marker PTF	5770-SS1 Resave level marker PTF
10/12/2012	IBM i 7.1 Technology Refresh 5 (TR 5)	RS-710-F RE12249	MF99005
05/18/2012	IBM i 7.1 Technology Refresh 4 (TR 4)	RS-710-E RE12066	MF99004
10/14/2011	IBM i 7.1 Technology Refresh 3 (TR3)	RS-710-D RE11221	MF99003
05/13/2011	IBM i 7.1 Technology Refresh 2 (TR2)	RS-710-C RE11067	MF99002
09/10/2010	IBM i 7.1 Technology Refresh 1 (TR1)	RS 710-B RE10187	MF99001

21.2 Upgrading from i5/OS 5.4 to IBM i 7.1

For V5R4 to IBM i 7.1 upgrades, unique conversions occur within the programs in libraries, Java programs in directories, spooled files, and integrated file system names in file systems that are not case-sensitive.

21.2.1 Object conversions

A requirement for program conversions was introduced with IBM i 6.1. If your system is on release 6.1, there is no additional program conversion requirements for upgrading to IBM i 7.1.

If you are upgrading from i5/OS 5.4 to IBM i 7.1, the same object conversion considerations apply as though the target release were IBM i 6.1. Read the “Program conversion” section in the *i5/OS Memo to Users* for V6R1, which is available at the following website:

<http://publib.boulder.ibm.com/infocenter/iserics/v6r1m0/topic/rzaq9/rzaq9.pdf>

The program conversions refresh programs to take advantage of the latest system enhancements. Program conversion includes the conversion of programs in libraries and conversion of Java programs in directories. The conversions of Java programs in directories, however, do not affect the actual IBM i upgrade time.

When you upgrade from IBM i 5.4, allow more time to analyze your system and adjust your programs for conversions. The length of time that is required to perform the analysis varies based on the individual system environment. Program conversion can also affect vendor software. These vendors are contacted as part of the upgrade planning, as they might need to verify their applications support IBM i 6.1 or 7.1.

The **ANZOBJCVN** command was introduced for i5/OS 5.3 and i5/OS 5.4 to help object conversion planning for upgrades to IBM i 6.1. This command can also be used for upgrades to IBM i 7.1. The command is available through a set of PTFs. Information APAR II14306 provides a brief description of **ANZOBJCVN** and PTF requirements. For IBM i 5.4, PTF SI39402 adds the option to specify a target release of V7R1M0. To review Information APAR II14306, see 21.1.4, “Program temporary fixes for upgrading to IBM i 7.1” on page 797.

For complete preparation and planning details, see *IBM i Program Conversion: Getting Ready for 6.1 and Beyond*, REDP-4293.

The Information Center has detailed options and instructions at the following website:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzahc/rzahcswsv5r1con.htm>

21.2.2 Spooled file conversions

When you upgrade from IBM i 5.4 to IBM i 7.1, spooled file operations are processed more efficiently than in previous releases because of conversions that can be done either during the release upgrade or after the upgrade.

By default, conversion occurs during the upgrade, which can add a significant amount of time. **ANZOBJCVN** can help by identifying the number of spool files and providing an estimate for conversion time, which can help you determine your best options.

Spooled files that are restored to the IBM i 7.1 release are automatically converted. It is possible that the time for the spooled file conversion process can be reduced by saving and deleting the spooled files before you upgrade from IBM i 5.4 and then restoring them after you have IBM i 7.1 installed.

More options are available for managing the spool file conversion after the upgrade. The Information Center has detailed options and instructions at the following website:

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzahc/rzahcspooledfileconversions.htm>

Important: Only newly created spooled files or spooled files that were converted can be seen and used after the upgrade. Until the conversion is complete, unconverted spooled files appear not to exist. If a data area is used to direct the conversion, delete the data area after the conversion occurs.

21.2.3 Integrated file system conversions

As of IBM i 6.1, file systems that are not case-sensitive in the integrated file system support Unicode Standard 4.0 for names that are stored in directories.

In IBM i 5.4, the integrated file system stores all names in Unicode and supported Unicode Standard 2.0. After you upgrade from IBM i 5.4, an automatic conversion runs to update the integrated file system directories, in file systems that are not case-sensitive, to support Unicode Standard 4.0. Unicode Standard 4.0 defines more characters and updated casing rules. The file systems included in this conversion are “root” (/) and user-defined file systems (UDFS). These file systems can exist in any auxiliary storage pool that is created with the parameter value CASE(*M0N0) on the Create User-defined File System (CRTUDFS) command.

Before you upgrade from IBM i 5.4, review Information APAR II14306 and *IBM i Program Conversion: Getting Ready for 6.1 and Beyond*, REDP-4293. These resources help you analyze your system and help identify objects that are going to be affected by the Unicode conversion. You can then decide whether you want to change the names of the affected objects before you upgrade or allow the automatic conversion to occur. To view the information APAR, see 21.1.4, “Program temporary fixes for upgrading to IBM i 7.1” on page 797.

The conversion of the directories automatically begins shortly after IBM i 7.1 is installed. This conversion runs in the background during normal operations and does not significantly affect your system activity.

The Information Center has detailed options and instructions at the following website:

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzahc/ifsconv.htm>

21.2.4 Backup Recovery and Media Services

If you are upgrading Backup Recovery and Media Services (BRMS), 5770-BR1, product initialization is required before you can use the BRMS functions. BRMS product initialization was removed from the installation process to improve the time that is required to complete software upgrades that include BRMS.

The Information Center has detailed options and instructions at the following website:

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzahc/br1conv.htm>

21.3 Media delivery changes

IBM i 7.1 provides a simplified ordering and software delivery process. Additional secondary language media can still be ordered, if wanted. There is no change for LPPs that are ordered and shipped individually; they continue to ship with all NLVs.

21.3.1 IBM i 7.1 multi-language distribution media

DVD installation media is consolidated. Previously, there were 51 sets of unique language version media. These sets were reduced to three pre-packaged sets of multiple language version media.

IBM i 7.1: Three language groups for a keyed media set

This package has the following characteristics. See Table 21-2.

- ▶ One language group delivers 11 - 21 NLVs, depending on the group ordered.
- ▶ US English is included in each language group.
- ▶ Install as many primary and secondary languages as wanted from the group.
- ▶ For a release upgrade, the upgrade installs the same language that is used with the current release.
- ▶ On a manual installation, select the language wanted and then proceed with the installation.
- ▶ Full automatic installation and upgrades from optical media that are changing the primary language must set the installation language using the QINSTLNG API.
- ▶ You can order more than one group, if wanted, using 5770-NLV.

Table 21-2 IBM i 7.1 media language groups

Group 1 - 5770-SS1 Feature 5817		Group 2- 5770-SS1 Feature 5818		Group 3- 5770-SS1 Feature 5819	
2924	English	2924	English	2924	English
2963	Belgian Dutch MNCS	2995	Albanian	2938	English DBCS
2909	Belgium English	2954	Arabic	2984	English DBCS
2966	Belgian French MNCS	2974	Bulgarian	2930	Japanese Universal
2980	Brazilian Portuguese	2912	Croatian	2962	Japanese Katakana
2981	Canadian French MNCS	2975	Czech	2986	Korean
2926	Danish	2903	Estonian	2906	Laotian
2923	Dutch Netherlands	2998	Farsi	2989	Simplified Chinese
2925	Finnish	2957	Greek	2987	Traditional Chinese
2928	French	2961	Hebrew	2972	Thai
2940	French MNCS	2976	Hungarian	2905	Vietnamese
2929	German	2904	Latvian		
2939	German MNCS	2903	Lithuanian		
2958	Icelandic	2913	Macedonian		

Group 1 - 5770-SS1 Feature 5817		Group 2- 5770-SS1 Feature 5818		Group 3- 5770-SS1 Feature 5819	
2932	Italian	2978	Polish		
2942	Italian MNCS	2992	Romanian		
2933	Norwegian	2979	Russian		
2922	Portuguese	2914	Serbian		
2996	Portuguese MNCS	2994	Slovakian		
2931	Spanish	2911	Slovenian		
2937	Swedish	2956	Turkish		

For more information, see *Media labels and their contents* in the IBM i 7.1 Information Center at the following website:

<https://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp?topic=/rzahc/rzahcswsmedialabel.htm>

21.3.2 Media charges

Effective April 13, 2010, IBM charges a \$50 USD fee for orders on physical media. This fee applies to all initial and upgrade orders for currently available releases of IBM i, which include IBM i 5.4, IBM i 6.1, and IBM i 7.1. Each additional order for physical media also incurs an extra \$50 USD fee. There is no media charge if you use Electronic Software Delivery (ESD).

Expanded ESD support

IBM continues to expand the countries for which electronic delivery is supported. For a complete list, and instructions for using the ESD website, go to the following website:

http://www-947.ibm.com/systems/support/software/delivery/en_US/downloadinfo.html

Additionally, a new API called Fill Image Catalog (QVOIFIMG) was made available for IBM i 5.4, IBM i 6.1, and IBM i 7.1 through PTFs. This API makes it easier to use image catalogs when you work with images that are downloaded through the ESD process. Information APAR II14482 *Required PTFs for Upgrading to V7R1MX* includes the specific PTF numbers for each of these releases

To review Information APAR II14482, see 21.1.4, “Program temporary fixes for upgrading to IBM i 7.1” on page 797.

21.4 IBM i network upgrade

In October 2009, IBM announced that an IBM i 6.1 environment on a POWER6 processor-based server can be upgraded to IBM i 6.1.1 or IBM i 7.1 remotely using installation images on a network file server. This enhancement extends the ability to perform not only upgrades, but installations and maintenance activities for remote systems.

Before this enhancement, you needed physical media or virtual media that are located locally on the system that is being upgraded. Using virtual media once required using FTP to manually download the virtual images across the network to the individual systems to be installed.

The Network Filer Server (NFS) system is the repository for the virtual images, and can be any NFS system that can meet the basic requirements. On the IBM i client system, this new function takes advantage of the 632B-003 virtual optical device that supports virtual image files on a remote system in a network. An image directory identifies a network path on the central system that contains the virtual image files that are prepared for use with a target system.

Before you upgrade to IBM i 7.1, review Information APAR II14482 (for APAR access, see 21.1.4, “Program temporary fixes for upgrading to IBM i 7.1” on page 797) and the “Installing, upgrading, or deleting IBM i and related software” and “Storage Solutions” topics in the IBM i Information Center.

For complete details, see *IBM i Network Install using Network File System* at the following web page:

ftp://ftp.software.ibm.com/systems/support/power/i/nfs_optical_upgrade.pdf

21.4.1 Client system requirements

The client system (the system to be installed) accesses virtual optical images through a network and must meet the following requirements:

- ▶ The system must be at IBM i 6.1 or greater.
- ▶ The system must be POWER6 or later.
- ▶ The installation media must be IBM i 6.1.1 or later.
- ▶ PTFs are required for IBM i 6.1. See Information APAR II14482 for the latest PTFs. To review Information APAR II14482, see 21.1.4, “Program temporary fixes for upgrading to IBM i 7.1” on page 797

A client partition with virtual optical device type 632B model 003 can access images on a server using the NFS. This device can be used to install Licensed Internal Code, operating system, licensed programs, and PTFs. See Figure 21-3. The client must meet the following requirements:

- ▶ Either a service tools server or a LAN console connection must be configured.
- ▶ The Internet Protocol (IP) must be Version 4.




Figure 21-3 Remote installation or upgrade using NFS

The 632B-003 optical device is created by using the Create Device Description Optical (**CRTDEVOPT**) command:

```
CRTDEVOPT DEVD(virtual_device_name)
RSRCNAME(*VRT)
LCLINTNETA(*SRVLAN)
RMTINTNETA('X.X.XXX.XXX')
NETIMGDIR('/catalog_directory ')
```

Parameter definitions:

- ▶ **RMTINTNETA** is the remote IP address of the NFS server where this virtual optical device looks for virtual image files.
- ▶ The **NETIMGDIR** parameter specifies the network path on the NFS server that contains the virtual image files that were prepared for use with this device.

Determine whether you must configure a service tools server. The type of system and configuration determines what type of setup might be required. If a LAN console is configured, no further setup is required. For more information about configuring the Service Tools Server, see the following website:

ftp://ftp.software.ibm.com/systems/support/power/i/nfs_optical_upgrade.pdf

21.5 More considerations for upgrading to IBM i 7.1

Upgrading to IBM i 7.1 allows the users to benefit from the enhancements that were made to the operating system. Users must ensure that the system is ready to be upgraded. The following sections detail the necessary changes.

21.5.1 IBM i Console changes

Operations Console Direct attached is not supported by IBM i 7.1. If you have an Operations Console that is directly attached to the system, you must change the console to one of the supported console options before you start the upgrade. Failure to do so causes the installation to fail with unpredictable results.

Supported console options for IBM i 7.1

On POWER5 or POWER6 systems, supported console interfaces include the following options:

- ▶ Operations Console LAN attached
- ▶ Twinax console (which is IOP-based)
- ▶ HMC managed console

POWER6 does not support any IOPs in the central processor complex. Therefore, any IOP-based interface, such as Twinax, must be placed in an HSL attached IO drawer and an HMC is required to tag the console location.

Support: Operations Console Direct attached and Twinax console are not supported on any POWER7 processor-based server. IBM i console options on POWER7 consist of either Operations Console LAN attached or HMC managed console.

For more information about changing consoles, see the following website:

<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/iphca/chgconsoles.htm>

The new enhancements on IBM i 7.1 for console support are as follows:

- ▶ Auto-create service tools device IDs remote panel (RCP) privilege

By default, IBM i 7.1 sets the default value of the Auto-create service tools device IDs remote panel (RCP) privilege to be revoked. To view or change this default, go to Work with Service Tools Security Data—Option 12.

- ▶ Console takeover / recovery status panel

The new default for IBM i 7.1, after you enter a Service Tool user ID and password, is that the Console Take Over Status panel is skipped and the previously displayed console panel is displayed. To view or change this default, go to Work with Service Tools Security Data—Option 13

- ▶ Console take over F18

In IBM i 7.1, you can take over a console type or console device type using the PF key 18. This key allows temporary switching of the console type from HMC console to LAN console without changing the tagging or resetting the operations console session. To view or change this default, go to Work with Service Tools Security Data—Option 14.

- Connecting LAN operations console for uninitialized Load Source (LS)

For a Manufacturing Default Configuration (MDC) system that does not have a preinstalled image, and is not HMC managed, the console type must be set to LAN by the Using the Console Service functions. For more information about changing consoles, see the “Changing consoles” topic in the IBM systems Hardware Information Center.

Additionally, if a LAN console uses the embedded Ethernet ports, then the Enable Ethernet embedded port (E1) function must be set through the console service functions.

Using the console service functions (65+21)

Console service functions (65+21) are the emergency console recovery functions. For more information, see the following website:

<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/iphbx/functions6521.htm>

Changes in DST/SST Console Security

Changes in the DST/SST console security are shown in Figure 21-4.

Work with Service Tools Security Data		System: BH080001
Select one of the following:		
1. Reset operating system default password		
2. Change operating system install security		
3. Work with service tools security log		
4. Restore service tools security data		
5. Save service tools security data		
6. Change password level	PWLV L 1	
7. Work with lock for device IDs from SST	Enabled	
8. Password expiration interval in days	<u>180</u>	
9. Maximum sign-on attempts allowed	<u>3</u>	
10. Duplicate password control	<u>18</u>	
11. Autocreate service tools device IDs	<u>10</u>	
12. RCP privilege on autocreated device IDs	Granted	
13. Display console status screen	Show	
14. Console F18 take over	Enabled	
Selection		
—		
F3=Exit F12=Cancel		

Figure 21-4 Work with Service Tools Security Data

21.5.2 Full automatic installation and upgrades from optical media

If you plan to change the primary language during the upgrade or installation, it is necessary to set the preferred installation language using the QINSTLNG API. This new API was introduced with IBM i 7.1. For details about using this API, see “Set Install National Language Version (NLV) (QINSTLNG) API” at the following website:

<https://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/apis/qinstlng.htm>

These changes are documented in the “Installing, upgrading, or deleting IBM i and related software” topic in the IBM i Information Center.

21.5.3 IBM i 7.1 5770-SS1 Option 1 installation actions

Several database files within QSYS2 might not be fully updated during a slip installation of Option 1 during an upgrade to IBM i 7.1. For the steps to complete before you do the upgrade, see *Memo to Users* at the following website:

<https://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzaq9/rzaq9.pdf>

21.5.4 Adjusting the cluster version of a cluster

The cluster version defines the level at which all the nodes in the cluster are actively communicating with each other. Cluster versioning is a technique that allows the cluster to contain systems at multiple release levels and fully interoperate by determining the communications protocol level to be used.

Before you upgrade to IBM i 7.1, ensure that the node has the appropriate cluster version. Clusters support only a one version difference. If all the nodes in the cluster are at the same release, upgrade to the new release before you change the cluster version. This upgrade ensures that all functions associated with the new release are available. For detailed actions for an upgrade to a new release, see the IBM i Information Center topic “Scenario: Upgrading operating system in a high-availability environment” at the following website:

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp?topic=/rzaig/rzaigscenupgrade.htm>

To verify and change the cluster version for a node, complete the following steps:

1. In a web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log in to the system with your user profile and password.
3. Click **Cluster Resource Services** on the IBM Systems Director Navigator for i5/OS window.
4. On the Cluster Resource Services window, select the **Display Cluster Properties** task.
5. On the Cluster Properties window, click the **General** tab.
6. Verify the cluster version setting or change the version to the wanted setting.

21.5.5 Java considerations

The following sections detail changes that were made to Java.

IBM Toolbox for Java - 5771-JC1

In IBM i 7.1, the IBM Toolbox for Java is now shipped with 5770-SS1 Option 3 of IBM i - Extended Base Directory Support. All functions and features that were available in the original (5761-JC1) product are available and supported in 5770-SS1 - Option 3. The installation process automatically removes 5761-JC1 if it is installed on the system.

Java Developer Kit 1.4 - 5761-JV1 Option 6

The upgrade to IBM i 7.1 automatically removes 5761-JV1 Option 6 and replaces it with 5761-JV1 Option 13 - J2SE 1.4 64 bit.

Java Developer Kit 5.0 - 5761-JV1 Option 7

The upgrade to IBM i 7.1 automatically removes 5761-JV1 Option 7 and replaces it with 5761-JV1 Option 8 - J2SE 5.0 32 bit and Option 9 - J2SE 5.0 64 bit.

Java SE Developer Kit 6 - 5761-JV1 Option 10

The upgrade to IBM i 7.1 automatically removes 5761-JV1 Option 10 and replaces it with 5761-JV1 Option 11 - Java SE 6 32 bit and Option 9 - Java SE 6 64 bit.

IBM Toolbox for Java - 5771-JC1

This function was moved to 5761-SS1 Option 3.

IBM Adopt Authority for Java for i5/OS - 5799-AAJ

This PRPQ provided support for Adopted Authority for Java for the Classic JVM. IBM i 7.1 supports only J9 VM and no longer supports Classic JVM.

21.5.6 Miscellaneous changes and withdrawals

This section covers changes and withdrawals to the following elements:

- ▶ Dynamic Host configuration Protocol
- ▶ System i Access for Wireless 5722XP1
- ▶ IBM Extended Integrated Server support for i 5761LSV
- ▶ IBM Secure Perspective for IBM i 5733PS1
- ▶ Performance Viewer open source Components
- ▶ IBM WebSphere Application Server Version 6
- ▶ Lotus Domino 7 and Lotus Domino 8
- ▶ IBM DataMirror® iCluster® SMB 5733ICS
- ▶ IBM Eclipse Help System

Dynamic Host Configuration Protocol (DHCP)

DHCP in IBM i 7.1 was moved from the base OS to 5770-SS1 option 31 (Domain Name System) and requires that 5771-SS1 Option 33 (Portable App Solutions Environment) be loaded. This feature includes the standard IBM i DHCP server and the new ISC DHCP server. For more information, see Chapter 10, “Networking enhancements” on page 413.

System i Access for Wireless - 5722-XP1

The product to replace the System i Access for Wireless is a browser-based alternative for System i Access for Web for 250 session or IBM System Director Navigator for system management functions.

IBM Extended Integrated Server Support for i - 5761-LSV

This product provided integrated server support for VMware ESX on iSCSI-attached BladeCenter and System x server, Linux on iSCSI attached BladeCenter, and System x server, and Linux running in IBM i hosted logical partitions.

VMware ESX on iSCSI attached BladeCenter and System x server are supported by IBM i 7.1 - 5770-SS1 Option 29 - Integrated Server Support.

Changes for Linux support

For Linux on iSCSI-attached BladeCenter and System x servers, 5761-LSV provides Linux operating system installation, save when active support for virtual disk (storage spaces), file level backup of files on the Linux server, dynamic linking and unlinking of virtual disk, and virtual optical and tape support. These functions are not supported in IBM i 7.1.

Suggested replacement

Use IBM i support for VMware ESX on iSCSI-attached BladeCenter or System x server to host Linux servers. However, save when active, file level backup, and virtual optical and tape are not supported.

Linux running in IBM i hosted partitions continue to be supported in IBM i 7.1, but save when active and file level backup are not supported.

IBM Secure Perspective for IBM i - 5733-PS1

Secure Perspective for IBM i (sold as 5733-PS1 and 5724-PS1) continues to be available from IBM System Lab Services. For more information, see the following website:

http://www-03.ibm.com/systems/services/labservices/platforms/labservices_power.html

Performance Viewer Open Source Components - 5761-SS1 option 46

This function was moved. It is included in IBM i 7.1 - 5770-SS1 option 3. No action is required.

IBM WebSphere Application Server Version 6 - 5733-W60

Replace this product with WebSphere Application Server V7.

Lotus Domino 7 for i5OS - 5733-LD7

Upgrade to Domino V8.5.2 (or later) 5733-LD8.

Lotus Domino 8 for System i - 5733-LD8

Upgrade to Domino V8.5.2 (or later) 5733-LD8.

IBM DataMirror iCluster SMB - 5733-ICS change

Use IBM DataMirror iCluster for System i - 5733-ICL. It has the same or lower price and more functionality than SMB.

IBM Eclipse Help System removed

The IBMHELP server was removed in IBM i 7.1. Previously, this TCP server was packaged with 5761-SS1 option 3 (Extended Base Directory Support).

Figure 21-5 gives a summary of the withdrawals of IBM features related to IBM i.

Function / Product	Notes	Recommended Replacement
Direct Attach Ops Console	Not supported with i 7.1	LAN Console or HMC
Java Developer Kit (Classic)	Not included with i 7.1	IBM Technology for Java
Extended Integrated Server Support	Not support on i 7.1 (5761-LSB) Linux not supported on iSCSI attached systems IBM i no longer supports save while active and File Level Backup for Linux partitions	IBM i in Option 29 - Integrated Server Support
IBM i Information Center Media	Physical Media no longer available	Use the online version of the Information Center
AnyNet	Still included in i 7.1, but no longer supported	Enterprise Extenders function of IBM i
System i Access for wireless	Not supported on i 7.1 (5722-XP1)	
Secure Perspectives	Not supported on i 7.1 (5733-PS1)	Lab Services offering
Rational Developer for i	Withdraw Feb 9 effective May 11 (5733-RDI)	Rational Developer for Power – RPG and COBOL Development for i
Zend Core for i	5639-ZC1 product renamed to Zend Server Community Edition for i	Zend Server Community Edition for i
Zend Platform for i	Withdraw 4/13 effective 7/17 (5619-ZP1, 5771-ZP1)	Zend Server Support for i 5771-ZC1
Zend and MySQL Value Pack	Withdraw 4/13 effective 7/17 (5639-MVP)	Zend Server Support for i 5771-ZC1

Figure 21-5 Withdrawal notes summary

21.6 IBM i network installation

IBM i now supports the installation of partitions from images in the network. An IBM POWER7 processor-based server can be installed without using physical media; you can use IBM i 7.1 optical images that are stored on a network file server.

21.6.1 Requirements for IBM i network installation

This section describes the requirements of installing partitions using network installation.

Source system requirements

Here are the source system requirements:

- ▶ IBM i 6.1 or greater.
 - 6.1: PTFs SI39390, SI35186, and SI35189
 - 7.1: PTFs SI39238
- ▶ The server must be able to share virtual optical images using Version 3 or later of the Network File System (NFS).

- ▶ Install image of IBM i 6.1.1 or later.
 - IBM i 6.1 required PTFs:
 - SI39400 (Lead PTF. This PTF causes the other PTFs to be ordered and installed.)
 - SI39390
 - SI35186
 - SI35189
 - SI35747
 - MF47284
 - MF47285

Client system requirements

Here are the client system requirements:

- ▶ IBM i 6.1.1 or later
- ▶ Server firmware
 - POWER7: 01Ax720 or later; 01Ax710_097 or later (for physical adapters)
 - POWER7: 01Ax730_049 or later (for virtual adapters)
 - POWER 6: 01Ex350_085 or later
- ▶ HMC 7.7.2 or later

Support for setting a Logical Host Ethernet Adapter (LHEA) or a Virtual Ethernet Adapter as the installation source requires HMC Version 7.7.3 and later.
- ▶ Available PCI Ethernet adapter, LHEA port (HMC 7.7.3 and later), or Virtual Ethernet Adapter (HMC 7.7.3 and later) with network connectivity to the installation source server

Network configuration requirements

Here are the network configuration requirements:

- ▶ The MTU value that is set on the TCP/IP interface (Server/IBM i) must always match the MTU value that is set by the **chsysstate** command (Client/HMC CLI).
- ▶ Because **chsysstate** supports only two different values for MTU (1500 and 9000), the Ethernet line description can be set only to MAXFRAMES 1496 or 8996 on the server side.
- ▶ MAXFRAMES 1496 (line description) translates to MTU 1500 (TCP/IP interface), and MAXFRAMES 8996 (line description) translates to MTU 9000 (TCP/IP interface).
- ▶ Only MAXFRAMES 1496 (thus MTU 1500) is supported for network installations. If the Ethernet line description that you are planning to use on the server to set up the network installation is set to MAXFRAMES 8996 (Jumbo frames), you must change it to use MAXFRAMES 1496.

For IBM i 7.1 for PTF and setup information, see APAR II14482.

For complete details, see *IBM i Network Install using Network File System* at the following website:

ftp://ftp.software.ibm.com/systems/support/power/i/nfs_optical_upgrade.pdf

Also, see *IBM Network Installation Using HMC* at the following website:

http://www-912.ibm.com/s_dir/slkbases.NSF/DocNumber/563261959

Related publications

The publications listed in this section are considered suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 818. A few of the documents referenced here might be available in softcopy only.

- ▶ *DS8000 Copy Services for IBM i with VIOS*, REDP-4584
- ▶ *End to End Performance Management on IBM i*, SG24-7808
- ▶ *Getting Started with DB2 Web Query for i*, SG24-7214
- ▶ *Hardware Management Console V7 Handbook*, SG24-7491
- ▶ *IBM BladeCenter JS23 and JS43 Implementation Guide*, SG24-7740
- ▶ *IBM i 6.1 Independent ASPs: A Guide to Quick Implementation of Independent ASPs*, SG24-7811
- ▶ *IBM i 6.1 Technical Overview*, SG24-7713
- ▶ *IBM i and Midrange External Storage*, SG24-7668
- ▶ *IBM i Program Conversion: Getting Ready for 6.1 and Beyond*, REDP-4293
- ▶ *IBM Power 520 and Power 550 (POWER6) System Builder*, SG24-7765
- ▶ *IBM Power 520 Technical Overview*, REDP-4403
- ▶ *IBM Power 550 Technical Overview*, REDP-4404
- ▶ *IBM Power 710 and 730 (8231-E2B) Technical Overview and Introduction*, REDP-4636
- ▶ *IBM Power 720 and 740 (8202-E4B, 8205-E6B) Technical Overview and Introduction*, REDP-4637
- ▶ *IBM Power 750 and 755 (8233-E8B, 8236-E8C) Technical Overview and Introduction*, REDP-4638
- ▶ *IBM Power 770 and 780 (9117-MMB, 9179-MHB) Technical Overview and Introduction*, REDP-4639
- ▶ *IBM Power 795 (9119-FHB) Technical Overview and Introduction*, REDP-4640
- ▶ *IBM PowerVM Virtualization Active Memory Sharing*, REDP-4470
- ▶ *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940
- ▶ *IBM System i5, eServer i5, and iSeries Systems Builder IBM i5/OS Version 5 Release 4 - January 2006*, SG24-2155
- ▶ *IBM System i5 V5R4 Technical Overview Redbook*, SG24-7271
- ▶ *IBM System i Security: Protecting i5/OS Data with Encryption*, SG24-7399
- ▶ *IBM Systems Director Navigator for i*, SG24-7789
- ▶ *Implementing IBM Systems Director 6.1*, SG24-7694
- ▶ *Implementing PowerHA for IBM i*, SG24-7405

- ▶ *PCI, PCI-X, PCI-X DDR, and PCIe Placement Rules for IBM System i Models*, REDP-4011
- ▶ *Security Guide for IBM i V6.1*, SG24-7680
- ▶ *VMware VI3 on BladeCenter and System x Integrated with System i*, SG24-7408
- ▶ *Windows-based Single Signon and the EIM Framework on the IBM eServer iSeries Server*, SG24-6975

Other publications

These publications are also relevant as further information sources:

- ▶ Droms, et al, *DHCP Handbook, 2nd Edition*, SAMS, 2002. 0672323273
- ▶ *Rational Development Studio for i ILE RPG Language Reference*, SC09-2508
- ▶ *Rational Development Studio for i ILE RPG Programmer's Guide*, SC09-2507

Online resources

These web pages are also relevant as further information sources.

- ▶ AFP Font Collection page
http://www-03.ibm.com/systems/i/software/print/afpfonthome_m_ww.html
- ▶ Application Runtime Expert for IBM i page
<http://www-03.ibm.com/systems/power/software/i/are>
- ▶ Backup Recovery & Media Services page
<http://www-03.ibm.com/systems/i/support/brms/index.html>
- ▶ *Connecting to IBM i - IBM Systems Director Navigator for i*
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzatg/rzatgdirect.pdf>
- ▶ IBM i Access page
<http://www-03.ibm.com/systems/i/software/access/>
- ▶ IBM i Access for Web page
<http://www-03.ibm.com/systems/i/software/access/web/>
- ▶ *IBM Advanced Job Scheduler for i*
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzaks/rzaksadjs.pdf>
- ▶ IBM Advanced Job Scheduler for i page
<http://www-03.ibm.com/systems/i/software/jscheduler/index.html>
- ▶ IBM Advanced Job Scheduler for i What's New page
http://www-03.ibm.com/systems/i/software/jscheduler/v7r1_feat.html
- ▶ IBM DB2 for i
<http://www-03.ibm.com/systems/i/software/db2/index.html>

- ▶ IBM DB2 Web Query for i
<http://www-03.ibm.com/systems/i/software/db2/webquery/index.html>
- ▶ IBM i 7.1 Information Center
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp>
- ▶ IBM i 7.1 Technical Overviews
<http://www-947.ibm.com/systems/support/i/library/techoverviews>
- ▶ IBM i 7.1 Upgrade Planning
<http://www-947.ibm.com/systems/support/i/planning/upgrade/v7r1/index.html>
 - Upgrade Planning - Software
<http://www-947.ibm.com/systems/support/i/planning/upgrade/v7r1/software.html>
 - Upgrade Planning - Hardware
<http://www-947.ibm.com/systems/support/i/planning/upgrade/v7r1/hardware.html>
 - Upgrade Planning - Statements of Direction
<http://www-947.ibm.com/systems/support/i/planning/upgrade/v7r1/direct.html>
 - Upgrade Planning - Planning Statements
<http://www-947.ibm.com/systems/support/i/planning/upgrade/v7r1/planstmts.html>
 - Upgrade Planning - Future Software / Hardware
<http://www-947.ibm.com/systems/support/i/planning/upgrade/future.html>
 - Upgrade Planning - Release Life Cycle
<http://www-947.ibm.com/systems/support/i/planning/upgrade/suptschedule.html>
 - Upgrade planning IBM i, i5/OS and OS/400 level mapping
<http://www.ibm.com/systems/support/i/planning/upgrade/osmapping.html>
 - Pre-upgrade verification tool
http://www-912.ibm.com/s_dir/slkbase.NSF/DocNumber/465353483
- ▶ IBM i integration with BladeCenter and System x page
<http://www.ibm.com/systems/i/advantages/integratedserver/library.html>
- ▶ *IBM i Memo to Users 7.1*
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzaq9/rzaq9.pdf>
- ▶ *IBM i on a POWER Blade Read-me First*
http://www-03.ibm.com/systems/resources/systems_power_hardware_blades_i_on_blade_readme.pdf
- ▶ IBM Power Blade Servers - IBM i
<http://www-03.ibm.com/systems/power/hardware/blades/ibmi.html>
- ▶ IBM PowerHA SystemMirror for i
<http://www-03.ibm.com/systems/power/software/availability/i/index.html>
- ▶ IBM Power Systems Hardware Information Center
<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp>

- ▶ *IBM PowerVM Active Memory Sharing: An Overview*
<http://www-03.ibm.com/systems/power/software/virtualization/whitepapers/ams.html>
- ▶ IBM Rational Developer for Power Systems Software page
<http://www-01.ibm.com/software/rational/products/rdp/>
- ▶ IBM Rational Development Studio for IBM i page
<http://www-01.ibm.com/software/awdtools/wds400/>
- ▶ IBM Rational Host Access Transformation Services page
<http://www-01.ibm.com/software/awdtools/hats/>
- ▶ IBM Rational RPG Cafe
<http://www-949.ibm.com/software/rational/cafe/community/rpg>
- ▶ IBM Rational Software Delivery Platform for Power
<http://www-03.ibm.com/systems/power/software/i/rational/index.html>
- ▶ IBM Rational Team Concert for Power Jazz project page
<http://jazz.net/projects/rational-team-concert-power/>
- ▶ IBM Rational Team Concert for Power Systems Software page
<http://www-01.ibm.com/software/rational/products/rtcp/>
- ▶ IBM Systems Director page
<http://www-03.ibm.com/systems/software/director/index.html>
- ▶ IBM Systems Director Active Energy Manager V4.2
http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/aem_420/frb0_main.html
- ▶ IBM Systems Management solutions for Power Servers
<http://www-03.ibm.com/systems/power/software/management/>
- ▶ IBM Systems Software Information Center
<http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp>
- ▶ IBM Systems Workload Estimator page
<http://www.ibm.com/systems/support/tools/estimator>
- ▶ *IBM i Virtualization and Open Storage Read-me First*
http://www-03.ibm.com/systems/resources/systems_i_Virtualization_Open_Storage.pdf
- ▶ InfoPrint AFP Utilities for IBM i page
http://www-03.ibm.com/systems/i/software/print/afputilhome_i_ww.html
- ▶ *Installing, upgrading, or deleting IBM i and related software 7.1*
<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzahc/sc415120.pdf>
- ▶ Interoperability Guide - IBM BladeCenter
<http://www-947.ibm.com/systems/support/supportsite.wss/docdisplay?Indocid=MIGR-5073016&brandind=5000020>

- ▶ Licensed program releases and sizes
<http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzahc/rzahcswslpsze.htm>
- ▶ Linux on integrated servers page
<http://www.ibm.com/systems/i/advantages/integratedserver/linux/>
- ▶ Lotus Enterprise Integrator for IBM i page
<http://www-03.ibm.com/systems/i/software/domino/related/lei/index.html>
- ▶ Lotus Products for IBM i page
<http://www-03.ibm.com/systems/i/software/domino/index.html>
- ▶ *Lotus Software for IBM i operating system Compatibility Guide*
<http://www-03.ibm.com/systems/resources/releasesupport.pdf>
- ▶ Online Lotus Server Backup page
<http://www-03.ibm.com/systems/i/support/brms/domino.html>
- ▶ Online Lotus Server Backup - DAOS page
<http://www-03.ibm.com/systems/i/support/brms/domdaos.html>
- ▶ Performance Management on IBM i page
<http://www-03.ibm.com/systems/i/advantages/perfmgmt/ptools.html>
- ▶ PM for Power Systems page
<http://www-03.ibm.com/systems/power/support/perfmgmt>
- ▶ *Power Systems Logical partitioning*
<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphathat/iphathat.pdf>
- ▶ Print Services Facility for IBM i page
http://www-03.ibm.com/systems/i/software/print/psfhome_i_w.html
- ▶ Profile Data Trace Visualizer (PDTV) page
<http://www.alphaworks.ibm.com/tech/ptdv>
- ▶ System Storage Interoperation Center page
http://www-03.ibm.com/systems/support/storage/config/ssic/displayesssearchwithoutjs.wss?start_over=yes
- ▶ Virtualization with IBM i, PowerVM and Power Systems page
<http://www-03.ibm.com/systems/i/os/index.html>
- ▶ WebSphere Application Server for IBM i
<http://www-03.ibm.com/systems/i/software/websphere/>
- ▶ Zend Products page
<http://www.zend.com/en/products/>

How to get Redbooks

The linked Redbooks publication material listed in the previous section is available in softcopy on the internet from the IBM Redbooks server.

Using the web browser of your choice, go to the following website:

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

At the upper right corner of the IBM Redbooks publication web page, there is a search window. You can search for a specific book by book ID or by title text using that search capability.

At the left is a navigation menu. Advance search, Software, Storage, and Systems & Servers are helpful options.

The easiest way to download a Redbooks publication is to complete the following steps:

1. Locate the book through one of the menu or search options,
2. Select the Redbooks publication. This action opens a summary page for the book.
3. Select the **Download PDF** link. You then have two options:
 - Right-click the link and click **Save target as**. This action downloads and saves the PDF.
 - Click the link. This action displays the PDF. Select the **Save file** icon.

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



IBM i 7.1 Technical Overview with Technology Refresh Updates

(1.5" spine)
1.5" x 1.998"
789 <-> 1051 pages



IBM i 7.1 Technical Overview with Technology Refresh Updates



Redbooks®

Boosted efficiency with enriched virtualization including Live Partition Mobility

Enriched database functionality and new application development tools

Reenergized management and performance tooling web interface

This IBM Redbooks publication introduces a technical overview of the new features, functions, and enhancements available in IBM i 7.1. It provides a summary and brief explanation of new capabilities and what changed in the operating system. This publication also describes many of the licensed programs and application development tools that are associated with IBM i. This publication includes all the enhancements that were delivered with the Technology Refresh 4 and 5 that were made available in 2012.

Many new and enhanced functions are described:

- ▶ IBM DB2 for i support for XML and column level encryption
- ▶ IBM PowerHA for i asynchronous Geographic Mirroring and LUN-level switching
- ▶ Virtualization enhancements for IBM i hosted partitions and IBM PowerVM
- ▶ Storage management enhancements for solid-state drives
- ▶ Systems management enhancements for Systems Director Navigator for i
- ▶ IBM Rational software development product enhancements
- ▶ Zend PHP Enhancements
- ▶ Web Application Serving enhancements

The information that is provided in this book is useful for clients, IBM Business Partners, and IBM service professionals that are involved with planning, supporting, upgrading, and implementing IBM i 7.1 solutions.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks