

IBM i 7.1


Technical Overview

Including Technology Refresh Updates - November 2011

Boost efficiency with enriched virtualization
and storage integration

Exploit enhanced database functions and
new application development tools

Benefit from centralized system
management and performance tooling



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

Redbooks



International Technical Support Organization

IBM i 7.1 Technical Overview

November 2011

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

Second Edition (November 2011)

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

© Copyright International Business Machines Corporation 2010. 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	xv
Trademarks	xvi
 Preface	 xix
The team who wrote this book	xix
Now you can become a published author, too!	xxiv
Comments welcome	xxiv
Stay connected to IBM Redbooks	xxiv
 Chapter 1. IBM i 7.1 introduction and summary	 1
1.1 Clarification of operating system terminology	2
1.2 What is new in IBM i 7.1	2
1.2.1 IBM i 7.1 on POWER based servers	2
1.2.2 Technology Refresh	2
1.2.3 Integrated DB2 database for IBM i	3
1.2.4 PowerHA SystemMirror for i	3
1.2.5 IBM i virtualization	3
1.2.6 IBM i storage management	4
1.2.7 Application Development on IBM i Enhancements	5
1.2.8 Rational Support for IBM i 7.1	6
1.2.9 Systems management enhancements	7
1.2.10 Printing	7
1.2.11 IBM i iSCSI integration with BladeCenter and System x	8
1.2.12 Security enhancements in IBM i 7.1	8
1.2.13 Web application serving	8
1.2.14 IBM i Access Family 7.1 enhancements	8
1.2.15 Networking enhancements	9
1.2.16 Upgrade for IBM i enhancements	9
1.2.17 IBM i network installation	9
1.2.18 IBM i Processor Entitlement Transfer	9
1.2.19 BRMS and basic backup and recovery	10
 Chapter 2. Technology Refresh	 11
2.1 Overview	12
2.2 What is in a Technology Refresh	12
2.2.1 When to install a Technology Refresh	13
2.2.2 How it all fits together	13
2.2.3 How to determine the Technology Refresh PTF Group level installed on a system	14
2.2.4 How to determine the Resave level installed on a system	15
2.2.5 How to determine the Technology Refresh PTF level installed on a system	16
2.2.6 Planning for an IBM i Technology Refresh update	17
2.2.7 How to install an IBM i Technology Refresh	17
2.2.8 How to order an IBM i Technology Refresh PTF Group	18
2.2.9 How to install an IBM i Technology Refresh PTF Group	18
2.2.10 How to order and install an IBM i Resave	18
2.2.11 How to order an IBM i Technology Refresh PTF	18
2.2.12 How to install an IBM i Technology Refresh PTF	19
2.2.13 How does a Technology Refresh PTF or PTF Group affect other PTFs?	19

2.2.14	How to prevent or reduce the impact of a 'double IPL'?	19
2.2.15	Why has IBM moved to this Technology Refresh process?	20
2.2.16	Will there still be new releases of IBM i?	20
2.2.17	Is an IBM i Technology Refresh the same as an AIX Technology Level?	20
2.2.18	Related publications	21
Chapter 3.	Security enhancements	23
3.1	ASP encryption enhancements	24
3.2	Column Encryption	24
3.3	User profiles enhancements	24
3.3.1	CRTUSRPRF and CHGUSRPRF commands	24
3.3.2	Display Expiration Schedule	25
3.4	Other security enhancements	26
3.4.1	Printing	26
3.4.2	DB2 for i	26
Chapter 4.	Backup and recovery	27
4.1	New and enhanced system save and restore functions	28
4.1.1	Fast restore using tape position	28
4.1.2	New ALWOBJDIF (*COMPATIBLE) restore option	29
4.1.3	Enhanced save and restore limits	29
4.1.4	Save While Active Support for Integrated VMware ESX Servers	30
4.1.5	Miscellaneous enhancements	30
4.2	New and improved backup and recovery and media services for IBM i functions	31
4.2.1	Support for Domino 8.5 Domino Attachment and Object Service (DAOS)	31
4.2.2	Link list enhancements	32
4.2.3	Support of more than 999,999 objects	32
4.2.4	STRRCYBRM (Start recovery using BRM) command enhancements	33
4.2.5	Recovery report enhancements	36
4.2.6	Output file support for BRMS restores	37
4.2.7	Ability to select Saved ASP on RSTLIBBRM and RSTOBJBRM commands	42
4.2.8	Distributed backup support	43
4.2.9	Maintenance enhancements	47
4.2.10	Planning Media Movement Report	49
4.2.11	Improved clearing of volumes in *ERR status	50
4.2.12	Protection of volumes marked for duplication	51
4.2.13	Improved recovery times with new media position function	52
4.2.14	BRMS support for special value *COMPATIBLE for ALWOBJDIF parameter	53
4.2.15	Improved control when running non-BRMS saves	53
4.3	BRMS enhancements to GUI and web interfaces	53
4.3.1	Added support for IBM Systems Director web browser environment	54
4.3.2	IBM Systems Director navigation to BRMS functions	54
4.3.3	IBM Systems Director Navigator for i Navigation to BRMS Functions	59
4.3.4	Enhancements to the BRMS initial panel	61
4.3.5	BRMS advanced functions panel	62
4.3.6	Scheduling Support for BRMS	62
4.3.7	Added option to BRMS log to filter messages by control groups	72
4.3.8	Ability to mark and unmark volumes for duplication	75
4.3.9	Multiple email address support	82
4.3.10	High availability support for independent ASPs in a BRMS network	84
4.3.11	Enhanced maintenance features in backup policy	88
4.4	Additional references	89

Chapter 5. High availability	91
5.1 PowerHA SystemMirror for i	92
5.1.1 New PowerHA Packaging.	92
5.1.2 PowerHA versioning	93
5.1.3 PowerHA SystemMirror for i Enhancements.	93
5.1.4 PowerHA SystemMirror for i - Graphical interfaces	96
5.1.5 Duplicate Library Error Handling.	100
5.1.6 N_Port ID virtualization support	100
5.1.7 Asynchronous geographic mirroring	100
5.1.8 LUN level switching.	103
5.1.9 Space-Efficient FlashCopy	104
5.1.10 Better detection of cluster node outages	105
5.1.11 Improved geographic mirroring full synchronization performance	107
5.1.12 Cluster administrative domain enhancements	107
5.1.13 IBM HA Assist for i	107
5.1.14 IPv6 Support.	108
5.1.15 New CL commands for programming cluster automation.	108
5.2 Journaling and commitment control enhancements	109
5.3 Additional information	109
Chapter 6. DB2 for i	111
6.1 Introduction: Getting around with data	112
6.2 SQL data description and data manipulation language	112
6.2.1 XML support	113
6.2.2 MERGE statement	118
6.2.3 Creating and using global variables	120
6.2.4 Support for arrays in procedures.	121
6.2.5 Result set support in embedded SQL.	122
6.2.6 FIELDPROC support for encoding and encryption	124
6.2.7 MQ Series integration	126
6.2.8 Miscellaneous	130
6.2.9 OVRDBF SEQONLY(YES, buffer length)	134
6.3 Availability and consistency	134
6.3.1 Journal management	134
6.3.2 Remote journaling.	136
6.3.3 DISPLAY_JOURNAL (easier searches of journal)	136
6.3.4 Commitment control and independent ASPs	138
6.3.5 System managed access path protection (SMAPP)	139
6.3.6 Journal management functions: IBM Systems Director Navigator for i.	140
6.4 Performance and query optimization	140
6.4.1 Methods and tools for performance optimization	141
6.4.2 What is query optimization?	142
6.4.3 Global Statistics Cache (GSC)	142
6.4.4 Adaptive query processing (AQP).	143
6.4.5 Sparse indexes	145
6.4.6 Encoded vector index (EVI)	145
6.4.7 Preserve EVI indexes on ALTER enhancement	147
6.4.8 Keeping tables or indexes in memory.	147
6.4.9 SQE optimization for indexes on SSD	148
6.4.10 SQE support of simple logical files	148
6.4.11 QSYS2.INDEX_ADVICE procedure	148
6.4.12 SKIP LOCKED DATA and NC or UR	149
6.4.13 SQL routine performance integer arithmetic (requires recreate).	149

6.4.14	Automatic cancel of QSQSRVR jobs when an application is ended.	150
6.4.15	QAQQINI properties	150
6.4.16	Alter table performance.	151
6.4.17	Avoid short name collisions for CREATE PROCEDURE, FUNCTION or TRIGGER 152	
6.4.18	CREATE PROCEDURE (SQL) PROGRAM TYPE SUB.	153
6.4.19	Referential integrity and trigger performance	153
6.4.20	QSQBIGPSA data area	154
6.4.21	Validate constraints without checking	154
6.4.22	Limit the amount of processing on a RGZPFM cancel	154
6.4.23	CPYFRMIMPF performance.	155
6.4.24	QJOSJRNE API option to force journal entries without sending an entry. . . .	155
6.4.25	QDBRTVSN API performance	155
6.4.26	Add Total DB Opens job level instrumentation to Collection Services	156
6.5	New Functionality for DB2 developers	156
6.5.1	QSYS2.SYSCOLUMNS2 view	156
6.5.2	QSYS2.SYSPARTITIONDISK and QSYS2.SYSPARTITIONINDEXDISK Views	157
6.5.3	QSYS2.Object_Statistics table function ease of use	159
6.5.4	EARLIEST_POSSIBLE_RELEASE	159
6.5.5	SIGNAL support for native triggers.	160
6.5.6	Hierarchical queries using CONNECT BY clause.	160
6.5.7	Additional parameter marker support (LAND, LOR, XOR, and TRANSLATE) . .	161
6.5.8	Support PROGRAM NAME on CREATE TRIGGER	161
6.5.9	TINYINT in CLI	162
6.5.10	CHGPFM and CHGLFM UNIT support.	162
6.5.11	SYSTOOLS procedures	162
6.5.12	Add client and server IP address and port number to QSYS2.TCPIP_INFO . .	163
6.5.13	QSYS2.GROUP_PTF_INFO view	163
6.5.14	QSYS2.DUMP_SQL_CURSORS procedure	165
6.5.15	QIBM_SQL_NO_RLA_CANCEL environment variable	166
6.5.16	QSYS2.FIND_AND_CANCEL_QSQSRVR_SQL and QSYS2.CANCEL_SQL procedures	166
6.5.17	QSYS2.FIND_QSQSRVR_JOBS() procedure	168
6.5.18	SQL Server Mixed Mode for Batch Processing.	169
6.5.19	QDBRPLAY() API - Disable or Enable Constraints option	170
6.5.20	SQL0901 logger education	170
6.6	DB2 security enhancements	171
6.6.1	Query Manager profile auditing.	171
6.6.2	Add SECURE column to QSYS2.SYSCOLUMNS2	172
6.6.3	QIBM_DB_SQLADM and QIBM_DB_SYSMON (*JOBCTL special authority no longer required)	173
6.6.4	STRDBMON Over a View (Enables faster SQL statement auditing)	174
6.6.5	SECURE Columns (Monitor and Plan Cache masking of variable values). . . .	178
6.6.6	Add QDDMDRDASERVER server authentication entry special value	179
6.6.7	Enhanced audit capability of RUNSQLSTM or STRSQL	182
6.6.8	Add FTRSQLCODE parameter to STRDBMON	183
6.6.9	Extend FTRINTNETA pre-filter on STRDBMON command to work with non-Database clients	184
6.6.10	Extend STRDBMON to support FTRUSER against group profiles.	184
6.7	DB2 Availability and Recovery Enhancements	184
6.7.1	Preserve SQE Cache Size Across IPLs	184
6.7.2	Pre check physical file size during restore	185
6.7.3	Prevent Index Rebuild On Cancel During Catch Up.	185

6.8 DB2 for Web Query for i (5733-QU2, 5733-QU3, and 5733-QU4)	185
6.8.1 DB2 Web Query for i - 5733-QU2.	185
6.8.2 DB2 Web Query Report Broker - 5733-QU3.	187
6.8.3 DB2 Web Query Software Developer Kit - 5733-QU4	187
6.8.4 DB2 Web Query for i Standard Edition	187
6.9 OmniFind Text Search Server for DB2 for i (5733-OMF)	188
6.9.1 OmniFind Updates Page.	188
Chapter 7. Performance tools	189
7.1 Introduction to performance tooling	190
7.1.1 Gathering data	190
7.1.2 Analyzing data	190
7.2 Overview of data collection in IBM i 7.1	191
7.2.1 Collection profiles for Collection Services.	192
7.2.2 Information repository: Performance tools tables	192
7.2.3 Additional Enhancements	202
7.3 IBM Systems Director Navigator for i Performance interface	203
7.4 IBM iDoctor for IBM i	204
7.4.1 Installation.	204
7.4.2 My Connections View	204
7.4.3 Main Window	206
7.4.4 Collection Options.	209
7.4.5 Data Viewer	210
7.4.6 Collection Services Investigator (CSI).	212
7.4.7 Job Watcher	229
7.4.8 PEX Analyzer	237
7.4.9 Additional Information	240
Chapter 8. Virtualization	241
8.1 PowerVM enhancements	242
8.1.1 Active memory sharing	242
8.1.2 Enhanced support for IBM System Storage	246
8.1.3 PowerVM Virtualization and I/O enhanced with NPIV.	247
8.1.4 Expanded HBA and switch support for NPIV on power blades.	250
8.1.5 PowerVM N_Port ID Virtualization attachment of DS5000	250
8.1.6 Enhanced mirroring algorithms	250
8.1.7 PCIe2 Riser Card (Gen2) (#5685) Direct Support	250
8.1.8 Redundant VIOS partitions support	251
8.2 Additional OS levels combinations of server and client logical partitions	253
8.3 Hardware Management Console virtual device information	255
8.4 Virtualizing an optical device to IBM i client partitions.	258
8.4.1 Support for embedded media changers	267
8.4.2 IBM i to IBM i virtual tape support	267
8.5 Virtual Partition Manager Enhancements	268
8.5.1 Ethernet Layer-2 Bridging	268
8.5.2 Preparing for Ethernet layer-2 bridging.	268
8.5.3 Recommended practices	269
8.5.4 Configuring Ethernet layer-2 bridging	269
8.5.5 Common errors	269
8.5.6 Managing Ethernet layer-2 bridging	270
8.6 Partition Suspend and Resume	270
8.6.1 Requirements for Suspend / Resume.	270
8.6.2 Restrictions for Suspend / Resume.	270

8.7 HEA Daughter cards	271
8.8 10 Gb FCoE PCIe Dual Port Adapter	272
Chapter 9. Storage and solid state drives	273
9.1 General and direct-attached storage management	
enhancements	274
9.1.1 Concurrent removal of disk units.	274
9.1.2 Hot spare for mirroring	276
9.1.3 Dual storage I/O adapters.	277
9.1.4 Encrypted ASP enhancements	282
9.1.5 Disk response time buckets enhancements	284
9.1.6 CEC node level mirroring	285
9.1.7 EXP24S SFF Gen2-bay drawer (#5887).	288
9.1.8 Higher Capacity 10K RPM SFF SAS disk drives	288
9.1.9 Tape performance instrumentation	289
9.1.10 Tape library resource name changes for IOP-less IOA attachment	289
9.1.11 Tape library unreadable barcode changes for IOP-less IOA attachment	289
9.1.12 DVD/Tape SAS External Storage Unit for Power 795 CEC Rack.	290
9.2 SAN storage management enhancements	290
9.2.1 Multipathing for virtual I/O	290
9.2.2 DS5000 native attachment	292
9.2.3 Level of protection reporting for multi-path disk units	294
9.2.4 Library control paths for IOP-less Fibre Channel IOA tape attachment	295
9.2.5 External disk storage performance instrumentation	296
9.2.6 Thin Provisioning for DS8700, DS8800 storage servers and for VIOS shared storage pools	301
9.3 SSD storage management enhancements	302
9.3.1 DB2 media preference	302
9.3.2 ASP balancer enhancements for SSDs	307
9.3.3 User-defined file system media preference.	310
9.3.4 177 GB SFF SSD with eMLC	312
9.3.5 IBM Disk Sanitizer PRPQ extended to include SSD devices	312
Chapter 10. Networking enhancements	315
10.1 TCP/IP enhancements summary	316
10.1.1 IPv6 support enhancements	316
10.1.2 ISC-based DHCP server supports IPv6 and failover	316
10.1.3 DHCPv6 Client	318
10.1.4 SNMP	319
10.1.5 SMTP	320
10.1.6 IPv6 support added for PPP connections	321
10.2 HTTP server supports PASE FastCGI PHP Processing.	325
10.3 Telnet client support for SSL encryption	326
10.4 Sockets programming: New user exits	327
10.4.1 Exit points are defined in the user registry	328
10.4.2 Example: User exit program for QIBM_QSO_ACCEPT	328
10.5 IKEv2 support	330
10.5.1 Enhanced Cryptographic Algorithms have also been added	330
10.6 IPv6 TCP/IP applications and V6R1 enablement PTFs	331
10.7 AnyNet is no longer supported on IBM i 7.1	332
10.8 Ethernet link aggregation	332
10.8.1 Overview of Ethernet link Aggregation	332
10.8.2 Prerequisites for Ethernet link aggregation.	332

10.8.3	Configuring Ethernet link aggregation.	333
10.8.4	Example - Configuring four aggregated network links	334
10.9	Sharing physical Ethernet connections across partitions (Ethernet layer-2 bridging)	335
10.9.1	Introduction to Ethernet layer-2 bridging.	335
10.9.2	How Ethernet layer-2 bridging works on IBM i	336
10.9.3	IBM i Prerequisites for Ethernet layer-2 bridging	336
10.9.4	Configuring a shared network connection via Ethernet layer-2 bridging.	337
10.10	IBM Portable Utilities for i (5733-SC1) supported versions.	339
Chapter 11.	Printing enhancements	341
11.1	Printing overview.	342
11.2	Overview of IBM i 6.1 print enhancements	343
11.3	Enhancements to base print	344
11.3.1	CPYSPLF command enhanced to copy to PDF or TIFF stream files	344
11.3.2	New QIBM_QSP_SECURITY exit point and formats	347
11.3.3	Encryption of PDF output (Infoprint Server LP).	349
11.4	Print Services Facility for IBM i (PSF) enhancements.	350
11.4.1	Disabling offset stacking and edge mark printing	351
11.4.2	Specify public data authority for directories created by PSF.	351
11.4.3	Improved PSF debugging capabilities.	352
11.5	Transform services enhancements	353
11.5.1	Generate PDFs from existing spooled files.	353
11.5.2	PDF transform enhancements	353
11.6	IBM Systems Director Navigator for i print enhancements	354
11.6.1	Viewing printer output in PDF format	354
11.6.2	Export printer output in PDF format using InfoPrint Server.	357
11.7	IBM i Access for web print enhancements	362
11.8	Host Print Transform enhancements	362
11.8.1	The Host Print Transform API now uses Transform Services.	362
11.8.2	Example: Host Print Transform API with Transform Services.	363
11.9	References	364
Chapter 12.	Integration with BladeCenter and System x.	365
12.1	iSCSI software targets	366
12.1.1	IBM i Integrated server object model with a Target iSCSI	366
12.1.2	IBM i Integrated server object model with a software target.	367
12.2	Defining the iSCSI software target support.	367
12.2.1	CRTDEVNWSH CL command interface	368
12.2.2	IBM Systems Director Navigator for i changes for iSCSI software target support	369
12.3	Service Processor Manager function	370
12.4	VMware support changes	370
12.4.1	New NWSD types	371
12.4.2	VMware ESX server management	371
12.4.3	SWA storage spaces for VMware ESX servers	372
12.5	New planning work sheets	372
12.6	IBM Systems Director Navigator for i	372
12.6.1	Create Server task	373
12.6.2	Clone Integrated Windows Server task.	375
12.6.3	Delete Server task	377
12.6.4	Launch web console	378
12.7	New IBM i CL commands	378
12.7.1	INSINTSVR (Install Integrated Server)	379

12.7.2 DLTINTSVR (Delete Integrated Server) command	379
12.8 IBM i changed CL commands	380
12.8.1 INSWNTSVR (Install Windows Server) CL command	380
12.8.2 CRTNWSCFG (Create NWS Configuration) and CHGNWSCFG (Change NWS Configuration) CL commands	380
12.8.3 INSLNXSVR (Install Linux Server) CL command	381
12.8.4 No new integrated Linux servers.	381
12.9 Fewer IBM i licensed programs required	381
12.10 Changes to IBM i integration with BladeCenter and System x documentation	381
12.10.1 A new IBM i iSCSI Solution Guide PDF	381
12.10.2 IBM i 7.1 Information Center	382
12.10.3 New IBM i integration with BladeCenter and System x group on developerWorks	382
12.10.4 New IBM i Technology Updates page on developerWorks.	383
12.10.5 IBM i integration with BladeCenter and System x Marketing Web site	383
12.11 Additional information	383
Chapter 13. IBM Advanced Job Scheduler for i enhancements	385
13.1 Advanced Job Scheduler capabilities	386
13.1.1 Scheduling a job	386
13.1.2 Job groups	386
13.1.3 Notification and report distribution.	386
13.1.4 Remote command jobs.	386
13.1.5 Multiple scheduling environments.	386
13.2 IBM Systems Director Navigator for i AJS support	387
13.2.1 Navigating to the AJS Menu	388
13.2.2 AJS pop-up menu	394
13.2.3 Start Scheduler action.	394
13.2.4 Stop Scheduler action.	395
13.2.5 AJS properties menu	395
13.2.6 Scheduled Jobs menu	402
13.2.7 Using the Scheduled Jobs action	404
13.2.8 Using the Scheduled Jobs Table	406
13.2.9 Adding a new scheduled job	412
13.3 Other AJS enhancements	425
13.3.1 Notification enhancements	425
13.3.2 Multiple Scheduling Environments function	426
13.3.3 Scheduling and running jobs.	427
13.4 References	430
Chapter 14. Web serving	431
14.1 Triggered cache manager removed	432
14.2 Web server search engine and web crawler support removed	432
14.3 Plug-ins and LoadModule directives	432
14.4 HTTP Server for i with HA IPv6 Support	438
Chapter 15. Web enablers.	441
15.1 IBM Integrated Web Services for i	442
15.1.1 Integrated Web Services Client for ILE.	442
15.2 Changes in IBM WebSphere Application Server	442
15.2.1 Packaging for WebSphere Application Server	442
15.2.2 Java for IBM i 7.1	443
15.2.3 Installation.	444
15.2.4 Upgrading to IBM i 7.1	444

15.2.5	IBM Installation Manager for WebSphere Application Server Version 8.0	445
15.2.6	Enabling IBM Technology for Java Virtual Machine	446
15.3	IBM Web Administration for i	447
15.3.1	Web Log Monitor	447
15.3.2	Permissions	458
15.3.3	Web Performance Advisor	466
Chapter 16.	Application development	469
16.1	High-level programming languages	470
16.1.1	C/C++	470
16.1.2	Predefined macros	470
16.1.3	Pragmas	472
16.1.4	Compiler options	472
16.1.5	ILE COBOL	473
16.1.6	ILE RPG	475
16.2	Control language CL	479
16.2.1	New workload capping commands	480
16.2.2	Retrieve CL source support for ILE CL	480
16.2.3	Longer Integer CL variables for ILE CL	480
16.2.4	Showing DO and SELECT nesting levels in compiler listing.	481
16.2.5	Encrypting the debug listing view for ILE CL	481
16.2.6	Nested INCLUDE support	481
16.3	PHP	482
16.3.1	Zend Server Community Edition for IBM i	482
16.3.2	Zend Server for IBM i	483
16.3.3	Zend Studio for IBM i	483
16.4	Lotus products for IBM i	484
16.5	Native archive and un-archive API support	484
16.5.1	QZIPZIP API	484
16.5.2	QZIPUNZIP API	485
16.5.3	QZIPUTIL Service program	487
16.6	IBM Toolbox for Java JDBC	487
16.6.1	JDBC 4.1	487
16.6.2	XML data type support	487
16.6.3	Database metadata updates	487
16.6.4	Currently committed support	487
16.6.5	Array type support	488
16.6.6	Long schema name support	488
16.7	Application Runtime Expert for i	488
16.7.1	Deployment template	489
16.7.2	Customize Plugins For Template	489
16.7.3	Application Runtime Expert Console	494
16.7.4	Application Runtime Expert Requirements	499
16.7.5	More Information	499
Chapter 17.	Rational Products	501
17.1	IBM Rational Developer for Power Systems Software	502
17.1.1	RPG and COBOL Development Tools	502
17.1.2	Rational Team Concert client integration for IBM i	506
17.1.3	Enhancements in Version 8.0.3	506
17.1.4	Migration to Rational Developer for Power System Software V8.0	508
17.2	Rational Team Concert	508
17.2.1	Integration with Rational Developer for Power system	509

17.2.2	What's New in Rational Team Concert V3.0.1	512
17.3	Rational Developer for i for SOA Construction	513
17.4	Rational Development Studio for i	513
17.4.1	ILE RPG Compiler	514
17.4.2	ILE COBOL Compiler	517
17.4.3	ILE C	518
17.4.4	ILE C++	518
17.5	Rational Open Access: RPG Edition	519
17.5.1	How to use Rational Open Access	520
17.5.2	Open Access Requirements	523
Chapter 18.	IBM Systems Director Navigator for IBM i 7.1	525
18.1	5250 emulation related enhancements	526
18.2	Set Target System enhancements	527
18.3	Database enhancements	529
18.3.1	On Demand Performance Center enhancements	529
18.3.2	Health center enhancements	536
18.3.3	Database management enhancements	537
18.3.4	Maintenance category enhancements	544
18.3.5	Long SQL schema names enhancements	547
18.3.6	Omnifind text search functions enhancements	547
18.4	New journal management enhancements	548
18.4.1	Show Journaled Objects function	548
18.4.2	Change Receivers function	549
18.4.3	Remote Journals function	549
18.4.4	Add Remote Journal function	550
18.4.5	Properties function	550
18.4.6	Activate function	550
18.4.7	Deactivate function	551
18.4.8	Remove function	551
18.4.9	Properties function	551
18.5	Integrated Server Administration enhancements	552
18.5.1	Create Server task	552
18.5.2	Cloning a Server	552
18.5.3	Delete Server task	554
18.5.4	Launch Web Console task	554
18.6	Printer output enhancements	555
18.6.1	View as PDF action	555
18.6.2	Export as PDF action	556
18.7	File systems enhancements	556
18.8	Networking enhancements	558
18.9	Disk management enhancements	559
18.9.1	Graphical view enhancements	559
18.9.2	Start/stop encryption on disk pools enhancements	561
18.9.3	Asynchronous delivery for Geographic Mirroring enhancements	561
18.10	Tape support enhancements	562
18.11	Performance enhancements	563
18.11.1	General health indicators	564
18.11.2	Viewing cross-partition CPU	574
18.11.3	QAPMCONF Perspective	585
18.11.4	Image and file exportation	588
18.11.5	Size the next upgrade	589
18.11.6	Java perspectives	593

18.11.7 Disk response time charts.	594
18.11.8 Interactive perspective development.	600
18.11.9 Metric finder	605
18.11.10 Performance Explorer content package	606
18.11.11 New metrics	606
18.11.12 Miscellaneous enhancements.	611
18.12 JS1- Advanced Job Scheduler for i enhancements	618
18.13 Backup Recovery Media Services: BRMS plugin enhancements	619
18.14 Additional information	619
Chapter 19. IBM i Access	621
19.1 IBM i Access for Windows.	622
19.1.1 Installation enhancements	622
19.1.2 .NET Data Provider enhancements	622
19.1.3 OLE Data Provider enhancements	622
19.1.4 Windows ODBC Driver enhancements.	623
19.1.5 Data Transfer enhancements	623
19.1.6 Personal Communications Emulator enhancements	623
19.1.7 Operations Console enhancements	623
19.2 IBM i Access for Web enhancements	623
19.2.1 AFP to PDF Transform	624
19.3 System i Navigator tasks on the web	626
19.4 System i Navigator Database administration functions.	628
19.5 System i Access for Wireless	630
Chapter 20. Miscellaneous enhancements	631
20.1 Licensed product program structures and sizes	632
20.2 Changed or new CL Commands and APIs	633
20.3 Temporary user-defined file systems	633
20.4 Watch for Event function (message enhancements)	635
20.4.1 Advantages of using watches	635
20.4.2 New Watch for Event function for IBM i 7.1	635
20.5 IBM Tivoli Directory Server for IBM i enhancements.	638
20.5.1 Create suffix entries automatically whenever necessary	638
20.5.2 Administrative roles.	638
20.5.3 User interface enhancements	638
20.5.4 Security enhancements	638
20.5.5 New password encryption options	638
20.5.6 Pass-through authentication	638
20.5.7 Enhanced password policy to use global date and time for initialization	639
20.5.8 Multiple password policies	639
20.5.9 Policy enforced for Digest-MD5 binds.	639
20.5.10 Persistent search	639
20.5.11 Replication configuration enhancements	639
20.5.12 Filtered replication.	639
20.5.13 Limit number of values returned by a search	640
20.5.14 Enhanced syntaxes and matching rules	640
20.5.15 IASP enablement for Directory Server on IBM i	640
20.5.16 Idapcompare utility	640
20.5.17 Provide re-entrant LDAP C client library	640
20.6 Automate Extra IPL for PTF Install	641
20.7 IBM i workload capping	642

20.7.1 Example of how workload capping works	642
Chapter 21. Upgrading IBM i	643
21.1 Installing or upgrading IBM i	644
21.1.1 IBM i 7.1 supports POWER5 and later hardware	644
21.1.2 Planning statements	645
21.1.3 Supported upgrade paths	645
21.1.4 Program temporary fixes for upgrading to IBM i 7.1	645
21.1.5 Pre-upgrade verification tool	646
21.1.6 License program releases and sizes.	646
21.1.7 Server firmware requirements.	646
21.1.8 IBM i 7.1 Resaves.	646
21.2 Upgrading from i5/OS 5.4 to IBM i 7.1	647
21.2.1 Object conversions	647
21.2.2 Spooled file conversions.	647
21.2.3 Integrated file system conversions	648
21.2.4 Backup Recovery and Media Services (BRMS)	648
21.3 Media delivery changes	649
21.3.1 IBM i 7.1 multi-language distribution media	649
21.3.2 Media charges	650
21.4 IBM i network upgrade	651
21.4.1 Client system requirements.	651
21.5 Additional considerations for upgrade to IBM i 7.1	653
21.5.1 IBM i Console changes.	653
21.5.2 Full automatic installation and upgrades from optical media	654
21.5.3 IBM i 7.1 5770-SS1 Option 1 install actions	655
21.5.4 Adjusting the cluster version of a cluster.	655
21.5.5 JAVA considerations.	655
21.5.6 Miscellaneous changes and withdrawals	656
21.6 IBM i network installation	659
21.6.1 Requirements for IBM i network installation	659
Related publications	661
IBM Redbooks	661
Online resources	662
How to get Redbooks.	665
Help from IBM	666
Index	667

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™	iSeries®	Print Services Facility™
AIX®	Jazz™	Quickr™
AS/400®	Language Environment®	Rational Team Concert™
BladeCenter®	Lotus Enterprise Integrator®	Rational®
DataMirror®	Lotus®	Redbooks®
DB2®	MQSeries®	Redpaper™
developerWorks®	OmniFind®	Redbooks (logo)  ®
Domino®	OS/400®	Sametime®
DRDA®	POWER Hypervisor™	System i®
DS6000™	Power Systems™	System p®
DS8000®	Power Systems Software™	System Storage®
EnergyScale™	POWER5™	System x®
Express Servers™	POWER6+™	System z®
FlashCopy®	POWER6®	Tivoli®
i5/OS®	POWER7™	WebSphere®
IBM®	PowerHA™	xSeries®
iCluster®	PowerVM™	
Integrated Language Environment®	POWER®	

The following terms are trademarks of other companies:

Catalyst, the AMD Arrow logo, and combinations thereof, are trademarks of Advanced Micro Devices, Inc.

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.

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.

SUSE, the Novell logo, and the N logo are registered trademarks of Novell, Inc. in the United States and other countries.

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

QLogic, and the QLogic logo are registered trademarks of QLogic Corporation. SANblade is a registered trademark in the United States.

SAP, and SAP logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries.

Emulex, and the Emulex logo are trademarks or registered trademarks of Emulex Corporation.

Red Hat, and the Shadowman logo are trademarks or registered trademarks of Red Hat, Inc. in the U.S. and other countries.

VMware, the VMware "boxes" logo and design are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

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

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

Preface

This IBM® Redbooks® publication introduces a technical overview of the main new features, functions, and enhancements available in IBM i 7.1. It gives a summary and brief explanation of new capabilities and what has changed in the operating system, and also discusses many of the licensed programs and application development tools associated with IBM i.

Many of the new and enhanced functions are described:

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

The information provided in this book is useful for clients, business partners and IBM service professionals 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 25 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.



Ahmed Mansour is an IT Specialist and a PMI Project Manager Professional (PMP) at IBM Lab in Cairo, Egypt. Ahmed's areas of expertise include software development for IBM i, globalization and multicultural support for IBM products, and project management activities. He currently works with IBM Content Collector for SAP development team. He is also an IBM Certified Solution Advisor for Cloud Computing. Ahmed holds a Bachelor's degree in Computer Science from Cairo University, and has co-authored other IBM Redbooks publications.



Michael Harper is a Senior IT Specialist working for ANZ Techline in Australia. He has 33 years of experience in IBM having started as a Hardware Engineer before then working for the Software Support organisation and the IBM Sales supporting S/36, S/38, AS/400 and its follow on products. He has a wide knowledge of the hardware and the software of the IBM i platform and is currently supporting IBM Business Partners and Distributors in Australia and New Zealand in the design and configurations of Power Systems™ for customers of all sizes.



Robert Jacobsen is an IT Architect specializing in IBM i systems management technologies at the GTS Service Delivery Data Center in Rochester, MN. He joined IBM in 1981, working in S/38, S/36 and AS/400 development for more than a decade before moving to GTS. He has broad knowledge of and hands-on experience with IBM i hardware and software products. His current responsibilities include designing and implementing IBM i solutions, and evaluating and integrating new technologies into the data center environment. He is a graduate of North Dakota State University with an MS in Computer Science and Applied Mathematics.



Jiri Sochr works as IBM i IT Specialist at IBM Delivery Centre Central Europe in Brno, Czech Republic. He provides the 3rd level support for IBM i customers. He has more than 20 years experience in IBM i hardware and software products, IBM i Development, workload management, performance and also has a broad knowledge of other IBM products even not related to IBM i. He specializes in IBM i complex problem determination, multiplatform problems solution. The other area he works in is a IBM i solution design including external storage and multiplatform integration based on customer requests. He holds a doctor degree (RNDr.) in natural sciences from the Faculty of Mathematics and Physics, Charles University, Prague, Czech republic from 1990. Jiri joined IBM in 2009. Prior joining IBM he worked for IBM Business Partner company as IBM i system engineer and the technology leader.



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 participated in previous ITSO residencies, preparing Technical Overview presentation material and writing TCP/IP Communications-related Redbooks for IBM i. He currently is an Enhanced Technical Support remote Account Advocate for several IBM Power Systems customers.



Matias Centeno Lozada is a SW Support Specialist at the Latin America SW Support Center in Buenos Aires, Argentina. He joined IBM in July 2004 and has been in his current role for the last seven years. During such time he has contributed to the Rochester Support Center with around 110 documents and 4 support tools. He has 22 years of experience working with AS/400, iSeries®, System i® and Power Systems. He had visited IBM Rochester five times to participate on the Maintenance and Technical (MTS) Systems Support Global Internship Programs. He also taught a 7 days workshop in IBM Chile, IBM Colombia and IBM Venezuela with focus on i5/OS® and IBM i (installation, upgrade, troubleshooting, etc...), and a 3 days workshop in IBM Argentina with focus on i5/OS and IBM i (Virtual Media Images). His areas of expertise cover a wide range of topics including, but not limited to, Work Management, Security, APIs, MI, SLIC Macros, HMC, VIOS, and VPM. As a curiosity he owns two AS/400 9401-150s (Processors 2269 w/V4R5M0 & 2270 w/V5R2M0, respectively) that he uses for research and to improve his skills on AS/400 Machine-Level programming and PL/1.



Figure 0-1 From left to right, Ahmed Mansour, Jiri Sochr, Matias Centeno Lozada, Michael Harper, Tom Vernailen and Robert Jacobsen

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

Ann Lund
 Jenifer Servais
 Linda Robinson
 International Technical Support Organization, Poughkeepsie Center

Al Levering
 Alan Smith
 Allan E Johnson
 Armin Christofferson
 Bob Seemann

Brian K. Bratager
Brian Warford
Bunny Chaney
Chris Beyers
Christopher Francois
Cindy Mestad
Curt Schemmel
Dan Boyum
Darcy Rucker
Dave Charron
David B Murray
David R Bhaskaran
David S Charron
Dawn May
Duane Wenzel
Ellen Stacey

Ellie B. Streifel
Eric Stec
Gerald Simon
Ginny McCright
Glen Nelson
Guy Vig
James Flanagan
Jasmeet Singh
Jason Clegg
Jenny Dervin
Jerry Evans
Jessica Erber-Stark
Jim Flanagan
Jim Lembke
Jim Tilbury
Joe Kochan
Joe Mulholland
Johnnie R Talamantes
Jon Rush
Joseph Kochan
Jossie McManus
Kathryn A Tri
Kent L Bruinsma
Kent Milligan
Kris Whitney
Kristi Harney
Lilo Bucknell
Margaret Fenlon
Marilyn Martin
Mark Goff
Mark J Anderson
Mark J Hessler
Mark R Vanderwiel
Mark Schroeder
Mary Ann Pierce
Mervyn Venter
Randy Johnson
Rick Saltness

Ricky M Peterson
Robert Gagliardi
Robert J Bestgen
Robert Padzieski
Rodney Traff
Rodney Klingsporn
Roger Guderian
Ron McCargar
Scott Forstie
Shauna Smith
Stacy L Benfield

Stacy L. Haugen
Stephen A Knight
Steve Ransom
Steve Sparrow
Steve Will
Steven M Howe
Terry D. Ackman
Terry Hennessy
Terry Luebbe
Tim Klubertanz
Tim Kramer
Tim Rowe
Tom Crowley
Wayne Holm
Wesley Varela
IBM Rochester

Gang Shi
Jian Sang
Sheng Li Li
Ting Ting Sun
IBM China

Alison Butterill
Barbara Morris
George G Farr
Phil Coulthard
Philip Mawby
IBM Toronto

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

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

Ivan Berrios
Ingo Dimmer
Robert Jacobsen
Allyn Walsh
Tom Vernailen
Prema Sathasivam
Jos Vermaere

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>



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 Redbooks publication *IBM i 6.1 Technical Overview*, SG24-7713 was published. We highlight those enhancements and the new capabilities in the operating system, the licensed programs, and application development tools associated with IBM i. 7.1.

The following chapters provide additional 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 relating to IBM i 7.1.

IBM i Memo to Users 7.1 is an valuable starting point for readers, and is available at the following web page:

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

More detailed information about the IBM i 7.1 enhancements can be located in a variety of other sources:

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

1.1 Clarification of operating system terminology

When referring to operating systems releases, it is important to understand the naming transition from OS/400® to i5/OS to IBM i. When IBM introduced POWER5™, OS/400 was renamed i5/OS. When the 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.

1.2 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 allow clients to reduce cost, improve service, and manage risk of their IT infrastructure.

1.2.1 IBM i 7.1 on POWER based servers

IBM i 7.1 is supported on the following platforms:

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

1.2.2 Technology Refresh

IBM i 7.1 includes an enhancement to make it easier to deliver hardware and firmware enhancements by using 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, etc.

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

Future deliveries of new hardware-related function 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.

Throughout the whole book, we assume you have IBM i 7.1 with the latest Technology Refresh as well as the latest Group PTFs installed in order to get advantage of the new enhancements.

1.2.3 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 OmniFind® Text Search Server, adding support for searching XML documents. It is available with IBM i for 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. Advanced SQE query optimizer now supports native logical files. Also included is improved management tooling 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).

1.2.4 PowerHA SystemMirror for i

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

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

Additional enhancements are as follows:

- ▶ New PowerHA packaging
- ▶ PowerHA versioning
- ▶ New CL commands for programming cluster automation
- ▶ Cluster administrative domain enhancements
- ▶ Better detection of cluster node outages
- ▶ Improved geographic mirroring full synchronization performance
- ▶ IPv6 support

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

- ▶ Support for managing IBM System Storage SAN Volume Controller (SVC) 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.

1.2.5 IBM i virtualization

IBM i virtualization support is enhanced to offer 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, AIX®, and Linux client partitions, or IBM i 7.1 can be the server and virtualize the storage 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 using 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 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 Active Memory™ Sharing environment when fully hosted by VIOS or fully hosted by IBM i.

The Virtual Partition Manager (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 are enabled in VPM.

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, thereby freeing server resources that were in use by that partition, and later resume operation of that partition and its applications on the same server.

1.2.6 IBM i storage management

In release 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 even further by using SSD technology to collect I/O performance data and move data that is accessed most frequently to SSD storage. This 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 complete removal of internal disk units or external storage LUNs without an IPL or outage, and is supported on any POWER® hardware running IBM i 7.1. Prior to 7.1, an IPL (or a vary off of the independent ASP) was required to complete the process of removing storage disk units or LUNs. For more details, see 9.1.1, “Concurrent removal of disk units” on page 274.

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 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 have been included with IBM i 7.1:

- ▶ CEC 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 CEC Rack
- ▶ Thin Provisioning for DS8700, DS8800 storage servers and for VIOS shared storage pools
- ▶ 177 GB SFF SSD with eMLC
- ▶ IBM Disk Sanitizer PRPQ extended to include SSD devices

1.2.7 Application Development on IBM i Enhancements

New enhancements were done for Application Development on IBM i. The following sections highlight the enhancements.

High-level programming languages

Enhancements in High-level programming languages include:

- ▶ For C/C++, the following additions and modifications are 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 is an enhanced version of open source PHP. It is a lightweight version of Zend Server and replaces Zend Core. It is preloaded on IBM i 7.1.
- ▶ Zend Server for IBM i

This is a robust PHP production environment that helps ensure that application written in PHP run smoothly at all time. It replaced Zend platform. It offers all the features provided in Zend Server CE for IBM i plus more additional features.
- ▶ Zend Studio for IBM i

This product includes all the development components necessary for the full PHP application life cycle 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 consolidation of Lotus® collaboration solutions. IBM i 7.1.

Supports Lotus Domino® 8.5.1, Lotus Sametime® 8.5.1, Lotus Quickr™ 8.5, Lotus Enterprise Integrator® 8.5.2, Lotus Forms Server 3.5.1 FP2, IBM Forms Server 4.0, Lotus Workflow 7.0

Native API support of Archive and un-archive of (.zip) files

New enhancement for API support of native archive and un-archive of (.Zip) files. This support includes the following native APIs and a service program to create (.zip) archive files:

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

IBM Toolbox for Java JDBC Enhancements

Java JDBC interface now supports several features of the latest JDBC 4.1 definitions with DB2 i.

Several other enhancements are added like:

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

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 allows you to collect and verify a customized set of information, systems settings, and attributes about:

- ▶ Applications
- ▶ IBM i System
- ▶ Runtime Environment

ARE collects the information and build 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.

1.2.8 Rational Support for IBM i 7.1

The following sections discuss Rational Products for IBM i:

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

Rational Developer for Power System Software

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

Rational developer for Power Systems adds new feature that combines IBM i development tools with IBM Rational Application Developer Standard Edition for WebSphere® Software

V8.0 (RAD SE). The new Power tools allows developers to have all their development tools integrated into one work environment. Power tools also reduces the operation costs by deploying a single development workbench image to all developers.

Rational Team Concert

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

RTC 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 RTC 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.

1.2.9 Systems management enhancements

IBM i 7.1 continues the transition by deploying new features and functions 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 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 has been added, and 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 BladeCenter® and System x® installation wizard

1.2.10 Printing

The CPYSPLF command has been enhanced to copy to PDF or TIFF stream files.

1.2.11 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 10Gb. It uses standard Ethernet adapters and the network to connect IBM i on Power Systems to BladeCenter and System x. Additional operating system support has also been 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 deleting of servers in the IBM i Integrated Server environment.

1.2.12 Security enhancements in IBM i 7.1

ASP encryption can now be turned off and on 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 are user-written exit programs that are executed every time 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.

1.2.13 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 the Classic JVM. It replaced this by adding support for J9 VM.

The integrated Web Application Server, which ships with IBM i, supports Java5 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, V7.0 and V8.0 when ordered for IBM i 7.1.

Integrated Web Services for i enables integrated language environment (ILE) applications to play in the Web services and Service Oriented Architecture (SOA) arena with very little effort, knowledge and resources. The convergence of Web service 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.

1.2.14 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 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:
 - Additional option for viewing spooled files as PDF documents

1.2.15 Networking enhancements

Highlights include the new IBM i DHCP server, which is based on the Internet Systems Consortium (ISC) DHCP server enabling support for IPv6 and DHCP failover. IPv6 is also supported for the DHCP client. IPv6 and Remote Authentication Dial In User Service (RADIUS) supported has been 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 to link up to eight Ethernet links together in a single line description for improved throughput and reliability.

1.2.16 Upgrade for IBM i enhancements

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

1.2.17 IBM i network installation

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

1.2.18 IBM i Processor Entitlement Transfer

IBM i Processor Entitlement Transfer was announced in April, 2010. This 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 further details, see the program summary at the following web page:

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

1.2.19 BRMS 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 is a new function that 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 provides additional management functions for an improved user interface. A full suite of BRMS functions have been 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 has been included with IBM i 7.1



Technology Refresh

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

2.1 Overview

Following the General Availability (GA) of IBM i 7.1, IBM i point / modification releases have been replaced by a new release delivery mechanism called a Technology Refresh. A Technology Refresh 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.

The Technology Refresh PTF Group contains a Technology Refresh PTF plus related PTFs, including the most recent fixes prior to 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

We have the following items which are all described later on in this chapter:

- ▶ TR PTF - Technology Refresh PTF
- ▶ TR PTF Group -Technology Refresh PTF Group
- ▶ TR Reqs - Technology Refresh Requisite
- ▶ GA Resave - General Availability Resave
- ▶ CUM - Cumulative Package
- ▶ Managed Add Function PTFs
- ▶ Other PTFs
- ▶ Other Function GAed at the same time

2.2 What is in a Technology Refresh

A Technology Refresh may provide the following:

- ▶ Support for new hardware and firmware (e.g. 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. It is recommended that customers 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.

2.2.1 When to install a Technology Refresh

A Technology Refresh should be installed in the following cases:

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

It is a best practice to keep the system up to date with the latest Technology Refresh installed, to avoid unnecessary down time. To determine the Technology Refresh level needed for a specific hardware feature code, refer to the IBM Prerequisite tool. The PTF Cover Letter or Preventive Service Planning (PSP) information will identify 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, please refer to Preventive Service Planning website for additional info:

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

2.2.2 How it all fits together

Please see Table 2-1 on page 13 for 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
10/14/2011	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

Technology Refresh release date

This 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 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 prior to 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 (e.g. HIPER PTF Group) before 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 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 (e.g. RS710-A, RS710-B).

Marker PTF

The marker PTF is a PTF that can only appear on a system 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 found can be used to determine what Resave level of the product is installed.

5770-999 Technology Refresh PTF

This 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 will supersede previous versions. The currently active version of a Technology Refresh on a system is the PTF with the highest MF99xxx PTF ID in applied status.

2.2.3 How to determine the Technology Refresh PTF Group level installed on a system

In order to determine the Technology Refresh PTF Group level installed on a system, you can use WRKPTFGRP (Work with PTF Groups). Then find the PTF group named SF99707 as shown in Figure 2-2 on page 15. There may be multiple different levels of the group installed on the system. The latest level (the one with the highest level number) with the status of Installed is the current 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 installed

2.2.4 How to determine the Resave level installed on a system

In order to determine the Resave level of 5770-999, use DSPPTF (Display PTF Status) for product 5770-999 as shown in Figure 2-3 on page 16. 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 13, 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

2.2.5 How to determine the Technology Refresh PTF level installed on a system

In order to determine the Technology Refresh PTF level of 5770-999, use DSPPTF (Display PTF status) for product 5770-999 as shown in Figure 2-4 on page 17. 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 13 above, 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
                                More...

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

```

Figure 2-4 Technology Refresh PTF level installed

2.2.6 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 needed for hardware refer to the IBM Prerequisite tool.

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 be dependent on it, and those PTFs cannot be loaded until their requisite Technology Refresh PTF has been permanently applied, which requires an IPL. Therefore, it is a best 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 one need only apply the latest Technology Refresh to keep the system current.

2.2.7 How to install an IBM i Technology Refresh

The Technology Refresh PTF itself is enablement for supporting the new function of the Technology Refresh – it alone is not sufficient for obtaining the complete support of the new function, and may not contain the latest fixes available. Therefore, it is recommended that

when updating a system to a new Technology Refresh level, a client uses 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.

2.2.8 How to order 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 SNDPTFORD (Send PTF Order) 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.

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

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

2.2.9 How to install 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 INSPTF (Install Program Temporary Fix) command or by taking 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.

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp?topic=%2Frzam8%2Frzam8fixinstall1.htm>

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

Please refer to 2.2.14, “How to prevent or reduce the impact of a ‘double IPL?’” on page 19 for more information on how to avoid or reduce the impact of a ‘double IPL’ during PTF installation.

2.2.10 How to order and install an IBM i Resave

Refer to IBM i Resaves for instructions and other information related to ordering and installing an IBM i Resave.

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

2.2.11 How to order 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 SNDPTFORD (Send PTF Order) 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.

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

Before ordering a Technology Refresh PTF you think you may need, check and verify the PTF is not already on your system, as a requisite of another PTF.

2.2.12 How to install 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.

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=%2Frzam8%2Frzam8fixinstall1.htm>

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

Please refer to 2.2.14, “How to prevent or reduce the impact of a ‘double IPL’?” for more information on how to avoid or reduce the impact of a ‘double IPL’ during PTF installation.

2.2.13 How does a Technology Refresh PTF or PTF Group affect other PTFs?

If a later PTF changes a part or module that is contained in the Technology Refresh PTF, the Technology Refresh PTF will become a special kind of prerequisite for that PTF, called a *TRREQ. The Technology Refresh PTF will have to be permanently applied on the system before the PTF that requires it can be loaded. It therefore is a best 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, if it were found to be a requisite of another PTF that is needed. PTFs that do not involve parts or modules contained in a Technology Refresh PTF will not require the Technology Refresh PTF to have been applied before they can be loaded.

2.2.14 How to prevent or reduce the impact of a ‘double IPL’?

Customers can avoid the double IPL by making sure the Technology Refresh PTF has been permanently applied. Then if a PTF requires it, the Technology Refresh PTF will already be on the system. This will not reduce the total number of IPLs required, but it will allow scheduling the IPLs when they will be most convenient for the operation of the system.

Ordering and installing the Technology Refresh Resave will also ensure that the Technology Refresh PTF is permanently applied and the “double IPL” avoided.

New function PTF SI43585 is available to automate, not eliminate, any extra IPL required during PTF install. When you’re installing PTFs, there are two conditions where we prompt you to perform an IPL to apply some of the PTFs, require you to restart the PTF install after the first IPL, and then make you perform yet 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, we’ll save your PTF install parameters and use them during the next IPL. Instead of seeing “Confirm IPL for Technology Refresh or Special Handling PTFs” panel, you’ll 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 won’t see any messages or panels, we’ll just power down. On the next normal IPL, we’ll

do your second “GO PTF” during the “PTF Processing” IPL step in the SCPF job, and then automatically do a second IPL of the partition. So when the system IPLs the second time up to sign-on, your PTFs are all activated and ready to go.

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

2.2.15 Why has IBM moved to this Technology Refresh process?

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

Moving up to a Technology Refresh should be simpler and cheaper than qualifying a point release, so that customers can take advantage of new function 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 have already been fixed, unnecessary downtime and calls to IBM support.

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

While “backing out” a point / modification release requires a scratch install 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.

2.2.16 Will there still be new releases of IBM i?

Yes, there will still be new releases of IBM i. With a Technology Refresh, only the changed parts 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 together. Developers are working on the next release of IBM i and architects are looking at possible content for the next release beyond that.

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

2.2.17 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, via 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.

2.2.18 Related publications

For more information related to Technology Refreshes please refer to the following:

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 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/slkbases.NSF/DocNumber/14854405

Guide to fixes

<http://www-947.ibm.com/systems/support/i/fixes/guide/>



Security enhancements

This chapter describes the following Security Enhancements for IBM i 7.1:

- ▶ “ASP encryption enhancements” on page 24
- ▶ “Column Encryption” on page 24
- ▶ “User profiles enhancements” on page 24

3.1 ASP encryption enhancements

In IBM i 7.1 enhancements were made to ASP encryption. ASP encryption can now be turned off and on 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 needs to be processed. This can affect system performance.

For more information about ASP encryption enhancements, see 9.1.4, “Encrypted ASP enhancements” on page 282.

3.2 Column Encryption

To enhance data security, column encryption can be accomplished by using a new database feature called *field procedures*.

Field procedures are user-written exit programs that are run every time a column is changed or new values inserted.

For more information about column encryption see 6.2.6, “FIELDPROC support for encoding and encryption” on page 124.

3.3 User profiles enhancements

This section covers enhancements that were made to user profile commands.

3.3.1 CRTUSRPRF and CHGUSRPRF commands

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

- User expiration date USREXPDATE

This 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	To be calculated based on the value entered in the user expiration interval parameter.
Date	Specifies a date when the user profile expires. Date must be in job date format.

► User expiration interval USREXPITV

This parameter specifies the number of days (1–366) before the user profile is automatically disabled. See Table 3-2.

Table 3-2 Value for the USREXPITV

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

Note: 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.

Note: The aforementioned parameters can only be seen when using the character based interface.

3.3.2 Display Expiration Schedule

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

User Profile Expiration Schedule				
User Profile	Expiration Date	Action	Owned Object Option	New Owner
CHUA	04/23/10	*DELETE	*CHGOWN	PREMA
MARIE	04/23/10	*DISABLE		
<p>F3=Exit F11=Primary group info F12=Cancel F17=Top F18=Bottom</p> <p>(C) COPYRIGHT IBM CORP. 1980, 2009.</p>				

Figure 3-1 Output from DSPEXPSCD

TIP: To change the expiration action for a specific user profile run the CHGEXPSCDE (Change Expiration Schedule Entry) command.

3.4 Other security enhancements

3.4.1 Printing

See 11.3.2, “New QIBM_QSP_SECURITY exit point and formats” on page 347

3.4.2 DB2 for i

See 6.6, “DB2 security enhancements” on page 171



Backup and recovery

This chapter discusses enhancements to the IBM i backup and recovery functions.

We start with basic 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 28.

After a discussion of the basics, this chapter moves on to the more advanced Backup Recovery and Media Services (BRMS) product in 4.2, “New and improved backup and recovery and media services for IBM i functions” on page 31.

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 53.

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

4.1 New and enhanced system save and restore functions

This section reviews 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” on page 28
- ▶ “New ALWOBJDIF (*COMPATIBLE) restore option” on page 29
- ▶ “Enhanced save and restore limits” on page 29
- ▶ “Save While Active Support for Integrated VMware ESX Servers” on page 30
- ▶ “Miscellaneous enhancements” on page 30

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, saving time by minimizing tape searching for the object. This enables a user 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 containing 1.1 million IFS objects reduced object restore time from 22 minutes to less than three 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 previously mentioned.

The following restore interfaces support the POSITION parameter:

- ▶ RSTLIB (Restore Library), RSTOBJ (Restore Object) and RST (Restore IFS Object) 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 *FIRST, which restores using the current search from the beginning mode. When using 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 RSTOBJ (Restore Object) 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 specifying 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 has been added to the ALWOBJDIF (Allow Object Differences) 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 leaves the user 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:

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

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

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

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

4.1.3 Enhanced save and restore limits

Prior to IBM i 7.1, database files containing more than 16 MB of descriptive information could not be saved. This restriction has been removed. This 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 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. Refer to the IBM i integration with BladeCenter and System x website for more information.

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

4.1.5 Miscellaneous enhancements

There are numerous miscellaneous enhancements:

- ▶ Supported TGTRLS (Save Target) releases are V7R1M0, V6R1M0, and V5R4M0.
- ▶ SAVLIB and SAVCHGOBJ of library QUSRSYS now includes performing RTVSYNINF and saves the results.

If recommended procedures are followed, this information is already saved. Often, this data is inadvertently not saved and cannot be recovered. This change ensures 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 spoolfiles.
- ▶ Save and restore menu options that bring the system to a restricted state have been enhanced to gracefully end TCP/IP servers, host servers, and TCP/IP before ending to 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 include 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 are logged and the ability to adjust the size of the data.

To change the flight recorder size to 10 MB, use 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 library QTEMP or QGPL.

4.2 New and improved backup and recovery and media services for IBM i functions

This section reviews new and improved functions that are part of the BRMS product. Topics are as follows:

- ▶ “Support for Domino 8.5 Domino Attachment and Object Service (DAOS)” on page 31
- ▶ “Link list enhancements” on page 32
- ▶ “Support of more than 999,999 objects” on page 32
- ▶ “STRRCYBRM (Start recovery using BRM) command enhancements” on page 33
- ▶ “Recovery report enhancements” on page 36
- ▶ “Output file support for BRMS restores” on page 37
- ▶ “Ability to select Saved ASP on RSTLIBBRM and RSTOBJBRM commands” on page 42
- ▶ “Distributed backup support” on page 43
- ▶ “Maintenance enhancements” on page 47
- ▶ “Planning Media Movement Report” on page 49
- ▶ “Improved clearing of volumes in *ERR status” on page 50
- ▶ “Protection of volumes marked for duplication” on page 51
- ▶ “Improved recovery times with new media position function” on page 52
- ▶ “BRMS support for special value *COMPATIBLE for ALWOBJDIF parameter” on page 53
- ▶ “Improved control when running non-BRMS saves” on page 53

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

Domino Release 8.5 contain a new enhancement 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's attachments.

Before DAOS, attachments were part of each Domino database (.nsf) file. If a large attachment were 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 described in the previous paragraph, rather than one occurrence of the attachment stored in each mail file, there is one NLO stored per Domino server, thus saving storage space.

BRMS has been 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 has been made available through PTFs for V6R1 (SI34918) and V5R4 (SI31916)

When configuring DAOS on Domino servers, be careful with the configuration of attachment sizes that are externalized into NLOs. If you select a small size, a large number of NLO objects can be created, each of which is an IFS object that can significantly lengthen the IFS backup time. The default is 4096, however 1000000 or larger is recommended.

DAOS references

The following references provide more information regarding DAOS:

- ▶ 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>

4.2.2 Link list enhancements

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

The QIBMLINK link list for IBM IFS directories is now automatically added to the supplied system backup control group *SYSGRP for new installs only. In V5R4 i5/OS and IBM i 6.1, QIBMLINK existed, but was not automatically added to *SYSGRP. It is advised that existing installations add QIBMLINK manually to *SYSGRP. 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 has been added in IBM i 7.1. QALLUSRLNK is used to save user IFS directories and files. QALLUSRLNK is to be used in conjunction with the QIBMLINK link list. QALLUSRLNK omits the following directories:

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

Use of QIBMLINK followed by QALLUSRLNK enables more granularity than the *LINK control group entry and ensures that IBM directories are restored prior to user directories in case a system restore is necessary. The use 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

In support of more than 999,999 objects, BRMS enhanced the WRKMEDIBRM panel and the BRMS recovery report QP1ARCY.

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 on page 33 shows a WRKMEDIBRM Object Detail panel. The circled field shows a saved item with more than 999,999 objects.

```

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	Item Type	-----Objects-----	Control	Obj Dt1
					Saved Not Saved	Group	
—	*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
—	Q1WWTCL	3/20/10	3:42:31	*LIB	80	0 BKUP6	*YES
—	QUSRBRM	3/20/10	4:01:17	*OBJ	19	0 *NONE	*YES

```

F3=Exit F5=Refresh F11=Duplication status F12=Cancel
F23=More options
Bottom

```

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										
File	QP1ARCY								Page/Line	20/49
Control	-1								Columns	1 - 130
Find										
*.....1.....2.....3.....4.....5.....6.....7.....8.....9.....0.....1.....2.....3										

STEP 019 : Recover Directories and Files										
Start date/time _____ Stop date/time _____ Duration _____										
You should restore the current version of your objects in directories.										
If you are performing a complete system restore, run the following										
command to continue.										
STRRCYBRM OPTION(*RESUME)										
otherwise, run the following command.										
STRRCYBRM OPTION(*LNKLIST) ACTION(*RESTORE)										
Type the command choice and press "Enter".										
Select the saved item(s) listed below from the "Select Recovery Items"										
display and press "Enter" to recover these saved items. Recovery of										
these saved items will require the volumes listed on the report or										
duplicate volumes.										

Saved	Save	----	ASP	-----	Save	Save	---	Sequence	Control	Volume
Item	Type	Name		Number	Date	Time	Not	Number	Group	Identifier
-----	-----	-----	----	-----	-----	-----	Saved	-----	-----	-----
*LINK	*FULL	*SYSBAS		00001	3/13/10	2:26:30	1222113	5	294 BKUP6	LT0899

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

4.2.4 STRRCYBRM (Start recovery using BRM) command enhancements

The STRRCYBRM command has been enhanced to override specific recovery elements to use another time period. This requires that the override recovery element (OVERRIDE) parameter be set to *YES. This 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, *SECDA and *SAVCFG.

- ▶ ***SAVSYS**

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

- ▶ ***SECDA**

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 ignore 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 these. Each set has the following parameters:

- Recovery element

- ***ALLSYS**

This element specifies that you want recover the entire system using an override. I 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.

- ***SECDA**

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 created before the specified time and date are not included in the items 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 have been saved. Any entries saved prior to 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 created after the specified time and date are not included in the recovery items 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 have been saved. Any recovery items 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 found regardless of time, except for the operating system restore items, which selects nothing newer than 6 p.m. on 03/01/2010 due to 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 noting the override, which is an indicator to the user that the latest saved data is not being used. See the red box in Figure 4-3.

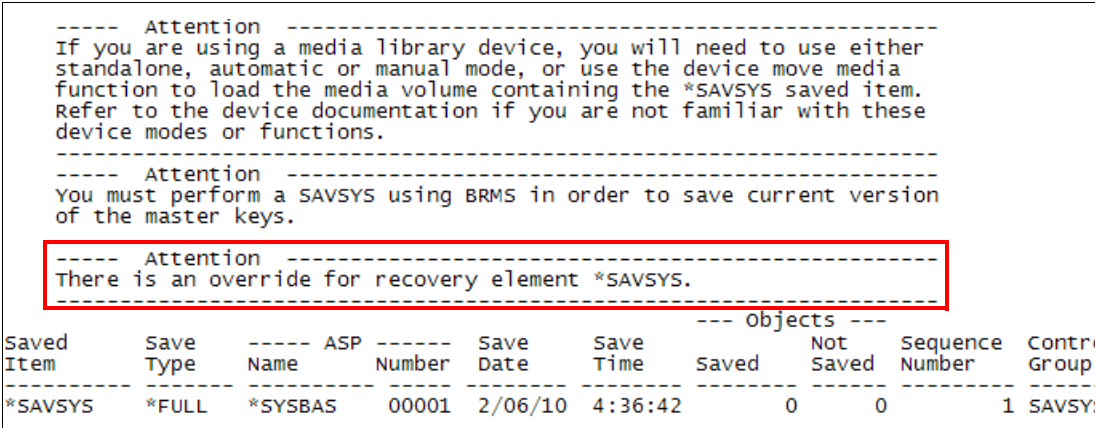


Figure 4-3 BRMS recovery report with override attention block

4.2.5 Recovery report enhancements

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

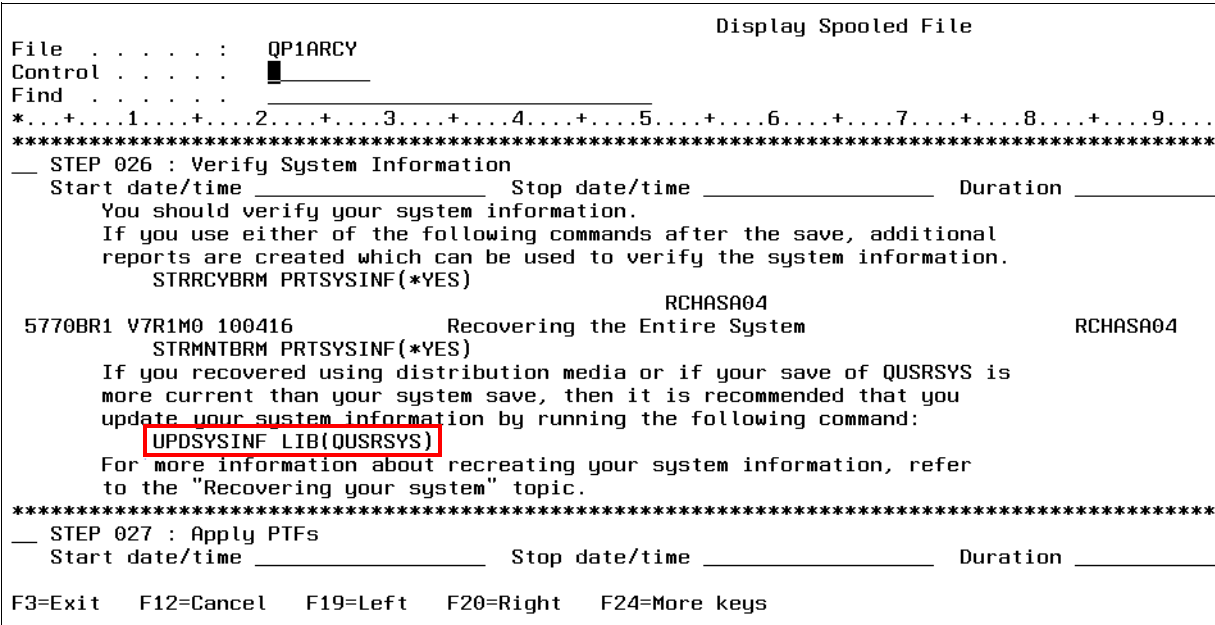


Figure 4-4 BRMS recovery report with new UPDSYSINF instructions

The “Restoring User Profiles” step has been updated to include the new Allow Object Differences *COMPATIBLE value, replacing the *ALL value of previous releases, as circled in red in Figure 4-5 on page 37.

Display Spooled File	
File	QP1ARCY
Control	<input checked="" type="checkbox"/>
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 saved through BRMS. The *OUTFILE value can be specified on the OUTPUT parameter in these functions:

- BRMS recovery policy

Figure 4-6 displays 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.	<input checked="" type="checkbox"/> *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	
F3=Exit F4=Prompt F5=Refresh F9=System policy F12=Cancel		Bottom

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:

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

- RSTLIBBRM (Restore Library using BRM) command
- RSTOBJBRM (Restore Object using BRM) command
- RSTBRM (Restore (IFS) Object using BRM) 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 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 file LIBS in library RESTORE.

```
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

Bottom
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 WRKMEDIBRM (Work with Media Information) command.

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

The following sequence of WRKMEDIBRM panels reaches the Recovery Defaults panel, which contains the new output parameters: From the WRKMEDIBRM (Work with Media Information) panel, select option 7 (Restore) for the library to be restored as circled 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
—	Q1WJTLCL	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, 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.

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
<u>1</u>	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, Step 2

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

Restore Command Defaults

Type information, press Enter.

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

F12=Cancel

Figure 4-10 WRKMEDIBRM recovery defaults panel

- STRRCYBRM (Start Recovery using BRM) 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 show 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.

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	Save 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
-	APSSD010	3/24/10	1:21:23	*FULL		LT0894	14	4/13/10
-	APSSP010.0	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...
F3=Exit F5=Refresh F9=Recovery defaults F11=Object View								
F12=Cancel F14=Submit to batch F16=Select								

Figure 4-11 STRRCYBRM Selected Recovery Items

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 discussed on page 37, 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: Restore options

4.2.7 Ability to select Saved ASP on RSTLIBBRM and RSTOBJBRM commands

This new function allows the user to specify the auxiliary storage pool (ASP) from which a library is saved when restoring on the RSTLIBBRM (Restore Library using BRM) command and the RSTOBJBRM (Restore Object using BRM) command.

Suppose a system has a system ASP and three independent ASPs. Each of the ASPs have library TOOLS and the entire system, including the IASPs has been saved. There are three saves of library TOOLS. This function allows selecting which of the saves is to be restored.

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

- ▶ ***ANY**
The library and objects saved is restored from any ASPs save. This is the default value, which works as did before IBM i 7.1.
- ▶ ***SYSTEM**
The saved library and objects are restored from the system ASP's save.
- ▶ **ASP number 1 through 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's save.

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

There are restrictions for which objects can be restored to non-system ASPs. These are objects that 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 . . . . . > TOOLS      Name, *MEDINF
      + for more values
Device . . . . . > MLD50      Name, *MEDCLS
      + for more values
Parallel device resources:
  Minimum resources . . . . . *SAV      1-32, *SAV, *NONE, *AVAIL
  Maximum resources . . . . . *MIN      1-32, *MIN, *AVAIL
  Save level . . . . . *CURRENT      1-99, *CURRENT, *SAVDATE
  Save level time reference:
    Save date . . . . .      Date
    Save time . . . . . *LATEST      Time, *LATEST
  Saved auxiliary storage pool . . > TESTASP      Name, 2-32, *ANY, *SYSTEM
  End of media option . . . . . *REWIND      *REWIND, *LEAVE, *UNLOAD
  Option . . . . . *ALL      *ALL, *NEW, *OLD, *FREE
  Database member option . . . . . *MATCH      *MATCH, *ALL, *NEW, *OLD
  Spooled file data . . . . . *NEW      *NEW, *NONE
                                          More...

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

```

Figure 4-13 SAVASP (Save ASP) keyword of 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. This option reduces the data on systems that have no need to know about the saved history from other systems. This part of 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) had no option 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	Network ID	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 at 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

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

-----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

► 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 will include 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 is 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 will include library level information only.

► Remote overrides field

This field represents media information overrides that are in place for the remote system for I-ASP high availability support discussed in 4.3.10, “High availability support for independent ASPs in a BRMS network” on page 84.

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 performing media moves through the STRMNTBRM (Start Maintenance for BRM) command, the user now receives 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.

When maintenance is running and a second job issues a BRMS command that attempts to use files in library QUSRBRM used by the maintenance job, a BRM6714 Job (job-name) is being held by job (maintenance-job-name) is issued to that second job’s message queue and is displayed as shown in 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/MSRJ/374662 is being held by job QPADEV000F/MSSSJ/274660.
```

Figure 4-17 BRMS6714 message displayed

When maintenance has used 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 executing 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 . . : JACOBSEN01      User . . : MSRJJ      System: RCHASA004
                                           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 STRMNTBRM job message queue

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

4.2.10 Planning Media Movement Report

Users can now print a report that can help them plan future moves of BRMS media.

The PRTMOVBRM (Print Media Movement) 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 the following panel (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. :
: :
: More... :
: 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 used by BRMS.

4.2.11 Improved clearing of volumes in *ERR status

Improvements have been made to simplify and speed the clearing of tape volumes in error status (*ERR). Prior releases required the user to remove the media volume from inventory, add the media volume back in inventory, 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 WRKMEDBRM (Work with Media using BRM) command. This 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      Creation Expiration  Move   Media   Dup
   Serial      Status   Date      Date      Location Date   Class   Sts
---
 1 JA3860      *EXP  03/09/10  *NONE    AVAIL    *NONE   AVAIL
 2 JA3861      *EXP  03/09/10  *NONE    AVAIL    *NONE   AVAIL
 3 JA3862      *EXP  03/09/10  *NONE    AVAIL    *NONE   AVAIL
 4 JA3863      *EXP  03/09/10  *NONE    AVAIL    *NONE   AVAIL
 5 JA3864      *EXP  03/09/10  *NONE    AVAIL    *NONE   AVAIL
 6 JA3865      *EXP  03/09/10  *NONE    AVAIL    *NONE   AVAIL
 7 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 *INZ status, the media volume can be reinitialized using one of the following commands

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

The user needs to verify 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 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 STRMNTBRM (Start Maintenance for BRM) command is executed, when the STREXPBRM (Start Expiration for BRM) command is executed, or when the WRKMEDBRM (Work with Media using BRM) 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

Control

Find

Page/Line 1/2

Columns 1 - 130

*...+...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 BRMS expiration report

When trying 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

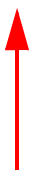
Position to _____
Starting characters

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 attempt to expire volume marked for duplication

4.2.13 Improved recovery times with new media position function

BRMS supports the new media position function mentioned in 4.1.1, “Fast restore using tape position” on page 28.

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 SAVLIBBRM (Save Library using BRM) 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),7 (specify object).
- ▶ WRKMEDIBRM option 9 (Work with Saved Objects). option 9 (Work with Saved Members), 7 (Restore), 7 (Specify object).

4.2.14 BRMS support for special value *COMPATIBLE for ALWOBJDIF parameter

BRMS restore functions support the *COMPATIBLE special value for the ALWOBJDIF parameter previously mentioned in 4.1.2, “New ALWOBJDIF (*COMPATIBLE) restore option” on page 29.

4.2.15 Improved control when running non-BRMS saves

When running 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 the user the option to load another non-BRMS volume.

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 have been 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 have been 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 notes which are also new to the System i Navigator product. This section describes the following enhanced functions:

- ▶ 4.3.1, “Added support for IBM Systems Director web browser environment” on page 54
 - 4.3.2, “IBM Systems Director navigation to BRMS functions” on page 54
 - 4.3.3, “IBM Systems Director Navigator for i Navigation to BRMS Functions” on page 59
- ▶ 4.3.4, “Enhancements to the BRMS initial panel” on page 61
- ▶ 4.3.5, “BRMS advanced functions panel” on page 62
- ▶ 4.3.6, “Scheduling Support for BRMS” on page 62
- ▶ 4.3.7, “Added option to BRMS log to filter messages by control groups” on page 72
- ▶ 4.3.8, “Ability to mark and unmark volumes for duplication” on page 75
- ▶ 4.3.9, “Multiple email address support” on page 82
- ▶ 4.3.10, “High availability support for independent ASPs in a BRMS network” on page 84
- ▶ 4.3.11, “Enhanced maintenance features in backup policy” on page 88

Note: To use the IBM i 7.1 enhancements, the user 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 has been added to IBM Systems Director and IBM Systems Director Navigator for i. See 18.4, “New journal management enhancements” on page 548 for more information.

4.3.1 Added support for 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 panel.

IBM Systems Director is intended for multiple systems and multiple system platforms. IBM Systems Director Navigator for i is intended to manage 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 discusses how this is done. This section also discusses the main BRMS pane functions.

To access the BRMS functions for a particular iSeries or IBM i, you must perform the following steps:

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

Having logged on the IBM Systems Director, you are greeted with the Systems Director Welcome panel. In the view pane, at the upper left corner is the “Navigate Resources” link as in Figure 4-23.

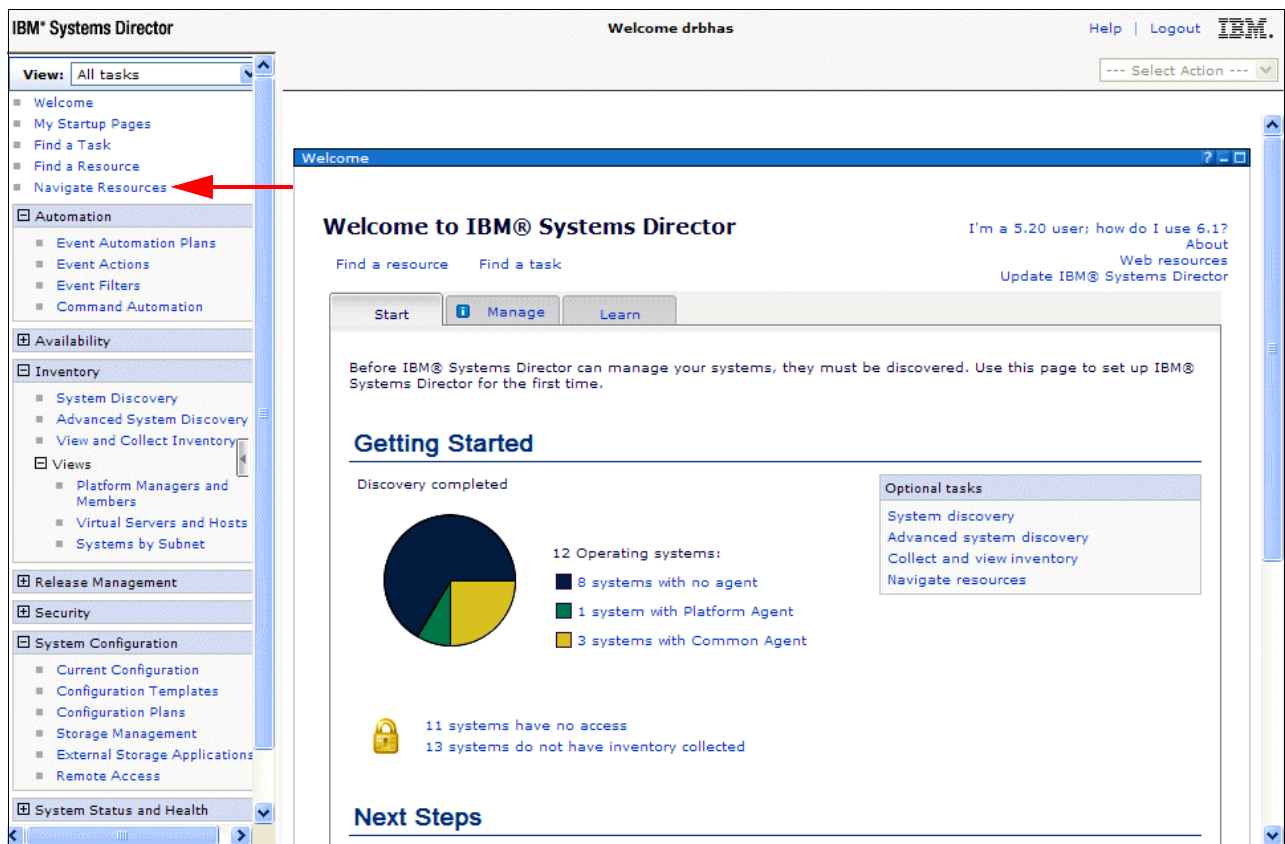


Figure 4-23 IBM Systems Director Welcome panel and Navigate Resources link

When you select the “Navigate Resources” link, a new panel listing the various resource groups is shown (Figure 4-24). 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.

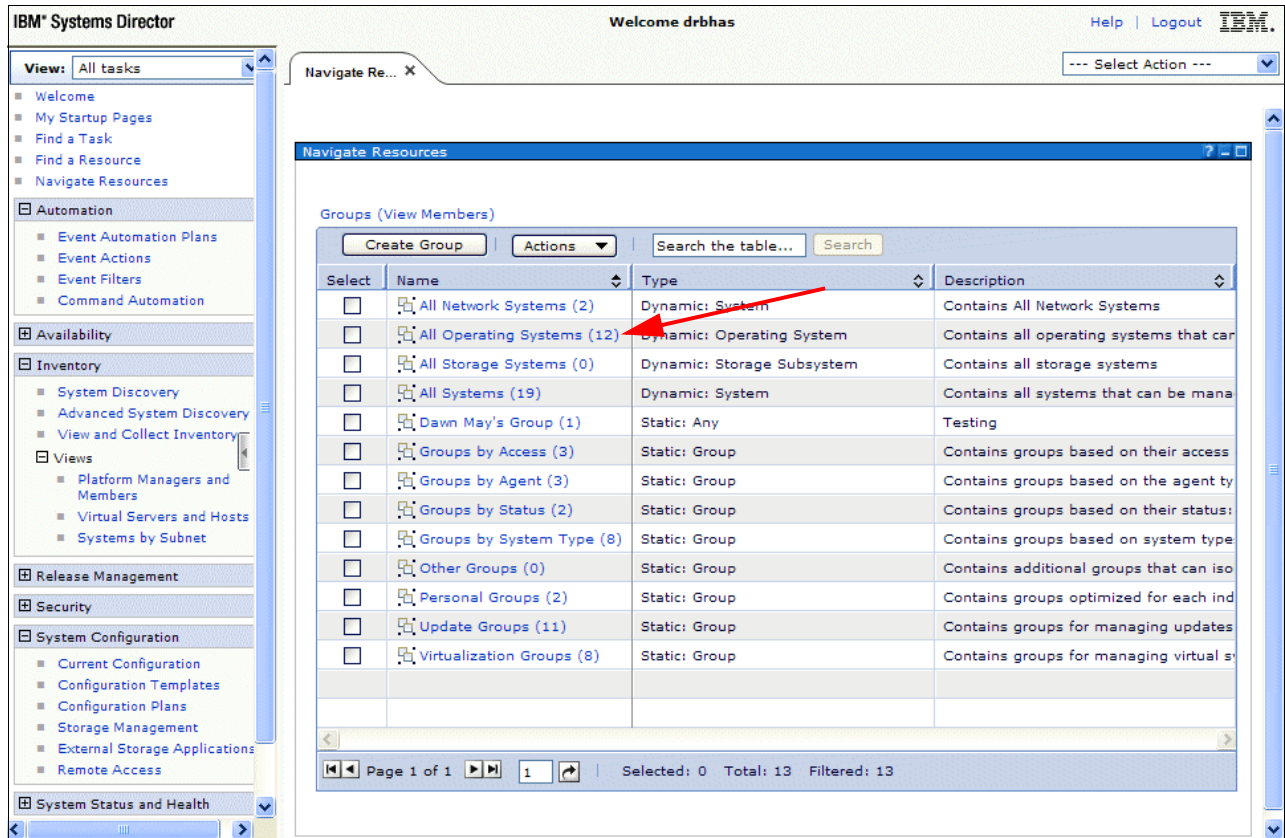


Figure 4-24 IBM Systems Director Navigate Resources group list panel

Having selected the “All Operating Systems” group’s link, the panel in Figure 4-25 is displayed 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 panel is titled 'Navigate Resources' and displays a table of member 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 link for 'lsz1p13.rchland.ibm.com' in the OS Type column.

Select	Name	Access	Problems	Compliance	IP Addresses	OS Type
<input type="checkbox"/>	elm8a105.beaverton.ibm.com	Offline	OK	OK	9.47.80.105	Unknown
<input type="checkbox"/>	elm8a155.beaverton.ibm.com	Offline	OK	OK	9.47.80.155, 19;	Windows
<input type="checkbox"/>	elm8a167.beaverton.ibm.com	Offline	OK	OK	9.47.80.167	Unknown
<input type="checkbox"/>	elm8a172.beaverton.ibm.com	Offline	OK	OK	9.47.80.172	Unknown
<input type="checkbox"/>	elm8a20.beaverton.ibm.com	Offline	OK	OK	9.47.80.20	Unknown
<input type="checkbox"/>	elm8a22.beaverton.ibm.com	Offline	OK	OK	9.47.80.22	Hypervis
<input type="checkbox"/>	lsz1p13.rchland.ibm.com	OK	OK	OK	9.5.215.103	IBM i
<input type="checkbox"/>	rchasyum.rchland.ibm.com	Offline	Information	OK	9.5.50.82	IBM i
<input type="checkbox"/>	x0716p3.rchland.ibm.com	OK	OK	OK	9.5.53.254, 127;	IBM i
<input type="checkbox"/>	x0716p7.rchland.ibm.com	OK	Minor	OK	9.5.52.190, 200;	Linux
<input type="checkbox"/>	z1014p26.rchland.ibm.com	OK	OK	OK	9.5.103.39	IBM i
<input type="checkbox"/>	z1014p28.rchland.ibm.com	No access	Information	OK	9.5.105.200	IBM i

Page 1 of 1 | Selected: 0 Total: 12 Filtered: 12

Figure 4-25 IBM Systems Director Navigate Resources Operating System Group list

Having selected the system, the resource panel in Figure 4-26 is displayed. When you select the Actions drop down menu, the first drop down box 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 4-26.

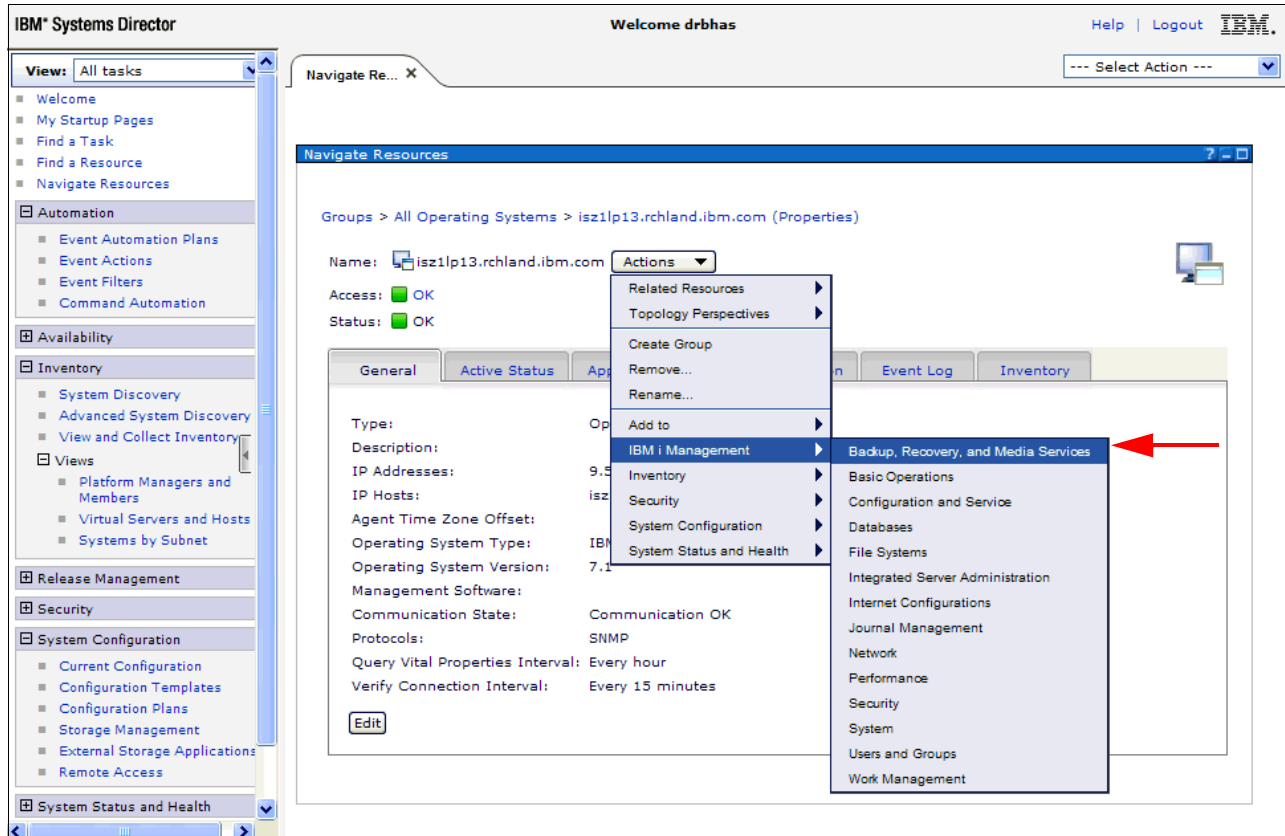


Figure 4-26 Navigation from selected system to the BRMS function

Having selected the Backup Recovery and Media Services link, the BRMS initial panel is displayed, as in Figure 4-27.

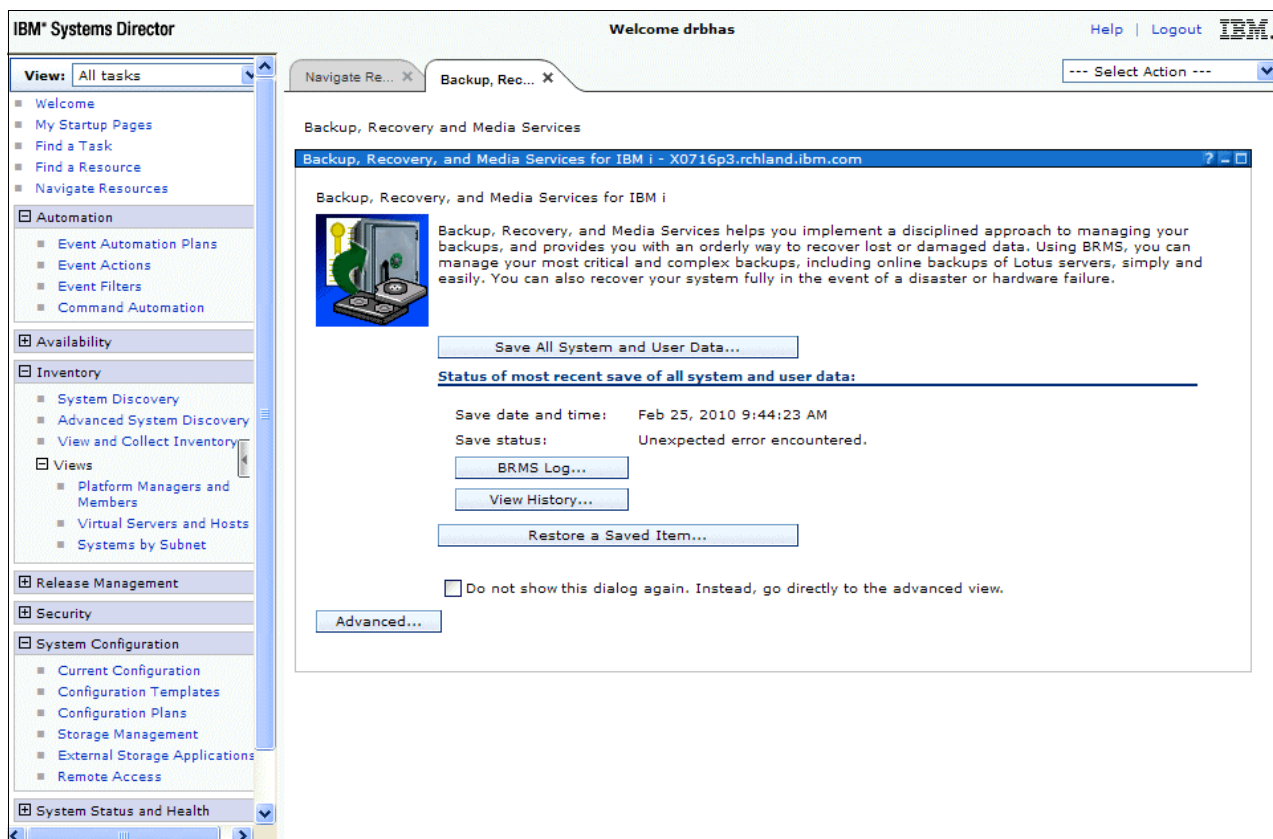


Figure 4-27 IBM Systems Director BRMS initial panel

The initial panel is discussed later in 4.3.4, “Enhancements to the BRMS initial panel” on page 61.

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 to the access the BRMS function than the IBM Systems Director product. Although this is not new function, it is included here for reference. This section reviews how this is done and discusses changes to the main panel of the BRMS functions.

To access the IBM Systems Director Navigator for i welcome page (Figure 4-28 on page 60), you must meet the following conditions:

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

Upon successful logging on, you are greeted with the Welcome pane shown in Figure 4-28. At the upper left corner of the navigation area is a minimized item labelled IBM i Management to which the red arrow points. Select this item to expand it.

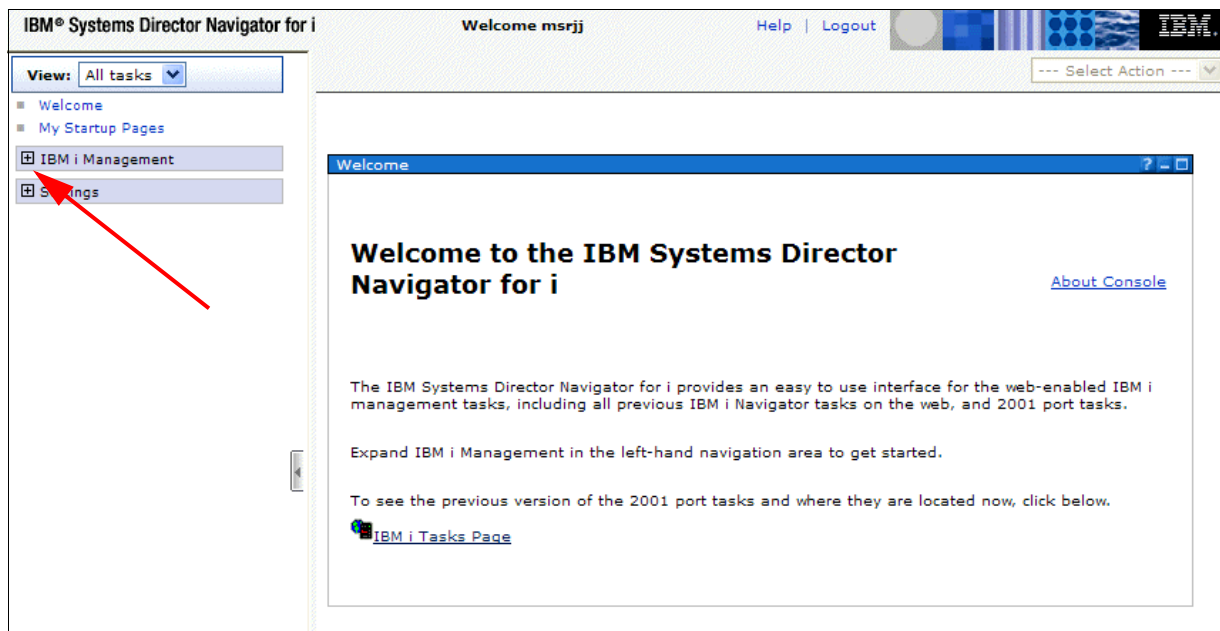


Figure 4-28 IBM Systems Director Navigator for i welcome panel

After expanding the IBM i Management item, you see a list of links to IBM i management functions (Figure 4-29).

At the bottom of the list is the **Backup, Recovery and Media Services** option. Click this link, a tab opens and the BRMS initial panel is shown.

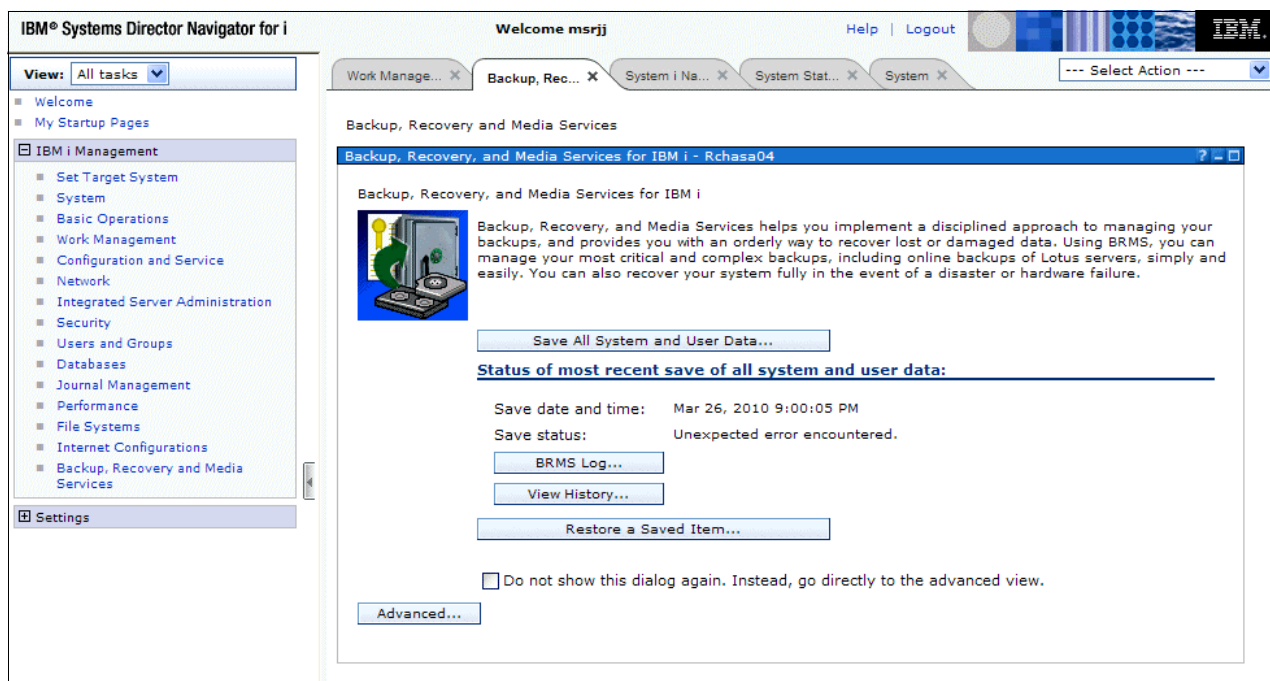


Figure 4-29 IBM System Director Navigator for i BRMS initial panel

4.3.4 Enhancements to the BRMS initial panel

This section discusses changes to the initial panel of the BRMS functions as accessed through IBM Systems Director and IBM Systems Director Navigator for i. The BRMS tab with the initial panel are nearly identical for both.

Here is what's new on the panel - see the number on Figure 4-30.

1. The date, time and status of the most recent save of all system and user data is shown.
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 aren't 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 panel and to go directly to the advanced panel.

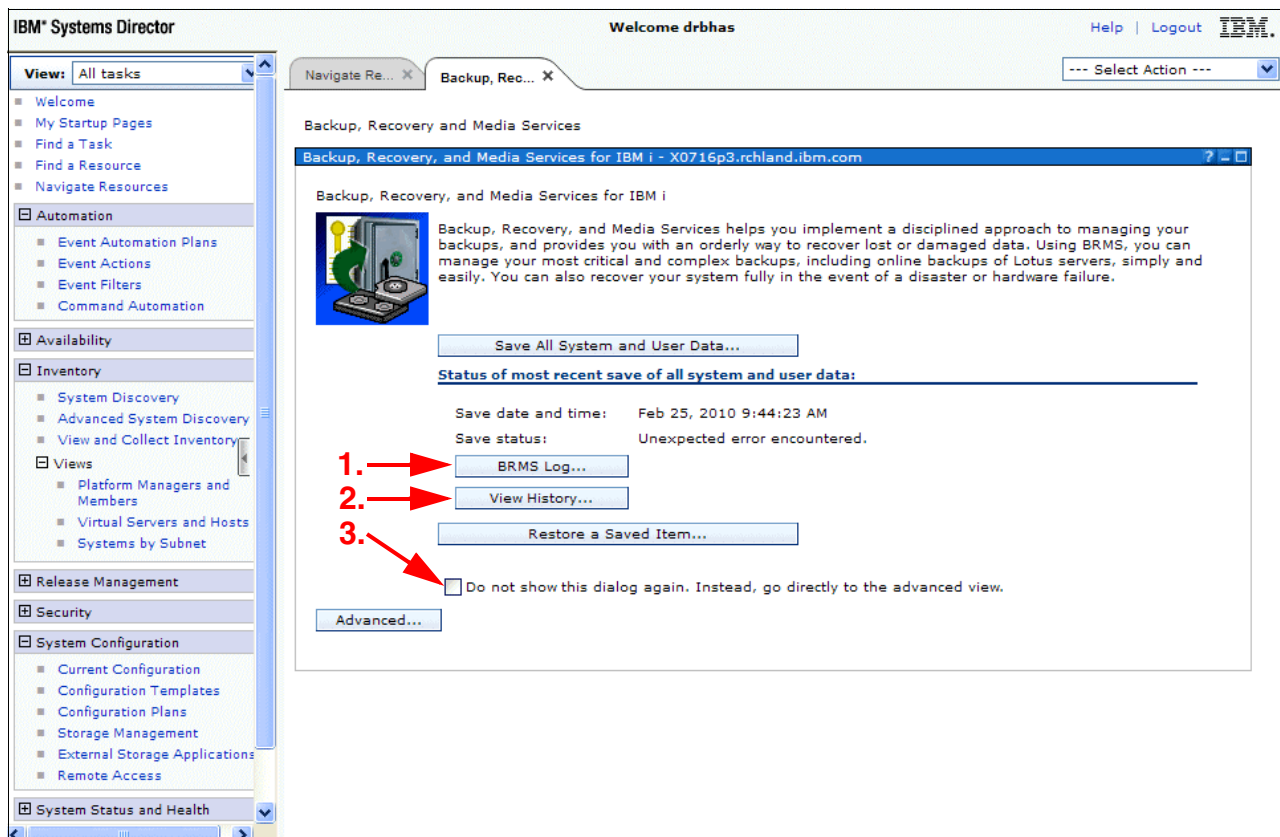


Figure 4-30 Enhancements to BRMS web initial panel

4.3.5 BRMS advanced functions panel

To access the BRMS advanced functions from the initial panel select the **Advanced** button. When the **Advanced** button is selected, the BRMS advanced functions panel is displayed, as shown in Figure 4-31.

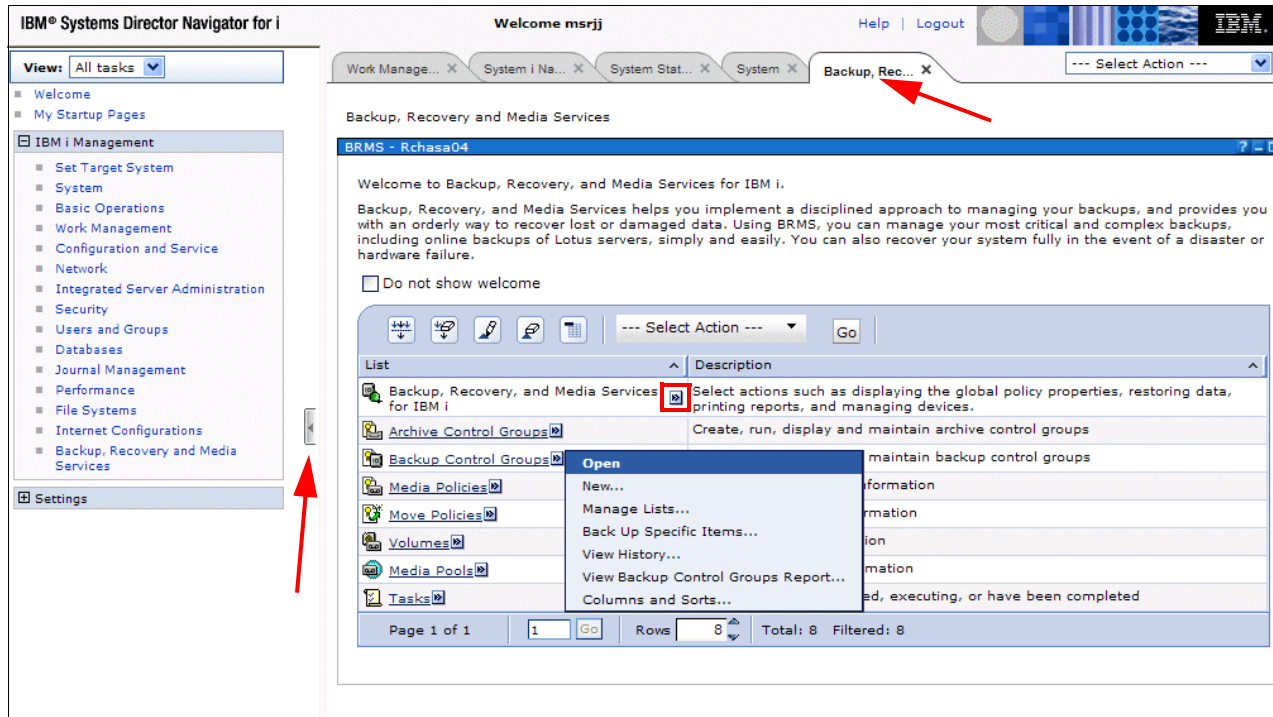


Figure 4-31 BRMS Advanced Function panel

The format of the panel is a page with a tab, selecting the tab (top arrow) brings that page to the foreground and places the others behind.

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

The small circled icon, when selected, opens a pop-up menu of actions. In Figure 4-31, the menu is shown for the BRMS Backup Control Groups field.

4.3.6 Scheduling Support for BRMS

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

- ▶ Backup 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 pop-up menu for Backup Control Groups from the BRMS advanced menu, as shown in Figure 4-32.

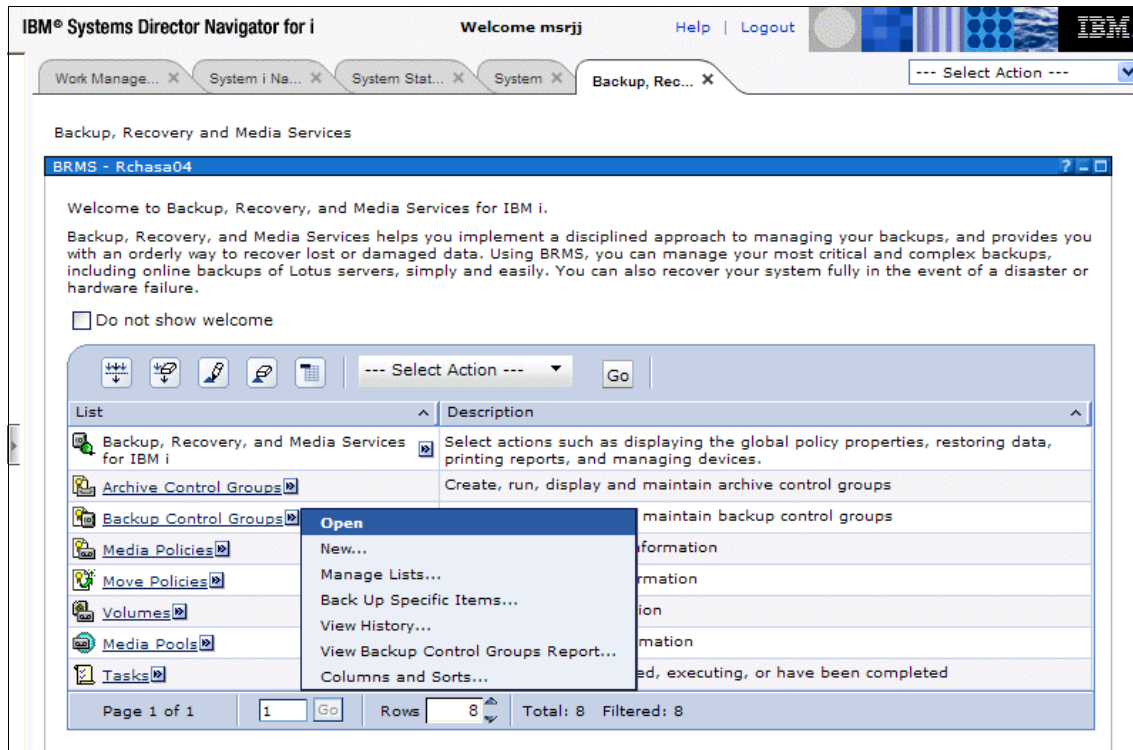


Figure 4-32 Open action of Backup Control Groups pop-up menu

Open Backup Control Groups

Perform the following steps to open backup control groups:

1. Select the Open action. A table of backup control groups (Figure 4-33) is displayed.

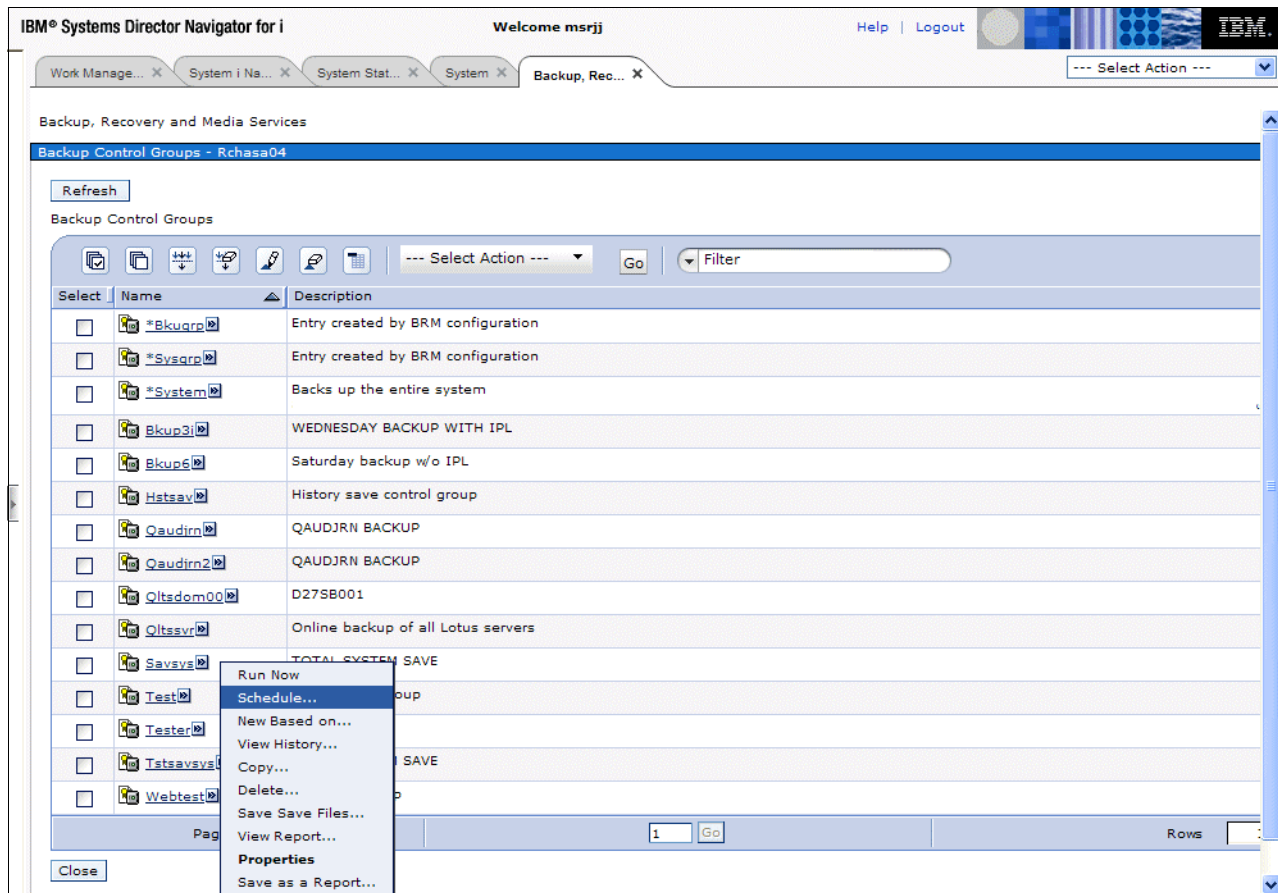


Figure 4-33 Selecting Schedule on a Backup Control Group

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

3. Select the Schedule action. The Schedule Information page is displayed, as shown in Figure 4-34.

The screenshot shows the 'Schedule Information' page for a task named 'QBRMBKUP'. The page 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' page has a title bar 'Schedule Information - Rchasa04'. The main content area includes fields for 'Task Name' (QBRMBKUP) and 'Task Description' (Run Control Group). Under the 'When to run:' section, the 'Once' radio button is selected. Below this, there are options for 'Absolute Day' and 'Relative Day' (with checkboxes for Last Week, First Week, Second Week, Third Week, Fourth Week, and Fifth Week). The 'Schedule day(s) or specific date to run' section has the 'Date to Start' radio button selected, with a date field set to '3/27/2010'. The 'Time to start:' field is set to '21:00:00 PM'. The 'Omit dates' section has an 'Omit Date' field set to '3/27/2010' and an 'Omit Dates' field set to '[Empty]'. At the bottom, there are 'OK' and 'Cancel' buttons.

Figure 4-34 Scheduling Information page

4. Schedule the backup control group to run once on March 27, 2010 at 21:00.

Scheduling BRMS maintenance

The following procedure discusses one way of scheduling BRMS maintenances.

1. From the BRMS advanced menu, click **Select Action** to display a drop down menu, as shown in Figure 4-35. From the drop down menu, click **Run Maintenance** and click **Go**.

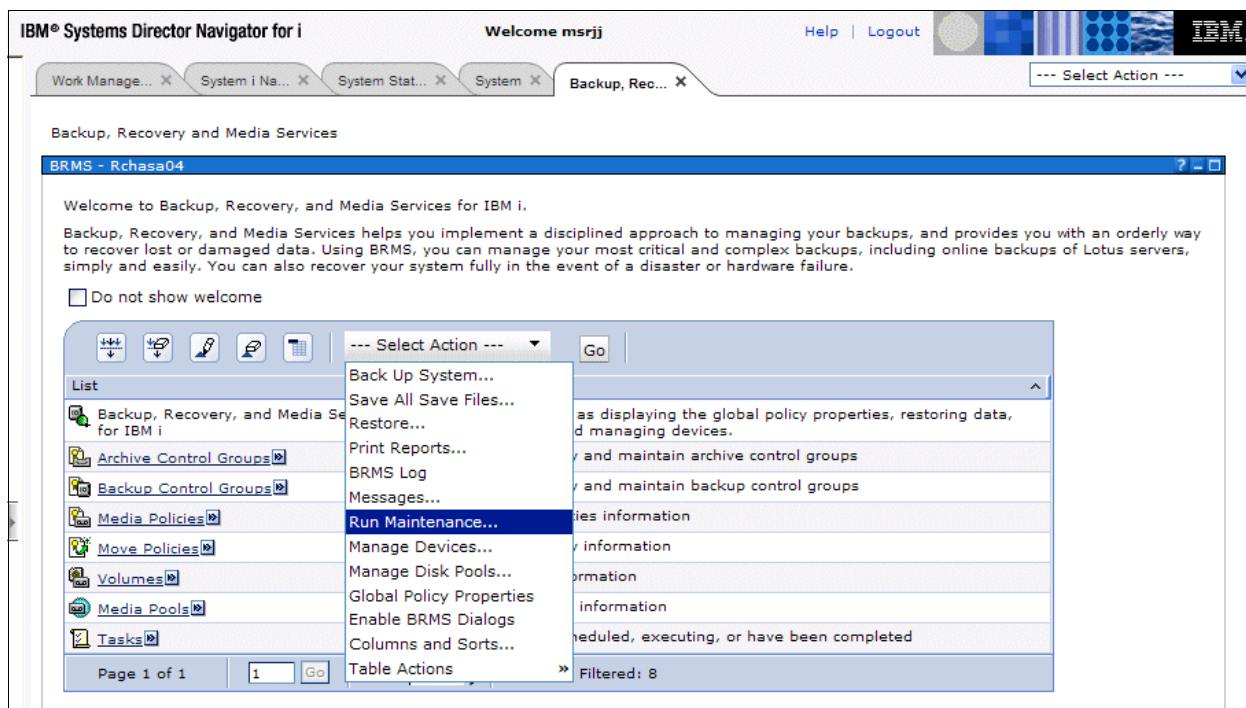


Figure 4-35 Selecting Run Maintenance from the Select Action drop-down menu

The Run Maintenance page (Figure 4-36) is displayed.

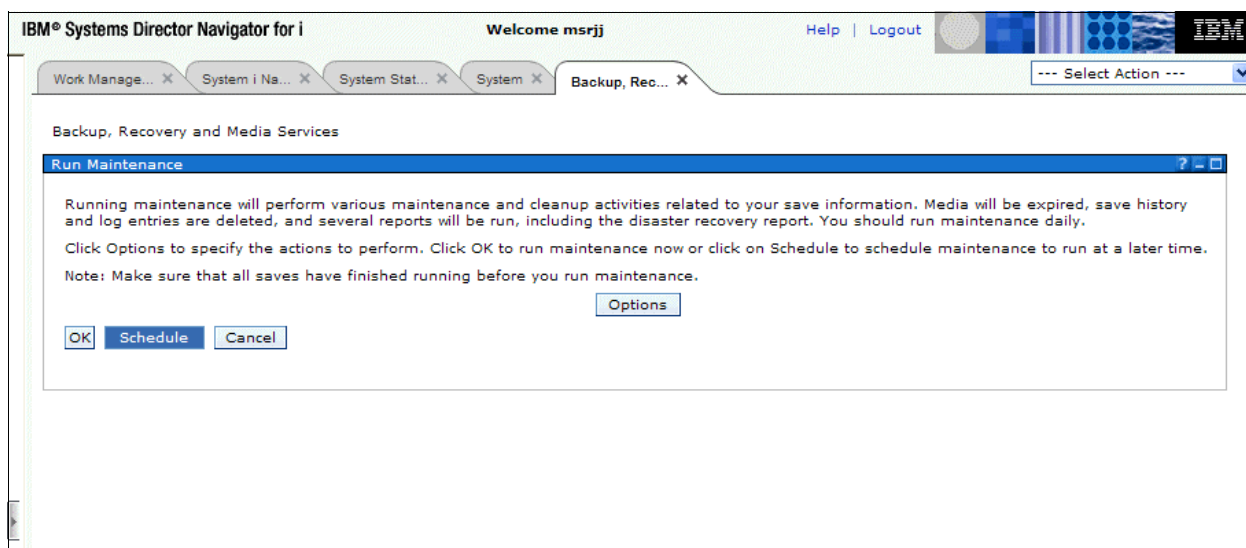


Figure 4-36 Run Maintenance page

2. Clicks **Options**. This displays the Run Maintenance Options window (Figure 4-37).




Figure 4-37 Run Maintenance Options: Media options page

The Run Maintenance Options window has three options to chose from:

- Media (shown in Figure 4-37)
- Print (shown in Figure 4-38 on page 68)
- File Management (shown in Figure 4-39 on page 68)

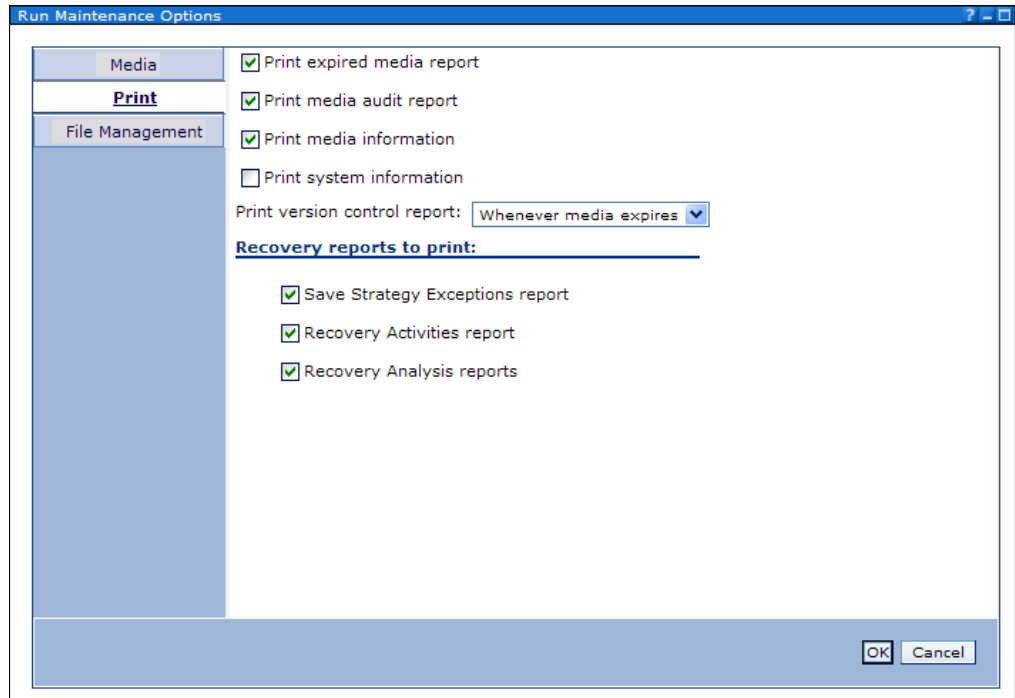


Figure 4-38 Run Maintenance Options: Print options page

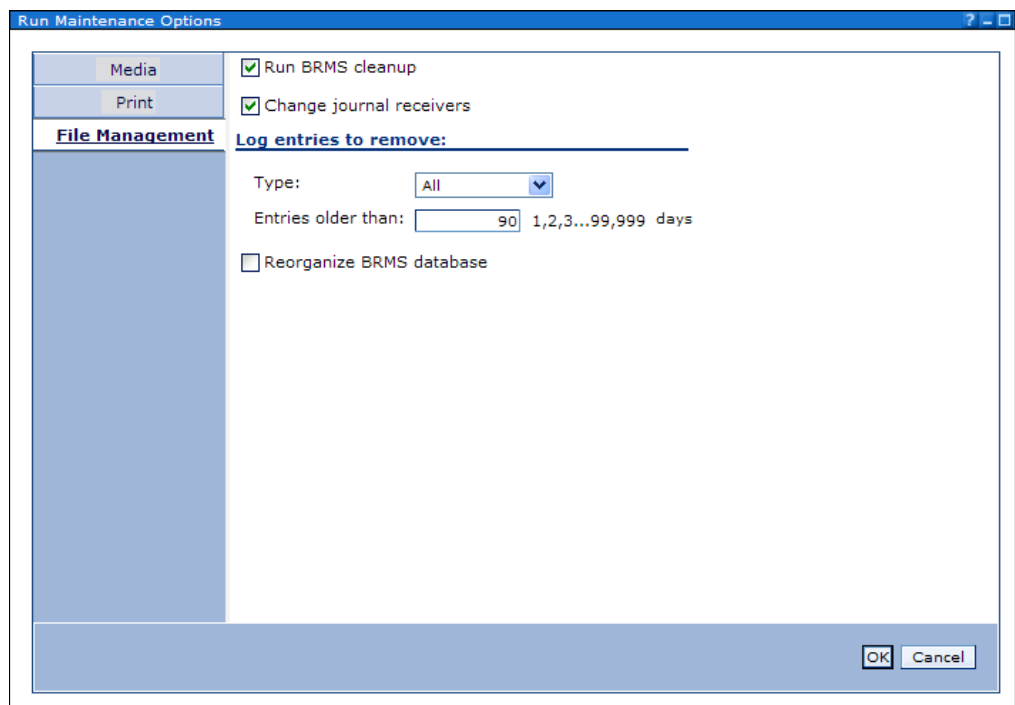


Figure 4-39 Run Maintenance Options: File Management options page

3. Select maintenance options as required
4. Click **OK** on any of the Run Maintenance options pages. The Run Maintenance page (Figure 4-36 on page 66) is displayed.

5. Click **Schedule**. The Schedule Information page (Figure 4-40) is displayed.

The screenshot shows the 'Schedule Information' page for a task named 'QBRMNT' with the description 'Run Maintenance Task'. The page 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 'When to run' section has 'Monthly' selected, with 'Relative Day' and 'Third Week' chosen. The 'Schedule day(s) or specific date to run' section has 'Days' selected, with 'Monday' through 'Friday' checked. The 'Time to start' is set to '4:45:25 PM'. The 'Omit dates' section shows '3/27/2010' as an omitted date. At the bottom are 'OK' and 'Cancel' buttons.

IBM® Systems Director Navigator for i

Welcome msrjj

Help | Logout

Work Manage... System i Na... System Stat... System Backup, Rec... Select Action

Schedule Information - Rchasa04

Task Name: QBRMNT

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

☐ Last Day of the Month

Time to start: 4:45:25 PM Example: 12:30:00 PM

Omit dates

Omit Date: 3/27/2010 Add

Omit Dates: [Empty] Remove

OK Cancel

Figure 4-40 Schedule Information page

This Schedule Information page looks identical and operates identically to the Schedule Information page seen when scheduling a backup control group. The differences are the Task name and Task Description field contents at the upper left of the page.

If the task has already been scheduled, the current values are shown. If the user makes no changes and clicks **OK**, the maintenance is not scheduled because it is already scheduled. An error page is displayed.

If the user changes the schedule, then clicks **OK**, a Run Maintenance confirmation page (Figure 4-41) is displayed, giving the user a chance to change any options, cancel the request, or accept the Run Maintenance as configured.

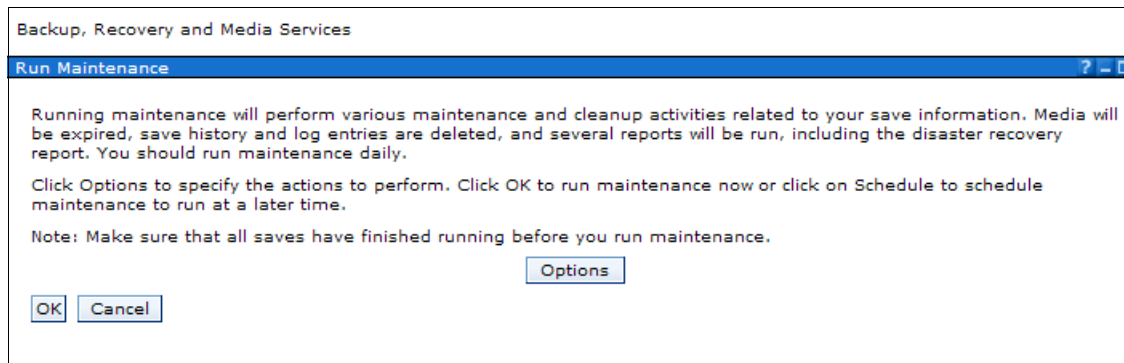


Figure 4-41 Run Maintenance Confirmation page

6. Click OK to complete scheduling maintenance.

Viewing scheduled tasks

Perform the following steps to view scheduled tasks:

1. Expand the pop-up menu for Tasks from the list column of the BRMS advanced menu.
2. From the pop-up menu, select Open, as shown in Figure 4-42.

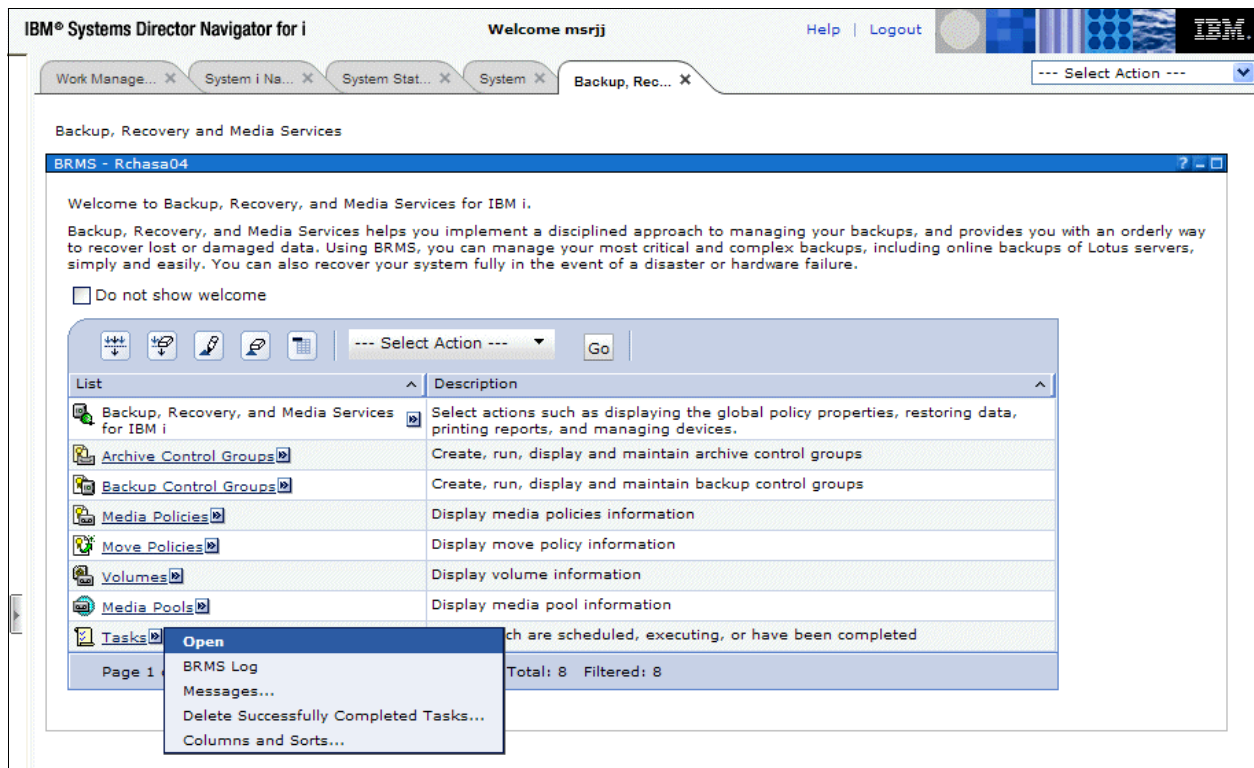


Figure 4-42 Task List pop-up menu with Open selected

Figure 4-43 shows the resulting page, a list of scheduled BRMS maintenance tasks.

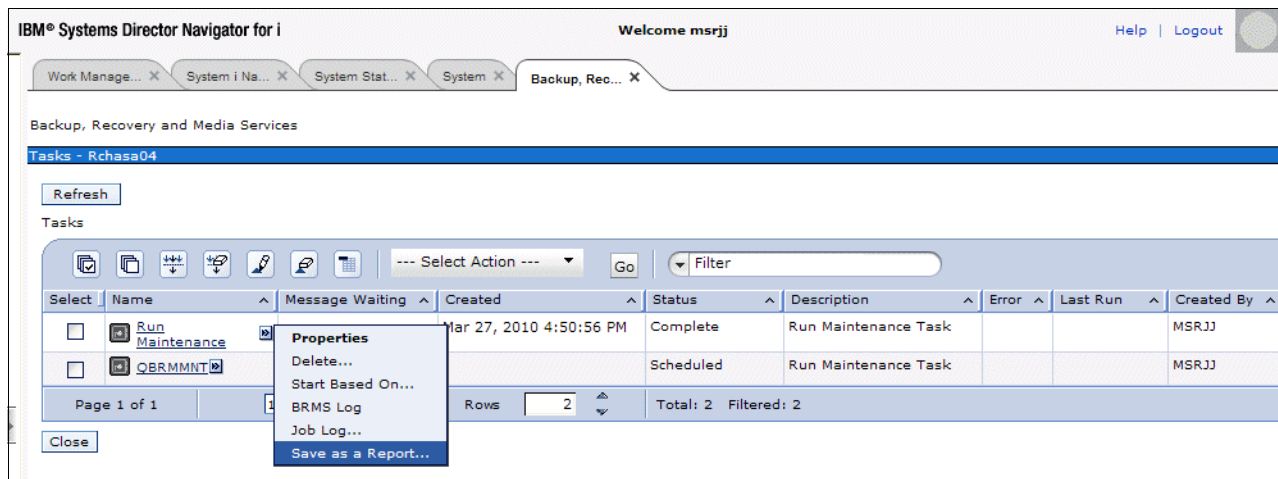


Figure 4-43 List of BRMS maintenance tasks

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 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 BRMS log to filter messages by control groups

In IBM i 7.1 the BRMS Log can now be filtered by control group. The user can filter by one control group at a time. Similar functionality is now available in the iSeries Navigator client.

The user can select BRMS Log from the BRMS initial menu shown in Figure 4-29 on page 60.

The user can also navigate to the Task list pop-up panel as shown in Figure 4-42 on page 70, except that instead of clicking **Open**, the user clicks **BRMS Log**.

Another way the user 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 4-44.

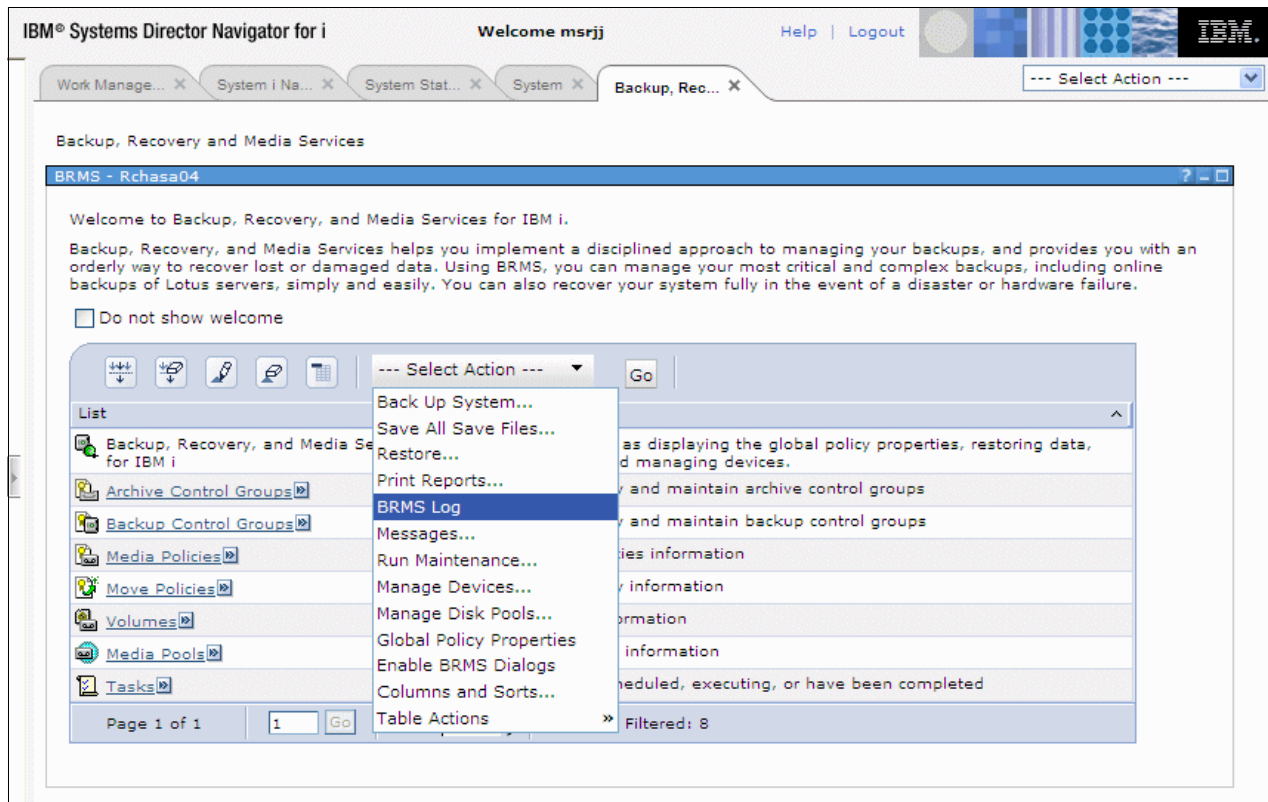


Figure 4-44 Selecting BRMS Log from the Select Action drop-down menu

Regardless of the method to access the BRMS log, the following BRMS Log - Include page is displayed in Figure 4-45.

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 4-45 New Control group selection of BRMS Log - Include page

The new control group selection parameters are shown in the circle. The **Browse** button displays a list of controls groups from which the user can make selections.

A partial view of the resulting control group for BRMS log listing follows in Figure 4-46.

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 ID
<input type="checkbox"/>	BRM1096	0	Backup	Job queue BRMMMAINT is released.	Feb 28, 2010	11:49:01 AM	Ms773	Secbkppaud	33
<input type="checkbox"/>	BRM1049	0	Backup	Control group QAUDJRN type *BKU processing is complete.	Feb 28, 2010	11:49:01 AM	Ms773	Secbkppaud	33
<input type="checkbox"/>	BRM1587	10	Backup	Save of BRM media information at level *OBJ complete.	Feb 28, 2010	11:49:00 AM	Ms773	Secbkppaud	33
<input type="checkbox"/>	CPC3701	20	Backup	19 objects saved from library QUSRBRM.	Feb 28, 2010	11:48:59 AM	Ms773	Secbkppaud	33
<input type="checkbox"/>	CPI6705	10	Backup	7485 blocks processed for sequence 333, volume LT0026, on device MLD50.	Feb 28, 2010	11:48:54 AM	Ms773	Secbkppaud	33
<input type="checkbox"/>	BRM1647	10	Backup	Starting save of media information at level *OBJ to device MLD50.	Feb 28, 2010	11:48:12 AM	Ms773	Secbkppaud	33
<input type="checkbox"/>	BRM4201	10	Backup	Change to media controls successful.	Feb 28, 2010	11:48:12 AM	Ms773	Secbkppaud	33
<input type="checkbox"/>	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 4-46 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 discussed.

In Figure 4-47, the Volume list pop-up menu is displayed with the Open option specified.

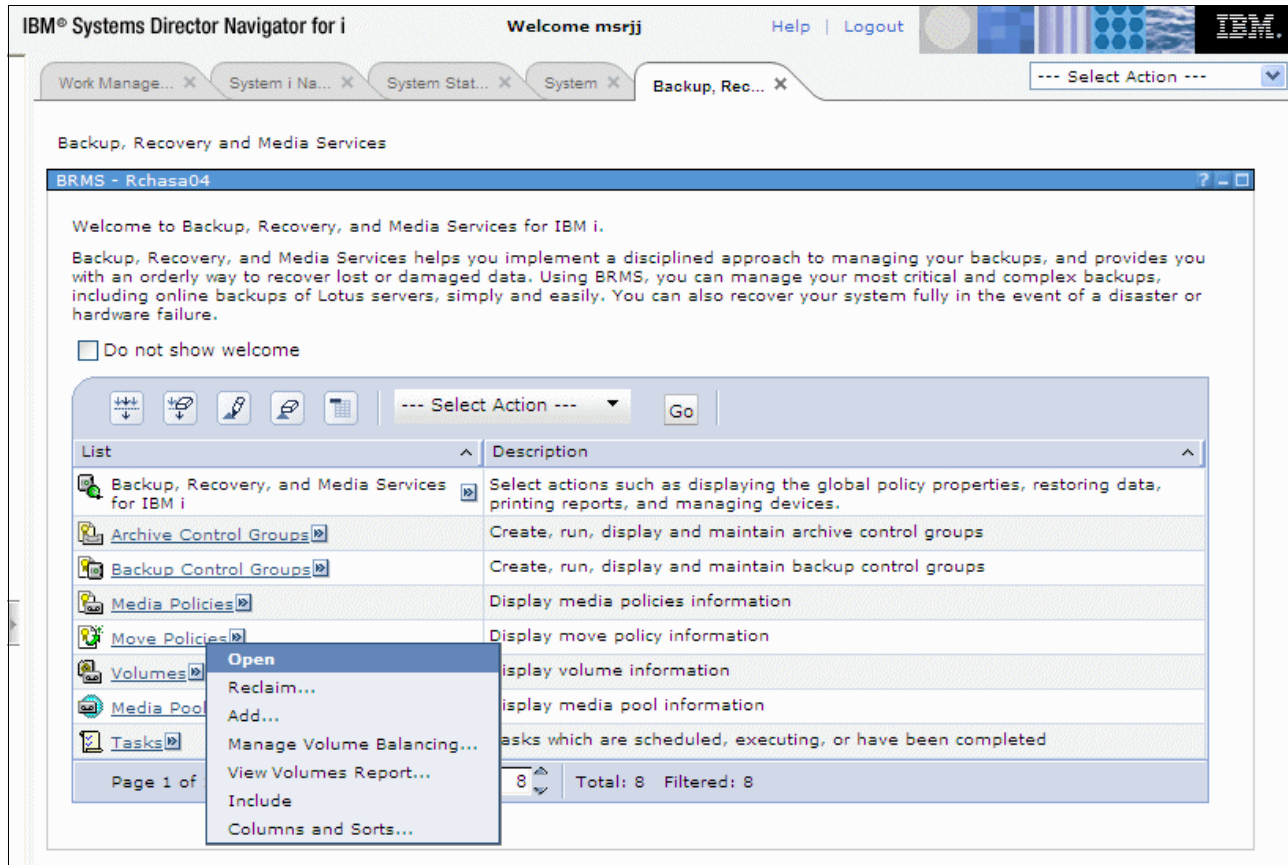


Figure 4-47 Opening the volume list

When the user clicks **Open**, the volumes table is displayed. From the volume table, the user opens the pop-up menu for a specific volume, as shown in Figure 4-48.

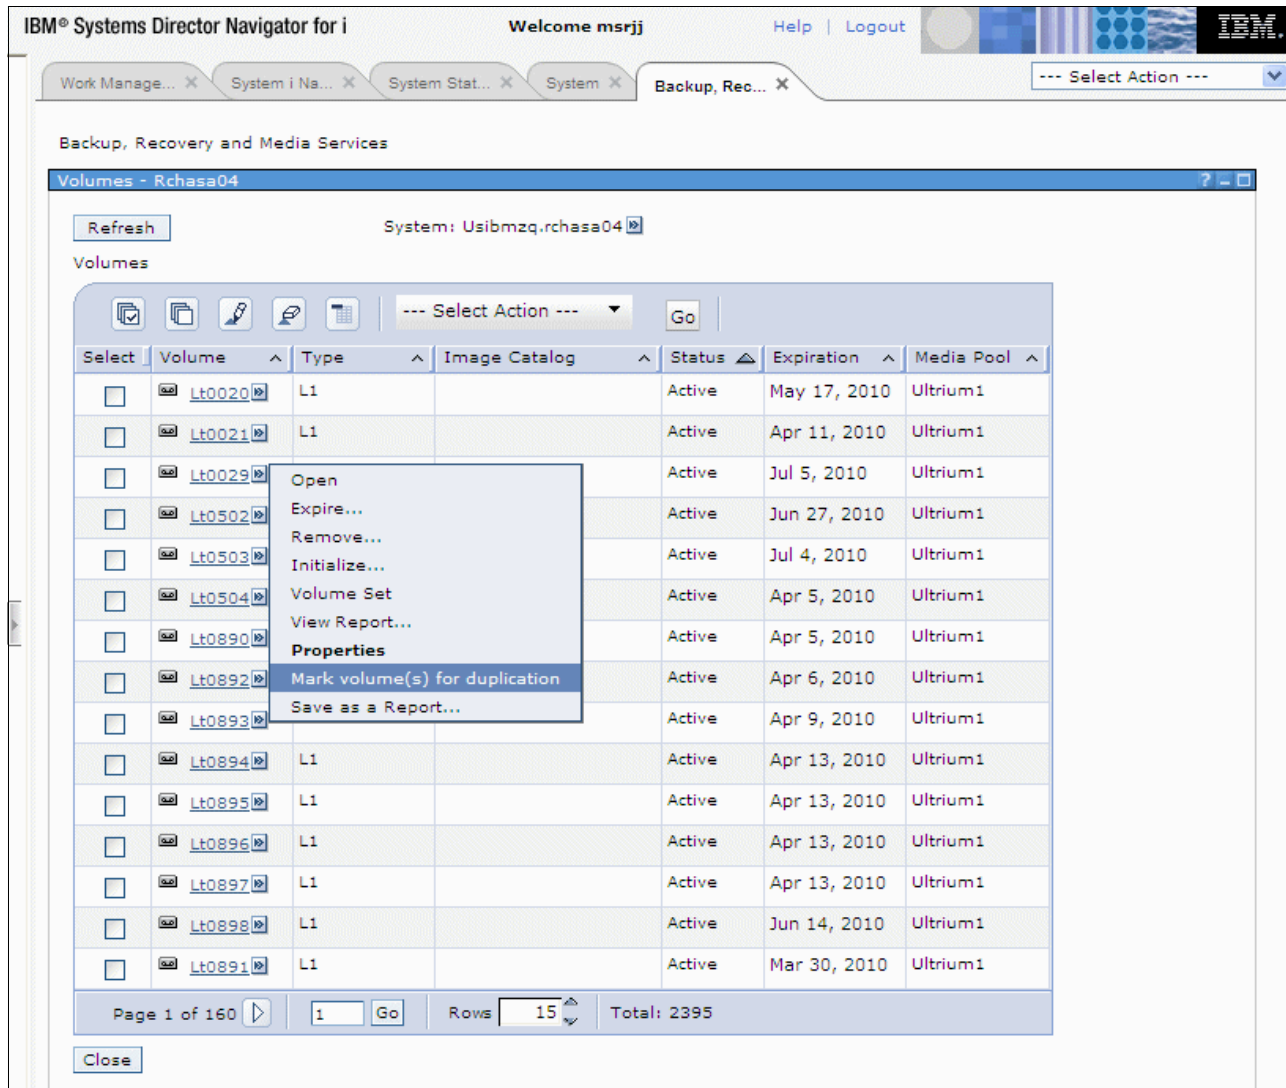


Figure 4-48 Volumes table with pop-up menu displayed for volume LT0029

The user can tell that the volume is not marked for duplication because the “Mark volume(s) for duplication” option exists in the menu. The user selects the mark option, and the Confirm Volume Action page (Figure 4-49) is displayed.

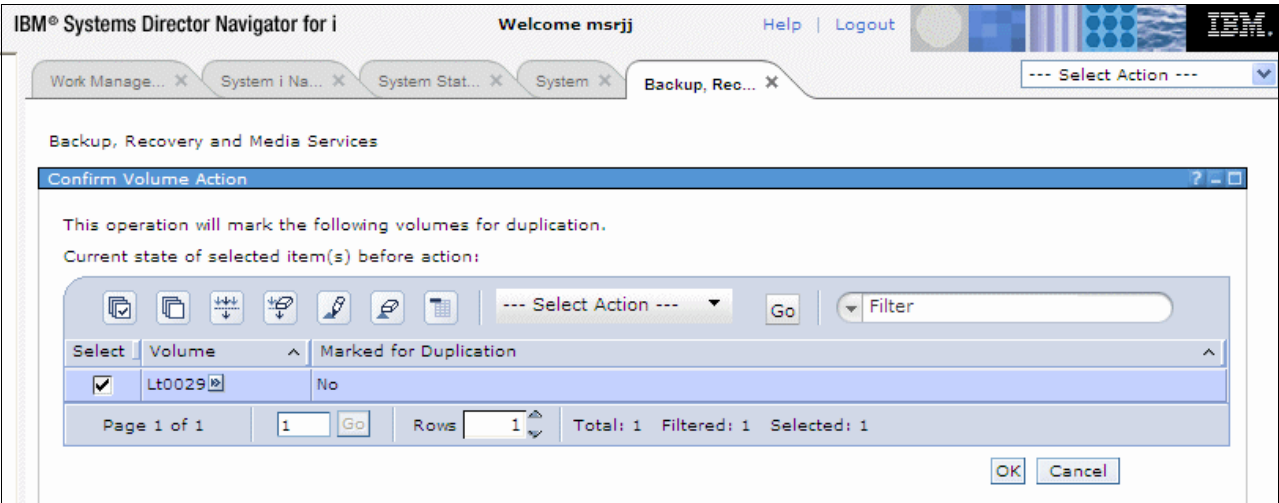


Figure 4-49 Confirm Volume Action page

The user clicks the **OK** button and the volume is marked duplication and the Volumes page is displayed again.

But the user cannot tell from the Volumes table page whether the volume is marked. One way to verify the volume is marked is to select the volume, open its pop-up menu, and look at its contents, as in Figure 4-50.

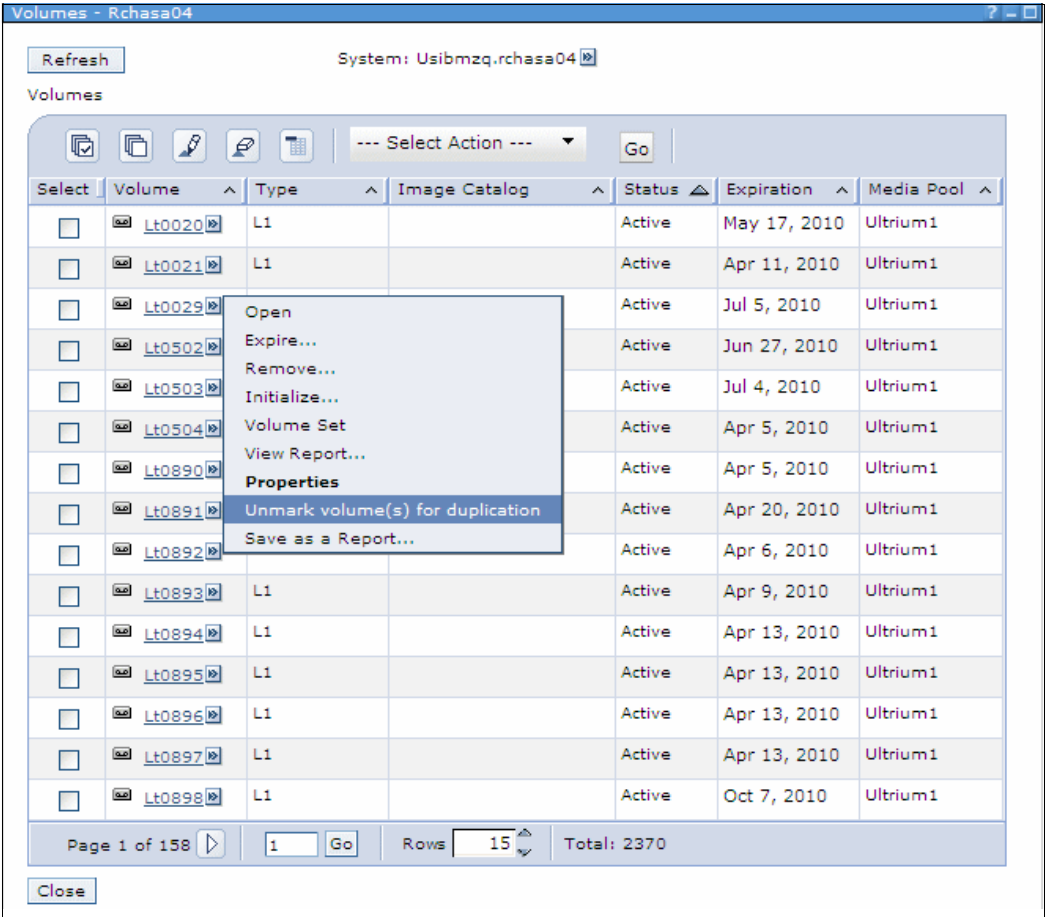


Figure 4-50 Volume pop-up menu showing Unmark volume(s) for duplication

Because the Unmark volume(s) for duplication option is shown, the user knows the volume is marked for duplication.

If the user wants to unmark the volume, clicking the unmark option will do so.

Suppose the user does a lot of marking and unmarking for duplication. The user can modify the volume table and add the “Marked for Duplication column” to the table view. In the following scenario, the user removes the “Image Catalog” column to keep the table narrow enough to fit on the panel.

To do this, open the Select Action drop-down menu as shown in Figure 4-51 and select the Columns and Sorts action.

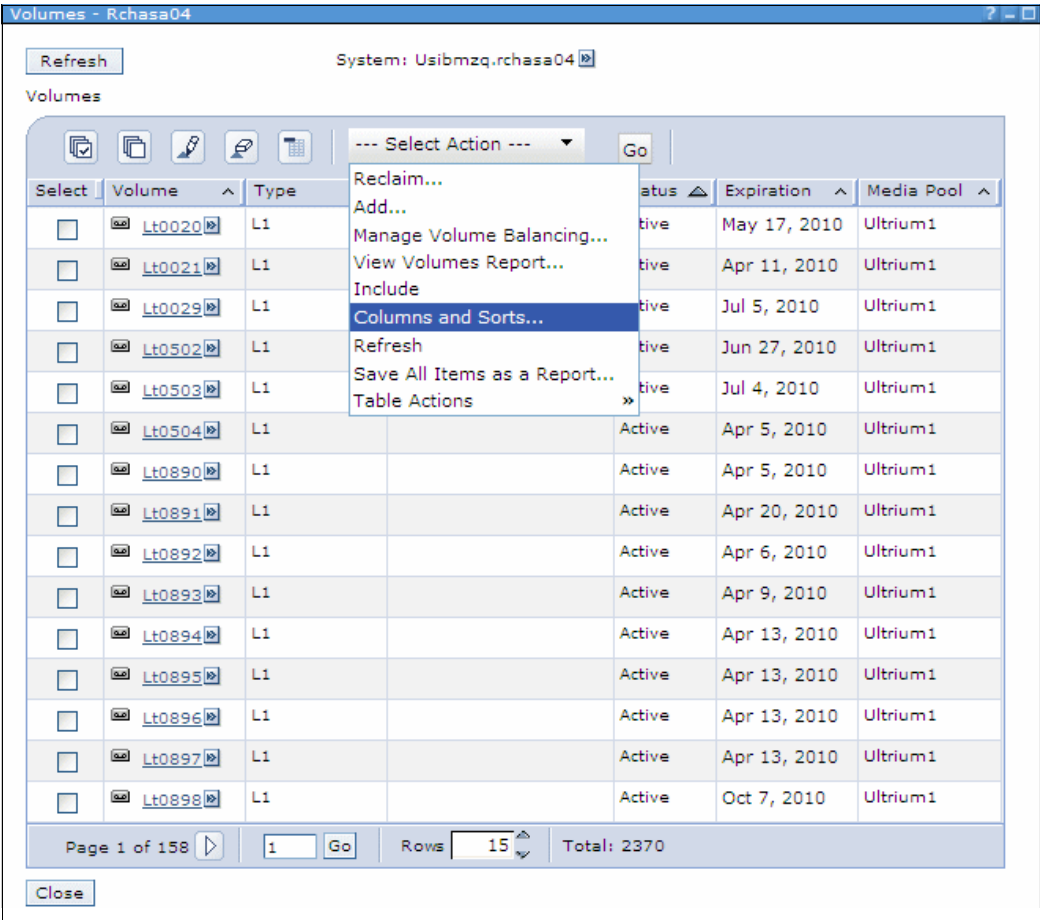


Figure 4-51 Select Action drop-down which Columns and Sorts action

When Columns and Sorts is selected, the Set Columns to Display and Sort page is displayed. Figure 4-52 displays a “Show Column” column in which the columns are selected for display. The two arrows point to columns the user wants to change.

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
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Volume		<input checked="" type="checkbox"/>
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Type		<input checked="" type="checkbox"/>
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Image Catalog		<input checked="" type="checkbox"/>
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Status	1	<input checked="" type="checkbox"/>
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Expiration		<input checked="" type="checkbox"/>
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Media Pool		<input checked="" type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	Container		<input checked="" type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	Date Added or Reused		<input checked="" type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	Description		<input checked="" type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	Duplicated Volume		<input checked="" type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	Expired		<input checked="" type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	Last Reused		<input checked="" type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	Location		<input checked="" type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	Marked for Duplication		<input checked="" type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	Moved Date		<input checked="" type="checkbox"/>

Page 1 of 2

1

Go

Rows

15

Total: 24

OK

Cancel

Figure 4-52 Set Columns to Display and Sort

The user does not want the Image Catalog column to display, so the user will clear that column. The user wants the Marked for Duplication column to display, so the user selects that column for display. When the user clicks the **OK** button, the volume display in Figure 4-53 is presented.

Refresh System: Usibmq.rchasa04

Volumes

--- Select Action --- Go

Select	Volume	Type	Status	Expiration	Media Pool	Marked for Duplication
<input type="checkbox"/>	Lt0020	L1	Active	May 17, 2010	Ultrium1	No
<input type="checkbox"/>	Lt0021	L1	Active	Apr 11, 2010	Ultrium1	No
<input type="checkbox"/>	Lt0029	L1	Active	Jul 5, 2010	Ultrium1	No
<input type="checkbox"/>	Lt0502	L1	Active	Jun 27, 2010	Ultrium1	No
<input type="checkbox"/>	Lt0503	L1	Active	Jul 4, 2010	Ultrium1	No
<input type="checkbox"/>	Lt0504	L1	Active	Apr 5, 2010	Ultrium1	No
<input type="checkbox"/>	Lt0890	L1	Active	Apr 5, 2010	Ultrium1	No
<input type="checkbox"/>	Lt0891	L1	Active	Apr 20, 2010	Ultrium1	No
<input type="checkbox"/>	Lt0892	L1	Active	Apr 6, 2010	Ultrium1	No
<input type="checkbox"/>	Lt0893	L1	Active	Apr 9, 2010	Ultrium1	No
<input type="checkbox"/>	Lt0894	L1	Active	Apr 13, 2010	Ultrium1	No
<input type="checkbox"/>	Lt0895	L1	Active	Apr 13, 2010	Ultrium1	No
<input type="checkbox"/>	Lt0896	L1	Active	Apr 13, 2010	Ultrium1	No
<input type="checkbox"/>	Lt0897	L1	Active	Apr 13, 2010	Ultrium1	No
<input type="checkbox"/>	Lt0898	L1	Active	Oct 7, 2010	Ultrium1	No

Page 1 of 158 1 Go Rows 15 Total: 2370

Close

Figure 4-53 Volume table with column changes

The Image Catalog column has been removed and the Marked for Duplication column is shown. The user can now see the Marked for Duplication status of each volume without having to select 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, GUI and IBM Systems Director and IBM Systems Director Navigator for i interfaces.

To configure this, access the Global Policy Properties Action from the BRMS advanced menu, as shown in Figure 4-54.

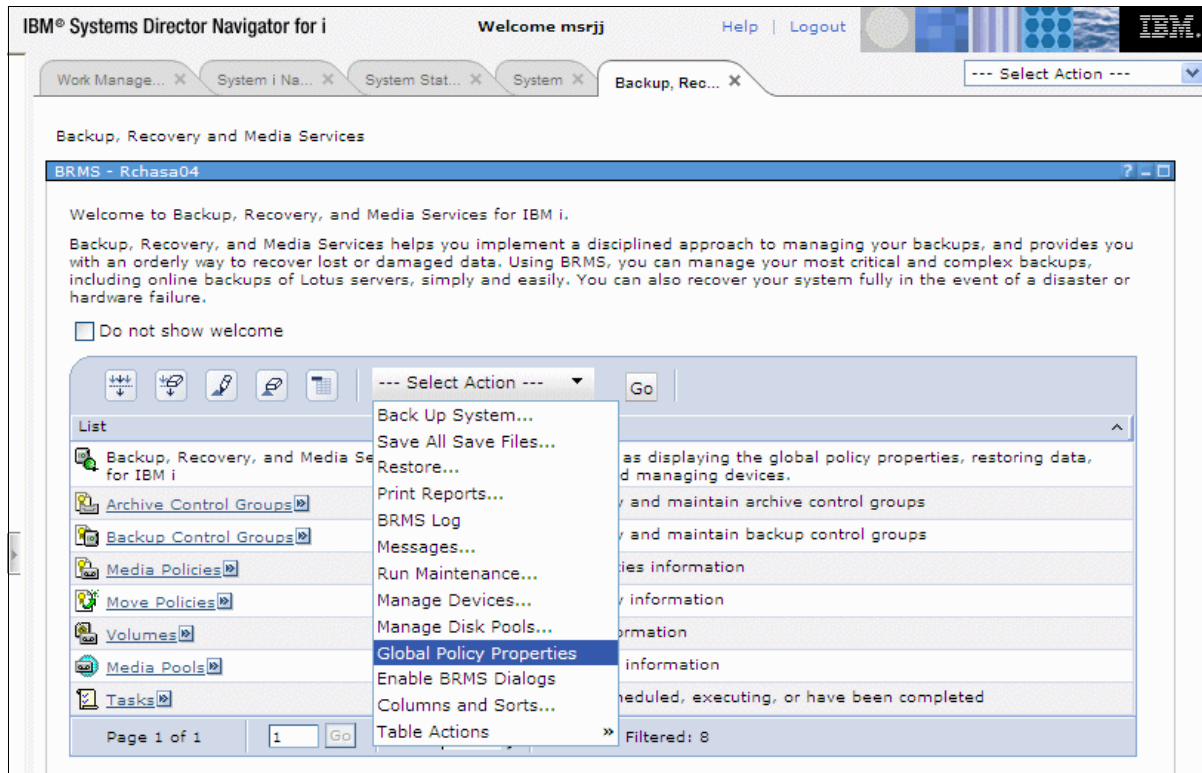


Figure 4-54 Selecting Global Policy Properties

When selected, the Global Policy Properties page shown in Figure 4-55 is presented.

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 4-55 Global Policy Properties page

The user selects the Network Properties page, shown in Figure 4-56.

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 ---

Backup, Recovery and Media Services

Global Policy Properties - Rchasa04

General

Notification

Signoff Exceptions

Power Down

Maintenance

Retrieval

Network

Logging

System update interval (30-9999):

Offline notification period (30-99,999):

☒ Use TCP/IP for networking

☒ Use SNA for networking

☐ Receive history information

Distribution:

Primary system: None

Secondary system: None

E-mail address: user@yourco.com, user2@yourco.com

Maximum message length: 5000 0,1,2...5,000

Network restricted state interfaces to start:

--- Select Action ---

Interface Address	Type
None	

1 Go Rows 0 Total: 0 Filtered: 0

BRMS network systems:

--- Select Action ---

System	System Status	Network Status	Network Identifier	APPC Name	TCP/IP Name	Relation
None						

Figure 4-56 E-mail address field with multiple entries on Network policy properties page

In the email address field, multiple email addresses can now be entered.

4.3.10 High availability support for independent ASPs in a BRMS network

This function provides the ability to run 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, IBM Systems Director web environment or System i Navigator.

To configure this function, perform the following steps:

1. Select **Backup, Recovery, and Media Services** → **Global Policy Properties**. See Figure 4-54 on page 82.

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
Distribution:
Primary system: None [Browse...]
Secondary system: None [Browse...]
E-mail address: [Text Field]
Maximum message length: 4999 0,1,2...5,000

Logging

Network restricted state interfaces to start:

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:

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 4-57 Global Policy Properties - Network Properties

2. On the Global Policy Properties panel, select the **Network** tab.
3. Click the **Manage Disk Pool History** button to which the arrow in Figure 4-57 points. The BRMS graphical interface refers to IASPs as disk pools.

- On the Manage Disk Pool History to Send panel, click **List actions** from the menu bar and select **New**. See Figure 4-58




Figure 4-58 Send New Disk Pool History

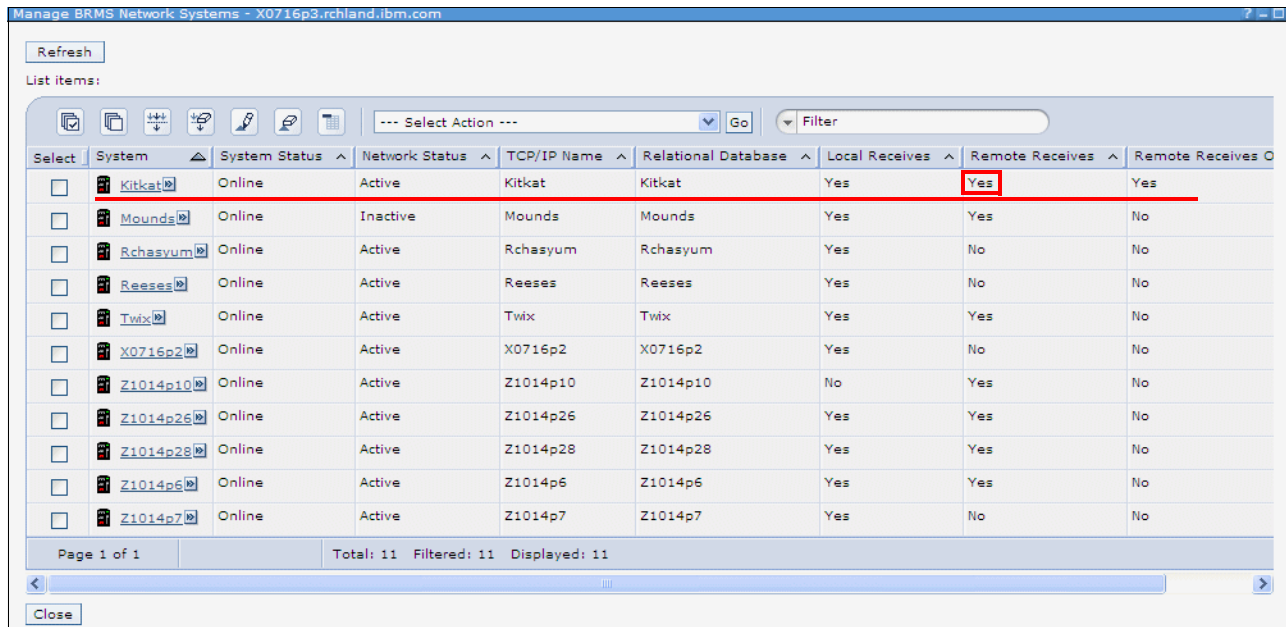
- Enter in the disk pool name, the system name, and network identifier of the system to which you want the history information to be sent.

Decide whether to use the BRMS media information remote receives value that is defined by the "Send to system". You can override that systems value by selecting the "Always send" option.



Figure 4-59 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, clicking **Network properties**, and clicking the **Manage Systems** button. The value is listed under the Remote Receives column for that remote system, as shown in the circle on Figure 4-60.



Manage BRMS Network Systems - X0716p3.rchland.ibm.com

Refresh

List items:

--- Select Action --- Go Filter

Select	System	System Status	Network Status	TCP/IP Name	Relational Database	Local Receives	Remote Receives	Remote Receives O
<input type="checkbox"/>	Kitkat	Online	Active	Kitkat	Kitkat	Yes	Yes	Yes
<input type="checkbox"/>	Mounds	Online	Inactive	Mounds	Mounds	Yes	Yes	No
<input type="checkbox"/>	Rchasyum	Online	Active	Rchasyum	Rchasyum	Yes	No	No
<input type="checkbox"/>	Reeses	Online	Active	Reeses	Reeses	Yes	No	No
<input type="checkbox"/>	Twix	Online	Active	Twix	Twix	Yes	Yes	No
<input type="checkbox"/>	X0716p2	Online	Active	X0716p2	X0716p2	Yes	No	No
<input type="checkbox"/>	Z1014p10	Online	Active	Z1014p10	Z1014p10	No	Yes	No
<input type="checkbox"/>	Z1014p26	Online	Active	Z1014p26	Z1014p26	Yes	Yes	No
<input type="checkbox"/>	Z1014p28	Online	Active	Z1014p28	Z1014p28	Yes	Yes	No
<input type="checkbox"/>	Z1014p6	Online	Active	Z1014p6	Z1014p6	Yes	Yes	No
<input type="checkbox"/>	Z1014p7	Online	Active	Z1014p7	Z1014p7	Yes	No	No

Page 1 of 1 Total: 11 Filtered: 11 Displayed: 11

Close

Figure 4-60 Checking a System's Remote Receives value

- Return to the Send Disk Pool History addition of a new disk pool and click **OK** to complete the addition.

4.3.11 Enhanced maintenance features in backup policy

BRMS maintenance features in the backup policy now allow the users to run movement, expire partial volume sets, and reorganize BRMS database files.

Figure 4-61 shows the Global Policy Properties - Backup Maintenance Options page. On the page the new Run move policies and Expire partial volume sets are circled.

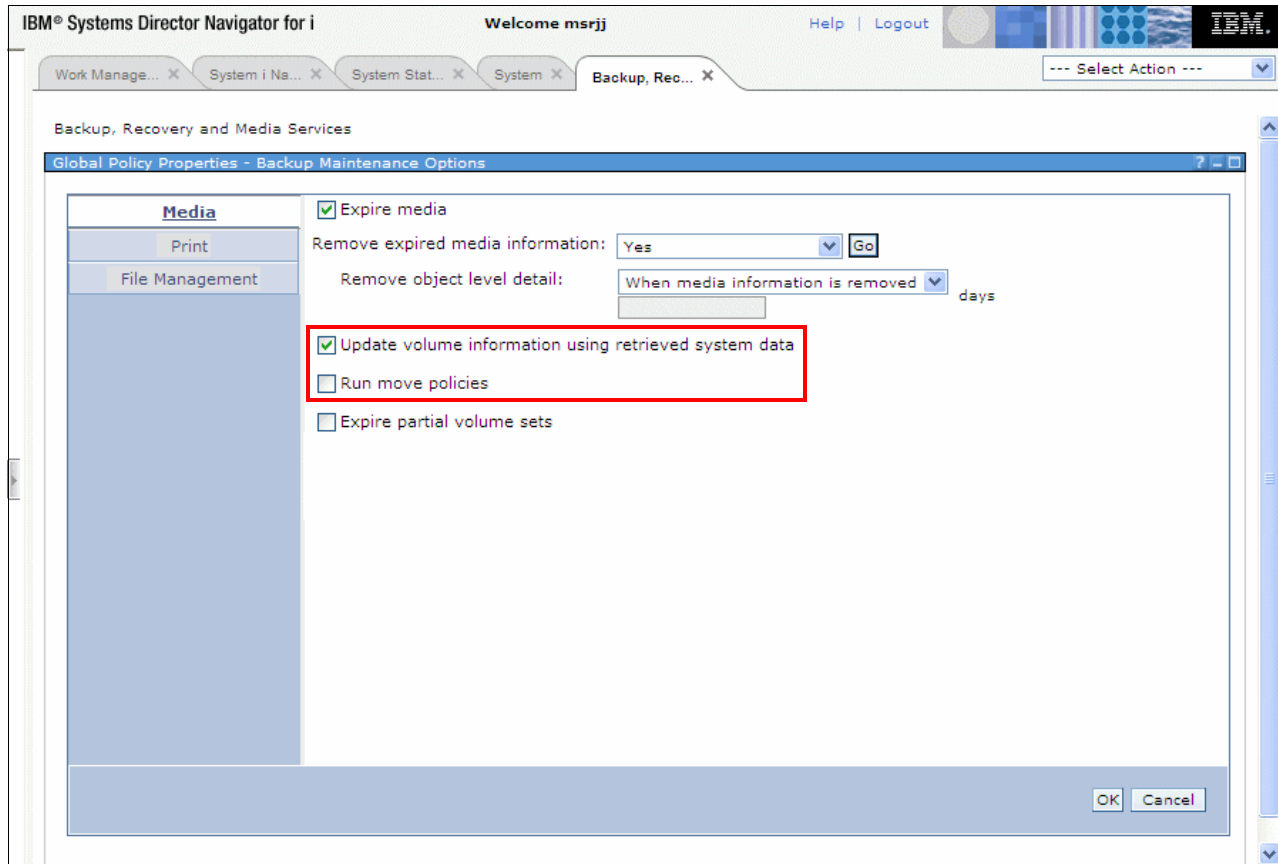


Figure 4-61 Backup policy properties -Run move policy and Expire partial volume sets options

By selecting the File Management page, the new Reorganize BRMS database option is available, as shown in the red circle of Figure 4-62.

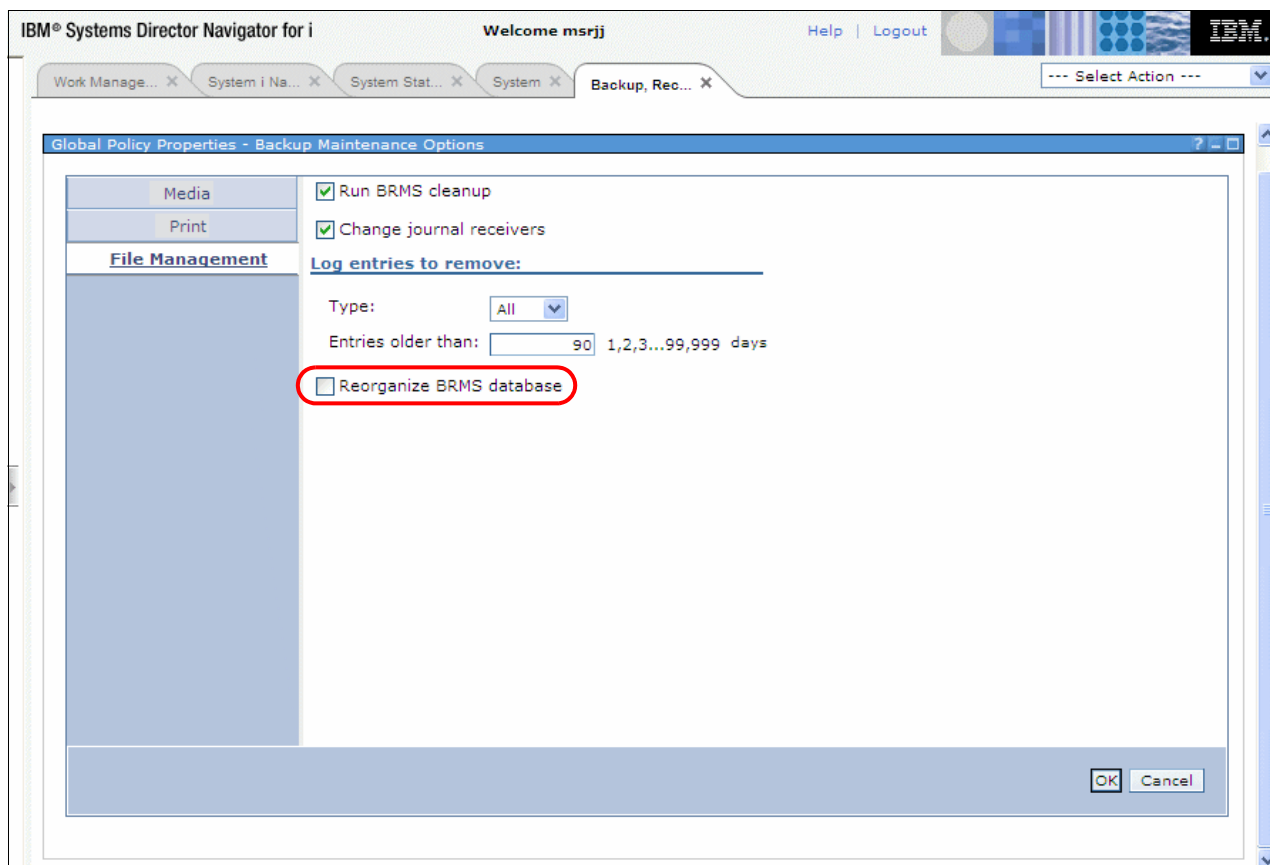


Figure 4-62 Backup control group Reorganize BRMS database option

4.4 Additional references

- ▶ 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
<http://publib.boulder.ibm.com/iserics/v7r1m0/topic/brms/sc415345.pdf>
- ▶ BRMS for IBM i web site
<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:

- ▶ Section 5.1, “PowerHA SystemMirror for i” on page 92
- ▶ Section 5.2, “Journaling and commitment control enhancements” on page 109

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:

- ▶ 5.1.1, “New PowerHA Packaging” on page 92
- ▶ 5.1.2, “PowerHA versioning” on page 93
- ▶ 5.1.3, “PowerHA SystemMirror for i Enhancements” on page 93
- ▶ 5.1.4, “PowerHA SystemMirror for i - Graphical interfaces” on page 96
- ▶ 5.1.5, “Duplicate Library Error Handling” on page 100
- ▶ 5.1.6, “N_Port ID virtualization support” on page 100
- ▶ 5.1.7, “Asynchronous geographic mirroring” on page 100
- ▶ 5.1.8, “LUN level switching” on page 103
- ▶ 5.1.9, “Space-Efficient FlashCopy” on page 104
- ▶ 5.1.10, “Better detection of cluster node outages” on page 105
- ▶ 5.1.11, “Improved geographic mirroring full synchronization performance” on page 107
- ▶ 5.1.12, “Cluster administrative domain enhancements” on page 107
- ▶ 5.1.13, “IBM HA Assist for i” on page 107
- ▶ 5.1.14, “IPv6 Support” on page 108
- ▶ 5.1.15, “New CL commands for programming cluster automation” on page 108

5.1.1 New PowerHA Packaging

IBM PowerHA for i has been 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 datacenter 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 the 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 on page 93.

PowerHA SystemMirror for i	Standard Edition	Enterprise Edition
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 Flash Copy	✓	✓
LUN level switching	✓	✓
GeoMirror sync delivery	✓	✓
GeoMirror async delivery		✓
Multi-Site HA/DR management		✓
DS8000/DS6000 Metro Mirror		✓
DS8000/DS6000 Global Mirror		✓

Figure 5-1 PowerHA SystemMirror for i editions

5.1.2 PowerHA versioning

To use any of the new PowerHA SystemMirror for i enhancements all nodes in the cluster need to be upgraded to IBM i 7.1. Prior to this both the cluster version and the PowerHA version need to be updated to the current cluster version 7 and PowerHA version 2 through 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 above command twice. As such a V5R4M0 system within a clustered environment can be upgraded towards i 7.1 by skipping i 6.1.

5.1.3 PowerHA SystemMirror for i Enhancements

The following new functions are delivered with the October 2011 announcement for PowerHA SystemMirror for i with 5770-HAS, PTF SI44148 and the new 5799-HAS Program request Pricing Quotation (PRPQ):

- Support for managing IBM System Storage SAN Volume Controller (SVC) and IBM Storwize V7000 Copy Services in IBM PowerHA SystemMirror for i.
- IBM i command line command 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.
- New PowerHA GUI

The PRPQ has the following characteristics:

- ▶ Product ID: 5799-HAS
- ▶ Product Name: PowerHA SystemMirror for i Enhancements.
- ▶ PowerHA PTF required: SI44148 for 5770-HAS.
- ▶ Language: English (2924) only.
- ▶ Secondary language support is not needed.

In the following we provide a brief overview of these new enhancements. For more detailed information please refer to the latest redbook IBM PowerHA SystemMirror for IBM i Cookbook SG24-7994.

Support for SVC / 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 those we use for DS8000 Copy Services, but some parameters are different:

- ▶ Add SVC ASP Copy Description (ADD**SVCCPYD**): this command is used to describe a single physical copy of an auxiliary storage pool (ASP) that exists within a SVC and to assign a name to the description.
- ▶ Change SVC Copy Description (CHG**SVCCPYD**): this command changes an existing auxiliary storage pool (ASP) copy description.
- ▶ Remove SVC Copy Description (RMV**SVCCPYD**): this command is used to remove an existing ASP copy description. It does not remove the disk configuration.
- ▶ Display SVC Copy Description (DSP**SVCCPYD**): this command displays an ASP copy description.
- ▶ Work with ASP Copy Description (WRKASPCPYD) shows both DS8000 and SVC / V7000 copy descriptions.
- ▶ Start SVC Session (STR**SVCSSN**): this command assigns a name to the Metro Mirror, Global Mirror or FlashCopy session which links the two ASP copy descriptions for the source and target IASP volumes and starts an ASP session for them.
- ▶ Change SVC Session (CHG**SVCSSN**): this command is used to change an existing Metro Mirror, Global Mirror or FlashCopy session.
- ▶ End SVC ASP Session (END**SVCSSN**): this command ends an existing ASP session.
- ▶ Display SVC Session (DSP**SVCSSN**): this command displays an ASP session.

Note: The new PowerHA GUI currently does not support the SVC / V7000 Copy Services.

Configure Device ASP (CFGDEVASP) command

The new Configure Device ASP (CFGDEVASP) command is part of the 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 using with the *CREATE action, it:

- ▶ 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 using with the *DELETE action, it:

- Deletes the IASP.
- Deletes the ASP device description if it was created by this command.

Please refer to Figure 5-2 for more information on creating/deleting an IASP by using the command line interface.

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 shown in Figure 5-3 on page 96 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 standalone IASP and create a geographic mirror copy.

In order to get this command, the 5770-HAS PTF SI44148 is needed on your system running 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

5.1.4 PowerHA SystemMirror for i - Graphical interfaces

With IBM i 7.1, we currently have two different Graphical User Interfaces within IBM Systems Director Navigator for i:

- Cluster Resource Services GUI
- High Availability Solutions Manager GUI

Cluster Resource Services GUI

You can get to it by performing the following as shown in Figure 5-5:

1. Expand **IBM i Management**.
2. Select **Cluster Resource Services**

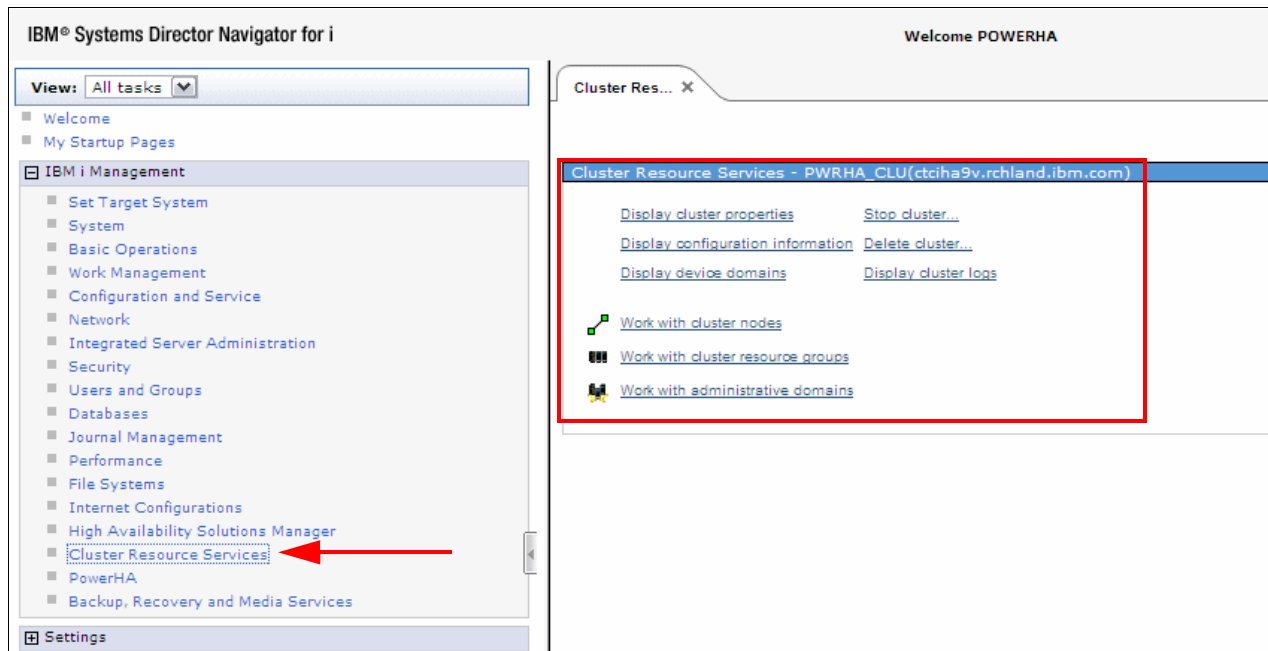


Figure 5-4 Cluster Resource Services GUI

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

Note: The clustering GUI plug-in for System i Navigator from High Availability Switchable Resources licensed program (IBM i option 41) has been removed in IBM i 7.1. The Clustering HA environments can continue to be configured and managed using the PowerHA for i licensed product (5770-HAS) CL commands.

High Availability Solutions Manager GUI

You can get to it by performing the following as shown in Figure 5-5:

1. Expand **IBM i Management**.
2. Select **High Availability Solutions Manager**

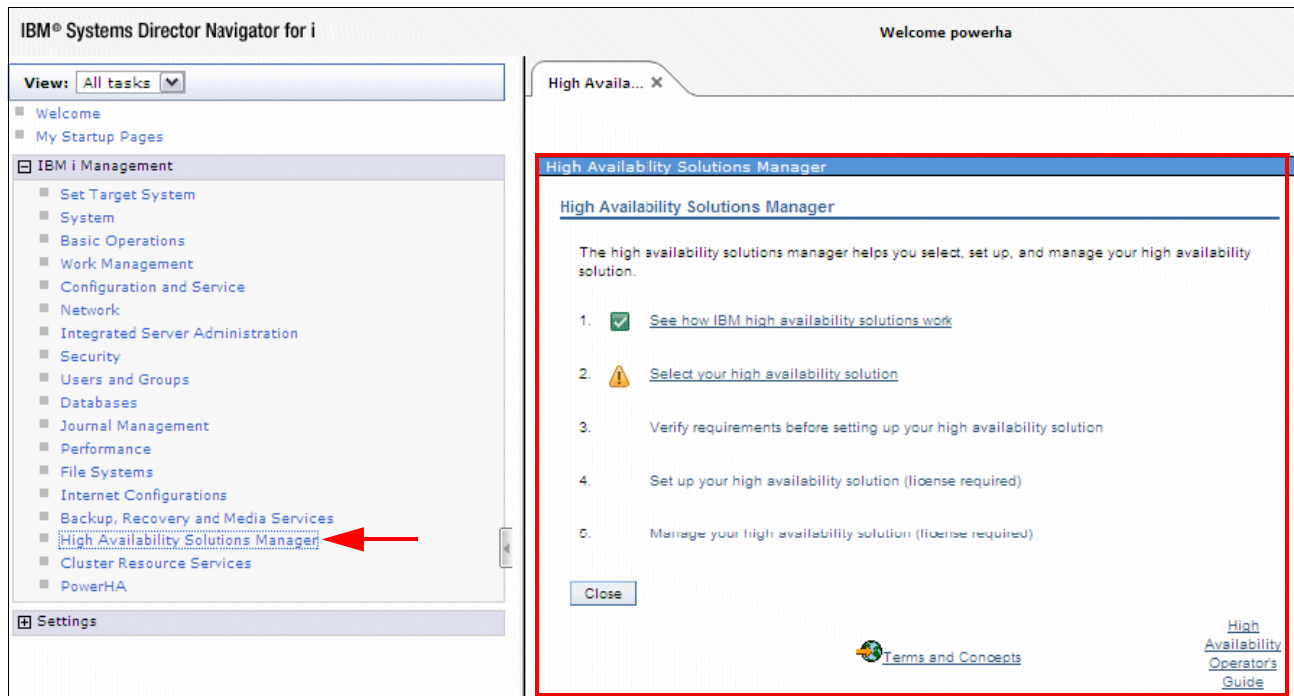


Figure 5-5 High Availability Solutions Manager GUI

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

With the October 2011 PRPQ announcement of IBM HA System Mirror for i, there is a completely new interface available for PowerHA within IBM Systems Director Navigator for i.

You can get to the new GUI by performing the following as shown in Figure 5-6 on page 99:

1. Expand **IBM i Management**.
2. Select **PowerHA**.

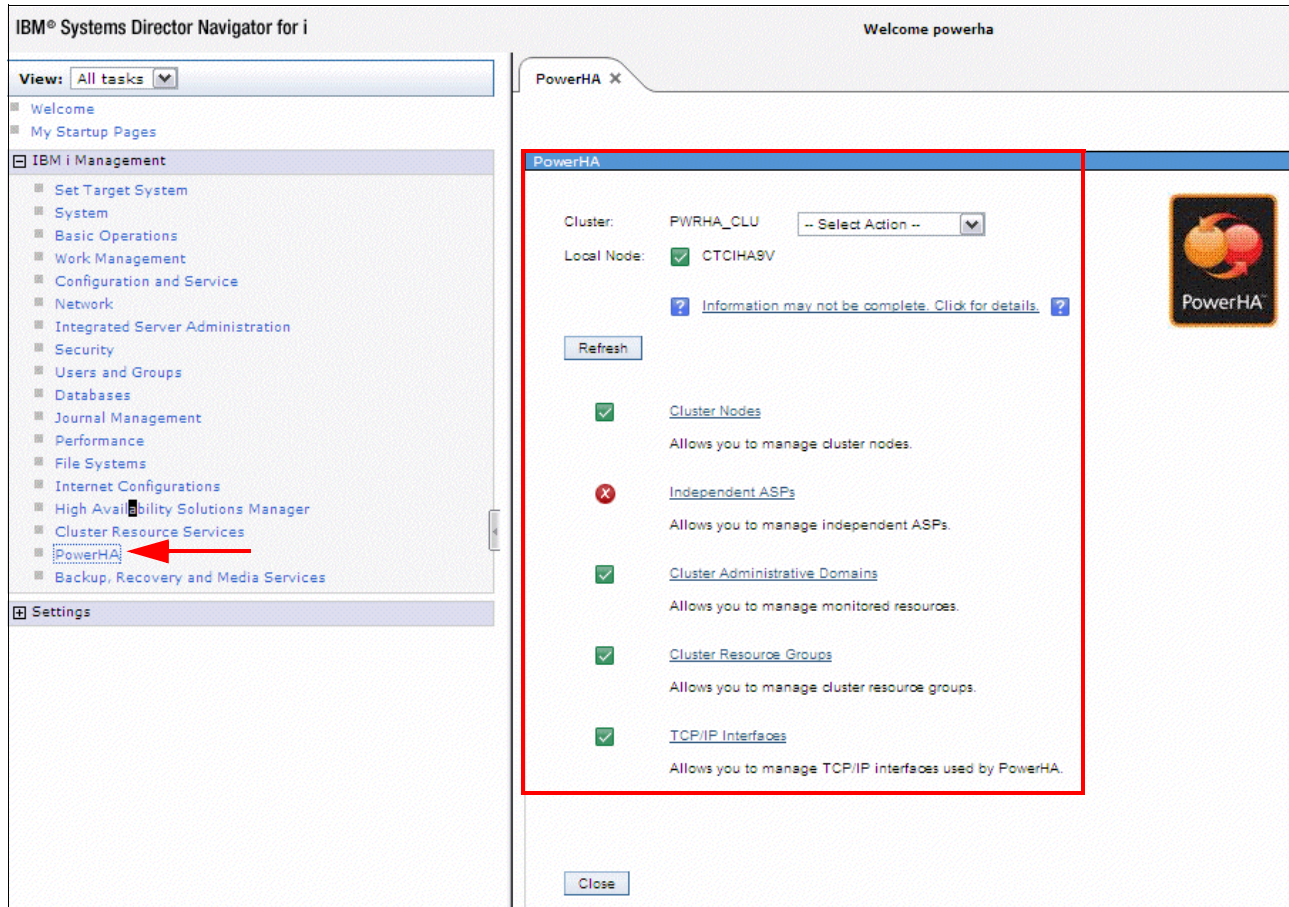


Figure 5-6 PowerHA SystemMirror for i - new GUI

The new PowerHA GUI provides the possibility to handle the high availability solution starting from one single screen. It currently supports the following:

- ▶ Geographic mirroring
- ▶ Switched disk (IOA)
- ▶ DS6000/DS8000 Metro Mirror
- ▶ DS6000/DS8000 Global Mirror
- ▶ DS6000/DS8000 FlashCopy

For more detailed information about the enhancements within the new GUI, please refer to **graphical interfaces** in the PowerHA SystemMirror for i Cookbook chapter 9.

In Figure 5-7 on page 100 we show 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 5-7 Main differences between the graphical interfaces

As the PowerHA GUI can be seen as a combination of the two other ones, those will be withdrawn in a later release.

5.1.5 Duplicate Library Error Handling

With this enhancement, the message ID CPDB8EB will be displayed in the QSYSOPR message queue at IASP vary on time when a duplicate library is found in SYSBAS and the IASP. The vary on of the IASP can be continued or canceled once the duplicate library issue is resolved.

5.1.6 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 when switching the IASP affects the virtual Fibre Channel client adapter only, instead of all ports of the physical Fibre Channel IOA, which can get reset in a native-attached storage environment.

For an overview of the new NPIV support by IBM i see Chapter 8, “Virtualization” on page 241.

For further information about NPIV implementation in an IBM i environment see *DS8000 Copy Services for IBM i with VIOS*, REDP-4584.

5.1.7 Asynchronous geographic mirroring

Asynchronous geographic mirroring is a new function supported by PowerHA SystemMirror for i Enterprise Edition with IBM i 7.1 extending the previously available synchronous geographic mirroring option, which for performance reasons, is practically 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 preserving the write-order to the target system. Therefore, at any given 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 being available, the currency of the target system, and memory overhead on the source system due to asynchronous geographic mirroring, can be checked with the DSPASPSSN (Display ASP Session) command showing the total data in transit, as shown in Figure 5-8.

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 5-8 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 CHGASPSSN (Change ASP Session) command's new DELIVERY(*ASYNCR) parameter as shown in Figure 5-9.

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 *ASYNCR

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 5-9 CHGASPSSN command - *ASYNCR Transmission delivery parameter

It can also be changed through the IBM Systems Director Navigator for i GUI by navigating to **Configuration and Service** → **Disk Pools**. Open the IASP drop down menu and choose **Properties** → **Geographic Mirroring**, as shown in Figure 5-10.

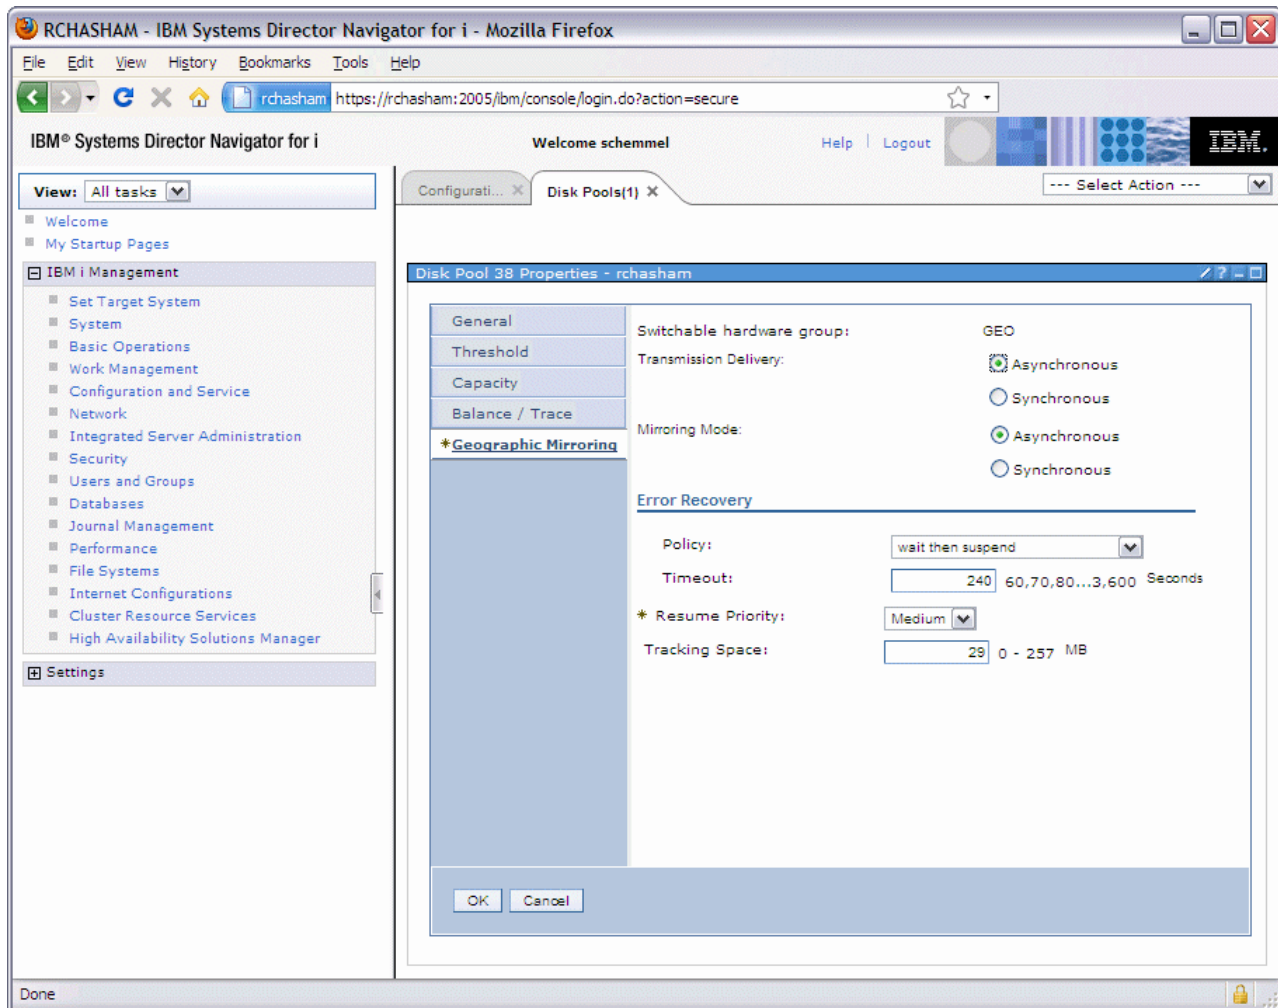


Figure 5-10 IBM Systems Director Navigator GUI GeoMirror asynchronous delivery change

5.1.8 LUN level switching

LUN level switching is a new function provided by PowerHA SystemMirror for i in IBM i 7.1 that allows for 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 managed by a cluster resource group device domain and located 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 MetroMirror or GlobalMirror is used for disaster recovery and protection against storage subsystem outages. When this happens, additional LUN level switching at the production site is used for local high availability protection eliminating the requirement for a site-switch in case of potential IBM i server outages.

For implementation of LUN level switching, an ASP copy description needs to be created for each switchable IASP using the ADDASPCPYD (Add ASP Copy Description) command, which has been enhanced with recovery domain information for LUN level switching, as shown in Figure 5-11.

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 5-11 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.

Note: Setting up an ASP copy description for LUN level switching is only supported from the green-screen 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.

5.1.9 Space-Efficient FlashCopy

PowerHA for SystemMirror for i with IBM i 7.1 newly supports space-efficient FlashCopy of the IBM System Storage DS8000 series.

The IBM System Storage DS8000 series FlashCopy SE licensed feature allows creation of space-efficient FlashCopy target volumes that can help to reduce the required physical storage space for the FlashCopy target volumes. These volumes are typically needed only for a limited time (such as for the duration of a backup to tape).

A space-efficient FlashCopy target volume has a virtual storage capacity reported to the host matching 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 is done on demand for host write operations from a configured repository volume shared by all space-efficient

FlashCopy target volumes within the same DS8000 extent pool as shown in Figure 5-12 on page 105.




Figure 5-12 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 no-copy option. The reason for this is PowerHA SystemMirror for i internally interrogates the DS8000 to determine the type of FlashCopy relationship and makes sure it uses the corresponding correct DSCLI command syntax. The syntax check is done for either traditional FlashCopy or FlashCopy SE when issuing the mkflash and rmflash commands.

For further information about using IBM System Storage DS8000 FlashCopy SE with IBM i see IBM Redbooks publication *IBM System Storage Copy Services and IBM i: A Guide to Planning and Implementation*, SG24-7103.

5.1.10 Better detection of cluster node outages

There are situations of a sudden cluster node outage such as a main storage dump, HMC immediate partition power-off or a system hardware failure which so far resulted in a partitioned cluster. In this case the user is alerted with a failed cluster communication message CPFBB22 sent to QHST and an automatic failover not started message CPFBB4F sent to QSYSOPR message queue on the first backup node of the CRG.

The operator needs to determine the reason for the cluster partition condition, which can have been caused either by a network problem or a sudden cluster node outage, and either solve the network communication problem or declare the cluster node as failed, which can be done using the CHGCLUNODE (Change Cluster Node Entry) command 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 is done by registering with a HMC or Virtual I/O Server (VIOS) management partition on IVM managed systems. This way clustering gets notified in case of severe partition or system failures to trigger a cluster failover event instead of causing a cluster partition condition.

For LPAR failure conditions it is the POWER Hypervisor™ (PHYP) that notifies the HMC that a LPAR has 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 an IBM 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 for 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 involved with CIM). On the IBM i cluster nodes the CIM event listener compares the events with available information about the nodes constituting the cluster to determine if 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 needs to be added to the IBM i cluster nodes, e.g. through the new ADDCLUMON (Add Cluster Monitor) command as shown in Figure 5-13. This enables communicating to the CIM server on the HMC or VIOS.

Add Cluster Monitor (ADDCLUMON)

Type choices, press Enter.

Cluster	HASM_CLU	Name
Node identifier	CTCV71	Name
Monitor type	*CIMSVR	*CIMSVR
CIM server:		
CIM server host name	HMC1	
CIM server user id	hmcuser	
CIM server user password . . .	password	
Bottom		
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display		
F24=More keys		

Figure 5-13 IBM i ADDCLUMOD command

For further information about configuring clustering advanced node failure detection see the IBM i 7.1 Information Center at the following web page:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=%2Frzaue%2Frzaueconceptnodefailureddetection.htm>

5.1.11 Improved geographic mirroring full synchronization performance

Performance improvements have been implemented in IBM i 7.1 for geographic mirroring full synchronization.

Changes within System Licensed Internal Code (SLIC) provide more efficient processing of data that is sent to the target system in the event of a full resynchronization. Even with the advent of source and target side tracking, there are still instances that require a full synchronization of the production copy, such as any time that the IASP cannot be normally varied off, for example, because of a sudden cluster node outage.

The achievable performance improvement varies based on the IASP data. IASPs with a large number of small objects see more benefit than those with a smaller number of large objects.

5.1.12 Cluster administrative domain enhancements

The IBM cluster administrative domain support has been 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 the following web page:

<http://public.boulder.ibm.com/infocenter/iserics/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 to allow adding MREs even if the object cannot be created. An inconsistent state of the MRE reminds the user in this case that the object still needs to be manually created.

5.1.13 IBM HA Assist for i

IBM HA Assist for i is a new licensed product (5733-HAA) for IBM i 6.1 and later which was announced with IBM i 6.1.1 as an extension for PowerHA only. IBM HA Assist for i is based on iCluster® code to replicate objects not supported for IASPs or by the cluster administrative domain. It is primarily targeted at customers with existing applications that cannot be fully migrated to an IASP environment. In these types of environments IBM HA Assist for i can be used to replicate those objects and keep them synchronized between various IBM i cluster nodes

For further information see the following:

- ▶ the IBM iCluster for i section of the IBM PowerHA web page:

<http://www-03.ibm.com/systems/power/software/availability/icluster/index.html>

- the High Availability Services for IBM Power Systems section on the IBM Systems Lab Services and Training - Systems Services - Power Services web page:

http://www-03.ibm.com/systems/services/labservices/platforms/labservices_power.htm
1

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 have been extended for field names holding either a 32-bit IPv4 or a 128-bit IPv6 address as shown for the CHGCLUNODE (Change Cluster Node Entry) command in Figure 5-14. An IPv6 internet address is specified in the form x:x:x:x:x:x:x with x being a hexadecimal number from 0 through FFFF, and "::" can be used once in the IPv6 address to indicate one or more groups of 16 bits zeros.

Change Cluster Node Entry (CHGCLUNODE)

Type choices, press Enter.

Cluster	> <i>HASM_CLU</i>	Name
Node identifier	> <i>CTCV71</i>	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 5-14 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:

- RTVCLU (Retrieve Cluster) command
- RTVCRG (Retrieve Cluster Resource Group) command
- RTVASPCPYD (Retrieve ASP Copy Description) command
- RTVASPSSN (Retrieve ASP Session) command
- PRTCADMRE (Print Cluster Administrative Domain Managed Resource Entry) command

For further information about the new PowerHA CL commands see the IBM i 7.1 Information Center at the following web page:

<http://public.boulder.ibm.com/infocenter/iseres/v7r1m0/index.jsp?topic=/rzaue/rzauewhatnew.htm>

5.2 Journaling and commitment control enhancements

Journaling and, ideally, commitment control are the base building blocks for any HA solution, as they guarantee database consistency and recoverability.

See 6.3, “Availability and consistency” on page 134 for information about new IBM i journaling and commitment control enhancements.

5.3 Additional information

For additional information and in-depth details about the latest enhancements in IBM PowerHA SystemMirror for i, please refer to the latest redbook PowerHA SystemMirror for IBM i Cookbook SG24-7994.



DB2 for i

In this chapter we describe what is new for DB2 for IBM i. It covers the following topics:

- ▶ SQL data description and data manipulation language
- ▶ Availability and consistency
- ▶ Performance and query optimization
- ▶ New functionality for DB2 developers
- ▶ DB2 security enhancements
- ▶ Enhancements to products related to DB 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.

Due to the unique 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 not only makes 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 shares this focus on SQL standards with DB2 for i, so this investment in SQL also enables DB2 for i to continue 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 will find many modifications and improvements as part of the new release. All of these are available to any of the development and deployment environments supported by the IBM POWER platforms on which IBM i 7.1 can be installed.

Note: Many DB2 enhancements for IBM i 7.1 are also available for 6.1 and some even for 5.4. In case you need to verify their availability, please see this link:

<https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/IBM%20i%20Technology%20Updates/page/IBM%20i%20Technology%20Updates>

This link is a part of developerWorks® i Zone.

6.2 SQL data description and data manipulation language

There are several changes and additions to the SQL language, which we describe in this chapter:

- ▶ 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
- ▶ MQSeries® 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

6.2.1 XML support

Extensible Markup Language (XML) is a simple, flexible text format 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, see the following web page:

<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 had to 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 forth).

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 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 representing an XML document using the XMLSERIALIZE (see “XML serialization” on page 115) 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 114) 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 containing 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 defining 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 6.4.15, “QAQQINI properties” on page 150). If this value has not been set, the default is 1208. The CCSID is established for XML data types 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 containing 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 containing 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
xmlnamespaces	Constructs namespace declarations from the arguments
xmlparse	Parses the arguments as an XML document and returns an XML value

Function	Description
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 having the input argument as the content
xmlvalidate	Returns a copy of the input XML value 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. With this you can either remove or maintain any whitespace 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 invoke the XMLSERIALIZE function to request XML serialization explicitly. The most common use 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. However, the XMLSERIALIZE scalar function lets you 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 on page 116) or DBCLOB before sending 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 containing information about XML schemas. XML instance documents might contain a reference to a Uniform Resource Identifier (Go) 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 overhead 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 begins registration of an XML schema.

The SYSPROC.XSR_ADDSCHEMADOC procedure adds additional XML schema documents to an XML schema that you are in the process of 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 registering or adding 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.

Note: Because an independent auxiliary storage pool (IASP) can be switched between multiple systems, there are additional 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 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 the XML data is to be stored in
- ▶ 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 inserting it, or applying conditions for insertion.

The following are items to consider when annotating your XML schema, along with pointers to related documentation:

- ▶ 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 it is being mapped to.
- ▶ Ensure complex types derived by restriction or extension are properly annotated.
- ▶ Confirm that no decomposition limits and restrictions are violated.
- ▶ Ensure that the tables and columns 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 XML document into relational database tables, such as `db2-xdb:defaultSQLSchema`, `db2-xdb:rowSet`, `db2-xdb:column`, etc.

In one XSR, multiple target tables can be specified, hence data in an XML document can be shredded to more than one target tables using one XSR. But the order of insertion into tables can not be specified with existing decomposition annotations, therefore, if the target tables have reference relationship, the insertion of dependent row fails if its parent row is not inserted before that.

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 MERGE statement

This statement enables the simplification of matching rows in tables and allows to 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 only be operated on once. A row in the target can only be identified 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 **upsert**.

Using the MERGE statement has a good potential 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 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 have been 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 must be specified. Each group has an activities table. For example, activities_groupA contains all activities Group A organizes, and the archive table contains all upcoming activities 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.

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

```

There is a difference in how many updates are done, depending on whether a NOT ATOMIC MERGE or an ATOMIC MERGE has been specified.

- In an ATOMIC MERGE, the source rows are processed as though a set of rows are processed by each WHEN clause. Thus, if five rows are to be updated, any row level update trigger is fired five times for each WHEN clause. This 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 was executed for each source row, meaning that, in the aforementioned case, the triggers are fired only 5 times.

After executing 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 operated on by the MERGE statement, excluding rows identified by the ELSE IGNORE clause.

The ROW_COUNT item and SQLERRD(3) does 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 already exist before the MERGE statement was executed. That is, 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 only be accessed 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 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. This 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 has been 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 hold 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. Note that 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` already 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 has been 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 on page 120 creates a global variable that defines a user class. This variable has its initial value set based on the result of invoking a function called `CLASS_FUNC`. This function is assumed to assign a class value such as administrator or clerk 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 will see 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 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 between 0 and 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 required to represent an array value is usually 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 on page 121, we use two array data types (intArray and stringArray), and a persons table with two columns (ID and name). The processPersons procedure adds three additional 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 only be created 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 forth) 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 to write, 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 has 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 used by the stored procedure to return the result set.

- The 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 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 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 executed 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 execute 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 has already been 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 big enough to hold descriptions of all columns in the result set.

You can use DESCRIBE CURSOR only for cursors for which you executed ALLOCATE CURSOR previously.

After you execute 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, perform 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.

Example 6-6 gives you an idea on how to implement this 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 allows for transparent encryption/decryption or encoding/decoding of data accessed through SQL or any other interface. It allows for transparent encryption or encoding of data accessed through SQL or native.

When values in the column are changed, or new values inserted, the field procedure is invoked 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 invoked for each value, which is encoded, and must decode it back to the original value. Any indexes 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. That 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. That 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 invoked during the processing of the CREATE TABLE or ALTER TABLE statement. That operation is called a **field-definition**. When so invoked, 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 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 having 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 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.

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 defining and using the parameters in the field procedure, take care to ensure that no more storage is referenced for a given parameter than is defined for that parameter. The parameters are all stored in the same space and exceeding a given parameter's storage space can overwrite another parameter's value. This, 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:

- ▶ 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 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 is dependent on function code.

- If function code 8, then the NULL value. This parameter is input only.
- If function code 0, then the data to be encoded. This parameter is input only.
- If function code 4, then the location to place the decoded data. 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 calling 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 MQ Series integration

IBM WebSphere MQ is a family of network communication software products, allowing independent and potentially non-concurrent applications on a distributed system to communicate with each other.

This implementation 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 MQSeries location, specified by *receive-service*, using the policy 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 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 MQSeries location, specified by *receive-service*, using the policy 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 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 MQSeries location, specified by receive-service, using the policy 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 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 MQSeries location, specified by receive-service, using the policy 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 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 aforementioned 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 MQ location specified by send-service, using the policy 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 MQ location specified by receive-service, using the policy 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 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 MQ location specified by receive-service, using the policy 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 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 MQ location specified by receive-service, using the policy 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 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 MQ location specified by receive-service, using the policy 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 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 aforementioned table functions, the receive-service is not specified or the null value is used, the DB2.DEFAULT.SERVICE is used.

DB2 MQ tables

The DB2 MQ tables contain service and policy definitions that are used by the DB2 MQ functions. The DB2 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 additional services and policies by inserting additional rows in the tables.

DB2 MQ CCSID conversion

When a message is sent, the message sent 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 defined to be 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 executed. The message data passed from DB2 is sent unchanged along with a CCSID which informs the receiver of the message how to interpret the message data. The CCSID that is sent depends on the value specified for the CODEDCHARSETID of the service 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 does not get 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 DB2 MQ scalar or table function, the msg-data return parameter (and the MSG result column for the DB2 MQ table functions) are also defined to be in job default CCSID. DB2 does no conversions and relies on WebSphere MQ to perform any necessary conversions. Whether WebSphere 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 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 read or received into that CCSID. Specifying something other than -3 for CODEDCHARSETID in a service used to read or receive messages is not recommended, because the msg-data return parameter and MSG result column are defined by DB2 to be in job default CCSID.

When reading or receiving 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 have been 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 whose 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 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

This new function allows simplifying 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

The support for three-part names allows you to bypass the explicit CONNECT or SET CONNECTION. Statements that use three-part names and see distributed data, result in 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 needs to be defined at the local site. An alias at a remote site can see yet another server as long as a referenced alias eventually refers to a table or view.

All object references in a single SQL statement must reside 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 on page 131

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 on 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 is the default. This value directs the database manager to wait for the commit or rollback when encountering locked data that is in the process of being updated or deleted. Locked rows that are in the process of being inserted are not skipped. This option does not apply for read-only queries 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 encountering locked data in the process of being updated or deleted. Locked rows in the process of being inserted can be skipped. This option applies where possible when 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 using these options because they might affect application functionality.

You can specify the use for concurrent access resolution in several ways:

- With 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 command

Specifying the CREATE OR REPLACE statement makes it easier to create an object without having to drop when it already exists. This 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 of course, 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 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 from -2^{112} to $2^{112} - 1$, 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 supporting fewer bits is cast to a value with the data type of the argument supporting more bits. This cast impacts 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.

Due to 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 of the BITXOR function is recommended 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 creating 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 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 creating 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 invoking 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.

6.2.9 OVRDBF SEQONLY(YES, buffer length)

OVRDBF will add support to allow the user to specify the buffer length rather than the number of records for OVRDBF SEQONLY(*YES N). N can be:

- ▶ *BUF32KB
- ▶ *BUF64KB
- ▶ *BUF128KB
- ▶ *BUF256KB

This means the number of records will be the number of records that will fit into a 32KB, 64KB, 128KB, or 256KB buffer.

6.3 Availability and consistency

Several enhancements have been made in the area of the integrity preservation and journaling. The main objectives are to provide easier interfaces on the setup and the monitoring of the database persistency, including in HA setups.

6.3.1 Journal management

Journals (more familiarly known as logs on other platforms) are used to keep track of changes to various types of 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 STRJRNLIB (Start Journal Library) 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, the STRJRNLIB (see Figure 6-1 on page 135) 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 can be specified with a specific or generic name. The default is to apply the rule to all objects that match the other criteria specified in the inherit rule regardless of the object name. This gives you the ability to start journaling on new production work files, but no journal temporary work files if they are named uniquely.

```

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 6-1 STRJRNLIB command prompt

If the library is already journaled and you want to define one of the new inherit rules, use the CHGJRNOBJ (Change Journaled Object) command. If the library is not already journaled, the new rules can be specified through the STRJRNLIB (Start Journal Library) command. To view the current inherit rules associated with a journaled library use the DSPLIBD (Display Library Description) command. Then click F10 - Display inherit rules.

There is an equivalent in the Systems Director Navigator for i to do the same task. Navigate to **Expand File Systems** → **Select Integrated File System** → **Select QSYS.LIB**. Select the library you want to journal and the Journaling action, as shown in Figure 6-2 and Figure 6-3 on page 136.

File System... x Integrated ... x --- Select Action ---

Journaling for ASAMPLE in ASP: System disk pool - 9.5.168.75

***Journaling Properties**

Recovery Information

Status: Never journaled

* Journal: qsqjrn Browse...

* Library: asample

Last journal start

Inherit rules

Select	Object type	Operation	Rule action	Omit journal entry	Images	Allow remote journal filtering	Name filter
<input checked="" type="checkbox"/>	All	All	Include	Object default	Object default	Object default	All Objects

Page 1 of 1 1 Go Rows 1 Total: 1 Filtered: 1 Selected: 1

Start End

Figure 6-2 Select a library for journaling




Figure 6-3 Setting a rule

6.3.2 Remote journaling

When a remote journal connection is ended with a recoverable error, you can now specify that the operating system try to restart the connection automatically. This is done by identifying the number of attempts and the time, expressed in seconds, between restart attempts. These settings can be set using the CHGRMTJRN (Change Remote Journal) command or with the QJOCHANGEJOURNALSTATE (Change Journal State) 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 the following web page:

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

The same command allows you to filter the remote journals. Filtering out journal entries that are not absolutely needed on the target system can decrease the amount of data sent across the communication line.

This remote journal filtering feature is available with option 42 of the IBM i Operating System, 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 activating a remote journal. Different remote journals associated with the same local journal can have different filter criteria. Remote journal filtering can only be specified for asynchronous remote journal connections. Because journal entries might be missing, filtered remote journal receivers cannot be used with the RMVJRNCHG (Remove Journalized Changes) command. Similarly, journal receivers that filtered journal entries by object or by program cannot be used with the APYJRNCHG (Apply Journalized Change) command or the APYJRNCHGX (Apply Journalized Change Extend) command.

The WRKJRNA (Work with Journal Attributes) command now has the ability to 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 occurring for a remote journal connection.

6.3.3 DISPLAY_JOURNAL (easier searches of journal)

Displaying a journal entry from a GUI interface today either requires using APIs or writing the journal entries to an outfile. The APIs are labor intensive and the outfile is somewhat restrictive and slower since a copy of the data required.

QSYS2.Display_Journal is a new table function that allows the user 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 of the special values, see the QjoRetrieveJournalEntries API in the Information Center.

Unlike many other UDTFs residing in QSYS2, this one has no DB2 for i provided view.

The following 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 a value of the Journal_Library.

► **Starting_Receiver_Library and Starting_Receiver_Name**

If the specified Starting_Receiver_Name is the null value, empty string, or blanks the *CURRENT will be 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 timestamp 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, or 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 containing 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' or 'R J' or 'R,J' or 'R, J'.

► **Journal_Entry_Types**

If the specified Journal_Entry_Types is the null value, or 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 containing 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' or 'R J' or 'R,J' or 'R, J'.

► **Object_Library, Object_Name, Object_ObjType, and Object_Member**

If the specified Object_Name is the null value, or 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 was not *FILE, the member name is ignored.

► **User**

If the specified User is the null value, or 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, or an empty string, or a blank string, *ALL is used.

Otherwise, the Job must identify a valid job name a specific job where the first 10 characters are the job name, the second 10 characters are the user name, and the last 6 characters are the job number.

► **Program**

If the specified Program is the null value, or an empty string, or a blank string, *ALL is used. Otherwise, the Program must identify a valid program name.

Example 6-20 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/qsqrjrn.
```

```
select * from table (  
Display_Journal (  
  'JSOCHR', 'QSQRJRN', -- 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 gives you a result table with data similar to what you get using DSPJRN (Display Journal Command).

6.3.4 Commitment control and independent ASPs

Commitment control allows you 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 have been either applied permanently or removed permanently. Furthermore, if any process or even a complete system performing 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 due to 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 has been deployed using independent ASPs (IASPs), you are actually using a database instance that resides in that IASP. This has an impact on how commitment control will work.

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 release IBM i 7.1 is that if you subsequently place independent disk pool resources under commitment control before system disk pool resources, the commitment definition is moved to the independent disk pool. This 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 reside in the current library name space. For example, they can reside in the independent ASP or *SYSBAS.

If the first resource that is placed under commitment control does not reside 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 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 reside 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 the data in both ASP groups is committed or rolled back atomically.

6.3.5 System managed access path protection (SMAPP)

System-managed access-path protection allows you 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 has effect on the overall system performance. The lower the target recovery time 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 built over files 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) shows the internal threshold used by SMAPP (see Figure 6-4). This panel was added in this release. 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 6-4 Display Details from Edit and Display Recovery for Access Paths

6.3.6 Journal management functions: IBM Systems Director Navigator for i

IBM Systems Director Navigator now supports additional journal management functions. With IBM i 7.1 the following functions were all added:

- ▶ Change journal receivers and attributes associated with a journal.
- ▶ View the properties associated with a journal receiver.
- ▶ View the objects journaled to a specific journal.
- ▶ Add and remove remote journals.
- ▶ View the list of remote journals associated with a specific journal.
- ▶ Activate and inactivate remote journals.
- ▶ View the details of a remote journal connection.

See Chapter 18, “IBM Systems Director Navigator for IBM i 7.1” on page 525 for more information.

6.4 Performance and query optimization

In the IBM i 7.1 release of DB2 for IBM 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 dictated by a large number of potential issues, typically associated with 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, the size and volumes to be processed, and so on. All of these influence largely the runtime, throughput, or response time, which is supported by the IT components and which is defined by the needs of the business. Performance optimization for database access therefore needs to address all the components that are used in obtained acceptable and sustainable results, covering the functional aspects and the technical components that support them.

In this section, we review the query optimization method. Then we describe what is behind the changes that are implemented in the database management components to relieve the burden, associated with the tools and processes a database administrator uses or follows to realize the non-functional requirements about performance and scalability:

- ▶ Global Statistics Cache

- ▶ 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

6.4.1 Methods and tools for performance optimization

Typically, the autonomous functions in IBM i, and the new ones in IBM i 7.1 addressed in the paragraphs that follow, 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 have taken the architectural decision for a new application to use a stateless runtime environment and that your detailed component model has designed the infrastructure for it. If the business processes it supports are changing and require a more statefull design, you might face an issue if you want to preserve information to keep track of the statefullness in your transactions. At that time, 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 factor this in, you have a serious issue. Just striving for being capable to handle high volumes with a low latency is good, but this needs to 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 defining components for a database support, develop a methodology and use best practices to obtain the best results. Any methodology has to be consistent, acceptable, measurable, and sustainable. You want to stay away from ad hoc measures or simple bypasses.

IBM i provides statistics on 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 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 need to consider more. The following paragraph provides an example.

Since the introduction of solid state drives (SSD) carrying a low latency, the IBM i storage manager has awareness about this kind of technology and uses it as appropriate. Since release 6.1, you can specify a preferred unit on the CREATE TABLE/INDEX and ALTER TABLE/INDEX commands (See 6.4.9, "SQE optimization for indexes on SSD" on page 148). The SYSTABLESTAT and SYSINDEXSTAT catalog tables provide additional I/O statistics (SEQUENTIAL_READS and RANDOM_READS) in 7.1 on these objects. Keep in mind that these statistics, generated by the database manager, only indicate possible candidates to be housed on SSD hardware. Further investigation on run time and the contribution to the performance and capacity of the infrastructure will reveal whether or not they are really eligible for that kind of settings.

More on the SSD topic can be found at Chapter 9, "Storage and solid state drives" on page 273.

Finally, and as a last resort, there is now a stored procedure available that allows you to cancel long running SQL jobs, using the QSYS2.CANCEL_SQL procedure.

6.4.2 What is 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 execute 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 (CPU, memory, 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 forth. 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 can, for instance, 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 is IPLed again. This index can be used by any query that might benefit of its existence. These autonomous indexes can be viewed and carry the information about which a database administrator can decide whether or not to make it a permanent object, 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, the existence of stale statistics on busy systems,
- ▶ Hidden correlations in the data, often due to a poor design, data skew, and data volatility
- ▶ Changes in the business or infrastructure environment.

In the last case, this is more likely to happen on variations in both memory and CPU allocations on partitioned systems, which are re-configured 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, read the IBM Redbooks publication *Preparing for and Tuning the SQL Query Engine on DB2 for i5/OS*, SG24-6598.

6.4.3 Global Statistics Cache (GSC)

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, up to 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 are typically post mortem actions and are labor intensive.

To reduce this labor intensive work, the DB2 Statistics Manager has been reviewed. 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 GSC, which is a system wide repository, containing those complex statistics. The Adaptive Query Processing (AQP) (See 6.4.4, “Adaptive query processing (AQP)” on page 143) inspects the results of queries and compares the estimated row counts with the actual row counts. All of the queries 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 re-optimize and re-start the query using the new query plan. Further, 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 allows generating faster query plans by the query optimizer.

Typically, observed statistics are for complex predicates such as with a join. A simple example is a query joining 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 DSPFD (Display File Description) command. This command only provides a limited number of I/O operations. Both counters (catalog tables and the object description) are reset at IPL time.

As a reminder: the statistics collection is defined by the system value QDBFSTCCOL (Data Base file statistics collection). The SM jobs that update the statistics, carry the same name.

6.4.4 Adaptive query processing (AQP)

The SQE uses statistics to build the mechanism on how to perform an SQL statement. These statistics come from two sources:

- ▶ Information contained in the indexes on the tables used in the statement
- ▶ Information 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 AQP. AQP analyzes actual query run time statistics and uses that information to correct previous estimates. These updated estimates can provide better information for subsequent optimizations. Secondly, it focuses on optimizing join statements to improve the join orders and minimizing the creation of large 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 two 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 due to inaccurate statistical estimates. This process is referred to as the *AQP Handler*.

After a query has completed, another task, the AQP Request Support starts, and runs in a system task so 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 re-optimize the query using partial 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 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 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 undertaken with by the AQP.

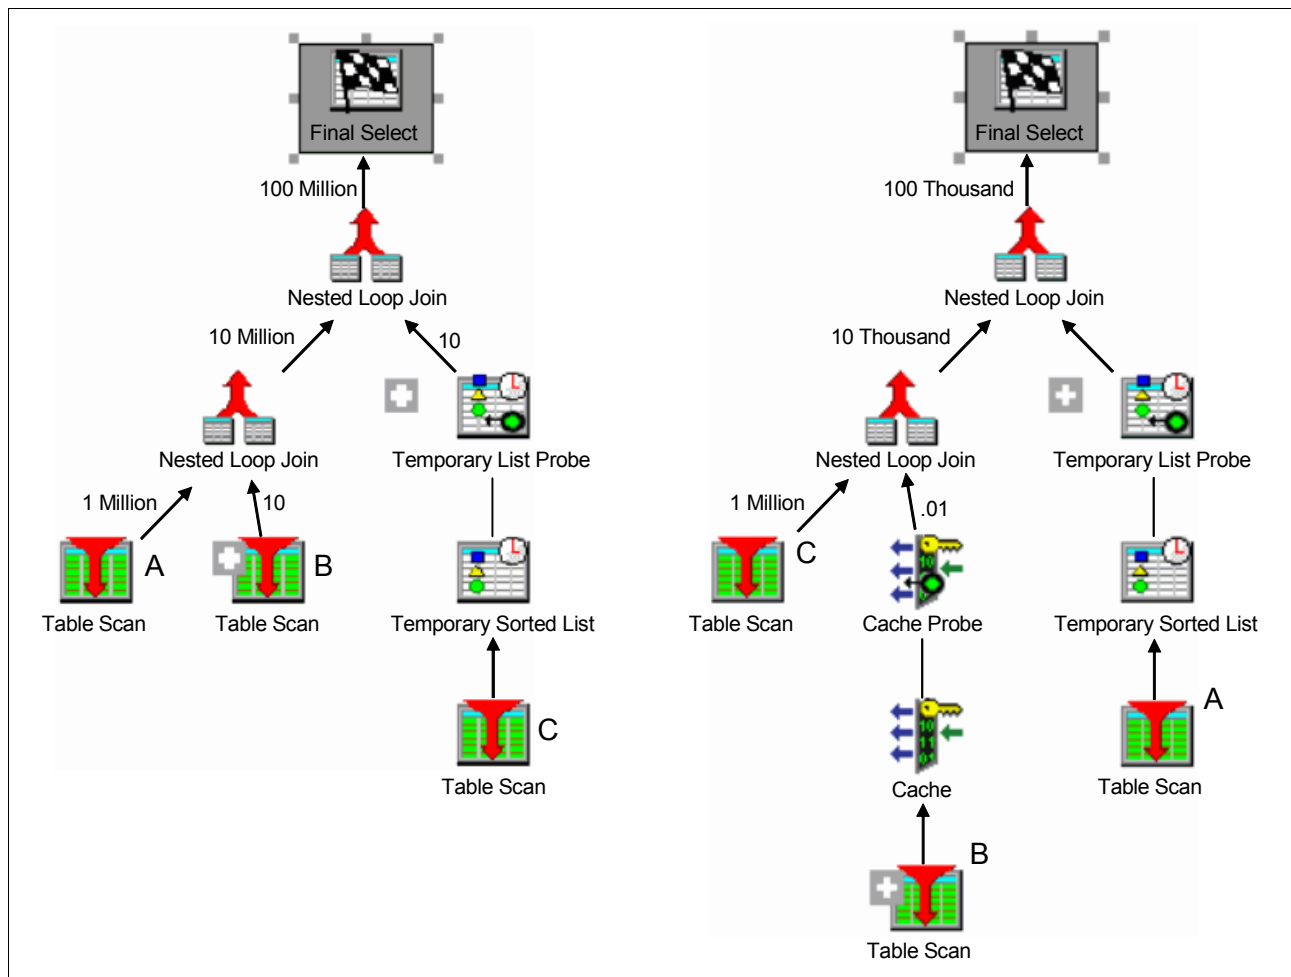


Figure 6-5 AQP Optimization on JOIN

Figure 6-5 on page 144 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 returned and dramatically reduced the size of the result set.

6.4.5 Sparse indexes

Starting from the IBM i 6.1 release, you can create a sparse index using a WHERE condition. In the 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.

For a sparse index to be used, the WHERE selection in the query must be a subset of the WHERE selection in the sparse index. That is, 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 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 6-21, 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 6-21 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, in the second one, all of them. 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 executed 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.

6.4.6 Encoded vector index (EVI)

Earlier in this chapter, we discussed the enhancements for the EVIs (see “Encoded Vector index” on page 133.) The EVI can be used for more than generating a bitmap or row number list to provide an asynchronous I/O map to the desired table rows. The EVI can also be used by two index-only access methods that can be applied specific 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 advantage for such an operation are that the pre-summarized results are readily available. It also only processes the unique values in the symbol table, avoiding processing table records. Similarly, it extracts all the data from the index unique key values or INCLUDE values, eliminating the need for a Table Probe or vector scan. With the INCLUDE statement, it provides ready-made numeric aggregates, eliminating the need to access corresponding table rows to perform the aggregation.

The advantages of this are quite 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 (see following section) 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 involved, making the symbol table large. You are likely to experience poor performance if a large portion of the symbol table has been put into the overflow area. On the other hand, you will 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 additional INCLUDE values. These are ready-made numeric aggregate results, such as SUM, COUNT, AVG, or VARIANCE values 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 addendum to the EVI keys for each symbol table entry. Because these are numeric results and finite in size, the symbol table is still a desirable compact size.

These 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. This 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 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 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.

6.4.7 Preserve EVI indexes on ALTER enhancement

Prior to 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 as long as the data type or other attribute of a key column of the index is not changed by the alter.

6.4.8 Keeping tables or indexes in memory

The KEEPINMEM parameter (KEEPINMEM) specifies whether the data for a file member or an access paths 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 QAQQINI (Query Options File) parameter MEMORY_POOL_PREFERENCE (see 6.4.15, “QAQQINI properties” on page 150) specifies the preferred main storage pool to be used.

This function applies only during the runtime of a query, and might therefore be substituted to the SETOBJACC (Set Object Access) command that puts the table or index in memory in a static function. After the query did complete, the memory might be freed again, contrary to the effects of the SETOBJACC (Set Object Access), where you need to clear it using the *PURGE option on the POOL (Storage Pool) parameter of the command.

Similarly, the DB2 database manager reduces the storage occupied by a table that does not contain any data, thus reducing the storage space needed for unused objects. This is also referred to as *deflated table support*.

6.4.9 SQE optimization for indexes on SSD

The query optimizer will now recognize that indexes may potentially be located on SSD and prioritize usage of those indexes higher than indexes located on spinning disk when ordering the indexes during optimization.

Indexes must have the SSD attribute specified via UNIT(*SSD) parameter on the CRTLF (Create Logical File) or CHGLF (Change Logical File) CL commands or by using the UNIT SSD clause on the SQL CREATE INDEX statement. See also 6.5.10, “CHGPFM and CHGLFM UNIT support” on page 162

6.4.10 SQE support of simple logical files

SQE supports simple logical files in 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 will only support read only queries. There is no SQE support of insert, update or delete using logical files.

The QAQQINI file option 'IGNORE_DERIVED_INDEX' will continue to be supported. If IGNORE_DERIVED_INDEX(*NO) is specified, and a select/omit logical file exists over the based on file of the simple logical file, then SQE will not process the query of the simple logical file.

6.4.11 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 IBM i index advice condenser is externalized through the QSYS2.CONDENSEDINDEXADVICE view. The view and underlying user defined table function are hard-wired to use the raw index advice stored within the QSYS2/SYSIXADV file. Some users need to have the ability to utilize the index advice condenser against a file that was saved and restored from a different machine.

A new database supplied procedure (QSYS2.INDEX_ADVICE) has been 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.

When the procedure is called with advice_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 will be altered to match the 7.1 advice format. This 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.

Example 6-22 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;
```

6.4.12 SKIP LOCKED DATA and NC or UR

SKIP LOCKED DATA clause allows a user 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 will be skipped. This behavior is typically desired 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.

Prior to this enhancement, SKIP LOCKED DATA was only allowed when the isolation level was CS or RS.

For more on this topic see:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/db2/rbafzskiplocked>

6.4.13 SQL routine performance integer arithmetic (requires recreate)

This improves the performance of the generated code within LANGUAGE SQL routines. When adding and subtracting 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 will not be observed when

building SQL routines for previous releases (e.g. SET OPTION TGTRLS=V6R1M0 or V5R4M0).

To achieve the improved code generation, SQL procedures, functions and trigger routine must be recreated after upgrading 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>
But these statements will not generate inline ILE C code:
SET v1 = v2 + <integer literal>
SET v2 = v1 + <bigint literal>
```

6.4.14 Automatic cancel of QSQRVR jobs when an application is ended

When an application using SQL Server Mode is terminated and the SQL Server Mode connections have not been ended, the termination of the application job may completely end while the QSQRVR job remains active performing a long running system operation.

Job termination has been improved to signal an SQL Cancel request to any QSQRVR jobs being used by the application. The cancel will interrupt some long running operations, allowing the QSQRVR job to observe that the application is ending.

SQL Server Mode users only need to apply the PTF to receive the improved cancel handling support.

6.4.15 QAQQINI properties

Table 6-3 lists the new parameters and their values which 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 IBM i.

Table 6-3 QAQQINI file

Parameter	Description
ALLOW_ADAPTIVE_QUERY_PROCESSING (6.4.4, “Adaptive query processing (AQP)” on page 143)	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 (6.2.6, “FIELDPROC support for encoding and encryption” on page 124)	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 uses. This option does not guarantee use 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	<p>This option 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 require a different plan to perform well, SQE will re-optimize the query.</p> <p>*DEFAULT – The default value is *NO.</p> <p>*NO – do not check host variable selectivity at pseudo-open time. This is compatible with previous behavior.</p> <p>*OPTIMIZE – The optimizer will determine when host variable selectivity should be checked. In general, the SQE engine will monitor the query and if after a certain number runs it determines that there is no advantage to checking host variable values, i.e. the selectivity isn't 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 will do the normal plan validation.</p> <p>*YES– Always check host variable selectivity at pseudo-open time.</p> <p>Note: If REOPTIMIZE_ACCESS_PLAN QAQQINI option is set to *ONLY_REQUIRED, the PSEUDO_OPEN_CHECK_HOST_VARS option has no effect.</p>
SQL_CONCURRENT_ACCESS_RESOLUTION (“Concurrent access resolution” on page 131)	Specifies the concurrent access resolution to use for an SQL query.
SQL_XML_DATA_CCSID (“XML data type” on page 113)	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 specified in an XML text search using the CONTAINS or SCORE function. The time zone is the offset from UTC (Greenwich mean time). It is only applicable when a specific time zone is not given for the value.

6.4.16 Alter table performance

ALTER TABLE can be a long running operation. The general performance of ALTER TABLE has been improved (though 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.

6.4.17 Avoid short name collisions for CREATE PROCEDURE, FUNCTION or TRIGGER

When SQL routines (CREATE PROCEDURE (SQL), CREATE FUNCTION (SQL) and CREATE TRIGGER) are created using a long name, the database will generate the system name of the routine. For long names, the first 5 characters of the long name are combined with '00001'. If an object with that system name already exists, the second half of the name is incremented by one and the create is retried.

If you have many SQL routines whose names begin with a common first 5 characters, the creation of the routines will be slowed down by all the name conflicts and rebuild attempts needed to find a system name which hasn't been used.

The QGENOBJNAM data area can be used to control the system name 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 have affect, the data area must be created as CHAR(10) and must reside within a library that is in the library list.

The user creating 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.

Example 6-23 Using 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
```

Will always use MNAME123 for the system name of the trigger program.

There are two ways to use the QGENOBJNAM data area:

1. Use question marks and a starting value ('?????xxxxx') where x is a number digit, the generated name will begin 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, the last 4 digits from the job number and '000'. For example, if the value of the data area was '*GEN00 ', the job number was 098435, 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.

Automatically assigned trigger system programs according to the value of QGETNOBJNAM:

Example 6-24 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 will show you the short names assigned to trigger programs.

6.4.18 CREATE PROCEDURE (SQL) PROGRAM TYPE SUB

There is a simple action which will improve the performance of SQL procedures is to use the PROGRAM TYPE SUB clause. When omitted or PROGRAM TYPE MAIN is used on 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 use of PROGRAM TYPE SUB is most relevant for procedures that are frequently called within a performance critical application. PROGRAM TYPE SUB procedures perform better due to the fact that 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 using 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.

6.4.19 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. Prior to this enhancement, if a large number of operations and nested transactions was performed as part of the outer transaction (perhaps due to 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.

6.4.20 QSQBIGPSA data area

For some DB2 for IBM 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 via the **QSQCSRTH** data area.

*DUMMY cursors exist when unique SQL statements are prepared using a statement name that isn't unique. The SQL cursor name is changed to '*DUMMY' to allow the possibility of the cursor being re-used in the future.

Prepared SQL statements are maintained within a thread scoped internal data structure 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, they will observe *DUMMY cursors being hard closed at each PSA compression.

This type of application is gaining little value from the PSA compression and has to endure the performance penalty of its *DUMMY cursors being hard closed.

A new data area control is being provided for this type of user. The **QSQBIGPSA** will indicate that the application wants to start with a large size for the PSA threshold.

By using this option, the application will skip all the PSA compressions 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 needs to exist within the library list for a job when the first SQL PREPARE statement is executed.

The data area merely needs to exist, it doesn't need to be set to any value.

6.4.21 Validate constraints without checking

In 7.1 a new CHECK parameter was added to the CHGPF CST (Change PF Constraint) 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 will verify that the table's data conforms to the constraint definition.

This can be a long running operation. CHECK(*NO) will enable the constraint without checking. If the data is not checked when the constraint is enabled, it is the responsibility of the user to guarantee that the data currently in the file is valid for the constraint.

Prior to 7.1, a data area can be created to enable a constraint without checking. When CHGPF CST (Change PF Constraint) is performed, DB2 will search for a data area in QTEMP called QDB_CHGPF CST. If the data area is found and its length is exactly 9 characters and contains the value 'UNCHECKED', DB2 will enable the constraint without validation.

6.4.22 Limit the amount of processing on a RGZPFM cancel

A RGZPFM (Reorganize Physical File Member) command with an ALWCANCEL(*YES) parameter can be canceled and then later restarted where it left off.

Prior to this enhancement, a possibly significant amount of processing was performed at the time of the cancel to allow the RGZPFM (Reorganize Physical File Member) to be restarted later and to return as much storage to the system as possible.

With this enhancement, the amount of time processing performed at cancel time is minimized, allowing the RGZPFM (Reorganize Physical File Member) to be canceled in a reasonable amount of time. The processing that is bypassed will be performed later when the RGZPFM (Reorganize Physical File Member) is restarted.

6.4.23 CPYFRMIMPF performance

Prior to this enhancement when issuing CPYFRMIMPF (Copy from Import File) 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 given 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 can drastically reduce the CPU utilized and the elapsed time for the CPYFRMIMPF (Copy from Import File). The more columns the table or physical file has the larger the performance benefit.

6.4.24 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.

Force journal entry. Whether 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 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 will result in an error.

6.4.25 QDBRTVSN API performance

Prior to this enhancement, finding the short name for a given long name of a table or view would always be 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 will now find the short name in most cases without enqueueing a request to the database cross reference.

6.4.26 Add 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 database full opens. (SQL and Native I/O)

Database has instrumented the number of full opens that occur within a job. This metric will be reported by Collection Services in the **QAPMJOBOS** 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.

Field Name - JBNUS

Description - The number of Native database (non-SQL) files and SQL cursors which have been full opened. Subtracting the value within field JBLBO from JBNUS will yield the number of non-SQL full opens.

Existing fields: (for SQL Cursors)

Field Name - JBLBO

Description - The cumulative number of SQL cursors which have been full opened.

Field Name - JBLBS

Description - The cumulative number of SQL cursors which have been pseudo-opened. Pseudo-opens are also known as reused SQL cursors.

6.5 New Functionality for DB2 developers

This section will cover the new functionality for DB2 for i developers.

6.5.1 QSYS2.SYSCOLUMNS2 view

QSYS2.SYSCOLUMNS2 is a view based on a table function that will return additional information not available in SYSCOLUMNS (such as the allocated length of a varying length column). Since it is based on a table function, it will typically return results faster if a specific table is specified when querying it.

Example 6-25 QSYS2.SYSCOLUMNS2 definition

```
SELECT * FROM qsys2.syscolumns2
WHERE system_table_schema = 'DBSCHM' and system_table_name = 'T1' ;
```

QSYS2.SYSCOLUMNS2 definition:

```
COLUMN_NAME FOR COLUMN NAME VARCHAR(128) ALLOCATE(100) DEFAULT NULL
TABLE_NAME FOR COLUMN TBNAME VARCHAR(128) ALLOCATE(18) NOT NULL
TABLE_OWNER FOR COLUMN TBCREATOR VARCHAR(128) ALLOCATE(100) NOT NULL
ORDINAL_POSITION FOR COLUMN COLNO INTEGER DEFAULT NULL
DATA_TYPE FOR COLUMN COLTYPE VARCHAR(8) ALLOCATE(8) DEFAULT NULL
LENGTH INTEGER DEFAULT NULL
NUMERIC_SCALE FOR COLUMN SCALE INTEGER DEFAULT NULL
IS_NULLABLE FOR COLUMN "NULLS" CHAR(1) DEFAULT NULL
```

IS_UPDATABLE FOR COLUMN UPDATES CHAR(1) DEFAULT NULL
 LONG_COMMENT FOR COLUMN REMARKS VARGRAPHIC(2000) ALLOCATE(100) CCSID 1200 DEFAULT NULL
 HAS_DEFAULT FOR COLUMN "DEFAULT" CHAR(1) DEFAULT NULL
 COLUMN_HEADING FOR COLUMN "LABEL" VARGRAPHIC(60) ALLOCATE(60) CCSID 1200 DEFAULT NULL
 STORAGE INTEGER DEFAULT NULL
 NUMERIC_PRECISION FOR COLUMN PRECISION INTEGER DEFAULT NULL
 "CCSID" INTEGER DEFAULT NULL
 TABLE_SCHEMA FOR COLUMN DBNAME VARCHAR(128) ALLOCATE(10) NOT NULL
 COLUMN_DEFAULT FOR COLUMN DFTVALUE VARGRAPHIC(2000) ALLOCATE(100) CCSID 1200 DEFAULT NULL
 CHARACTER_MAXIMUM_LENGTH FOR COLUMN CHARLEN INTEGER DEFAULT NULL
 CHARACTER_OCTET_LENGTH FOR COLUMN CHARBYTE INTEGER DEFAULT NULL
 NUMERIC_PRECISION_RADIX FOR COLUMN RADIX INTEGER DEFAULT NULL
 DATETIME_PRECISION FOR COLUMN DATPRC INTEGER DEFAULT NULL

 COLUMN_TEXT FOR COLUMN LABELTEXT VARGRAPHIC(50) ALLOCATE(50) CCSID 1200 DEFAULT NULL
 SYSTEM_COLUMN_NAME FOR COLUMN SYS_CNAME CHAR(10) DEFAULT NULL
 SYSTEM_TABLE_NAME FOR COLUMN SYS_TNAME CHAR(10) NOT NULL
 SYSTEM_TABLE_SCHEMA FOR COLUMN SYS_DNAME CHAR(10) NOT NULL
 USER_DEFINED_TYPE_SCHEMA FOR COLUMN TYPESCHEMA VARCHAR(128) ALLOCATE(100) DEFAULT NULL
 USER_DEFINED_TYPE_NAME FOR COLUMN TYPENAME VARCHAR(128) ALLOCATE(100) DEFAULT NULL
 IS_IDENTITY FOR COLUMN "IDENTITY" VARCHAR(3) ALLOCATE(3) DEFAULT NULL
 IDENTITY_GENERATION FOR COLUMN "GENERATED" VARCHAR(10) ALLOCATE(10) DEFAULT NULL
 IDENTITY_START FOR COLUMN "START" DECIMAL(31 0) DEFAULT NULL
 IDENTITY_INCREMENT FOR COLUMN "INCREMENT" DECIMAL(31 0) DEFAULT NULL
 IDENTITY_MINIMUM FOR COLUMN "MINVALUE" DECIMAL(31 0) DEFAULT NULL
 IDENTITY_MAXIMUM FOR COLUMN "MAXVALUE" DECIMAL(31 0) DEFAULT NULL
 IDENTITY_CYCLE FOR COLUMN "CYCLE" VARCHAR(3) ALLOCATE(3) DEFAULT NULL
 IDENTITY_CACHE FOR COLUMN "CACHE" INTEGER DEFAULT NULL
 IDENTITY_ORDER FOR COLUMN "ORDER" VARCHAR(3) ALLOCATE(3) DEFAULT NULL
 COLUMN_EXPRESSION FOR COLUMN EXPRESSION DBCLOB(2097152) CCSID 1200 DEFAULT NULL
 HIDDEN CHAR(1) DEFAULT NULL
 HAS_FLDPROC FOR COLUMN FLDPROC CHAR(1) DEFAULT NULL
 INLINE_LENGTH FOR COLUMN "ALLOCATE" INTEGER DEFAULT NULL
 NORMALIZE CHAR(1) DEFAULT NULL
 DATALINK_LINK_CONTROL FOR COLUMN DL_LINKC CHAR(1) DEFAULT NULL
 DATALINK_INTEGRITY FOR COLUMN DL_INTEG CHAR(1) DEFAULT NULL
 DATALINK_READ_PERMISSION FOR COLUMN DL_READP CHAR(1) DEFAULT NULL
 DATALINK_WRITE_PERMISSION FOR COLUMN DL_WRITEP CHAR(1) DEFAULT NULL
 DATALINK_RECOVERY FOR COLUMN DL_RECOVER CHAR(1) DEFAULT NULL
 DATALINK_UNLINK_CONTROL FOR COLUMN DL_UNLINKC CHAR(1) DEFAULT NULL
 DDS_TYPE CHAR(1) DEFAULT NULL
 SECURE CHAR(1) DEFAULT NULL

6.5.2 QSYS2.SYSPARTITIONDISK and QSYS2.SYSPARTITIONINDEXDISK Views

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 SSD.

See the SQL Reference for more detail:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/db2/rbafzcatsyspdisk.htm>

Example 6-26 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;
```

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	AGDTDL	AGDTDL	-	4096

Figure 6-6 Results from Example 6-26 on page 158

Example 6-27 Return allocation information for DB2 indexes (i.e., 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;
```

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 6-7 Results from Example 6-27

6.5.3 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 6-28 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 ;
```

6.5.4 EARLIEST_POSSIBLE_RELEASE

IBM DB2 for i now provides a way to see the earliest IBM i release which could be used for any SQL statement or program.

The SQL statement level detail is available via Database Monitor.

The program level detail is available via the **QSYS2.SYSPROGRAMSTAT** and **QSYS2.SYSPACKAGESTAT** catalogs. In both cases, you need to capture the DBMON or rebuild the program after applying the latest DB Group PTFs.

Database Monitor and the QSYS2.SYSPROGRAMSTAT & 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 are capable of being deployed on earlier IBM i releases or whether they are using SQL functionality unique to IBM i 6.1 or 7.1.

This field only applies if the SQL statement is dynamic. (QQC12= 'D')

Possible values for QQC82:

' ' - The statement release level has not been 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

QSYS2.SYSPROGRAMSTAT and QSYS2.SYSPACKAGESTAT column name
EARLIEST_POSSIBLE_RELEASE.

System column name is MINRLS

6.5.5 SIGNAL support for native triggers

INSERT/UPDATE/DELETE SQL statements have been changed to recognize when system triggers have used the SIGNAL SQL statement to communicate failure detail with the application.

If the system trigger executes the SIGNAL statement and sends an escape message to its caller, the SQL INSERT/UPDATE/DELETE statement will fail with MSGSQL0438 (SQLCODE=-438) instead of MSGSQL0443.

The SQLSTATE, MSG, and other values within the SQL diagnostics area or SQLCA will contain the values passed into the SIGNAL statement.

For details see the Database Programming book:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/dbp/rbaforzahftrm.htm>

which contains recommendations for native trigger programs, which includes the following:

Signal an exception if an error occurs or is detected in the trigger program. If an error message is not signalled 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 which caused the trigger to be fired through the use of SQL.

The SIGNAL SQL statement does not signal an exception, so be sure to use QMHSNDPM() API to send an escape message, after executing the SIGNAL statement.

Example 6-29 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.....
```

6.5.6 Hierarchical queries using CONNECT BY clause

DB2 for i has had recursive query support since V5R4. Another recursive query technique called a hierarchical query has been added. This technique is a more concise method of representing a recursive query.

For details see IBM i 7.1 Information Center supporting material:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/sqlp/rbafyrecursivequeries.htm>

and

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/db2/rbafzhierquery.htm>

Example 6-30 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
```

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 6-8 Result of hierarchical query above

6.5.7 Additional parameter marker support (LAND, LOR, XOR, and TRANSLATE)

Prior to release 7.1, there were many restrictions on where a parameter marker was allowed in an SQL statement. Many of these restrictions were removed by general availability of 7.1.

The LAND, LOR, XOR, and TRANSLATE scalar functions have been enhanced by removing similar restrictions.

Example 6-31 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'
```

6.5.8 Support PROGRAM NAME on CREATE TRIGGER

Provides the ability to specify a short name for the created trigger program.

When this is not supplied, the database will determine the system name, which could lead to differences in the system name for the trigger program across different machines.

Example 6-32 Usage of PROGRAM NAME in CREATE TRIGGER definition

```
CREATE TRIGGER trg21 AFTER UPDATE OF c1 ON TR1
REFERENCING OLD AS o NEW AS n
```

```
FOR EACH ROW MODE DB2ROW
PROGRAM NAME trg21
BEGIN ATOMIC
INSERT INTO tr2 VALUES(default, o.c1, n.c1);
END
```

6.5.9 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, SQL_C_UTINYINT, Unsigned TINYINT, which represents a one-byte unsigned integer value with a range of values from 0-255.

Refer to the IBM i 7.1 SQL CLI documentation for details:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/cli/rzadpwhatsnew.htm>

6.5.10 CHGPFM and CHGLFM UNIT support

CHGPFM (Change Physical File Member) and CHGLFM (Change Logical File Member) commands can now be used to move an individual member to or from 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 will 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.

CHGPFM t1 **UNIT(*SSD)**

CHGLFM v1 **UNIT(*SSD)**

Note: If the user is using logical replication, you will need the PTFs on the target as well as the source system!

6.5.11 SYSTOOLS procedures

SYSTOOLS is a set of DB2 for IBM 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 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 best practice for customers who are interested in such tools would be to periodically review the contents of SYSTOOLS.

For more detail on this topic see IBM Infocenter:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/rzajq/rzajqsystools.htm>

6.5.12 Add client and server IP address and port number to QSYS2.TCPIP_INFO

Client/server identification, logging and other instrumentation is possible using SQL.

The QSYS.TCPIP_INFO view returns detailed information for the current connection.

Note: The SERVER_IP_ADDRESS and SERVER_IP_ADDRESS_TYPE are only available on IBM i 7.1.

If you issue following SQL command:

Example 6-33 Usage of QSYS2.TCPIP_INFO

```
SELECT * from QSYS2.TCPIP_INFO
```

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 6-9 Results of QSYS2.TCPIP_INFO

6.5.13 QSYS2.GROUP_PTF_INFO view

This view allows you to use SQL to retrieve IBM i Group PTF information for the server. Data returned is similar to the Work with PTF Groups (WRKPTFGRP) command.

On IBM i 7.1, the Technology Refresh (TR) level can be determined using this

Example 6-34 QSYS2.GROUP_PTF_INFO

QSYS2.GROUP_PTF_INFO view definition:

```
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)
```

Try to issue command:

```
select * from QSYS2.GROUP_PTF_INFO order by PTF_GROUP_LEVEL DESC;
```

COLLECTED_TIME	PTF_GROUP_NAME	PTF_GROUP_DESCRIPTION	PTF_GROUP_LEVEL
2011-09-30 11:04:19...	SF99710	...	CUMULATIVE PTF PACKAGE C1116710
2011-09-30 11:04:19...	SF99709	...	GROUP HIPER
2011-09-30 11:04:19...	SF99701	...	DB2 FOR IBM I
2011-09-30 11:04:19...	SF99362	...	710 BACKUP RECOVERY SOLUTIONS
2011-09-30 11:04:19...	SF99368	...	IBM HTTP SERVER FOR I
2011-09-30 11:04:19...	SF99701	...	DB2 FOR IBM I
2011-09-30 11:04:19...	SF99369	...	IBM I INTEGRATION WITH BLADECENTER AND SYSTEM X
2011-09-30 11:04:19...	SF99708	...	GROUP SECURITY
2011-09-30 11:04:19...	SF99363	...	WEBSPHERE APP SERVER V7.0
2011-09-30 11:04:19...	SF99572	...	JAVA
2011-09-30 11:04:19...	SF99617	...	DB2 WEB QUERY FOR IBM I V1.1.1
2011-09-30 11:04:19...	SF99359	...	WEBSPHERE MQ FOR I5/OS -V7.0.1
2011-09-30 11:04:19...	SF99627	...	7.1 ELECTRONIC SERVICES GROUP PTF
2011-09-30 11:04:19...	SF99364	...	WEBSPHERE APP SERVER V6.1
COLLECTED_TIME	PTF_GROUP_NAME	PTF_GROUP_DESCRIPTION	PTF_GROUP_LEVEL
2011-09-30 11:04:19...	SF99710	...	CUMULATIVE PTF PACKAGE C1116710
2011-09-30 11:04:19...	SF99709	...	GROUP HIPER
2011-09-30 11:04:19...	SF99701	...	DB2 FOR IBM I
2011-09-30 11:04:19...	SF99362	...	710 BACKUP RECOVERY SOLUTIONS
2011-09-30 11:04:19...	SF99368	...	IBM HTTP SERVER FOR I
2011-09-30 11:04:19...	SF99701	...	DB2 FOR IBM I
2011-09-30 11:04:19...	SF99369	...	IBM I INTEGRATION WITH BLADECENTER AND SYSTEM X
2011-09-30 11:04:19...	SF99708	...	GROUP SECURITY
2011-09-30 11:04:19...	SF99363	...	WEBSPHERE APP SERVER V7.0
2011-09-30 11:04:19...	SF99572	...	JAVA
2011-09-30 11:04:19...	SF99617	...	DB2 WEB QUERY FOR IBM I V1.1.1
2011-09-30 11:04:19...	SF99359	...	WEBSPHERE MQ FOR I5/OS -V7.0.1
2011-09-30 11:04:19...	SF99627	...	7.1 ELECTRONIC SERVICES GROUP PTF
2011-09-30 11:04:19...	SF99364	...	WEBSPHERE APP SERVER V6.1

Figure 6-10 Results from QSYS2.GROUP_PTF_INFO call

There are the PTF group status messages:

PTF group status values:

- ▶ **UNKNOWN** - The PTF 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 currently 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 already temporarily or permanently applied.
- ▶ **RELATED GROUP** - The PTF group does not have any PTFs for products 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 will cause 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 will be delivered on either physical or virtual media.

6.5.14 QSYS2.DUMP_SQL_CURSORS procedure

There is now a possibility how to capture the list of open cursors for a job.

QSYS2.DUMP_SQL_CURSORS(

Job_Name VARCHAR(28),
Library_Name CHAR(10),
Table_Name CHAR(10),
Output_Option integer)

Job_Name is a qualified job name or a special value of '*' to indicate the current job.

Library_Name is a optional library name for the procedure output.

Table_Name is a optional table name for the procedure output.

Output_Option has these choices:

1. Ignore Library_Name and Table_Name inputs and return a result set
2. Ignore Library_Name and Table_Name inputs and place the results in table QTEMP/SQL_CURSORS (no result set)
3. Place the results in table in Library_Name and Table_Name (no result set). If the table doesn't exist, the procedure will create it. If the table does exist, the results will be appended to the existing table.
4. 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 6-35 Possible invocation

```
-- populate QGP/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)
```

6.5.15 QIBM_SQL_NO_RLA_CANCEL environment variable

The SQL Cancel support includes logic to insure that DB2 for i programs are active on the stack of the initial thread in the target job for the cancel request. Applications which utilize 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 will be accessed. If the environment variable exists, the cancel request will not be honored when RLA is the only database work on-going within the initial thread at the time the cancel request is received.

The environment variable is SQL Cancel operational switch. The variable can be created at the job or system level. Creating it once at the system level will affect how SQL Cancel's are processed for all jobs.

```
ADDENVVAR ENVVAR(QIBM_SQL_NO_RLA_CANCEL)
```

or

```
ADDENVVAR ENVVAR(QIBM_SQL_NO_RLA_CANCEL) LEVEL(*SYS)
```

6.5.16 QSYS2.FIND_AND_CANCEL_QSQSRVR_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 deciding how to deal with an orphaned or runaway process. End job immediate is a hammer, where SQL cancel is more like a tap on the shoulder. Prior to this improvement, the SQL cancel support was only available to 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's safe to cancel the statement. When it's determined to be safe to cancel the statement, an SQL0952 escape message is sent, causing the statement to terminate.

If it isn't safe to end the SQL statement, or if there's no active SQL statement, the request to cancel is ignored. The caller of the cancel procedure will observe a successful return code which only indicates that 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 cancelled SQL statement.

Procedure Definition

```
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 6-36 CALL of 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 has *JOBCTL special authority.

Description

The procedure has a single input parameter, the qualified job name of the job that should be cancelled. The job name must be upper cased. If that job is executing an interruptible SQL statement or query, the statement will be cancelled. The application will most likely receive an SQLCODE = SQL0952 (-952). In some cases, the failure returned could be SQL0901 or the SQL0952 could contain an incorrect reason code.

This procedure takes advantage of the same cancel technology 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 will fail with a descriptive SQL0443 failure if the target job isn't found.

The procedure will fail with SQL0443 and SQL0552 if the caller doesn't have *JOBCTL user special authority.

Commitment Control

When the target application is running without commitment control (i.e. COMMIT = *NONE or *NC), the cancelled SQL statement will terminate without rolling back the partial results of the statement. If the cancelled statement is a query, the query merely ends. However, if the cancelled statement was a long-running INSERT, UPDATE or DELETE SQL statement, the changes made prior to cancellation remain intact.

If the target application is using transaction management, the SQL statement will be running under the umbrella of a transaction save point level. When those same long running INSERT, UPDATE or DELETE SQL statements are cancelled, the changes made prior to cancellation are rolled back.

In both cases, the application receives control back with an indication that the SQL statement failed. It's up to the application to determine the next action.

Useful Tool

The QSYS2.CANCEL_SQL() provides a useful tool to database administrators for IBM i machines. Once you have the latest DB Group PTF installed, you can start calling this procedure to stop long-running or expensive SQL statements. Leave the hammer in the toolbox and try calling QSYS2.CANCEL_SQL() instead.

QSYS2.FIND_AND_CANCEL_QSQRVR_SQL procedure

The QSYS2.FIND_AND_CANCEL_QSQRVR_SQL() procedure uses the QSYS2.FIND_QSQRVR_JOBS and QSYS2.CANCEL_SQL() procedures derive the set of jobs which have active SQL activity, given a target application job. Each job found is made a target of an SQL cancel request.

Example 6-37 Using QSYS2.FIND_AND_CANCEL_QSQRVR_SQL procedure

```
CALL (QSYS2.FIND_AND_CANCEL_QSQRVR_SQL('564321/APPUSER/APPJOBNAME'));
```

6.5.17 QSYS2.FIND_QSQRVR_JOBS() procedure

Anyone responsible for administering, tuning or explaining the SQL Server Mode (e.g. QSQRVR jobs) activity might find the QSYS2.FIND_QSQRVR_JOBS() procedure a useful tool. This procedure has been added to QSYS2 after application of PTFs (see the Service Information section for details). The procedure is passed a single parameter, 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 QSQRVR 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: (1) Summary information and (2) Detailed SQL Server Mode job information.

How is this procedure useful? When you have an important application instance (job) that uses QSQRVR jobs, it can be quite 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 QSQRVR job that is consuming a lot of CPU or holding onto object locks? The FIND_QSQRVR_JOBS() procedure provides some of these answers by tying together the application and its SQL Server Mode job use.

Example 6-38 Invocation of QSYS2.FIND_QSQRVR_JOBS

```
call QSYS2.FIND_QSQRVR_JOBS('566463/EBERHARD/QPOZSPWP ')
```

Procedure definition:

```
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 invoke QSYS2.FIND_QSQRVR_JOBS the user needs *JOBCTL special authority.

On IBM i 7.1, to invoke QSYS2.FIND_QSQRVR_JOBS the user needs *JOBCTL special authority, QIBM_DB_SQLADM Function usage or QIBM_DB_SYSMON Function usage.

Otherwise you will get:

```
call QSYS2.FIND_QSQRVR_JOBS('650261/SCOTT/ QP0ZSPWP')
```

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 STRSQL (Start SQL Interactive Session) or elsewhere, the user needs to query the temporary tables to see the data.

Example 6-39 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;
```

6.5.18 SQL Server Mixed Mode for Batch Processing

SQL Server Mode users need to be able to fire SQL trigger programs from the Server Mode client job, via Native database I/O operations and have any SQL statements within the trigger program execute within the client job instead of being rerouted to a QSQRVR server job.

SQL Server mode has been extended to allow an application to direct the database to execute SQL statements within an SQL trigger program within the SQL Server mode client job instead of rerouting the SQL to a QSQRVR server job.

The change only affects SQL triggers fired via Native database I/O operations.

To enable the new function, an environment variable must exist before any SQL statements are executed within the client job. An easy way to deploy the environment variable would be

to define it at the system level.

```
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. Once 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 execute 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 will serialize the execution of the triggers. Only one trigger would be allowed to execute at a time.

The solution does not apply to native triggers - ADDPFTRG (Add Physical File Trigger) - built over programs which use SQL. The solution does not include SQL triggers which call procedures, fire User Defined Functions or cause nested triggers to execute.

6.5.19 QDBRPLAY() API - Disable or Enable Constraints option

A new option has been 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 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 new option is defined as follows

Disable constraints. 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.

- ▶ **0** - Do not disable constraints.
- ▶ **1** - Disable constraints.

6.5.20 SQL0901 logger education

SQL0901 is a message which is usually issued by DB2 for i when SQLCODE = -901 indicates that an unexpected error was encountered. Sometimes (for instance when the job has no joblog) it is very difficult to find the cause of such failure. That's why just before SQL0901 message is issued, DB2 for i creates a new record for this incident at QRECOVERY.QSQ901S table. A unique instance of job will log first three SQL0901 failures. Any subsequent failures for the job are not logged since 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 definition:

1. **SERVERNAME - VARCHAR(18)** - The server name.
2. **FAILTIME - TIMESTAMP** - Time when the failure occurred.
3. **FAILRSN - INTEGER(9)** - The unique failure reason that appeared in the SQL0901 message. This reason code is usually necessary for IBM service and is not documented externally.

4. **CURUSER - VARCHAR(18)** - The user who encountered the SQL0901 failure.
5. **JOBNAME - CHAR(28)** - The qualified job name which encountered the SQL0901 failure.
6. **MSGS - VARCHAR(3000)** -

And others to help you to identify the cause of the failure.

Try:

```
SELECT * FROM qsys2.syscolumns WHERE TABLE_SCHEMA = 'QRECOVERY' and
TABLE_NAME = 'QSQ901S' ORDER BY ORDINAL_POSITION;
```

and/or

```
SELECT * FROM qsys2.syscolumns2 WHERE TABLE_SCHEMA = 'QRECOVERY'
and TABLE_NAME = 'QSQ901S' ORDER BY ORDINAL_POSITION;
```

to get the information on contents of QRECOVERY.QSQ901S table.

The records in the QRECOVERY.QSQ901S table very likely show the internal failures inside DB2 for i and data from this table should be used when reporting problem to IBM. This will help with searching for PTFs for DB2 for i problems.

6.6 DB2 security enhancements

This chapter shows some DB2 for i enhancements related to security.

6.6.1 Query Manager profile auditing

You can now audit changes made to a Query Manager profile if auditing is enabled for AUDLVL(*SECURITY). A new **journal entry type of X2** will contain the old and new Query Manager profile values. An outfile is not provided for this journal entry. Instead view QSYS2.SQLQMPProfilesAudit can be queried.

Example 6-40 Creating a permanent table that contains the new journal entries

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

There are two main parts to this solution:

Part 1

A new journal entry (X2) would be sent to the QAUDJRN any time a QM (Query manager) profile changes. The journal entry would have the before and after profile information so we could tell exactly what changed in the profile.

Part 2

Normally, the way audit journal entries are externalized is that each has an associated output file. A particular audit entry type is output into that output file which is tailored to the information in that journal entry. However, output files are expensive because of language translation costs and are much more difficult and expensive to PTF. Also, the Query Manager

profile changes each release as IBM adds more SQL statements that one can authorize. That's why IBM has no plans to create an output file that maps this journal entry.

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 who and when made the change:

- ▶ Journal entry timestamp
- ▶ Current user
- ▶ Job name, job user, job number
- ▶ Thread

Most of the values stored in the QM (Query Manager) profile only have two possible values. For example the values for 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 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

6.6.2 Add SECURE column to QSYS2.SYSCOLUMNS2

QSYS2.SYSCOLUMNS2 is a view based on a table function that will return additional information not available in **QSYS2.SYSCOLUMNS** (such as the allocated length of a varying length column). For details see 6.5.1, "QSYS2.SYSCOLUMNS2 view" on page 156.

Since it is based on a table function, it will typically return results faster if a specific table is specified when querying it.

```
SELECT * FROM qsys2.syscolumns2
WHERE system_table_schema = 'MJATST' and system_table_name = 'T1';
```

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

SECURE column values:

NULL - This column has not been the target of a call to SYSPROC.SET_COLUMN_ATTRIBUTE.

'0' - This column does not contain data that needs to be secured in a database monitor or plan cache. This column was target of a call to SYSPROC.SET_COLUMN_ATTRIBUTE with parameter 'SECURE NO'.

'1' - This column contains data that needs to be secured in a database monitor or plan cache. This column was target of a call to SYSPROC.SET_COLUMN_ATTRIBUTE with parameter 'SECURE YES'.

Example 6-41 Values with 'SECURE YES'

```
call qsys.create_sql_sample('PRODLIB');

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' ;
```

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

Figure 6-11 Results of the previous SELECT command showing SECURE = '1' for 'SALARY'

What the database performance analyst will see for this query:

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

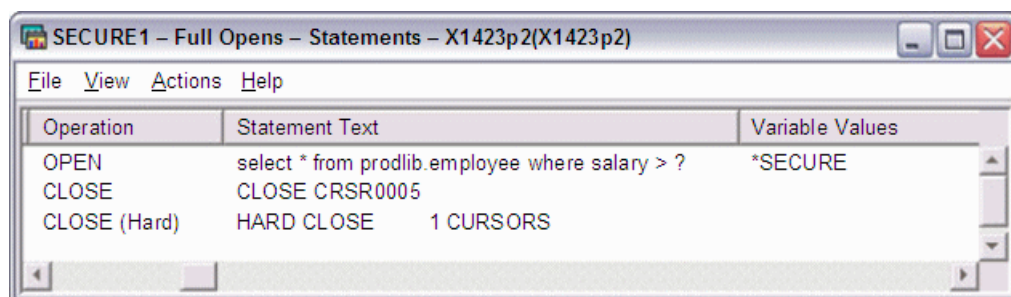


Figure 6-12 Results from Performance Analysis

6.6.3 QIBM_DB_SQLADM and QIBM_DB_SYSMON (*JOBCTL special authority no longer required)

Prior to this enhancement the administrator of the IBM i system had to allow *JOBCTL special authority for DB2 administrators and DB2 performance analysts. It was not ideal because the *JOBCTL special authority allows also possibility to access work management, work with jobs, work with all spooled files and so on. This enhancement is based on two functions - QIBM_DB_SQLADM and QIBM_DB_SQLMON. The first is for DB2 administrators, the second one is a subset of QIBM_DB_SQLMON which allows performance monitoring. Usage of these new functions allows also settings where even people with *ALLOBJ special

authority cannot do the DB2 administration and/or DB2 monitoring tasks. Also group profile can be specified in these functions setting. If user profile is associated with several supplemental group profiles the access is granted if at least one of these group profiles is set to *ALLOW in particular function. Adopted authority from program owner profiles has no effect on allowing access to DB2 administration and DB2 monitoring. Access is always granted according user profile under which program with adopted authority runs.

6.6.4 STRDBMON Over a View (Enables faster SQL statement auditing)

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

Utilize input only columns to only capture a subset of the monitor data into an underlying table.

In DB2 for i, the database performance monitor table has 276 columns. Assume an auditing application is only interested in collecting 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).

Example 6-42 View mechanism One

Step 1: Create a table with the 23 columns that are desired (QAQQDBMN is the model monitor file shipped with DB2 for i):

```
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, QQSMINT2,
  QQRID, QQI5, QQJFLD
FROM QAQQDBMN) WITH NO DATA;
```

Step 2: Next create a view that has 276 columns that match the database monitor table columns. Only the 23 desired columns will be input/output columns, the others are input only columns (those that are just CAST as NULL). The columns must have exactly the same attributes and be in exactly the same order as in the base database monitor table:

```
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 QSYS, 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 QQDEF ,
CAST(NULL AS DECIMAL(15,0)) AS QQSTN , CAST(NULL AS DECIMAL(15,0)) AS QQDTN ,
CAST(NULL AS DECIMAL(15,0)) AS QQDTL , 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 QQINTO1 , CAST(NULL AS INTEGER) AS QQINTO2 ,
 CAST(NULL AS INTEGER) AS QQINTO3 , CAST(NULL AS INTEGER) AS QQINTO4 ,
 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 QQQDBMN;

```

Step 3: Lastly, start the database monitor using the view.

```
STRDBMON mjatst.loopview4;
```

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

Utilize an INSTEAD OF trigger on the view to immediately process a row of monitor data without storing ANY data.

Example 6-43 View mechanism Two

Step 1: Create a view using the same technique as above or it can be simplified since in this example, no data will actually be inserted into a table.

```

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

```

Step 2: Next 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;

```

Step 3: Lastly, start the database monitor using the view.

```
STRDBMON mjatst/loopview5;
```

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

Comparison of the View Mechanisms and the Traditional Database Monitor

The following charts provide a comparison of the elapsed time and the storage consumed using the traditional database monitor and the two view techniques.




Figure 6-13 Comparison of methods One and Two

Note that in the case of the view with an instead of trigger, the elapsed time and the amount of storage consumed is completely under the control of the instead of trigger. For example, in the above the example, the instead of trigger simply sent the data to another system. This takes some CPU, but no persistent storage is used on the system that is being monitored.

6.6.5 SECURE Columns (Monitor and Plan Cache masking of variable values)

This procedure can be used to obscure host variable values in SQL Performance Monitors (aka Database Monitors) and from the SQL Plan Cache.

When an SQL statement refers to any column which has the SECURE attribute set to YES (6.6.2, “Add SECURE column to QSYS2.SYSCOLUMNS2” on page 172), all host variable values will appear as “*SECURE” when examined from the monitor and plan cache, unless the user is the QSECOFR user.

Example 6-44 The procedure definition

```
CREATE PROCEDURE SYSPROC.SET_COLUMN_ATTRIBUTE (
  Table_Schema VARCHAR(10),
  Table_Name VARCHAR(10),
  Column_Name VARCHAR(10),
  Attribute VARCHAR(10) )
LANGUAGE C PARAMETER STYLE SQL
MODIFIES SQL DATA
SPECIFIC QDBCATTR
EXTERNAL NAME 'QSYS/QDBSSUDF2(QDBCATTR)';

-- You mayt want to grant EXECUTE to public
GRANT EXECUTE ON PROCEDURE SYSPROC.SET_COLUMN_ATTRIBUTE TO PUBLIC;
```


A description of the parameters follow:

- ▶ **Table_Schema - VARCHAR(10)** - This will contain the system name of a table's schema (library).
- ▶ **Table_Name - VARCHAR(10)** - This will contain the system name of a table (file).
- ▶ **Column_Name - VARCHAR(10)** This will contain the system column name that is secured.
- ▶ **Attribute - VARCHAR(10)** This will contain the attribute being set for the column.

The valid values are:

SECURE NO - This column does not contain data that needs to be secured in a database monitor or plan cache.

SECURE YES - This column contains data that needs to be secured in a database monitor or plan cache.

All variable values for any query that references this column will not be visible in a database monitor or plan cache unless the security officer has started the database monitor or the security officer is accessing the plan cache.

Issue an SQL CALL to secure column CCNBR in table LIB1.Orders:

```
CALL SYSPROC.SET_COLUMN_ATTRIBUTE('LIB1', 'ORDERS', 'CCNBR', 'SECURE YES');
```

6.6.6 Add QDDMDRDASERVER server authentication entry special value

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

The dilemma faced when choosing to enforce password validation on the servers is that every user who needs to connect to the server needs admin work done on their behalf using **ADDSVRAUTE** (Add Server Authentication Entry), for every target server or require every user to supply a user id and password on each **CONNECT** statement.

This improvement allows to use special value **QDDMDRDASERVER** has been added to the **ADDSVRAUTE** (Add Server Authentication Entry) command **SERVER** parameter for DDM & DRDA connections. This special value allows an administrator to configure a user to work with all possible DDM or DRDA connections to any system in the TCP/IP network via a common user id and password. Once 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 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.

What are server authentication entries?

A server authentication entry is a way to define a user id & password to send via 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 **ADDSVRAUTE** (Add Server Authentication Entry) 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.

Usage details:

The special value QDDMSERVER already exists which allows the ability for 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 specifying the system name exists, and a user ID and password are **not passed** in on CONNECT statement, the user ID and password associated with the server authentication entry will take precedence over the server authentication entry for QDDMDRDASERVER.

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

For RDB DDM file connection requests, the server authentication entry specifying the system name will take precedence over the server authentication entry for QDDMDRDASERVER. For non-RDB DDM file connection requests, the server authentication entry QDDMSERVER will take precedence over the server authentication entry for QDDMDRDASERVER.

Example 6-45 Environment: Three systems (SYSA, SYSB, SYSC)

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

On SYSA:

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

At this point you have two connections made with shared 'youruid' and password 'yourpwd'.

Example 6-46 Environment: Three systems (SYSA, SYSB, SYSC)

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

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
```

At this point you have two connections made with user id 'yourotheruid' and password 'yourotherpwd'. This occurs because server authentication entries specifying the real system name take precedence over server authentication entries specifying QDDMDRDASERVER.

Example 6-47 Environment: Three systems (SYSA, SYSB, SYSC)

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

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
```

At this point you have two connections. The connection to SYSB would be made with user id 'testuserid' and password 'testpassword'. This occurs because specifying user id and password on a CONNECT statement takes precedence over server authentication entries. The connection to SYSC would be made with user id 'youruid' and password 'yourpwd' since it will use the QDDMDRDASERVER authentication entry when no other server authentication entry exists specifying the system name.

Example 6-48 Environment: Three systems (SYSA, SYSB, SYSC)

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

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)

SBMRMTCMD CMD('DSPLIB YOURLIB') DDMFILE(QTEMP/DDMF)
```

The connection to SYSB would be made with user id 'youruid2' and password 'yourpwd2'. This occurs because non-RDB DDM files will use QDDMSERVER for user id and password at connection time. If QDDMSERVER didn't exist, QDDMDRDASERVER would be used.

```
SBMRMTCMD CMD('DSPLIB YOURLIB') DDMFILE(QTEMP/DDMF2)
```

The connection to SYSB would be made with user id 'youruid' and password 'yourpwd'. This occurs because RDB DDM files will use QDDMDRDASERVER for user id and password at connection time if a server authentication entry doesn't exist specifying the system name.

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

The connection to SYSC would be made with user id 'yourotheruid' and password 'yourotherpwd'. This occurs because RDB DDM files will not use QDDMDRDASERVER for user id and password at connection time if a server authentication entry exists specifying the system name.

6.6.7 Enhanced audit capability of RUNSQLSTM or STRSQL

The database monitor (started by STRDBMON (Start Database Monitor) has been enhanced to be able to track RUNSQLSTM (Run SQL Statement) or STRSQL (Start SQL Interactive Session) activity connected to a specific user profile, application, program ID, accounting code or client workstation name.

There are five registers which appear on STRDBMON (Start Database Monitor) command:

- ▶ **CLIENT PROGRAMID** = 'STRSQL' (parameter **FTRCLTPGM**)
- ▶ **CLIENT APPLNAME** = 'START SQL INTERACTIVE SESSION' (parameter **FTRCLTAPP**)
- ▶ **CLIENT USERID** = The current user's name (parameter **FTRCLTUSR**)
- ▶ **CLIENT WRKSTNNAME** = The DB2 for i database name (parameter **FTRCLTWS**)
- ▶ **CLIENT ACCTNG** = The current user's accounting code
(i.e. the ACGCDE parameter on CRTUSRPRF (Create User Profiles) and CHGUSRPRF (Change User Profiles) commands)

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

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 6-14 STRDBMON - new filter parameters

6.6.8 Add FTRSQLCODE parameter to STRDBMON

The STRDBMON (Start Database Monitor) command has been enhanced to have a pre-filter parameter for the result of SQL statements on the FTRSQLCODE parameter.

There are several filter by SQLCODE (FTRSQLCODE) 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>**

Example 6-50 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)
FTRSQLCODE(-913)
```

For SQLSTATES and SQLCODEs for DB2 for IBM i 7.1 see this link to infocenter:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=%2Frzala%2Frzalaccl.htm>

6.6.9 Extend FTRINTNETA pre-filter on STRDBMON command to work with non-Database clients

Prior to this enhancement the STRDBMON (Start Database Monitor) command parameter Filter by internet address (FTRINTNETA) support was restricted to work against Database initiated connections.

This enhancement allows Filter by internet address to work as a Database monitor pre-filter against many non-Database initiated connections.

Note: TELNET connections are not recognized by this support.

6.6.10 Extend STRDBMON to support FTRUSER against group profiles

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

When determining whether the current user's SQL should be captured in the SQL Performance Monitor (aka Database Monitor) output, it is determined whether the user is a member of the group.

Wildcard group profile names are allowed.

For example, if the customer specifies FTRUSER(ADMIN*) and both ADMINGRP and ADMINGRP2 are group profiles, any SQL executed by users in either group will be captured.

System i Navigator's SQL Performance Monitor interface for "User" can be used to specify the group profile.

6.7 DB2 Availability and Recovery Enhancements

In this section we will cover the availability and recovery enhancements to DB2 for i.

6.7.1 Preserve SQE Cache Size Across IPLs

DB2 for i on IBM i 6.1 and 7.1 have been enhanced to preserve the SQL Plan Cache size across IPLs and slip installs of the operating system.

A scratch install of the OS will reset the SQL plan cache size to the default size.

Customers can explicitly increase the size of the SQE plan cache to allow more plans to be saved in the plan cache. This can improve performance for customers that have a large number of unique queries.

During an IPL, the SQL plan cache is deleted and recreated. Prior to this enhancement, when the plan cache was recreated, it was recreated to the default size of 512 megabytes even if the customer had explicitly specified a larger plan cache. Now, the size specified by the customer will be preserved and used when it is recreated during any subsequent IPL.

Once the latest DB2 Group PTFs are installed, the customer will need to change the plan cache size one more time (even if its 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; the desired SQL plan cache size in megabytes. If the value passed in is zero, the plan cache is reset to its default value.

```
CREATE PROCEDURE QSYS2.CHANGE_PLAN_CACHE_SIZE ( IN SIZE_IN_MEG INTEGER )
```

Example 6-51 Usage of QSYS2.CHANGE_PLAN_USAGE procedure

```
CALL qsys2.change_plan_cache_size(1024);
```

It also possible to get information about Plan Cache properties using the following procedure:

Example 6-52 QSYS2.DUMP_PLAN_CACHE_PROPERTIES

```
call qsys2.dump_plan_cache_properties('QGPL', 'SQECACHE1');
```

```
select heading,value from qgpl.sqecache1  
where heading like '%Time%';
```

6.7.2 Pre check physical file size during restore

Prior to this enhancement, when performing 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 could use all available storage and cause a system crash. A pre check will now be performed to ensure that enough storage exists on the target ASP.

6.7.3 Prevent 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 till the index is next opened.

This improves the performance of row change operations while the index maintenance is delayed.

Prior to this enhancement, if a job was canceled it was opening a delayed maintenance index, the entire index would be invalidated and have to be rebuilt from scratch. On large indexes this can be a lengthy operation. This enhancement will ensure that in this case, the cancel will not cause the entire index from being invalidated.

6.8 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 additional optional products 5733-QU3, and 5733-QU4. The latest version of these products now delivered is 1.1.2. IBM has also 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.

6.8.1 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 with drag and drop interface for reporting functions. Build dashboards, integrate data with new or existing spreadsheets, or choose to deliver reports in an analytical form that allows the user to slide and dice through the data interactively. DB2 Web Query for i can either be ordered in ala-carte, modular approach, or choose a **Standard Edition** for a pre-defined package of components that provide a very robust, recommended set of components.

Here are the modules:

1. **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 into the web based Report Assistant tool. DB2 Web Query is designed in a way that allows you to leave your data in DB2 for i and leverage all security and availability features of the IBM i platform. This base module is priced per processor tier and included a minimum number of users. Most clients will be able to upgrade into the base product at no charge if they own Query/400 and are current on IBM i software.
2. **Active Technologies** can combine the results of queries and create HTML reports which can be available to users without any needing to be connected to the DB2 for i server. Designed for users “on the go”, the reports contain a query results but the data can be viewed in a variety of ways from a browser, including functions to sort, filter the data by different criteria, and a calculated field, and chart information for visual impact.
3. The **DB2 Web Query Developer Workbench** feature is an open and intuitive environment that allows for rapid development of more customized Web-based reports and meta data. It includes an HTML layout painter for building dashboards, combining multiple reports onto a single view. It is a critical component for developing and managing meta data used to shield the complexities of the database from report authors and end users.
4. **OLAP Module** takes reporting to a whole new level by providing and interactive visualization of the data that allows end users to drill down or slice and dice for 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 pre-requisite to using an OLAP report.
5. **Run Time User Enablement.** Without this module each individual user needs to be licensed to the base product. With the Run Time User Enablement feature, one or more user licenses can now be defined as a group of run-time only users. If you are familiar with Group Profiles, this is similar in concept. Each member of the group is able to run reports concurrently, and each group can contain thousands of users, thereby providing an almost unlimited run-time user licensing model. Users defined as a run time user do not have the ability to create or edit report definitions, but have full functionality in executing reports, including parametrized dashboards, OLAP reports, and more.
6. **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 below noted adapter feature). End users with appropriate authority can invoke 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.
7. **DB2 Web Query Adapter for Microsoft SQL Server** provides connectivity from DB2 Web Query to remote SQL Server databases. Many IBM i customers have the majority 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 desired, and provides seamless access to this data for report authors and end users.
8. **The DB2 Web Query Adapter for Oracle JD Edwards** allows DB2 Web Query to report on data 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 meta data and security layers.

6.8.2 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, or 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 e-mail distribution list.

6.8.3 DB2 Web Query Software Developer Kit - 5733-QU4

This product is targeted toward application developers. The DB2 Web Query SDK provides a set of Web Services that allow you 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 execute DB2 Web Query reports and more. Simplify the programming effort by leveraging 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 developing with the SDK, the DB2 Web Query BASE product is required and Developer Workbench feature is highly recommended. Deployment (Run Time) environments will require the BASE product and the Run Time User Enablement feature of DB2 Web Query.

6.8.4 DB2 Web Query for i Standard Edition

IBM introduces the DB2 Web Query, Standard Edition. This new ordering vehicle simplifies the decision process of which features to get by including the most popular features into a single package. You can order additional features, such as the SQL Server or JDE adapter to the Standard Edition if desired, but the intent is to provide 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)
- ▶ 4 Additional User Licenses (that can be used as individual users or as a group of run time users)
- ▶ 1 PC license of Developer Workbench
- ▶ Active Technologies
- ▶ OLAP
- ▶ Runtime User Enablement
- ▶ Spreadsheet Client
- ▶ DB2 Web Query Report Broker
- ▶ DB2 Web Query Software Development Kit

6.9 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 has been enhanced to include additional 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 will 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, and/or an object set for all stream files in directory '/home/alice/text_data'.

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 have been 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.

For more details on this topic and generally OmniFind for i product, refer to the article **“Searching Spool Files and IFS Stream Files”** at developerWorks:

<https://www.ibm.com/developerworks/ibmi/library/i-omnifind/omnifind.html>

6.9.1 OmniFind Updates Page

More details plus additional technical resources you will find at:

<https://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/IBM%20i%20Technology%20Updates/page/OmniFind%20for%20IBM%20i?lang=en>



Performance tools

The IBM i operating system and its licensed products include a variety of applications for collecting, analyzing, and reporting performance data generically called *performance tools*¹ 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, modelling, aggregating and analyzing data. Although these tools still support a fixed function terminal (i.e. 5250 or a telnet client), most of them are using either a browser based or a specific client software.

These performance tools actually cover three areas of systems management:

- ▶ Performance and capacity management
- ▶ Diagnostic analysis of processes within a system context
- ▶ Specific diagnostics for the IBM support organization

With these functions you can set up practices for monitoring and managing your system's performance to ensure your IT infrastructure is aligned with the changing demands of your business.

In this chapter, we describe how both the Collection Services and the Analysis Tooling have changed.²

¹ Not to be confused with the licensed product 5770-PT1 Performance Tools

² In order to take advantage of all the Performance Tools enhancements discussed in this chapter the system will need to have the latest levels of PTFs installed

7.1 Introduction to performance tooling

We start with a brief review 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 the IBM i³:

- ▶ **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 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. In this chapter, we will refer 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 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.

7.1.2 Analyzing data

There are several tools and options to view, understand, and analyze data from the collection services. These are either built in 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:

- ▶ The Display Performance Data graphical user interface allows you to view performance data generated by the collection services through a 5250 interface.
- ▶ The Reports organize collection services performance data and trace data in a logical and useful format. In this release, there are no changes to the reports

³ All of the functions that allow the configure how to collect data, to start and end data collection and to manage the collection objects are part of the operating system.

- ▶ The Performance Tools Graphics function allows you 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.
- ▶ The Investigate data of the IBM System Director Navigator for IBM i allows you to perform deep analyses of the data from within a browser based interface. For more details, see 18.11, “Performance enhancements” on page 563.
- ▶ System i Navigator provides not only an interface to view and analyze data, it also gives you the possibility to define monitoring functions for resource use by individual, all, or selected jobs.
- ▶ The Job Watcher function of the IBM Systems Director Navigator for i Performance interface is also included in Performance Tools as an option.
- ▶ The Disk Watcher function of the IBM Systems Director 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 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 central processing unit (CPU) 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

In this section, we discuss the changes in the collection profiles and in the data kept when collecting performance data. These items have been adapted to allow an in-depth analysis of the activities on the system to set up the best practices to analyze the behavior of applications on a system, to provide valuable information about how the workload is built up and on 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.

The IBM i performance tooling allows you to keep track of how the performance and capacity objectives are met in an end-to-end approach. The detailed information these collection services generate allows you to shape not only the workload management of a system, but also to analyze issues that can be due to a large number of 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 *ENHCPCLN⁴ profiles) added to them:

- ▶ The Standard profile now gets additionally:
 - External storage, containing non-standardized data for disk units externally attached to an IBM i partition (see QAPMXSTGV).
 - System internal data, containing internal data for the system (see QAPMSYSINT).
 - Removable storage, containing data about removable storage devices connected to the system, more specifically tape device data.
- ▶ The Standard Plus profile gets information about the logical partition, containing performance data 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 customizing the categories, you can now also specify the intervals. This means that you can have different categories of data collected at different intervals. Figure 7-1 is an example of how you can set this.

Collection profile

☐ Select predefined collection profile
Standard plus protocol

☒ Customize collection profile

Available categories

Select	Category
<input type="radio"/>	Memory pool
<input type="radio"/>	Memory pool tuning
<input type="radio"/>	Hardware configuration
<input type="radio"/>	Subsystem
<input type="radio"/>	System CPU
<input type="radio"/>	System-level data
<input type="radio"/>	Jobs (MI tasks and threads)
<input type="radio"/>	Jobs (operating system)
<input type="radio"/>	SNADS
<input type="radio"/>	Disk storage

Page 1 of 3 | 1 | Go | Rows: 10 | Total: 29 | Selected: 0

Frequency to collect

☐ 15 seconds
☐ 1 minutes
☒ Use default collection interval

Categories to collect

Select	Category	Frequency
<input checked="" type="radio"/>	System bus	Default interval
<input type="radio"/>	IBM HTTP Server (powered by Apache)	

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 data base file contains data for external system buses.

Support for a particular bus as well as what metrics are supported for that bus is dependent on the type of bus, how that bus connects to the system, and if assigned to the partition.

⁴ IBM no longer recommends using this profile

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 had activity within the collecting partition

For newer technologies, the hardware may provide additional instrumentation. Hardware metrics represent bus usage by all partitions. Consequently, the collecting partition must be authorized to obtain this data (reference 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 Contents of the QAPMBUS table

Column	Description
INTNUM	Interval number: The nth sample database interval based on the start time in the create performance data (CRTPFRTDA) 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 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.
BUQSAS	Queue space available sent
BUTYPE	Bus type. Supported values: 'S' = SPD bus (no longer supported) 'P' = PCI bus 'V' = Virtual bus

Column	Description
BUCAT	Bus category. This field indicates if this bus record has some special characteristics, which may 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 <i>(Note: the following may or may not be used depending on what happens with switches)</i> X'02' = This record represents a switch. The data reported is the sum of all buses under the switch. The bus number reported is the first bus under the switch
BUHUB	Hub number. If this bus is associated with an I/O hub, this is the number of that hub. (Note: an I/O hub may be imbedded in the backplane).
BUMAXRATE	Maximum byte rate. When available from hardware, this is the estimated maximum rate that data may be both sent and received in bytes per second through the hardware port.
BUCBSND	Command bytes sent. When available from hardware, this is the number of command bytes sent through the hardware port.
BUDBSBD	Data bytes sent. When available from hardware, this is the number of data bytes sent through the hardware port.
BUCBRCV	Command bytes received. When available from hardware, this is the number of command bytes received through the hardware port.
BUDBRCV	Data bytes received. When available from hardware, this is the number of data bytes received through the hardware port.

QAPMBUSINT

This file contains data for internal system buses.

The metrics supported are dependent on the instrumentation within the hardware chips. Support for a particular bus is dependent on both the type of bus as well as the chip family.

There may be one or more records each interval for a reported bus. The number of records as well as the metrics supported are dependent on both bus type and chip type.

These metrics are instrumented in the hardware and represent bus usage by all partitions. Consequently, the collecting partition must be authorized to obtain this data (reference the "Allow performance information collection" option within the HMC partition configuration).

Table 7-2 Contents of the QAPMBUSINT table

Column	Description
BUNBR	The hardware assigned number associated with the bus or hub.
BUTYPE	Bus type. The supported bus types are <ul style="list-style-type: none"> ▶ 4 - 12X loop ▶ 6 - I/O hub (may be imbedded in backplane)

Column	Description
BUDFMT	Bus data format. This field is being provided to help 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: <ul style="list-style-type: none"> ▶ Type 4: Port identifier. One record will be present for each supported port. <ul style="list-style-type: none"> – 0 = even port – 1 = odd port ▶ Type 6: Category
BUPKTSND	Packets sent. Note: not supported for type 6.
BUPKTRCV	Packets received. Note: not supported for type 6
BUBYTESND	Data bytes sent. Note: not supported for type 6.
BUBYTERCV	Data Bytes received. Note: 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: <ul style="list-style-type: none"> ▶ Type 4: Reserved ▶ Yype 6: Command bytes sent
BUDATA2	The meaning of this field depends on the type (BUTYPE) field: <ul style="list-style-type: none"> ▶ Type 4: Reserved ▶ Yype 6: Command bytes received

QAPMDISK

You will find new entries in this table, detailing per path the total read and write operations and World Wide Node names for external disks. Table 7-3 shows the columns added.

Table 7-3 New columns in the QAPMDISK table

Column	Description
DSPTROP	The path total read operations reports the number of read requests 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 received by internal machine functions, which is not the same as the device write operations reported in the DSDWOP field.
DSWWNN	The World wide node name is unique identifier representing the external storage subsystem the disk belongs to. This is null for non-external disks.

QAPMDISKRB

Up to release 6.1, the QAPMDISK table contained a detailed set of data about the performance of the disk unit. This design has been kept, but complemented with a new table (QAPMDISKRB), containing only the disk operations per interval. At the same time, it increases the number of bucket definition boundaries reported from six 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 QAPMDISKRB 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 affinity⁵ status changes of a thread or process.

Table 7-5 lists the new columns in the QAPMJOBMI table:

Table 7-5 New columns in the QAPMJOBMI table

Column	Description
JBNFHN	An identifier of a resource affinity domain this software thread or task is associated with. Thread or task is associated with resource affinity domain at a create time, but operating system can decide to move it to another resource affinity domain at a later time.

⁵ On the 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 executes and the “nodal” location where that task will find its data is called “Memory Resource Affinity”.

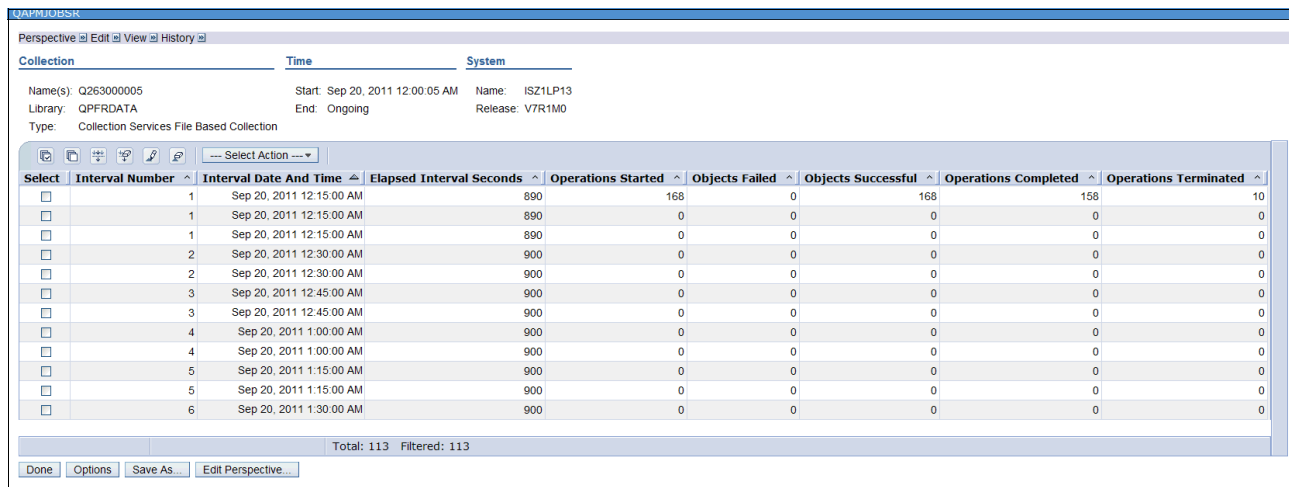
Column	Description
JBNFLVL	Resource affinity level specifies the relative strength of the binding between a thread and the internal machine resources with which it has affinity (processors and main storage). The strength is expressed as: <ul style="list-style-type: none"> ▶ X'00' = Processor normal, 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	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	Amount of CPU time used by the thread on the resource affinity domain this thread is associated with. The time is reported in internal model-independent units. This is called the <i>Local Dispatch Time</i> .
JBNFFNC	Amount of CPU time used by the thread on resource affinity domains other than the one this thread is associated with, but within the same group. The time is reported in internal model-independent units. This is called the <i>Non-Local Dispatch Time</i> .
JBNFHNP	Number of 4K page frames allocated for this thread during this interval from the resource affinity domain this thread is associated with. These are called <i>Local page frames</i> .
JBNFFNP	Number of 4K page frames allocated for this thread during this interval from resource affinity domains other than the one this thread is associated with, but within the same group. These 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 used to report data 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 they were associated with and short lifespan secondary threads are reported for the job to which they belong.
JBSACPU	Accumulated job scaled CPU time charged (in microseconds). The accumulated scaled interval CPU time charged for all threads of the job since the job started. This field is provided for primary threads only.
JBINDCPU	The amount of unscaled CPU time (in μ s) that represents the work done solely within this thread without regard for how server task work is charged.
JBSINDCPU	Thread scaled CPU time used (in microseconds). The amount of scaled CPU (in μ s) time that represents the work 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 executes.
JBVPDLY	The elapsed delay time (in microseconds) due to virtualization for a task when it was executing. 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 seizes held by this thread at the time the data was sampled.
JBPSLCKCNT	The number of process scoped locks held by this thread at the time data was sampled.
JBTSLCKCNT	The number of thread scoped locks held by this thread at the time data was sampled.

Column	Description
JBTSRCDLCK	The number of thread scoped database record locks held by this thread at the time data was sampled.
JBNFOGDT	Amount of CPU time used by the thread in a resource affinity group other than the one this thread is associated with. The time is reported in internal model-dependent units.
JBNFOGMA	Number of 4K page frames allocated for this thread during this interval from a resource affinity group other than the one this thread is associated with
JBFLDR2	Workload capping group delay time (in microseconds). The amount of time this thread could not be dispatched due to workload capping.
JBFLDR3	Workload capping group. The identifier for the workload capping group this thread belonged to at the time this data was sampled. A value of zero is reported when no group was assigned.

QAPMJOBSR

This file contains data for jobs that have performed save or restore operations. It contains one record per job for each operation type it has performed.

By selecting Collection Services, Collection Services Database Files, then QAPMJOBSR in Performance Data Investigator (PDI), you will 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
<input type="checkbox"/>	1	Sep 20, 2011 12:15:00 AM	890	168	0	168	158	10
<input type="checkbox"/>	1	Sep 20, 2011 12:15:00 AM	890	0	0	0	0	0
<input type="checkbox"/>	1	Sep 20, 2011 12:15:00 AM	890	0	0	0	0	0
<input type="checkbox"/>	2	Sep 20, 2011 12:30:00 AM	900	0	0	0	0	0
<input type="checkbox"/>	2	Sep 20, 2011 12:30:00 AM	900	0	0	0	0	0
<input type="checkbox"/>	3	Sep 20, 2011 12:45:00 AM	900	0	0	0	0	0
<input type="checkbox"/>	3	Sep 20, 2011 12:45:00 AM	900	0	0	0	0	0
<input type="checkbox"/>	4	Sep 20, 2011 1:00:00 AM	900	0	0	0	0	0
<input type="checkbox"/>	4	Sep 20, 2011 1:00:00 AM	900	0	0	0	0	0
<input type="checkbox"/>	5	Sep 20, 2011 1:15:00 AM	900	0	0	0	0	0
<input type="checkbox"/>	5	Sep 20, 2011 1:15:00 AM	900	0	0	0	0	0
<input type="checkbox"/>	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 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 caused by all partitions using the pool. You need to have a POWER6 system and firmware level xx340_075 or later for this data to be available. See Table 7-6 on page 199 for the data that is kept in this table, following 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	Description
SMPOOLID	Shared memory pool identifier. The identifier of the shared memory pool which this partition is using.
SMWEIGHT	Memory weight. Indicates the variable memory capacity weight assigned to the partition. Valid values are hex 0 -255. The larger the value, the less likely this partition is to lose memory.
SMREALUSE	Physical real memory used. The amount of shared physical real memory, in bytes, that was being used by partition memory at sample time.
SMACCDLY	Real memory access delays. The number of partition processor waits that have occurred because of page faults on logical real memory.
SMACCCWAIT	Real memory access wait time. The amount of time, in milliseconds, that partition processors have waited for real memory page faults to be satisfied.
SMENTIOC	Entitled memory capacity for I/O. The amount of memory, in bytes, currently assigned to the partition for use 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 can 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 has been mapped by I/O requests since the partition was last IPLed or the value was reset by an explicit request.
SMIOMDLY	I/O memory mapping delays. The cumulative number of delays that have occurred because insufficient entitled memory was available to map an I/O request since the partition was last IPLed.
MPACCDLY	Pool real memory access delays. The number of virtual partition memory page faults within the shared memory pool for all partitions.
MPACCCWAIT	Pool real memory access wait time. The amount of time, in millisecond, that all partitions processors have 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 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 served by the shared memory pool.
MPIOCUSE	Pool entitled I/O memory in use. The summation, in bytes, of I/O memory mapped by I/O requests from all active partitions 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.
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 used by this partition.
SYFRMSCPU	Processor scaled firmware time used (in microseconds). The amount of scaled processor firmware time used by this partition.
SYFRMSCPU	Processor scaled firmware time used (in microseconds). The amount of scaled processor firmware time used by this partition.
SYPFOLDSW	Identifies the current state of the processor folding ^a switch, where <ul style="list-style-type: none"> ▶ <i>blank</i> means data not available ▶ "0" off ▶ "1" on ▶ "2" system controlled
SYPFOLDST	The current state of processor folding <ul style="list-style-type: none"> ▶ <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 use of idle virtual processors (VP). This 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 amount of VPs on your configuration without a performance overhead. 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 Columns in QAPMSYSWLC

Column Name	Description
INTNUM	Interval number: The nth sample database interval based on the start time 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 assigned to the workload group when allocated by License Management
SWPRCAS	Processors assigned. The maximum number of processors which may be used concurrently by all threads of all processes which are associated with the workload group. This is the value 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 assigned to the group over time.
SWPRCUSE	Processor unscaled time used (in microseconds). The amount of unscaled processor time used within threads assigned to this group. Note: This does not include time charged to a thread by server tasks.
SWSPRCUSE	Processor scaled time used (in microseconds). The amount of scaled processor time used within threads assigned to this group. Note: This does not include time charged to a thread by server tasks.
SWDELAY	Dispatch latency time. The amount of time ready to run threads could not be dispatched due to 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 collected in the Removable storage (*RMVSTG) collection category. It contains one record per interval per tape device 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

Column Name	Description
TPWRTS	Number of writes
TPBRD	Bytes read
TPBWRT	Bytes written
TPWREQ	Time spent waiting for a request from the client (in milliseconds)
TPWRESP	Time spent waiting for a response from the drive (in milliseconds)
TPSFCMD	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. It should be noted that the support for *EXTSTG is disabled when shipped. Refer to the *Memo to Users* for more information and APAR SE41825 for PTF information. For more information, see Chapter 9, "Storage and solid state drives" on page 273.

7.2.3 Additional Enhancements

There have been several enhancements made to Collection Services component of Performance Tools, these include:

- Collection Services has been updated to support reporting system wide usage data for workload capping groups as well as TDE level data to assist in understanding performance issues related to capping actions.
 - The *JOBMI data category and QAPMJOBMI file has been modified to support additional TDE metrics that identify the group a TDE (thread) was associated with at sample time along with how much time that thread was not able to run due to workload capping constraints.

- The *SYSLVL collection category has been modified to collect WLC group data for groups that are in use. A new file QAPMSYSWLC has been created in the target performance database library to contain this data. The QAPMSYSWLC file and/or member will be created only if the source *MGTCOL collection contains data for workload capping groups.
- ▶ A page size flag metric has been added to the Dump main Memory Information (DMPMEMINF) command to indicate whether a page is 4K or 64K
- ▶ A new metric has been implemented by Collection Services to report the number of full opens that occur within a job and is being reported in the QAPMJOBOS file using the following fields
 - JBLBO - The cumulative number of SQL cursors which have been fully opened
 - JBLBS - The cumulative number of SQL cursors which have been pseudo-opened also known as reused SQL cursors
 - JBNUS - The number of Native database (non-SQL) files and SQL cursors which have been full opened
 - Subtracting the value within field JBLBO from JBNUS will yield the number of non-SQL full opens

7.3 IBM Systems Director 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, configurable for user-defined health thresholds.
- ▶ Collection services has the ability to collect high-level cross-partition processor performance metrics for all logical partitions on the same single physical server regardless of operating system. This 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 found under Physical System.
- ▶ 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 large number of 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 under the **Search** button when you launch Investigate Data.
- ▶ A simple Performance Explorer content package is included to start analysis of Performance Explorer data.
- ▶ Many new perspectives have been added to Investigate Data, and many new metrics have been 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 more-complete History data.⁶

For a full overview of all the functionality of the Performance Investigator and Collection management interface, see Chapter 18, “IBM Systems Director Navigator for IBM i 7.1” on page 525.

7.4 IBM iDoctor for IBM i

This section describes the changes in IBM iDoctor for IBM i for the IBM i 7.1 release. These changes are also available for the 5.4 and 6.1 releases unless otherwise noted.

We cover the changes for Job Watcher, Disk Watcher, PEX Analyzer and the Collection Services Investigator. This section also covers the changes to the various parts of the iDoctor GUI.

7.4.1 Installation

The following functions have been added to the installation process:

- ▶ Validation checks have been added for each partition specified to ensure that the partition will be able to install iDoctor
- ▶ The default install 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 has been added when installing Jobwatcher at 6.1 and 7.1 to ensure that the Job Watcher definitions file (QAPYJWDFN) exists in QUSRSYS and contains the IBM-supplied definitions

An example video of installing iDoctor is available here:

http://www.youtube.com/watch?v=pURHegIt0TQ&feature=channel_video_title

The iDoctor GUI now requires the Visual Studio 2008 SP1 redistributable package to be installed. More information regarding this requirement can be found at the following URL:

https://www-912.ibm.com/i_dir/idoctor.nsf/downloadsClient.html

7.4.2 My Connections View

The My Connections View, as 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 desired partitions
- ▶ Added menus to Load and Remote all iDoctor Stored Procedures
- ▶ Added *Uninstall iDoctor* option
- ▶ Added an option to *Edit* an existing connection
- ▶ Delete obsolete analysis files for each component

⁶ Most of these new functions are available for systems running IBM i 6.1 with the current IBM HTTP Server for i group PTF (SF99115)

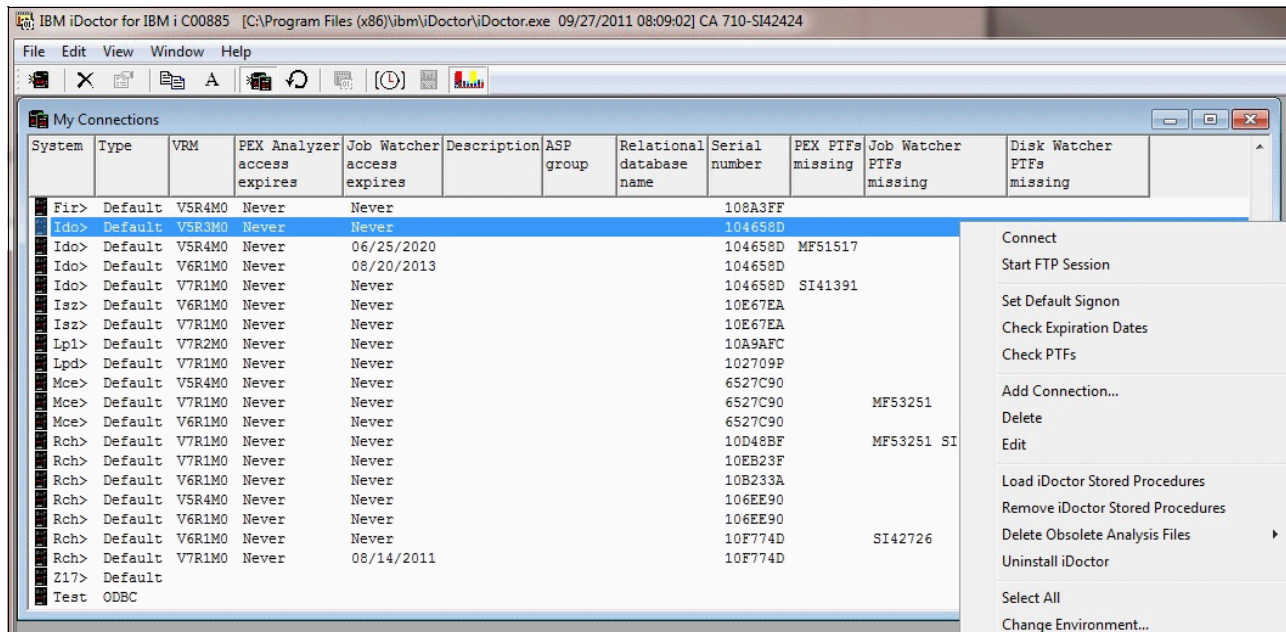


Figure 7-3 Options on My Connections view

When signing on to a system, iDoctor uses the configured sign-on setting defined in System i Navigator (**Properties** → **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. It also makes use of 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, have the **Prompt every time** option set within System i Navigator.

Support has been added to view collections stored in libraries 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. These values cause the QZRCRSRVs and QZDASOINIT jobs created by iDoctor to recognize data stored in the Independent ASP.

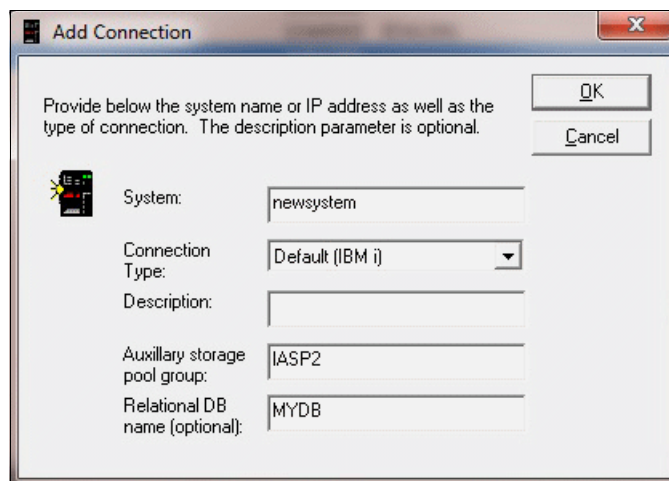


Figure 7-4 Add Connection

7.4.3 Main Window

On the main window toolbar a button has been added that enables/disables the display of situational analysis background colors in graphs. A simple click on the button will turn it on/off for all graphs (even open ones). Another click on the graph or legend will redraw the graph with/without the situations (if found in the data). See Figure 7-5




Figure 7-5 Button to enable/disable situations analysis background colors

On the main window, as shown in Figure 7-6 on page 206, 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 very small groupings are useful in PEX Analyzer and the very large groupings are useful in Collection Services. Additional checking has been added to the GUI to ensure that only relevant time grouping options are shown for the current data.

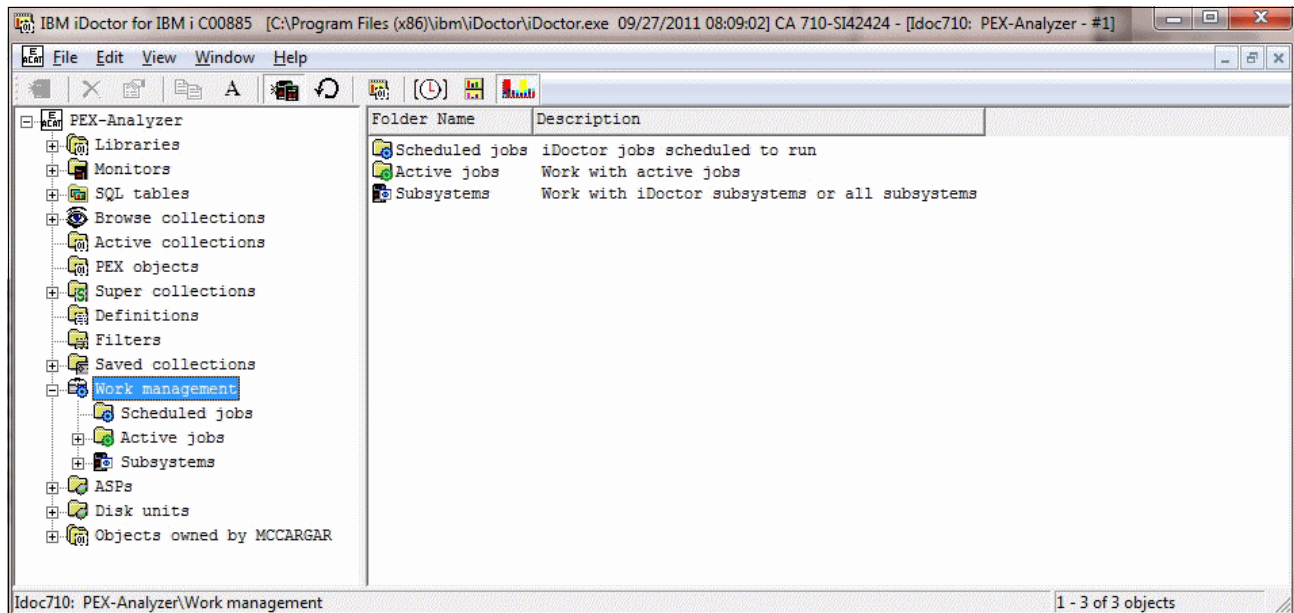


Figure 7-6 Main window

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 1 SQL table.

A new Browse Collections option has been added which provides alternate ways of looking at the collections stored on the system. This function is built from a repository that must be updated periodically by the user using the options found when right-clicking the Browse Collections folder.

Browse Collections gives the user the ability to find collections on the system in several ways:

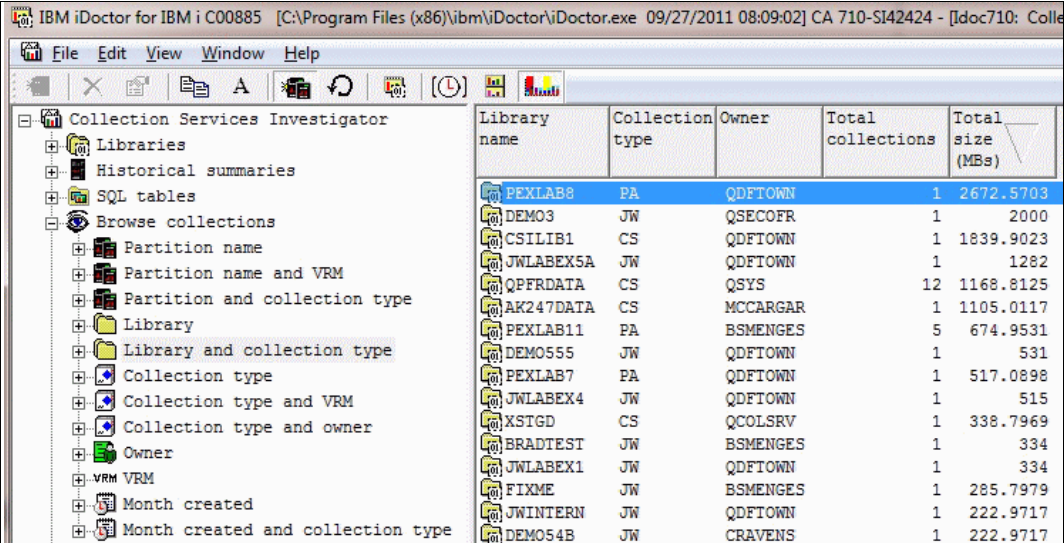
1. Partition Name
2. Partition Name and VRM
3. Partition and collection type
4. Library
5. Library and collection type
6. Collection type
7. Collection type and VRM
8. VRM
9. Month created
10. Month created and collection type

Each of the above views gives the total size of all collections (in MBs) in each group and the total number. See Figure 7-7 on page 207.

Added a Super collections folder that lets you work with the super collections that exist on the system. These contain a folder for each collection type collected within the super collection.

The Saved collections folder allows you to work with any save files found on the system that contain iDoctor collections saved previously using the iDoctor GUI.

The Work Management folder allows a user 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. See Figure 7-7.



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 showing the ASPs and disk units configured on the current system. The ASPs folder allows 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. (also includes total I/O, total sizes). See 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	0
48	4326	1	30.05	82.8	5.9	0	0	0	0	0	0	0
45	4326	1	30.05	82.8	5.9	0	0	0	0	0	0	0
44	4326	1	25.76	82.8	5.9	0	0	0	0	0	0	0
53	4326	1	30.05	82.8	5.6	.56	101.60	0	.56	0	101.60	0
51	4326	1	30.05	82.8	5.6	.11	16	.11	0	16	0	0
39	4328	1	120.59	82.8	1.3	.67	85.33	0	.67	0	85.33	0
40	4328	1	120.59	82.8	1.2	.56	92	.11	.44	8	113	0
33	4328	1	103.37	82.8	1.2	.56	80.80	.11	.44	12	98	0
37	4328	1	103.37	82.8	1.1	.67	76.67	0	.67	0	76.67	0
38	4328	1	120.59	82.8	1.1	.44	83	0	.44	0	83	0
41	4328	1	120.59	82.8	1	.44	81	0	.44	0	81	0
55	4326	1	30.05	82.8	0	0	0	0	0	0	0	0
56	4326	1	30.05	82.8	0	.11	8	0	.11	0	8	0
52	4326	1	25.76	82.8	0	.33	6.67	.11	.22	12	4	0
54	4326	1	30.05	82.8	0	.11	68	0	.11	0	68	0

Figure 7-8 Overview of the disk status in ASP1

Use the right-click menu option on the Disk Units or ASP folder and perform **Reset Statistics** to restart the collection of disk statistics. You can also use the **Select fields** menu option when right-clicking the folder to rearrange fields or add additional fields. The status bar of the Main window shows the times for first disk statistics snapshot, and the last one.

Similarly, you will find an *Active Jobs* (see Figure 7-6 on page 206) folder on the same panel that 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 CPU users, and so forth. There is also a filter option to filter the list by name, user, number, current user or minimum CPU percentage. Use the **Select fields** menu option when right-clicking the folder to rearrange fields or add additional fields. Expanding a job shows the threads and the thread statistics available for each. You can start Job Watcher or PEX Analyzer collections or add JW/ 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
QZRCRSRV	QUSER	092109	1	RUN	MCCARGAR	Interactive - Server			.2	38	20	560
QUMEPRVAGT	QSECOFR	686828	4	DEQW	QSECOFR	Interactive	cimprovagt		.1	34	50	0
TESTJOB	NICKI	937224	1	DEQW	NICKI	Interactive	TESTRECV		0	0	50	0
TESTJOB	NICKI	937280	1	DEQW	NICKI	Interactive	TESTRECV		0	0	50	0
TESTJOB	NICKI	937121	1	DEQW	NICKI	Interactive	TESTRECV		0	0	50	0
TESTJOB	NICKI	937213	1	DEQW	NICKI	Interactive	TESTRECV		0	0	50	0
TESTJOB	NICKI	937871	1	DEQW	NICKI	Interactive	TESTRECV		0	0	50	0
TESTJOB	NICKI	937620	1	DEQW	NICKI	Interactive	TESTRECV		0	0	50	0
TESTJOB	NICKI	937986	1	DEQW	NICKI	Interactive	TESTRECV		0	0	50	0
TESTJOB	NICKI	937956	1	DEQW	NICKI	Interactive	TESTRECV		0	0	50	0
TESTJOB	NICKI	937379	1	DEQW	NICKI	Interactive	TESTRECV		0	0	50	0
TESTJOB	NICKI	937361	1	DEQW	NICKI	Interactive	TESTRECV		0	0	50	0
TESTJOB	NICKI	937570	1	DEQW	NICKI	Interactive	TESTRECV		0	0	50	0

Figure 7-9 Overview of Active Jobs (WRKACTJOB)

The Objects owned by user folder allows you to manage (view/delete) the objects on the system that are owned by the current user. This is primarily intended to aid in disk space cleanup and options are provided to view only iDoctor objects, or all objects owned by the current user.

7.4.4 Collection Options

The collection menu (right-click) now contains an *Analyses* menu for all components.

Choosing an option under this menu will run a program that creates SQL tables needed for further analysis. In most cases additional reports will become available after the analysis completes and the collection is refreshed (F5.)

The Summarize menu option in CSI and JW has moved to *Analyses - Run Collection Summary*. Choosing this option will now display a window that allows a user 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 will be accessible using the graphing options provided under the collection.

The Create Job Summary menu option in CSI and JW has moved to *Analyses - Run Create Job Summary*.

There is a new iDoctor Report Generator for all collections. To access it, right-click a collection and click the *Generate Reports* menu. The default web browser is opened to show the HTML report after the reports have been captured to JPG files. As reports are running you can switch to other windows, but right before screen capture are taken the data viewer must move to the front of all windows to get the screen capture. This happens automatically but might look strange the first time you use it.

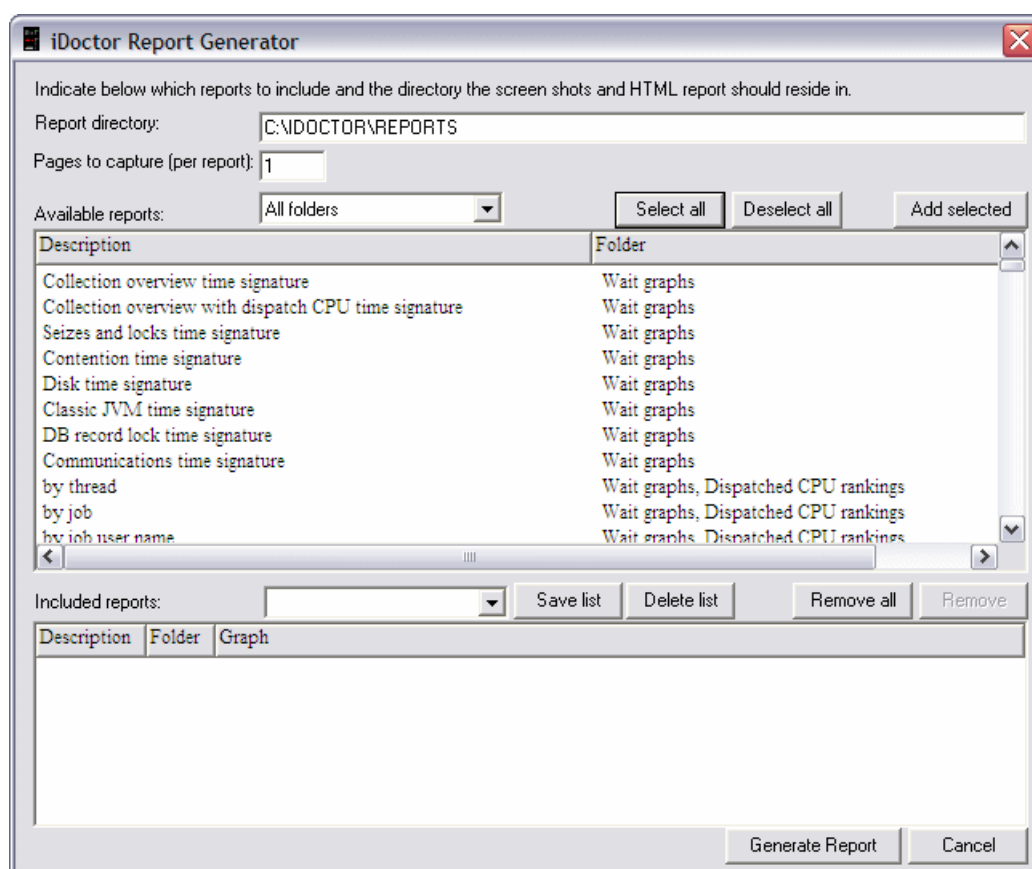


Figure 7-10 Report Generator

With the *Save* option (see Figure 7-11) you can save multiple collections or monitors as well. After using this option, the new Saved collections folder shows a record identifying the save file which you can use to restore the data, or distribute it.

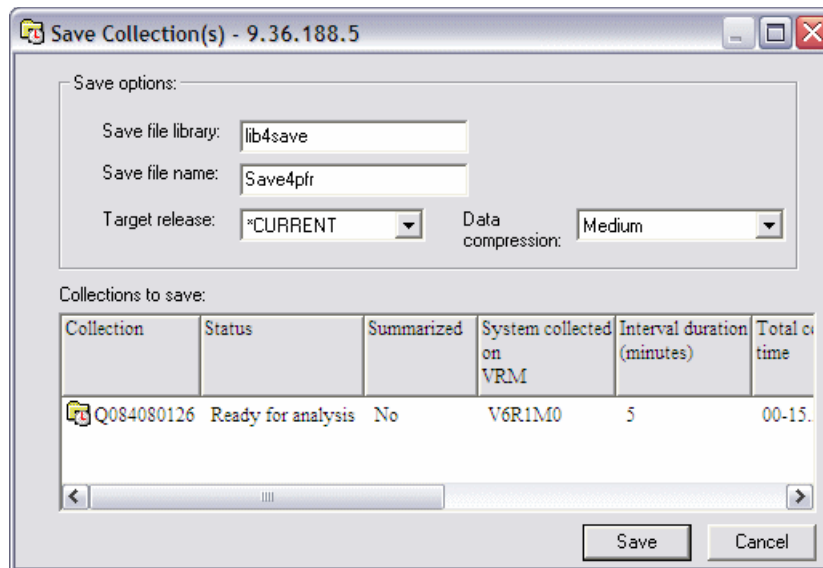


Figure 7-11 Saving a collection

The *Transfer to* option now supports all collection types. It has been modified as follows:

- ▶ The usage of commands FTPJWCOL/FTPPACOL by the GUI has been replaced with a combination of SAVPFCOL/FTPFILE.
- ▶ When transferring multiple collections, they are saved/sent using the same save file instead of different ones.
- ▶ Monitors (or multiple monitors) can now be transferred.
- ▶ Path/filenames has increased to 100 chars from 50 previously.
- ▶ The user has complete control to set the filename to whatever they want, but the recommended format is given.
- ▶ The user now has an option in the action list to FTP the collections to Ecurep.
- ▶ Mixed case passwords, user names, filenames are now supported which fixes problems when sending data to/from AIX.

7.4.5 Data Viewer

The Data Viewer toolbar has a toolbar that shows all the idle waits (include all buckets) for wait bucket *job* and larger grouping graphs in CSI and Job Watcher. This is a toggle button that allows you 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 brings up 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 the **Save** button on the Data Viewer toolbar to save the current graph and legend to a JPG image.

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 Max
- ▶ Percent of, Delta (current - prior)
- ▶ Delta (prior - current)

There are also changes to the design of graphs and reports.

Side by Side Comparisons:

Side-by-side comparisons allow you to sync up the scrolling and Y-axis scaling of any two graphs (or tables) in a Data Viewer.

Once two or more graphs/tables exist in a Data Viewer, the buttons will be ready for use. See Figure 7-13




Figure 7-13 Side by side comparisons

An example video of using side-by-side comparisons may be found here:

http://www.youtube.com/watch?v=0qkg_cmglAo&feature=channel_video_title

Table Filtering Options:

Right-click the desired column to filter on. Set the desired operator and value and press Apply to refresh immediately. See Figure 7-14 on page 212

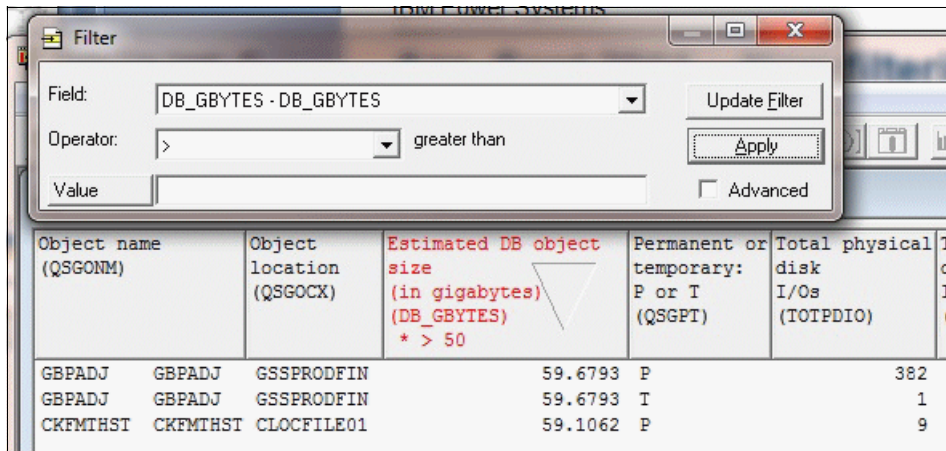


Figure 7-14 Filter

Additional options are found by right-clicking a column: Sort, remove filter, hide column and unhide all columns.

Graph Filter Options:

New graph filtering options allow you to define a filter over the desired column shown in the legend. Right-click the desired column description and use the *Add Filter* menu.




Figure 7-15 Example of Graph Filter Options

7.4.6 Collection Services Investigator (CSI)

The folders available in the CSI component have significantly changed. Instead of showing libraries containing CS data, new folders are available as shown in Figure 7-16:

- ▶ Libraries containing CS database file collections (filterable).
- ▶ Historical summaries containing 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 covered in section 7.4.3.

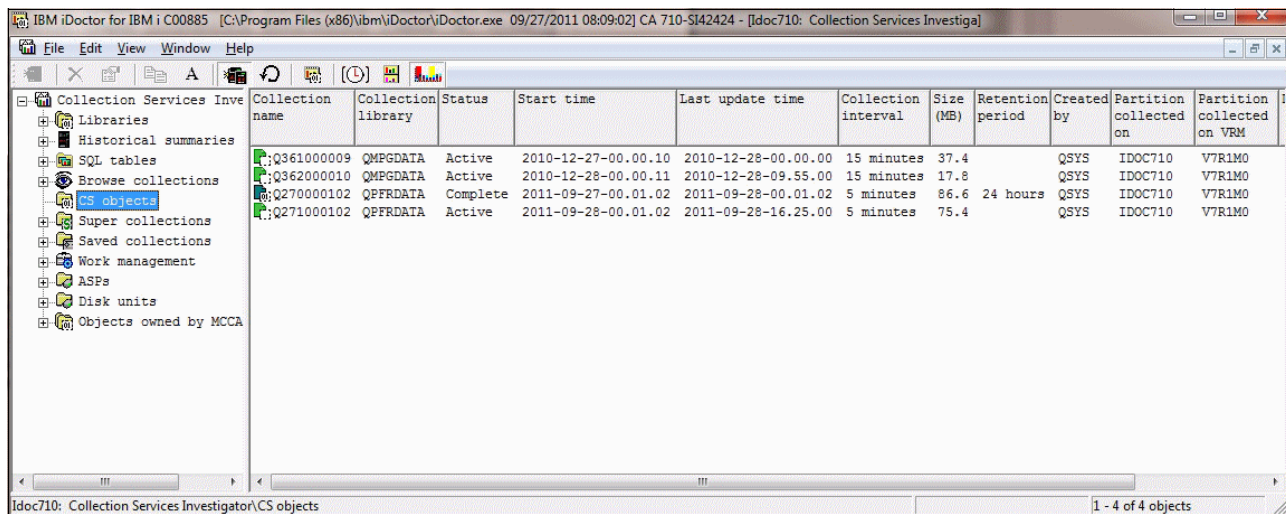


Figure 7-16 Folders in the CSI component

Historical Summaries

Historical Summaries consist of consolidated/reduced sets of Collection Services data for the purpose of graphing many days/weeks/months of data at one time.

See Figure 7-17 for 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. Use the *STRCSMON* command (or *Start Monitor* menu option from the Historical Summaries folder) that will summarize new Collection Services data every day at the desired time and add it to the system's default Historical Summary repository
2. Right-click one or more libraries and use the Analyses -> Run Historical Summary menu.

For best results, it's recommended that the 1 hour grouping option is used when creating the Historical Summary.

Historical Summaries provide a complete set of graphs very similar to the graphs provided under normal collections. A full set of “average day” and “average week” graphs are also supplied. Much more information about Historical Summaries can be found here:

[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 new menu on a collection called *Launch Workload Estimator* as shown in Figure 7-18. A webpage is shown with the average CPU and disk statistics for the collection.




Figure 7-18 Launch the Workload Estimator

Upon pressing *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 the following information:

System Name: S659F730
 Operating System: IBM i
 Version: 6.1
 Library Name: ANALYSIS
 Collection: DS6000

CPU Information
 Average CPU utilization 29.1187
 Average interactive utilization 6.0163

Disk Information

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

At the bottom left of the form is a "Submit" button.

Figure 7-19 Submit to WLE

Managing collections

On the pull-down as in Figure 7-18 on page 214, you see that a Copy function has been added. You can also use the CPYPFRDTA (Copy Performance Data) to obtain the same result. The Delete function now uses the DLTPFRDTA (Delete Performance Data) command.

The import data to WLE option is accessible from a CSI graph if a time range has been selected. The numbers provided to WLE is based on the selected time period.

A search function, similar to the one in Job Watcher is now available in this panel, allowing a user to generate a report showing job records 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 desired job over time. You can also search over multiple collections at one time by selecting the desired collections in the CSI component view's list side and then using the Search pop-up menu. After the search results are shown, the drill downs can be done on the same set of collections provided for the desired job/thread.

You can create graphs over multiple collections at one time in the same library. Select the desired 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, the drill downs for rankings and the single object over time graphs apply to this same set of collections.

In CSI, but also in PEX, you now have a Wait Buckets tab to collection properties that shows the wait buckets/ENUM mapping. Right-click *Collection* → *Properties* → *Wait Buckets* as shown in Figure 7-20.

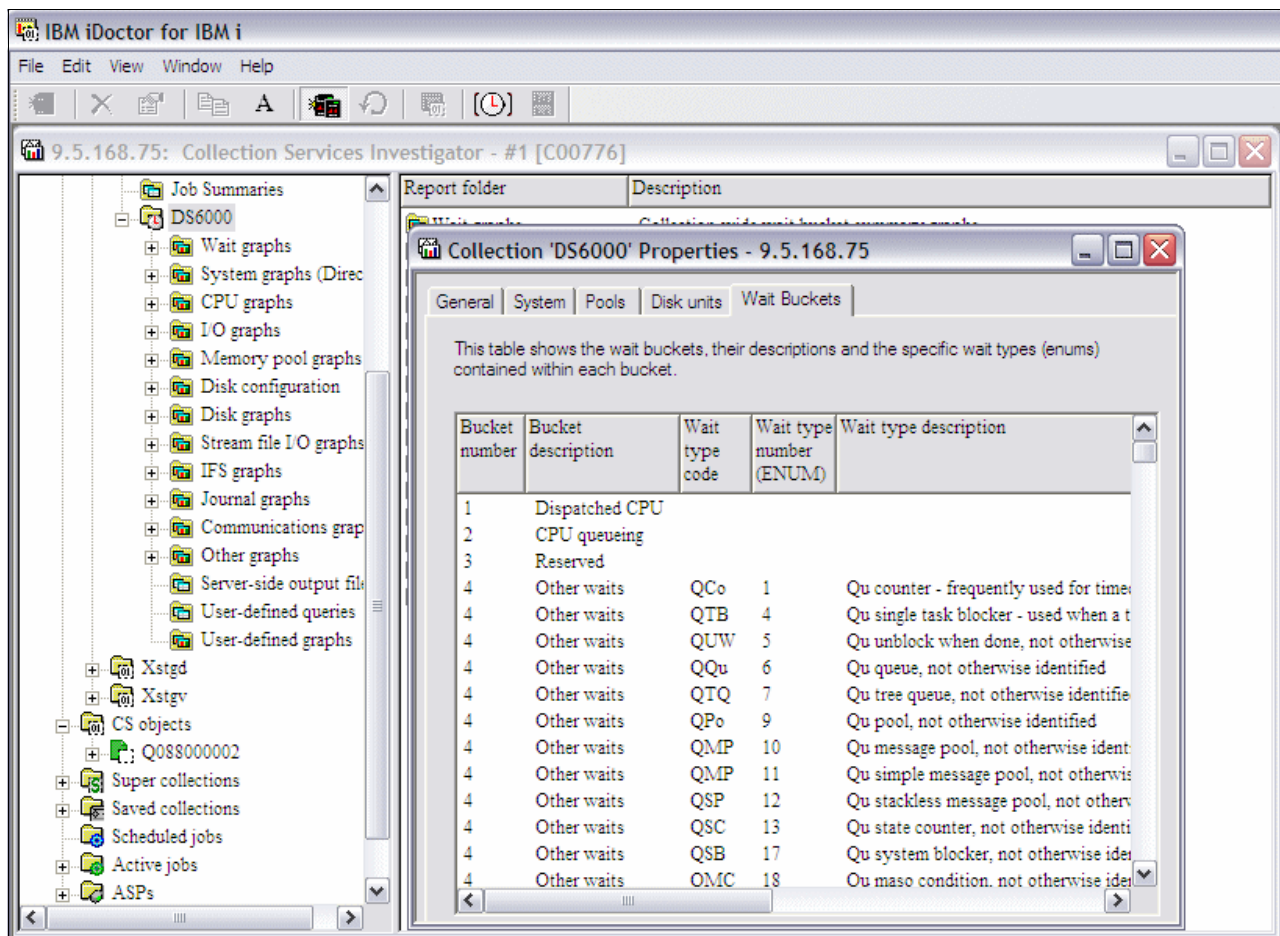


Figure 7-20 Wait and ENUM display

Situational analysis

To initiate a situational analysis, right-click the collection and pick the *Analyses → Run Situational Analysis* option. You can also right-click the collection and pick *Analyses → Analyze Collection* and from there click the *Situations* button to configure the default situations to be used by CSI.

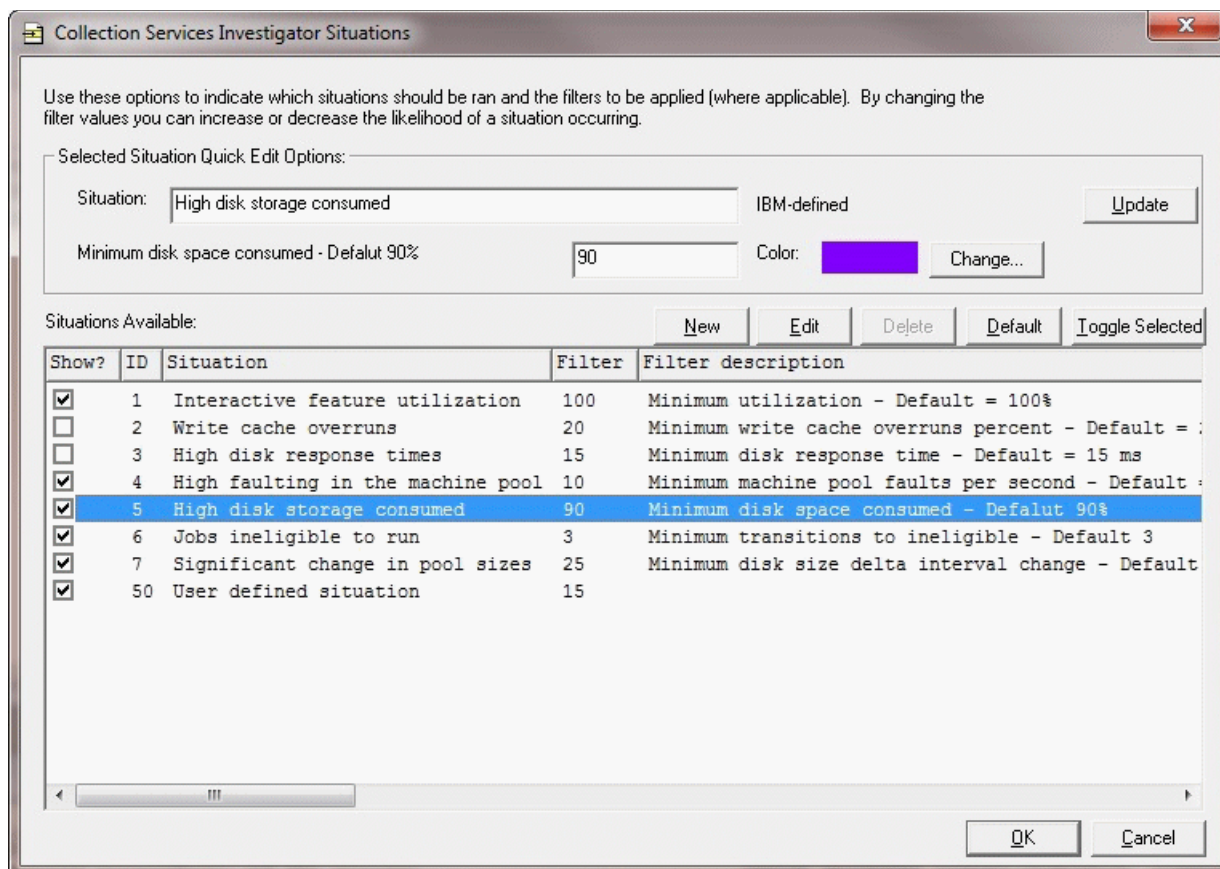


Figure 7-21 Collection Services Investigator Situations

All the time interval-based Wait graphs include different background colors, each identifying a situation as shown in Figure 7-22.

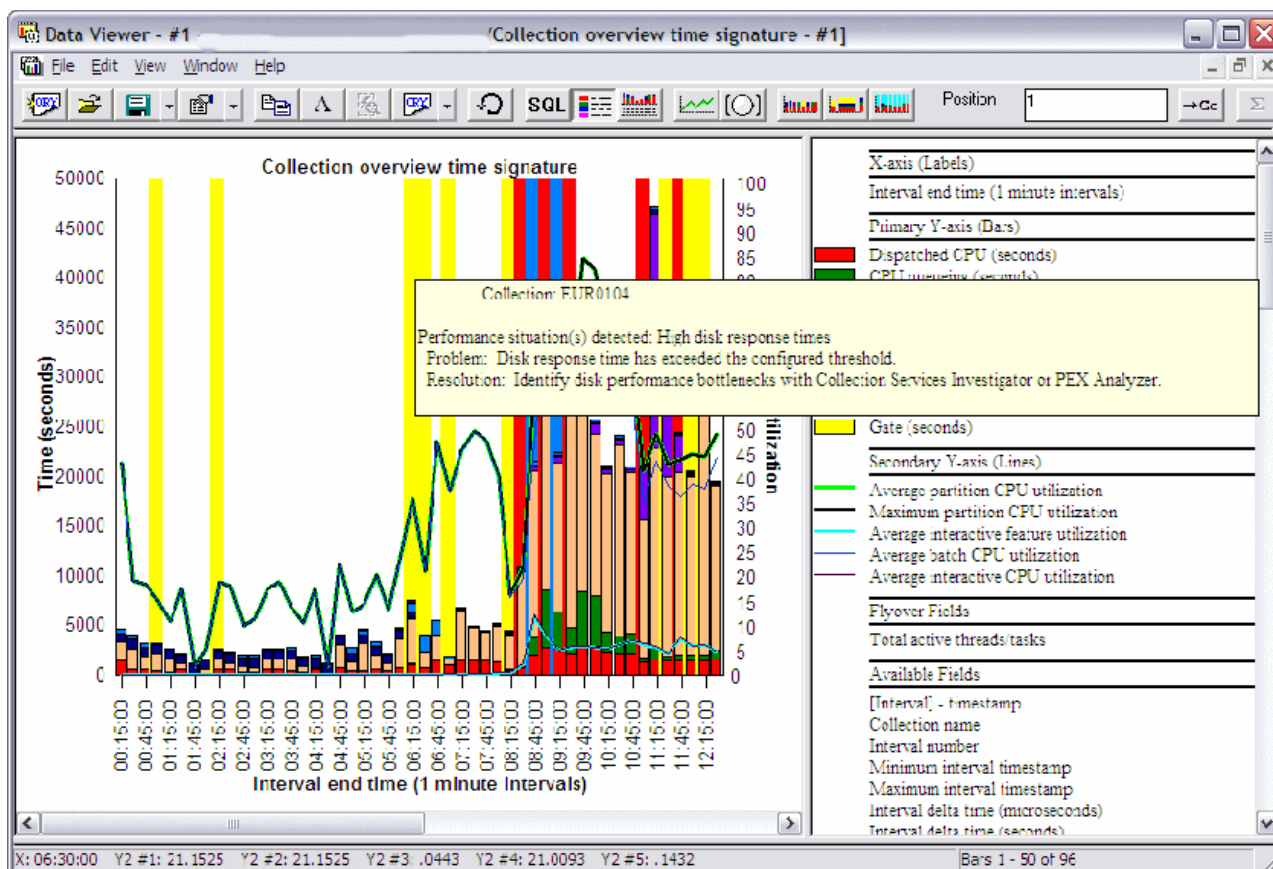


Figure 7-22 Situational Analysis example

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 sec
- ▶ High disk storage consumed - 90%
- ▶ Jobs ineligible to run - 3 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 296.

Physical system graphs

You can now view what has been collected by the IBM System Director tool for all partitions on which it has been running, as shown in Figure 7-23. This supports all POWER systems as of V5R4.

- ▶ CPU use
- ▶ I/O counts
- ▶ I/O rates

- Disk space
- Memory

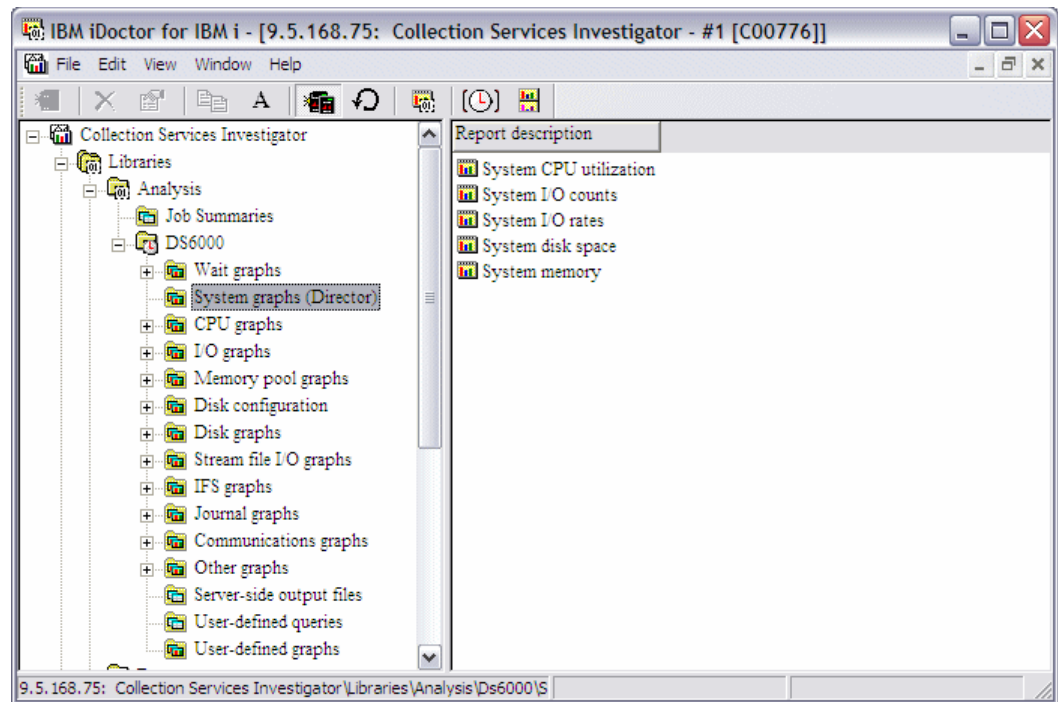


Figure 7-23 Physical system statistics

The information from the hypervisor for systems running release at 6.1 or higher, is now showing up in a new System Graphs HMC folder (see Figure 7-24).

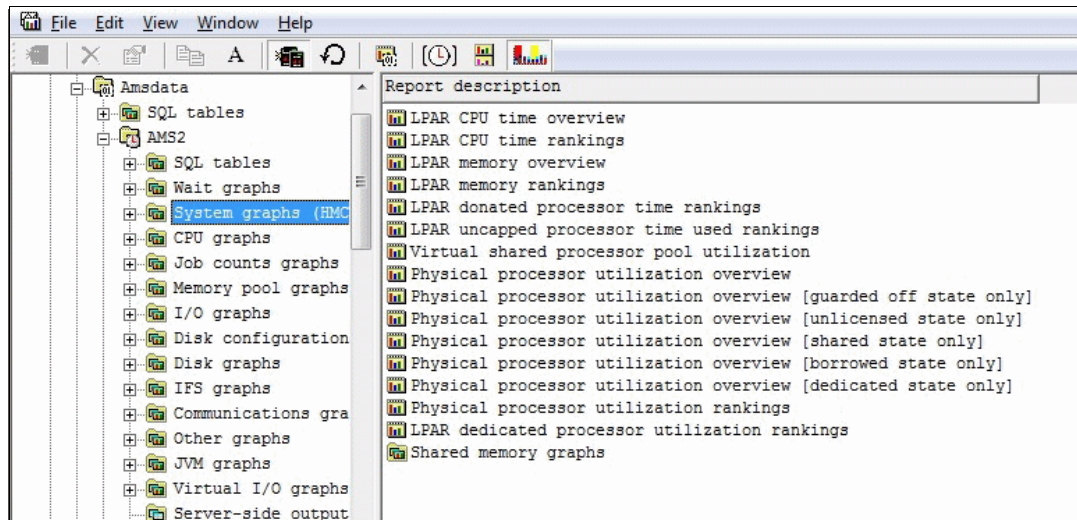


Figure 7-24 System Graphs (HMC) folder

It includes the following graphs

LPAR CPU time overview shows CPU time used for all partitions Figure 7-25 on page 220:

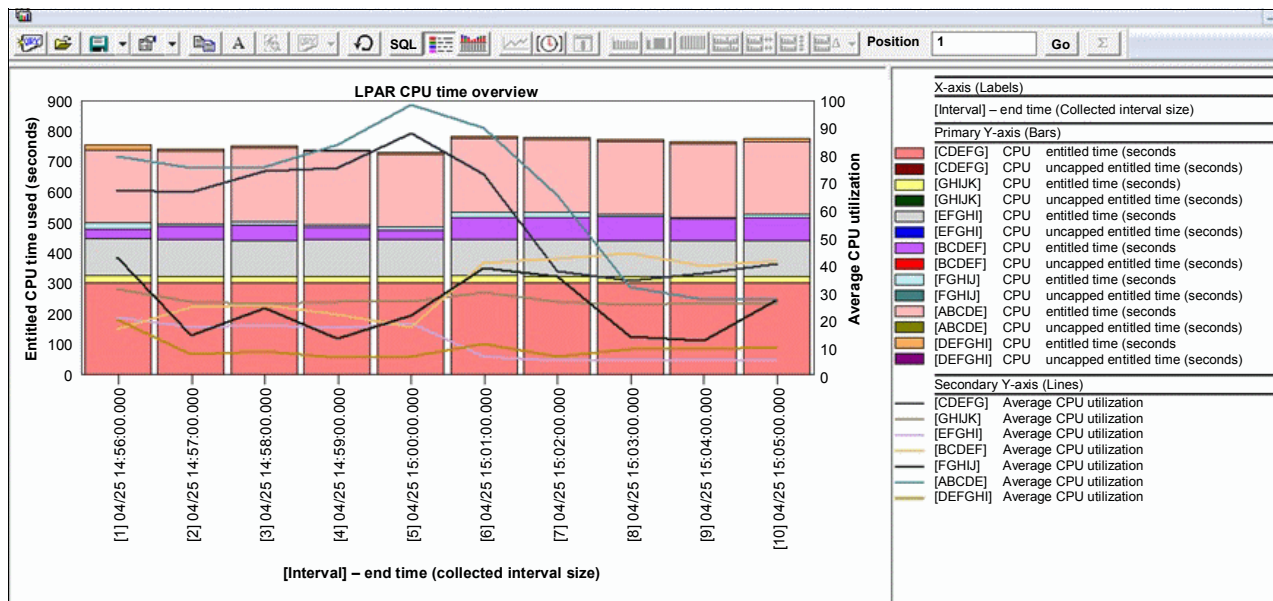


Figure 7-25 LPAR CPU time overview graph

LPAR cpu time rankings: ranks the partitions by CPU used

LPAR memory overview: shows memory consumption for all partitions Figure 7-26




Figure 7-26 LPAR memory overview

- ▶ LPAR memory rankings: ranks the partitions by memory used
- ▶ LPAR donated time rankings: ranks the partitions by donated CPU time
- ▶ LPAR uncapped processor time used rankings: ranks the partitions by uncapped CPU time
- ▶ Virtual shared processor pool utilization
- ▶ Physical processor utilization overview: shows the average CPU utilization for each physical processor over time see Figure 7-27 on page 221
- ▶ LPAR dedicated processor utilization rankings




Figure 7-27 Physical processor utilization overview

Shared Memory Graphs

- ▶ If file QAPMSHRMP file is available a *shared memory graphs* subfolder will be available containing the following additional graphs:
- ▶ Shared memory overview (Figure 7-28)

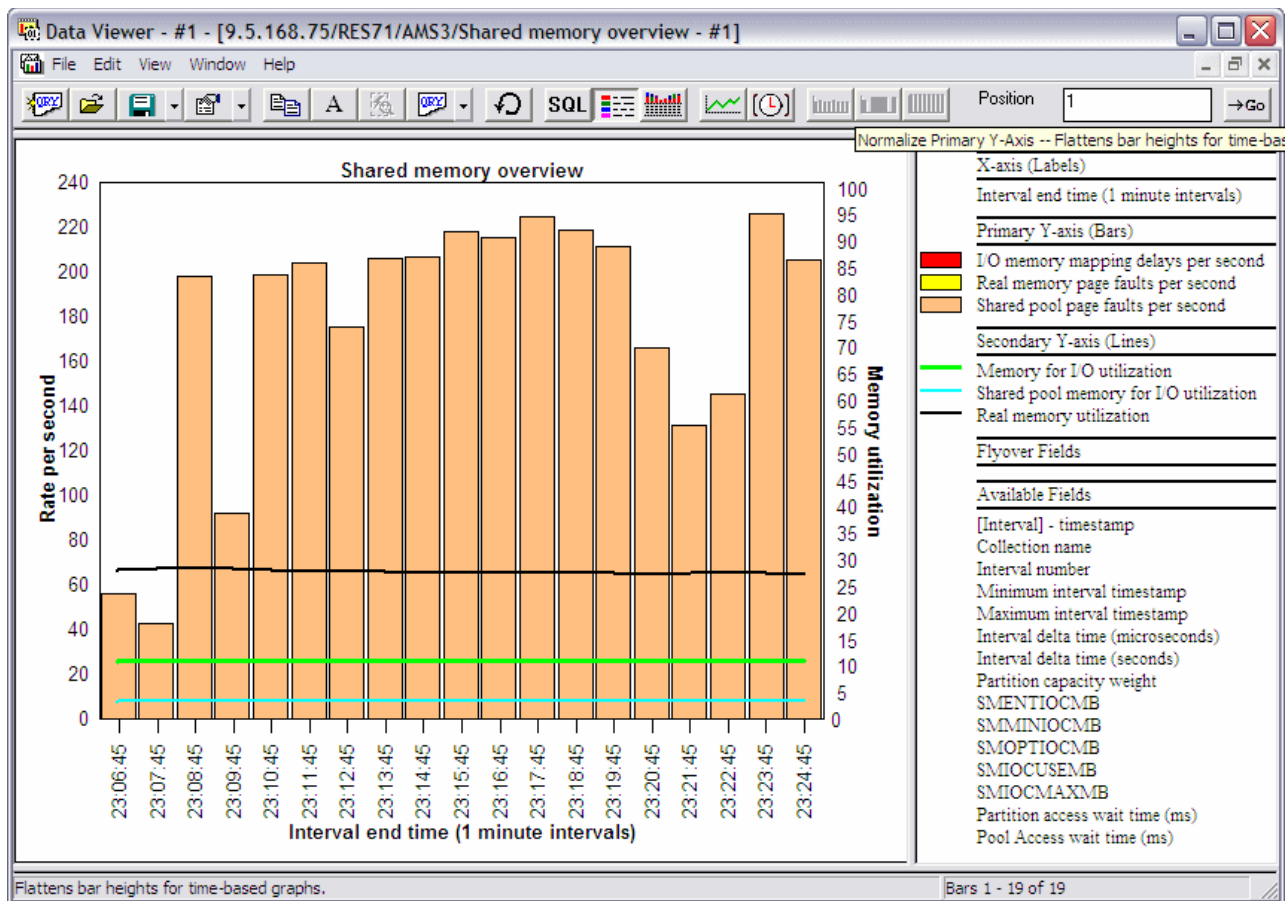


Figure 7-28 Shared memory overview

- Shared memory pool page faults (Figure 7-29)

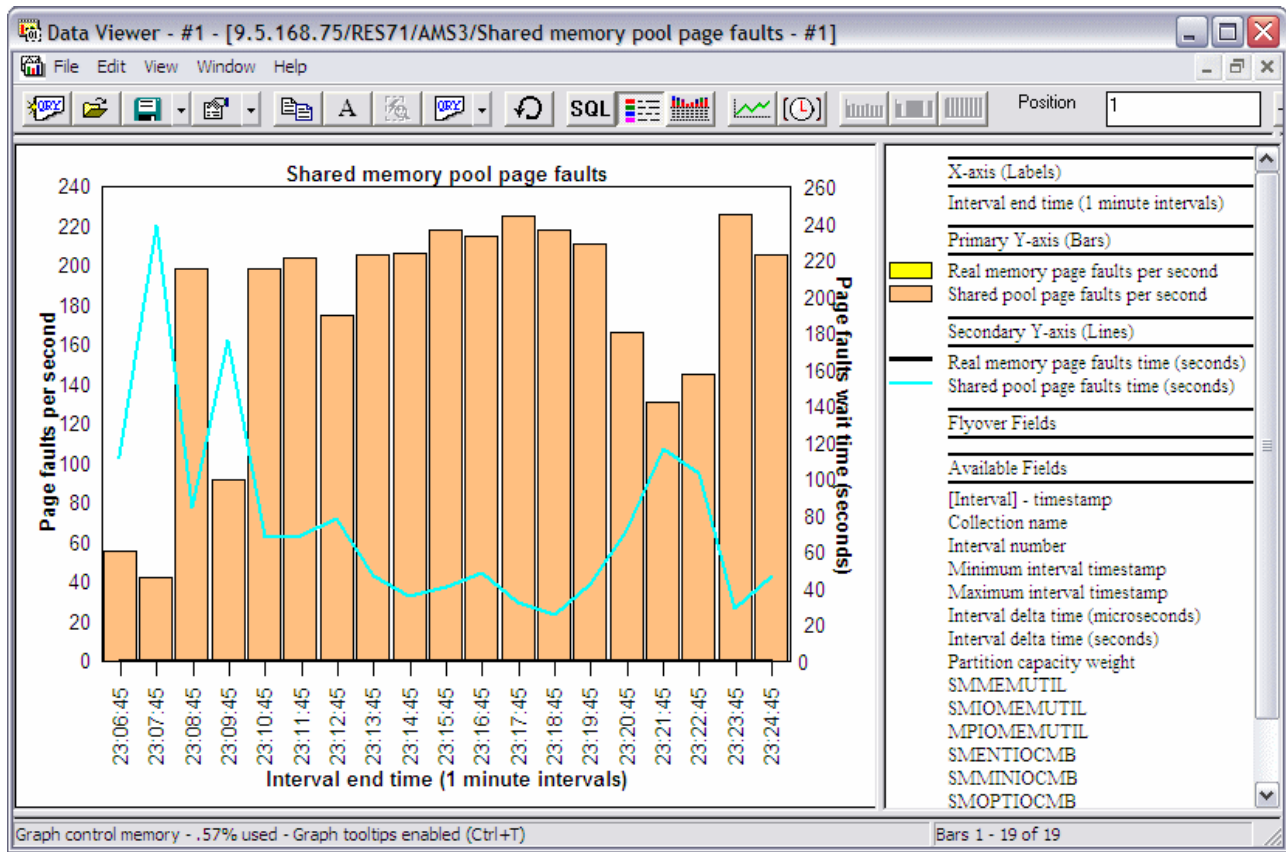


Figure 7-29 Shared memory page faults

- Memory for I/Os overview (Figure 7-30)

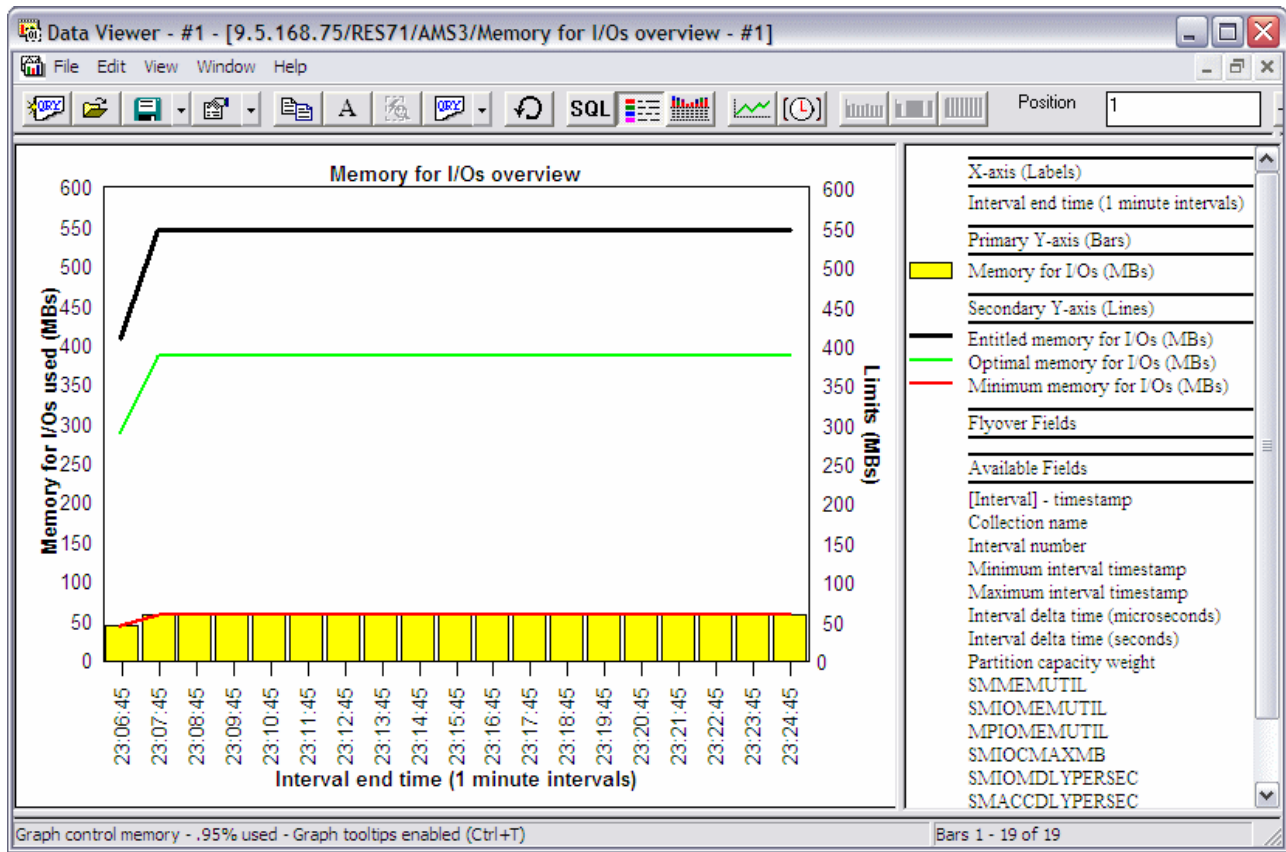


Figure 7-30 Memory for I/Os

Job Counts Graphs

If you right-click *Collection - Job counts* graphs the following options are available:

- ▶ Job Counts Figure 7-31
- ▶ Net jobs created
- ▶ Net jobs breakdown
- ▶ Job created/destroyed
- ▶ Job counts rankings (by job grouping, see Figure 7-32 on page 225)
- ▶ Net jobs breakdown rankings (by job grouping)




Figure 7-31 Job Counts

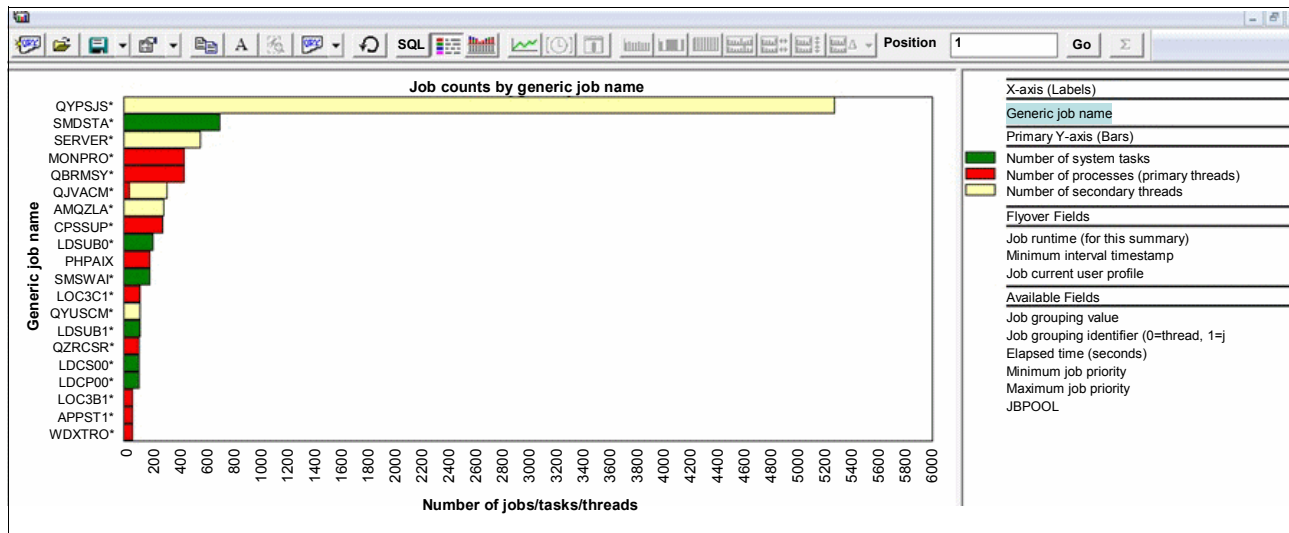


Figure 7-32 Job Counts by generic job name

Memory Pool Graphs

If you right-click **Collection** → **Memory pool graphs** you can generate graphs containing the following information:

- ▶ Memory pool consumption (by percentage), shown in Figure 7-33
- ▶ 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 vs non-DB faulting rates
- ▶ DB vs non-DB paging (pages read/written) rates




Figure 7-33 Memory pool consumption

The memory pool graphs also allow you to right-click the desired pool and time range in order to drill down and see the jobs within the desired 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 desired memory pool graph or utilize the *Historical Summary* analysis to graph multiple collections more easily.

Disk Configuration

A new *Disk configuration* folder under the collection contains information about the ASPs, IOPs, IOAs, and units on the system. This 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 additional reports showing the same disk configuration data, one is a flat table, the other is a tree. The tree provides counts and percentages of the units/IOAs/IOPS/ASPs within each prior level grouping. To access these reports, right-click *Collection - Disk configuration*. Figure 7-34 gives you an idea on how this looks.

Data Viewer - #1 - [9.5.168.75/ANALYSIS/DS6000/Disk configuration (tree) by ASP/IOP/IOA/Unit - #1]

File Edit View Window Help

Position 1

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>

Rows 1 - 25 of 70

Figure 7-34 Disk configuration by tree

Advanced Disk Graphs

As explained in “QAPMDISKRB” on page 195, 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 *Collection - Disk graph - Advanced*. An example is shown in Figure 7-35.

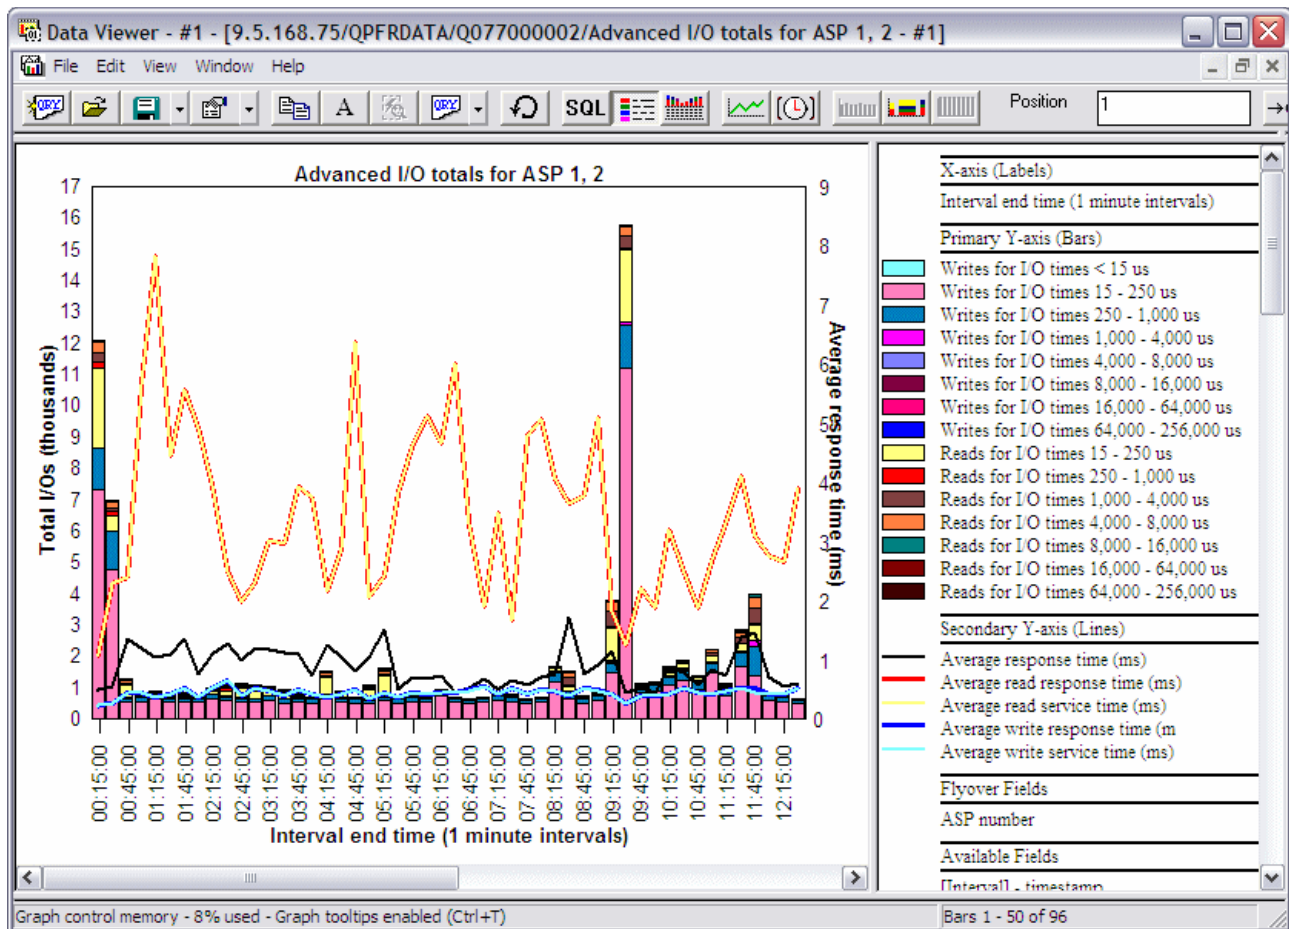


Figure 7-35 Advanced disk graphs

RIO HSL 12X Loop Graphs

Under the Communication graphs -> RIO HSL 12x loops folder a set of 12 graphs are now available to show HSL read/write throughput in various ways. (see Figure 7-36 on page 229) Options of using a filter of 0 MB/sec, 1 MB/sec, 100 MB/sec or 250 MB/sec are provided.




Figure 7-36 HSL Loop overview with 100MB/sec filter

After any of these graphs are opened, the current filter can be modified by right-clicking the legend. In graphs where there are large number of loops, the usage of the filters is recommended to group smaller values together.

Additional Graphs

Changes have been applied to a number of graphing capabilities, supporting release 5.4 and up:

- ▶ Communication graphs folder shows the following information:
 - Average IOP uses
 - Maximum IOP uses
 - SSL authentications
- ▶ Under *Disk graphs* a new graph *I/O size and Ethernet rates* is now available
- ▶ The collection overview wait graphs now show batch and interactive CPU use 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.

Starting with release 6.1, you now find 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 showing net pages allocated and net page frames requested. Net pages allocated are shown in megabytes and assumes the page size is 4 KB. Both sets of graphs includes the usual rankings graphs to graph the data by thread, job, generic job, user, and so forth.

7.4.7 Job Watcher

The folders available in the Job Watcher component have significantly changed. Instead of showing libraries containing Job Watcher data, new folders are available as shown in Figure 7-37 on page 230

- *Libraries* containing JW database file collections (filterable)
- A definitions folder providing a list of JW definitions on the system
- The rest of the folders are covered in section 7.4

Folder Name	Description
Scheduled jobs	iDoctor jobs scheduled to run
Active jobs	Work with active jobs
Subsystems	Work with iDoctor subsystems or all subsystems

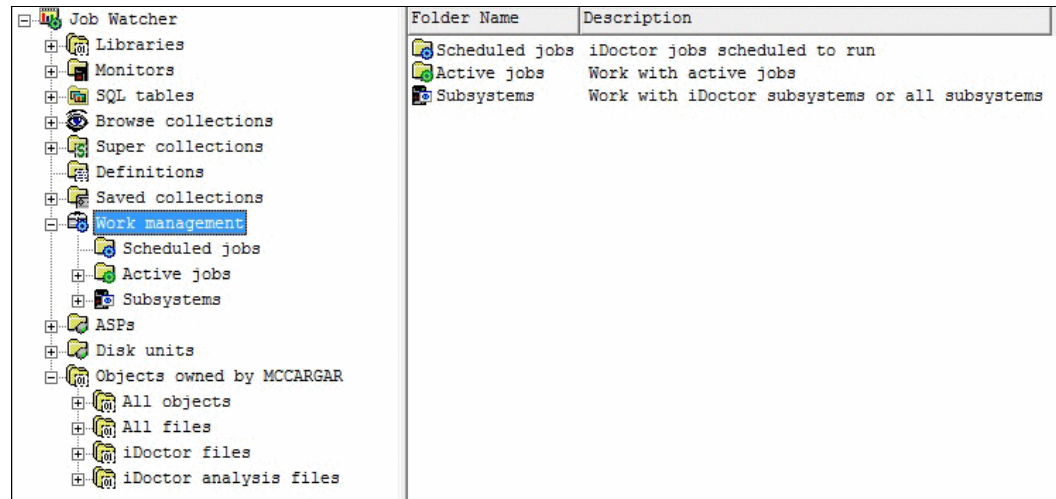


Figure 7-37 Job Watcher folders

Monitors

On the panel to start a Job Watcher (or Disk Watcher) monitor, you can specify the maximum collection size (in megabytes) for each collection running in the monitor.

The next set of changes apply to the following monitor commands: STRJWMON, STRPAMON, STRDWMON. These options can also be found in the GUI when starting a monitor.

The *Collection Overlap (OVLAP)* parameter is no longer used. The monitor will now detect that a new collection has started before ending the previous one.

The *Collection Duration (MAXSIZE)* parameter can now be specified in minutes with a decimal point (i.e. 59.5 minutes.)

When restarting a monitor if the *Maximum Historical Collections (COLNS)* parameter has been reduced added checks to delete the extra ones.

Deleting collections in a monitor is now done in a submitted job.

The following changes only apply 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 will be resubmitted if the RESUBMIT parameter is set to *YES and the current collection has ended prematurely.

Create Job Summary Analysis

Right-click *Collection - Analysis - Create Job Summary* to produce job totals for the desired jobs based on any filters given as shown in Figure 7-38.

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 7-38 Create Job Analysis

Collection Summary Analysis

As discussed previously a collection is summarized using right-click *Collection - Analyses - Analyze Collection* (for full options), or use *Collection - Run Collection Summary*. This new analysis is greatly simplified and many options that were previously on the *Summarize Collection* window have been removed. The only option that remains is the *Update Wait Bucket Actives + Idles* analysis and by default this analysis is not executed.

The Analyze Collection(s) window now has a new *Situations* button, which allows the user to customize which situational analysis option to execute, and the limits used for the situations (see Figure 7-39 on page 232). The selected analysis can now run as a batch job. A check box on the Analyze Collection(s) window allows you to indicate if this is to be done instead of running them in the Remote SQL statement status view (a separate GUI QZDASOINIT job). The same analysis and similar options are found in the CSI component.

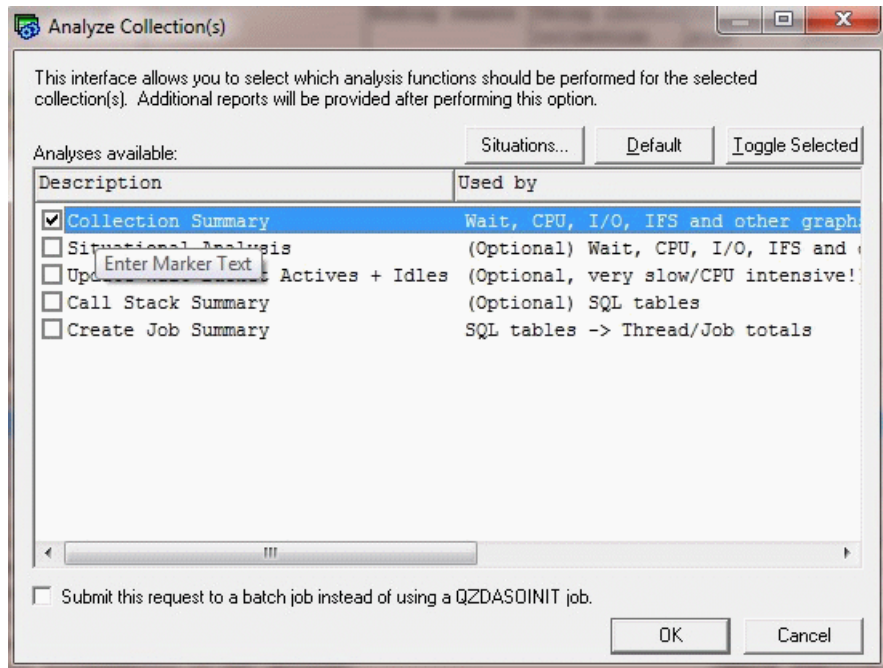


Figure 7-39 Analyze Collection(s) window

Advantages to running in batch are as follows:

- ▶ You can run start many analyses simultaneously
- ▶ You can start the analysis then kill your GUI session and it will keep on running
- ▶ You can start multiple analyses on multiple systems without waiting on the remote SQL statement status view to execute them in order.

Within the list of collections the status field indicates which files are not yet created yet.

Situational analysis

This option can be found under *Collection - Wait graphs* and has new situations:

- ▶ Concurrent write support not enabled
- ▶ Journal caching 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 due to 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

Top threads/tasks graphs

These graphs can be found under *Collection - Wait graphs - Top threads over time* and displays the threads/tasks that spent the most time in the desired wait bucket (such as CPU) over time

Objects waited on

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 CPU in the interval. Previously only jobs that used CPU in the interval were shown. Also added a checkbox to show segments waited on.

SQL server mode job information

For JW 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 7-40

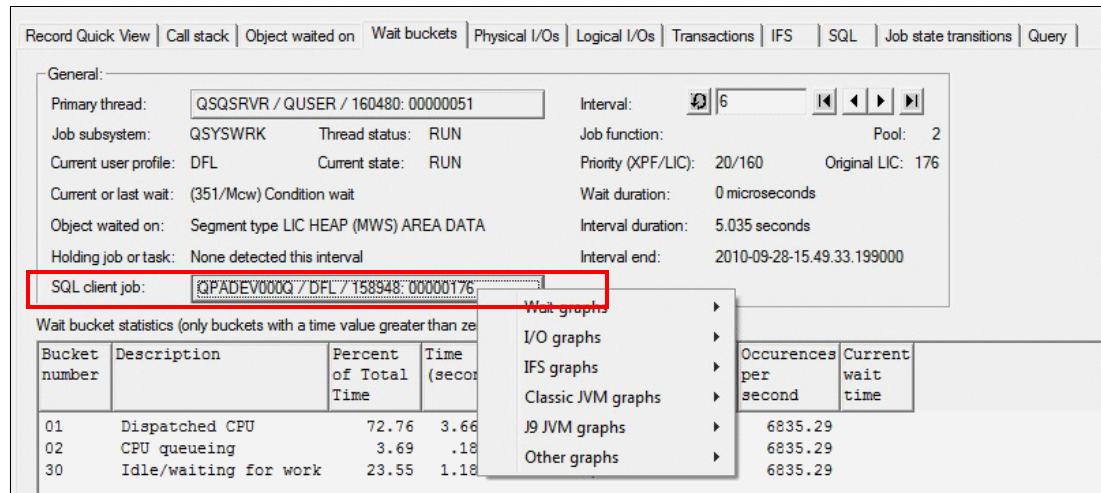


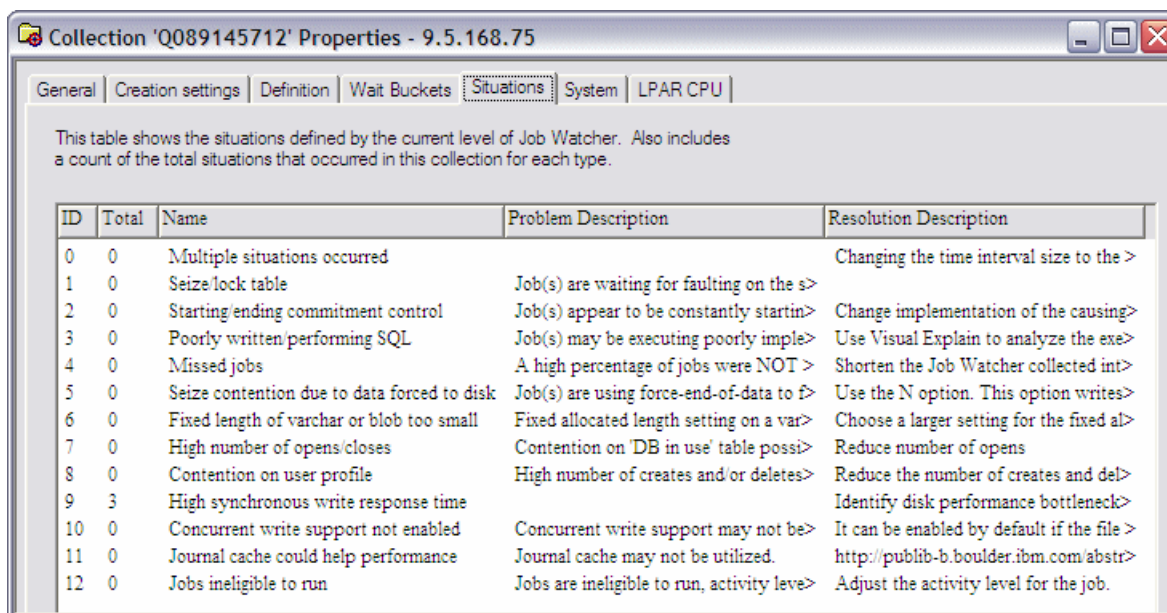
Figure 7-40 SQL client job drill-down options on the Interval Details - Wait Buckets window

Additional reporting options

There are several new reporting options for the situational analysis, similar to the call stack reports based on the selected situation. From these you can double-click a job/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 JW call stacks. J9 Java call stack entries are not usable with the call stack reports.

You will find a Situations tab in the Collection Properties (Figure 7-41), 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 7-41 Collection Situations

In the Interval Details interface, a button has been added to go to primary thread from a secondary thread. An example of this can be seen at the following URL:

http://www.youtube.com/watch?v=A_PZFK9Id18&feature=related

A new Call Stack Summary analysis has been added to identify the call stacks, waits and objects associated with the most frequently occurring call stacks found in the collection. The following URL has an example of this:

http://www.youtube.com/watch?v=NyycbX6eCL4&feature=more_related

Disk Watcher (DW)

The folders available in the Disk Watcher component have significantly changed. Instead of showing libraries containing Disk Watcher data, new folders are available as shown in Figure 7-42

- ▶ Libraries containing DW database file collections (filterable)
- ▶ A definitions folder providing a list of DW definitions on the system
- ▶ The rest of the folders are covered in 7.4.3

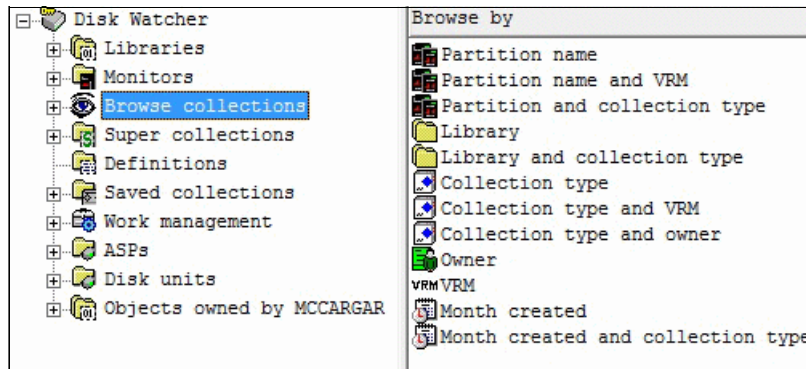


Figure 7-42 Disk Watcher

Collections

In the Start Disk Watcher Wizard (see Figure 7-43) you now find the option to collect the hardware resource file, to schedule the collections, and to check if there are any PTFs. Another parameter on this panel allows you to set the maximum collection size (in MB) for each collection.




Figure 7-43 Start a DW monitor

A copy function for Disk Watcher collections has been added. The CPYPFRCOL command (shipped with the OS) may also be used for this purpose.

Reporting

The graph titles match the naming convention used by the trace graphs. The word *pool* has been changed to *disk pool*, and *disk unit* to *disk path*.

A new trace DW drill down menu shows the top 25 I/O rates. The graphs in DW trace mode include the following (see Figure 7-44) 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

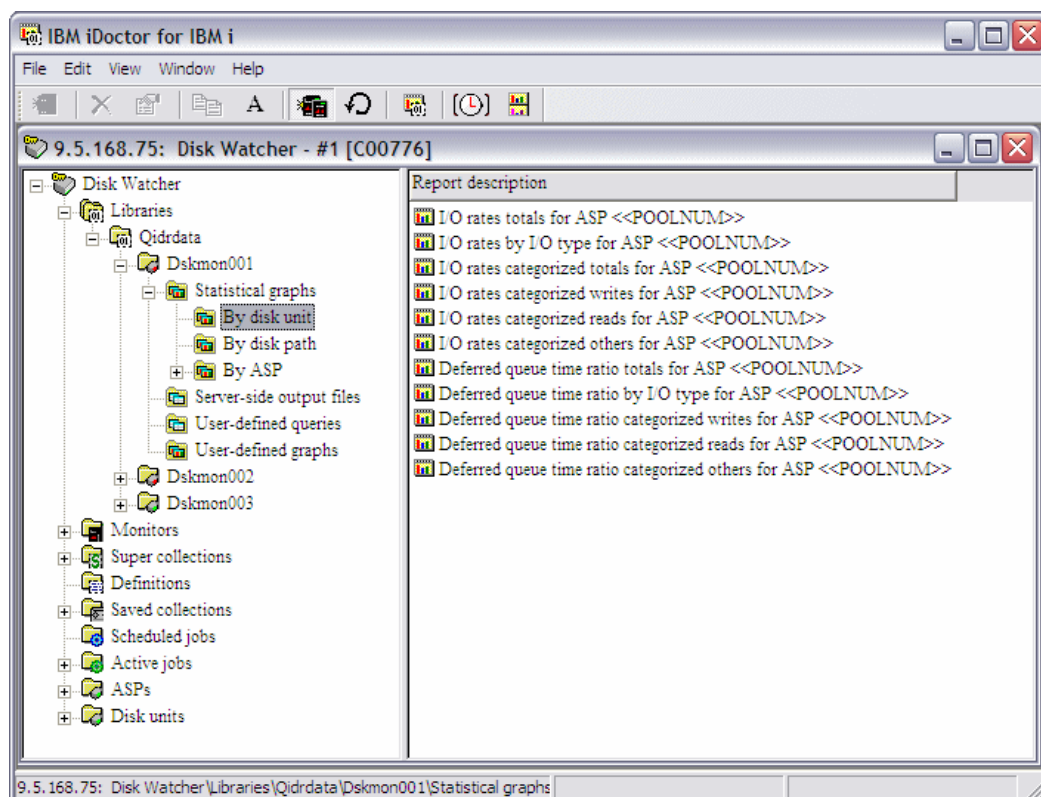


Figure 7-44 Statistical graphs - By disk unit

Figure 7-45 gives you an idea of how these charts look.

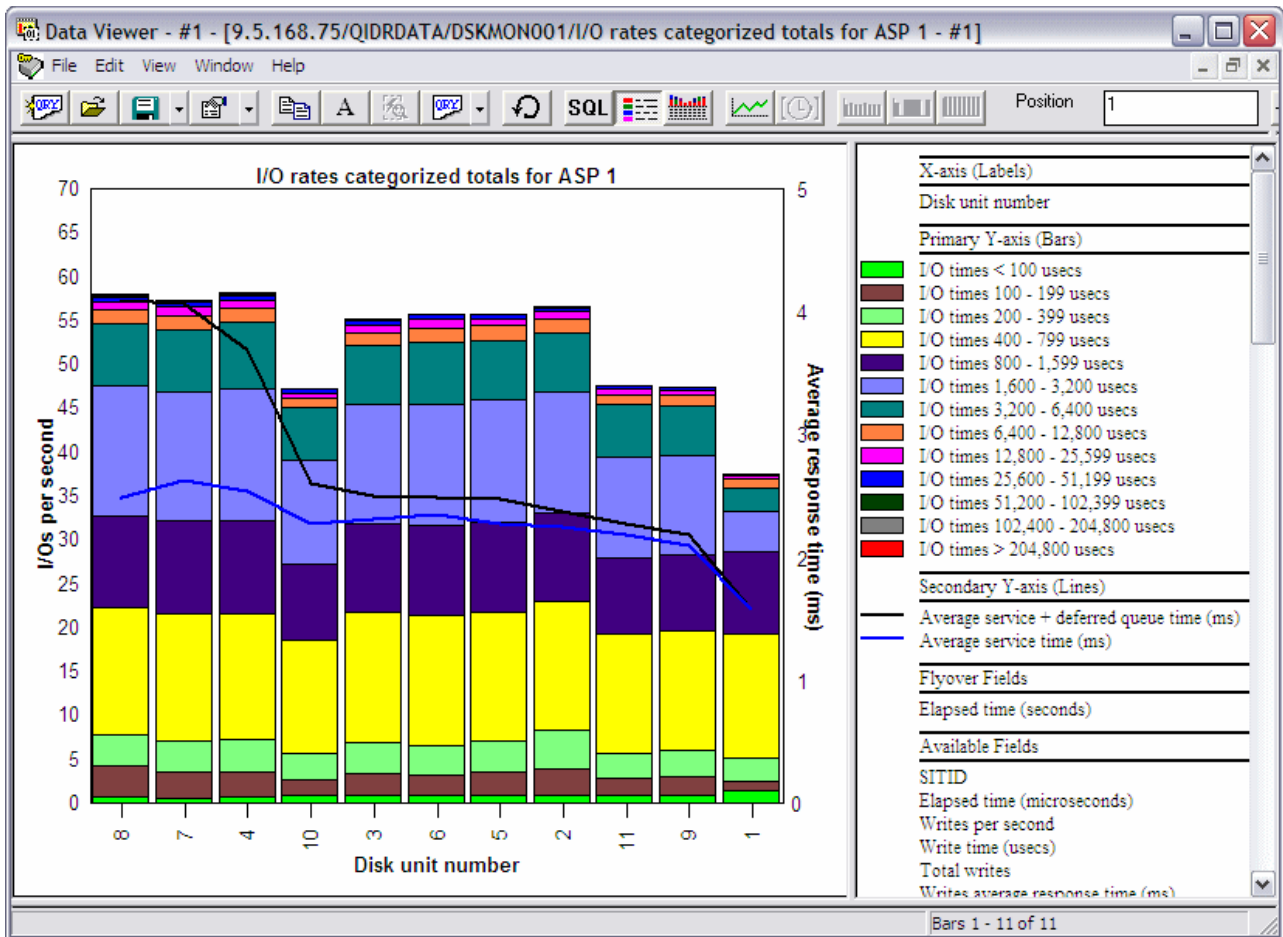


Figure 7-45 Categorized I/O totals

7.4.8 PEX Analyzer

The folders available in the PEX Analyzer component have significantly changed. Instead of showing libraries containing PEX Analyzer data, new folders are available as shown in Figure 7-46 on page 238:

- ▶ Libraries

This folder displays libraries containing PEX collections or libraries where active PEX collections created with the STRPACOL command (or the Start Collection Wizard) are currently running.

- ▶ Active collections

This folder allows you to work with any active PEX sessions on the system. This is similar to the ENDPEX command that lists the active PEX sessions

- ▶ PEX objects

This folder allows you to work with the PEX *MGTCOL objects on the system

- ▶ Definitions

This folder allows you to work with PEX definitions

- ▶ Filters

This folder allows you to work with PEX filters

- The rest of the folders are covered in section 7.4.3

Collection name	Collection library	Size (MB)	Created by	Partition collected on	Partition collected on VRM	Descrip...	Owner
CKRUN	QUSRSYS	512.1	CKOUR	IDOCEDU	V5R4M0	PEX data	QSYS
COL05	QUSRSYS	3.3	IDOCTOR05	IDOCEDU	V5R4M0	PEX data	QSYS
ID10COLL	QUSRSYS	3.3	IDOCTOR10	IDOCEDU	V5R4M0	PEX data	QSYS
ID10TPROF	QUSRSYS	1.8	IDOCTOR10	IDOCEDU	V5R4M0	PEX data	QSYS
MBTIPROF	QUSRSYS	3.3	IDOCTOR01	IDOCEDU	V5R4M0	PEX data	QSYS
MYCOLL02	QUSRSYS	3.3	IDOCTOR02	IDOCEDU	V5R4M0	PEX data	QSYS
SHIMETPROF	QUSRSYS	1.8	IDOCTOR01	IDOCEDU	V5R4M0	PEX data	QSYS
SUPCKOUR	QUSRSYS	513.1	CKOUR	IDOCEDU	V5R4M0	PEX data	QSYS
SUP01	QUSRSYS	1.8	CKOUR	IDOCEDU	V5R4M0	PEX data	QSYS
TPROF01	QUSRSYS	1.8	IDOCTOR01	IDOCEDU	V5R4M0	PEX data	QSYS

Figure 7-46 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 removed as of 6.1+: *MIPRECALL, *MIPOSTCALL, *JVAPRECALL and *JVAPOSTCALL
- Base event *CPUSWT added as of 6.1+
- Base events added as of 7.1: *PRCFDLSUSPEND, *PRCFLDRESUME, 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 the PDIO analysis, the STRPACOL (Start Pex Analyzer Collection) 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 forth) and the Netsize problem type. Not collecting with Format 2 now requires you to create your own PEX definition.

In PA 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. This allows the user 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 gets created using ADDPEXDFN by the GUI under the covers before STRPACOL is executed.

Analyses

Several changes have been implemented in this menu.

Classic analyses and support for the green panel QIDRPA/G* analysis commands have been removed and are 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 7-47). The menu also contains the *Analyze Collection* option which allows a user to kick off several analyses at once

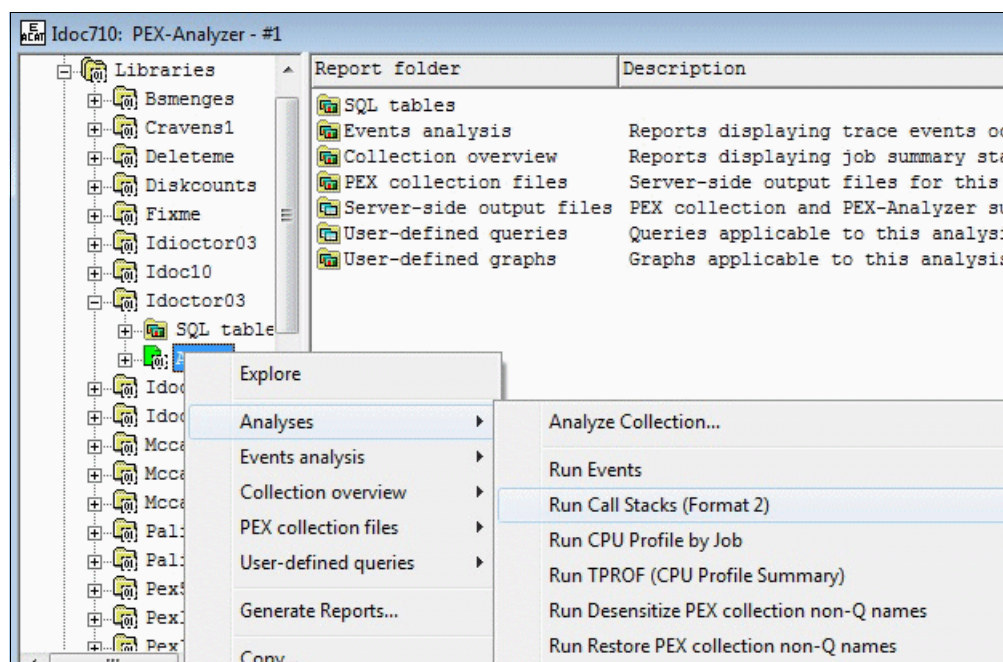


Figure 7-47 Analyses menu for a PEX collection

The TPROF analysis now has the tree table views that display the percentage of CPU hits in various ways.

PEX Analyzer has a new analysis called *Hot Sectors*. This SQL-based analysis is only available if the PDIO analysis has been run previously. It allows disk activity to be measured by portions of the disk address of the I/O in MB 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 have been collected. It shows estimated CPU consumption during the collection over time and CPU thread rankings for the desired time periods.

The MI user event analyses (LDIO, data area) now resolve the user program if Format 2 events were collected. This allows for MI entry/exit events to be excluded.

A database opens analysis, similar to the database LDIO analysis, provides statistics about the user program associated with the DB open events and reports 16 call level stacks, provided that 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 added graphs that show what the wait bucket time signature looks like for the desired thread/task (also known as TDE). See Figure 7-48. Additional drill-downs and reporting options are also provided.




Figure 7-48 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 CPU percent of total and the inline elapsed time percent of total information.

7.4.9 Additional Information

For additional information about the new features in iDoctor visit the iDoctor website.

http://www-912.ibm.com/i_dir/idoctor.nsf/

Presentations are created every few months with in-depth explanations of the latest features.



Virtualization

This chapter discusses the following topics:

- ▶ 8.1, “PowerVM enhancements” on page 242
- ▶ 8.2, “Additional OS levels combinations of server and client logical partitions” on page 253
- ▶ 8.3, “Hardware Management Console virtual device information” on page 255
- ▶ 8.4, “Virtualizing an optical device to IBM i client partitions” on page 258
- ▶ 8.5, “Virtual Partition Manager Enhancements” on page 268
- ▶ 8.6, “Partition Suspend and Resume” on page 270
- ▶ 8.7, “HEA Daughter cards” on page 271
- ▶ 8.8, “10 Gb FCoE PCIe Dual Port Adapter” on page 272

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 Power 6 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 (here after referred to as shared memory partitions), the hypervisor constantly 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 currently being used by shared memory partitions to other shared memory partitions that need to 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 known as a *paging space device*. Access to the paging space device is provided by a Virtual I/O Server (VIOS) logical partition known as the *paging service partition*. When the operating system of a shared memory partition attempts to access data that is located 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. See Figure 8-1 for an illustration of these AMS concepts.

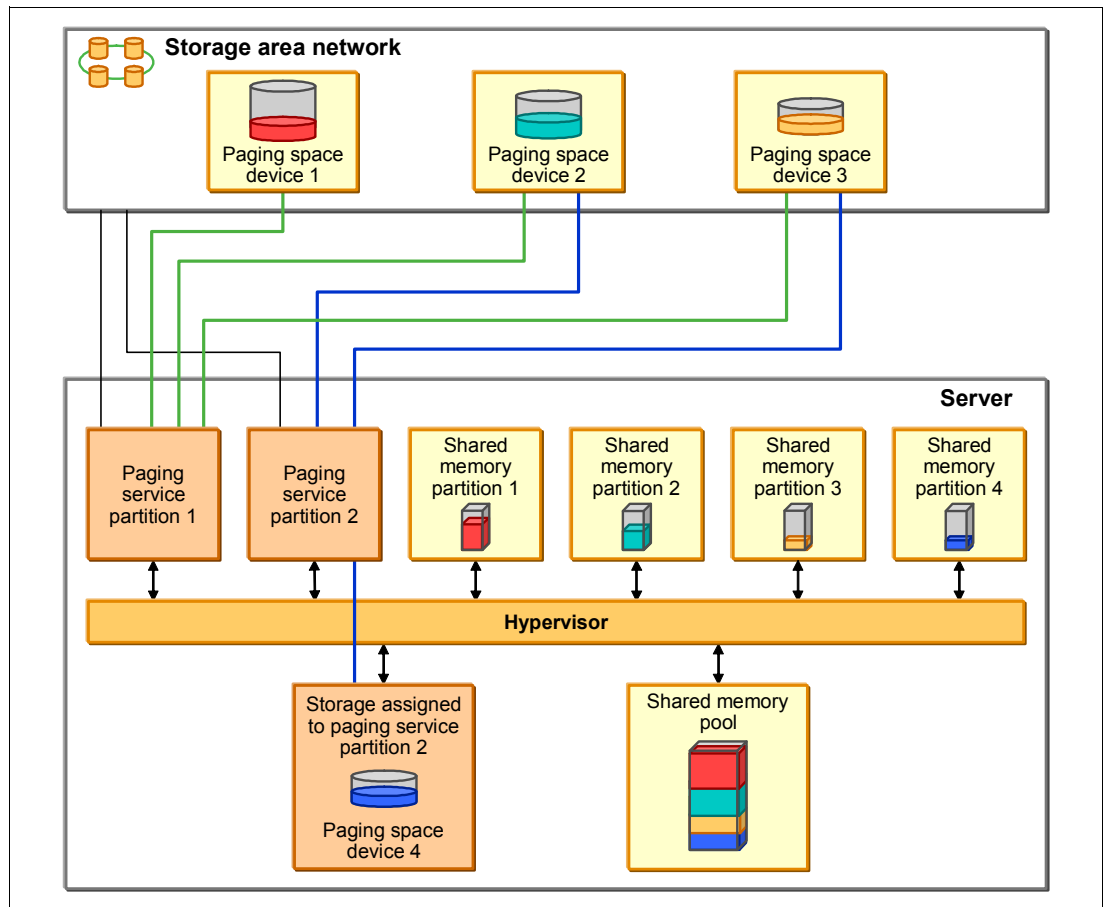


Figure 8-1 AMS concepts

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.

Note: 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, select a paging service partition, select a paging space device, and changing the IBM i client partition profile to use shared memory pool. See Figure 8-2.

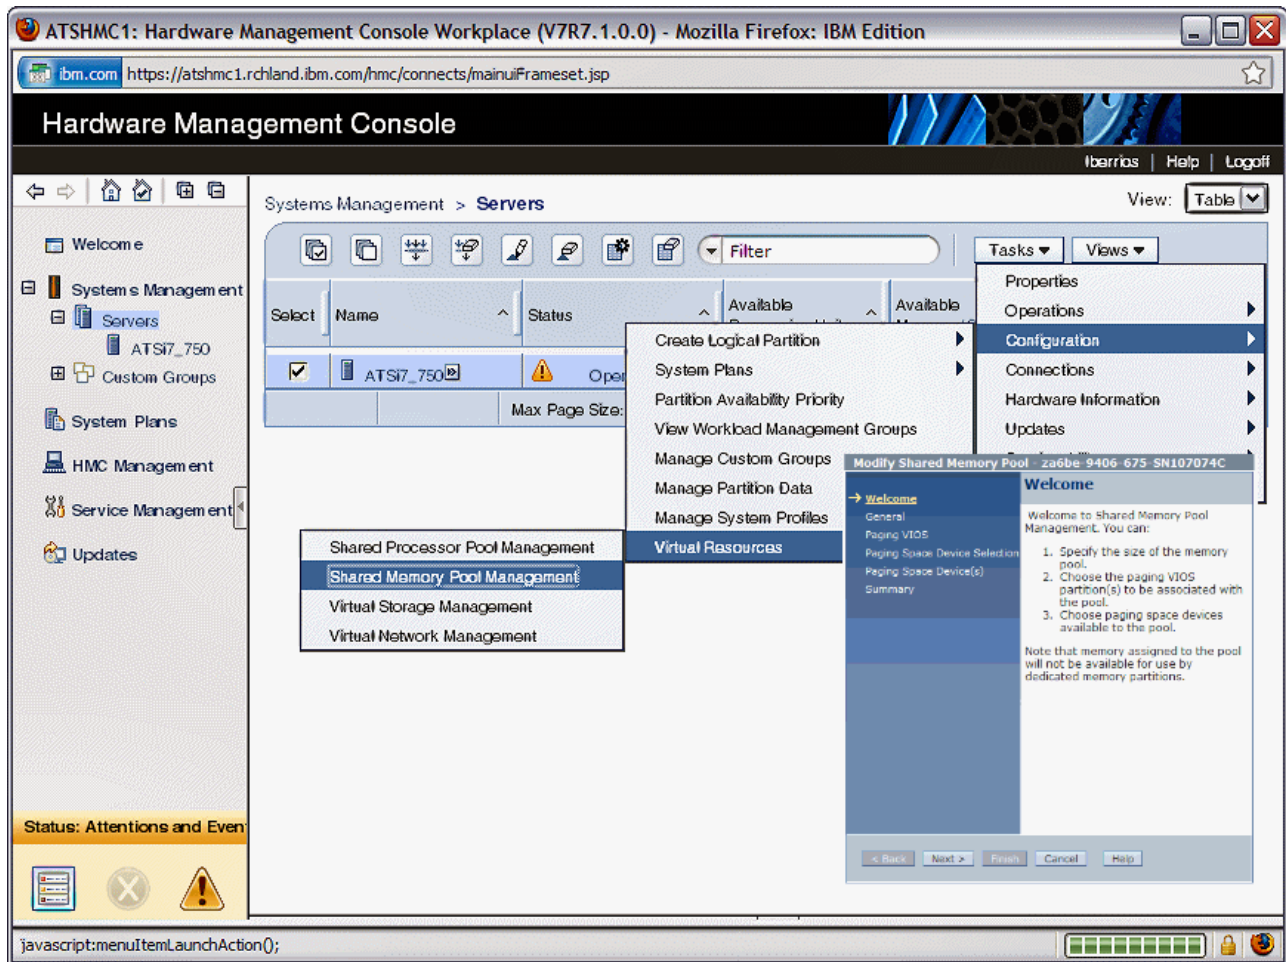


Figure 8-2 ASM setup using the HMC

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 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 251.

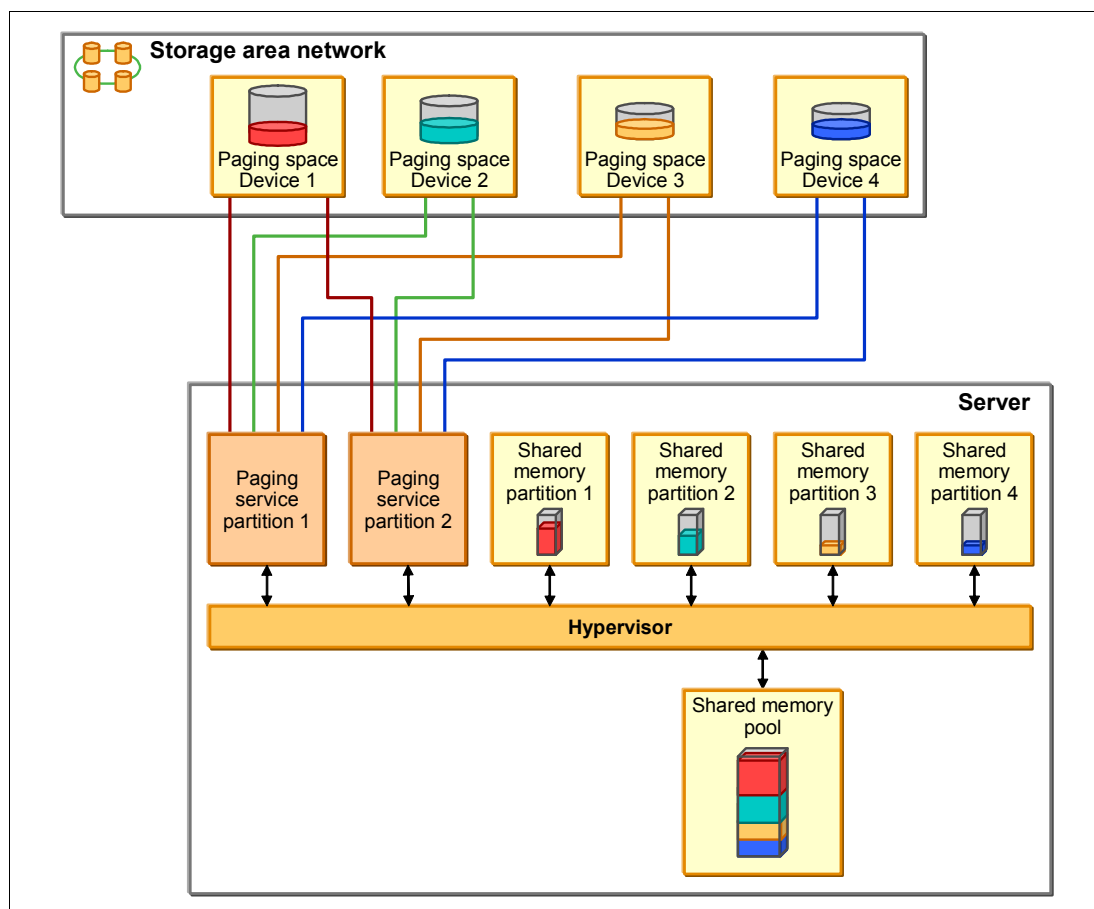


Figure 8-3 Redundant Paging Service Partitions

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 247. For multipath I/O for IBM i client partitions see 9.2.1, “Multipathing for virtual I/O” on page 290.

Note: When using redundant paging service partitions, common paging space devices must be located on SAN Disk Storage to enable symmetrical access from both paging service partitions.

System requirements for AMS is as follows:

- ▶ IBM Power Systems server or blade with POWER6 processors
- ▶ Virtual I/O Server (VIOS) 2.1.0.1Fixpack 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 6.1 TL3
- ▶ SUSE Linux Enterprise Server 11

Note: 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 302.

For an overview of AMS see the following web page:

<http://www.ibm.com/systems/power/software/virtualization/whitepapers/ams.html>

For more detailed information about AMS, see *PowerVM Virtualization Active Memory Sharing*, REDP-4470, at the following web page:

<http://www.redbooks.ibm.com/redpapers/pdfs/redp4470.pdf>

8.1.2 Enhanced support for IBM System Storage

On October 2009 IBM announced support for IBM System Storage DS5100 and 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 DS5020 are supported with IBM i 6.1 partitions through PowerVM VIOS attached to POWER6 processor-based servers and blades. See Figure 8-4 for a summary of SAN Disk Storage alternatives by IBM i. For more information about Storage Area Networks and IBM i see Chapter 9, “Storage and solid state drives” on page 273.

		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

Note: 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 via NPIV. See 8.1.3, “PowerVM Virtualization and I/O enhanced with NPIV” on page 247 for more information about NPIV.

For an overview on IBM i System Storage solutions see the *IBM i Virtualization and Open Storage Read-me First* at the following web page:

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 the following web page:

<http://www-03.ibm.com/systems/support/storage/config/ssic/>

For more information about available SAN Storage solutions for Bladecenter and IBM i see the *Bladecenter interoperability guide* at the following web page:

<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 in the following web page:

<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 LPARs. For NPIV the VIOS server acts as a FC pass-through instead of a SCSI emulator such as when using Virtual SCSI (see Figure 8-5).

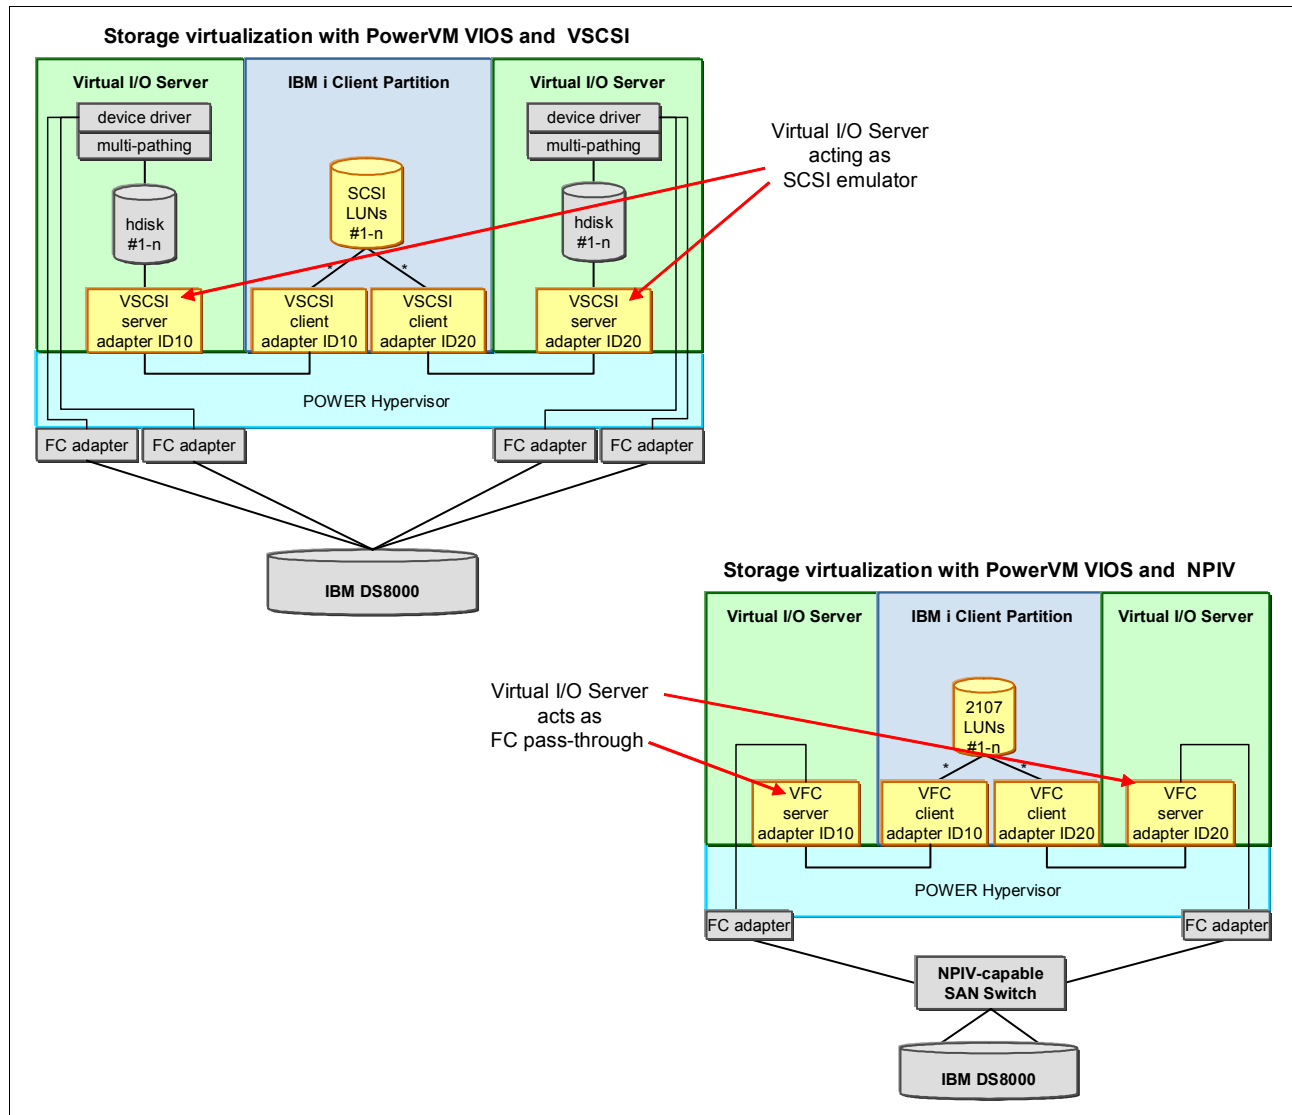


Figure 8-5 Comparing PowerVM storage virtualization with VSCSI and NPIV

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 LPAR, as shown in Figure 8-6.

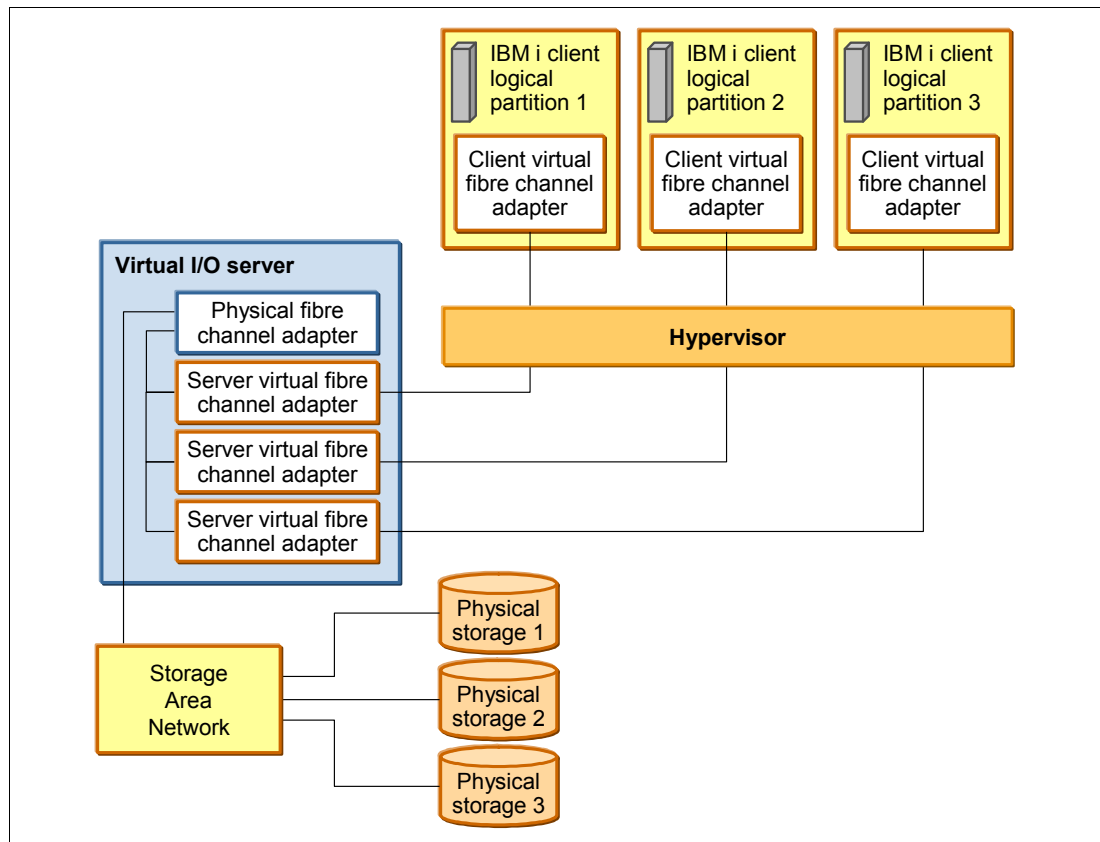


Figure 8-6 VIOS VFC server adapter and IBM i VFC client adapter

Two unique World-Wide 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 using NPIV:

- ▶ One VFC client adapter per physical port per partition
 - Intended to avoid single point of failure
- ▶ Maximum of 64 active VFC client adapter per physical port
 - Can be less due to other VIOS resource constraints
- ▶ 32,000 unique WWPN pairs per system platform
 - Removing adapter does not reclaim WWPNs
 - Can be manually reclaimed through CLI (mksyscfg, chhwres...)
 - or “virtual_fc_adapters” attribute
 - If exhausted, need to purchase activation code for more

Note: 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 using 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 LTO4 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 LTO3, LTO4, 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 via NPIV. This 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, flash copy, 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 use 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 take into consideration 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 this previously announced 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/adapters. It is an optional feature, containing four PCIe Gen2 low-profile (LP) slots for Gen1 or Gen2 adapter cards. It is physically very 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* article at the following web page:

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 the following web page:

<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 at the following web page:

<http://www-947.ibm.com/systems/support/supportsite.wss/docdisplay?lnodocid=MIGR-5073016&brandind=5000020>

For more information about IBM i on Power Blades see the *Read me first* and *Supported Environments* documents in the following link:

<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 with BladeCenter and Power Blades.

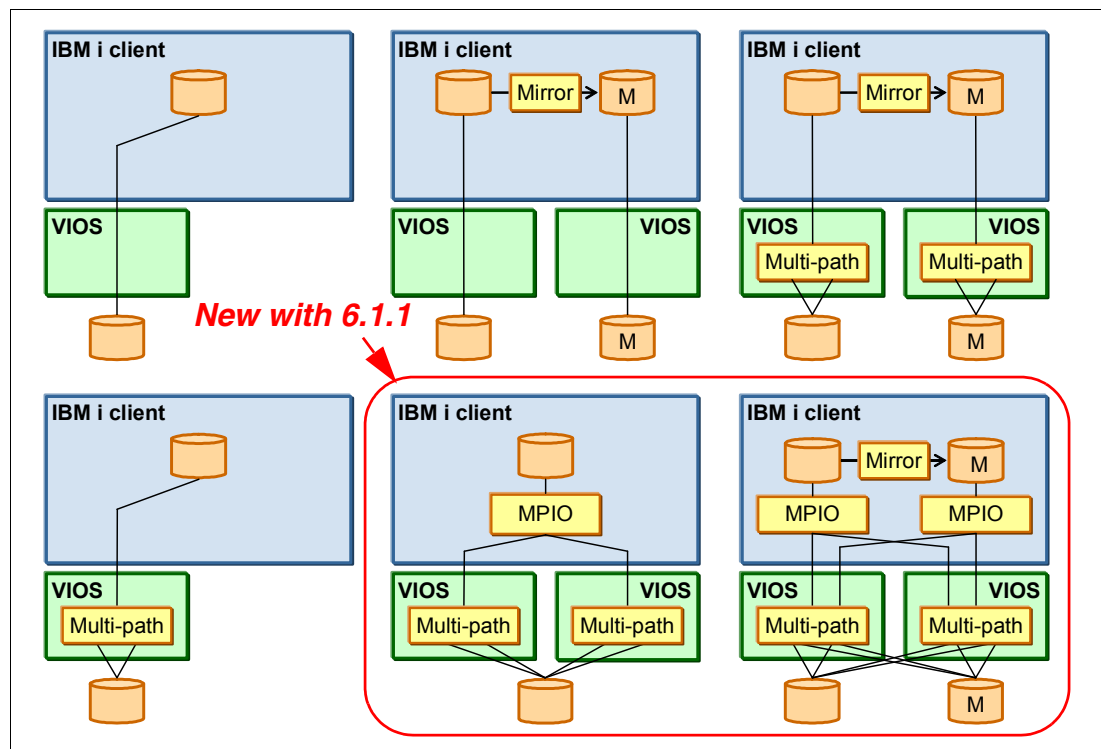


Figure 8-9 Redundant VIOS using VSCSI

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* article at the following web page:

http://www-03.ibm.com/systems/resources/systems_i_Virtualization_Open_Storage.pdf

8.2 Additional OS levels combinations of server and client logical partitions

IBM PowerVM continues to enable Power Systems with IBM i to achieve higher resource use by supporting additional OS levels combinations of server and client logical partitions, as shown in Figure 8-11.

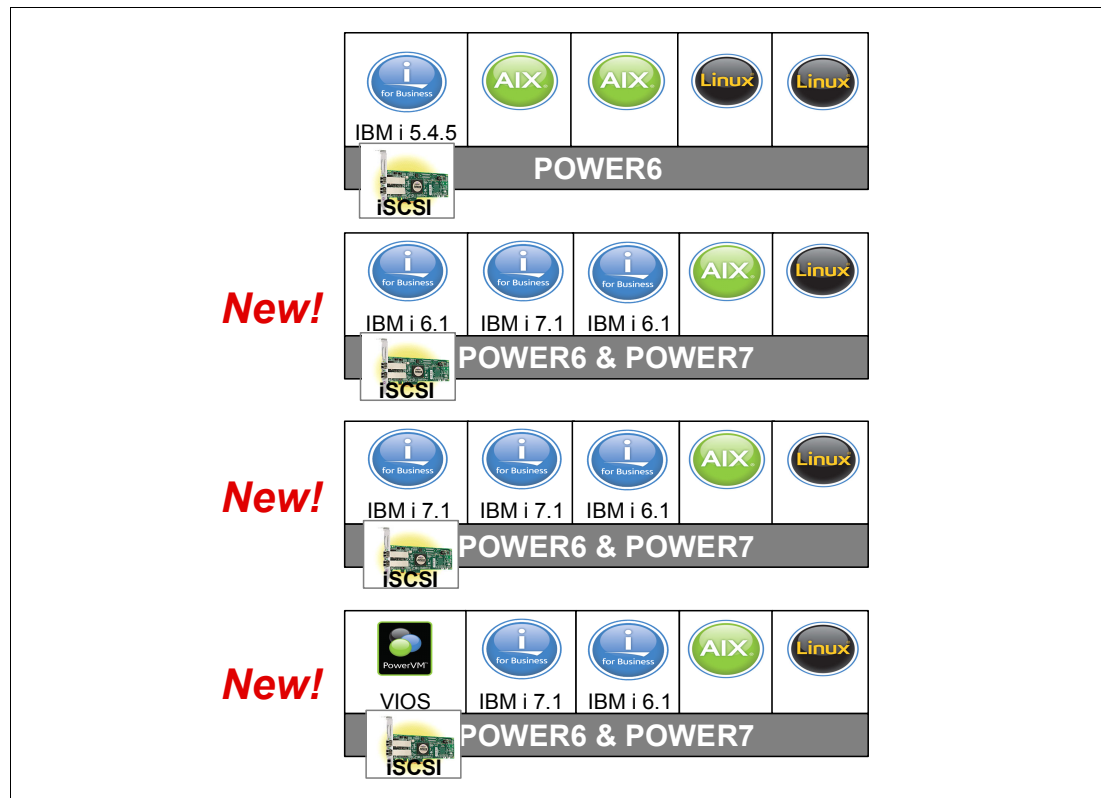


Figure 8-11 Additional OS level combination of server and client for IBM i and VIOS

- ▶ IBM i 6.1 server partition can provide virtual I/O resources to the following elements:
 - IBM i 7.1 and 6.1 or later client partitions
 - AIX 5.2, 5.3, 6.1 and SLES and Red Hat Linux client partitions
 - iSCSI attached System x and BladeCenter
- ▶ IBM i 7.1 server partition can provide virtual I/O resources to the following elements:
 - IBM i 7.1 and 6.1 or later client partitions
 - AIX 5.2, 5.3, 6.1 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 7.1 and 6.1 or later client partitions
 - AIX and Linux client partitions

The following list details benefits of using IBM i hosting:

- ▶ Same technology as IBM i hosting AIX, LINUX, and iSCSI x86 servers
- ▶ Use existing hardware investment
 - Create new IBM i 6.1 LPARs using only virtual hardware (No IOAs, IOPs, disk units, I/O slots necessary for client partitions), but can also use physical I/O.
- ▶ Rapidly deploy new workloads
 - Virtual disk created with 1 command or several clicks in System i Navigator
 - New LPAR, virtual resources deployed dynamically
 - Create test environments without hardware provisioning
 - Virtual resources allow new test environments of exact size to be created, deleted without moving hardware
 - Test new applications, tools, fixes in virtual test LPAR
 - Test the next release in the client partition

For more information about **PowerVM** see *IBM i 6.1 Technical Overview*, SG24-7713 available at the following web page:

<http://www.redbooks.ibm.com/abstracts/sg247713.html>

For more information about IBM i client partitions visit the IBM i Information Center at the following web page:

<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 on Figure 8-12.

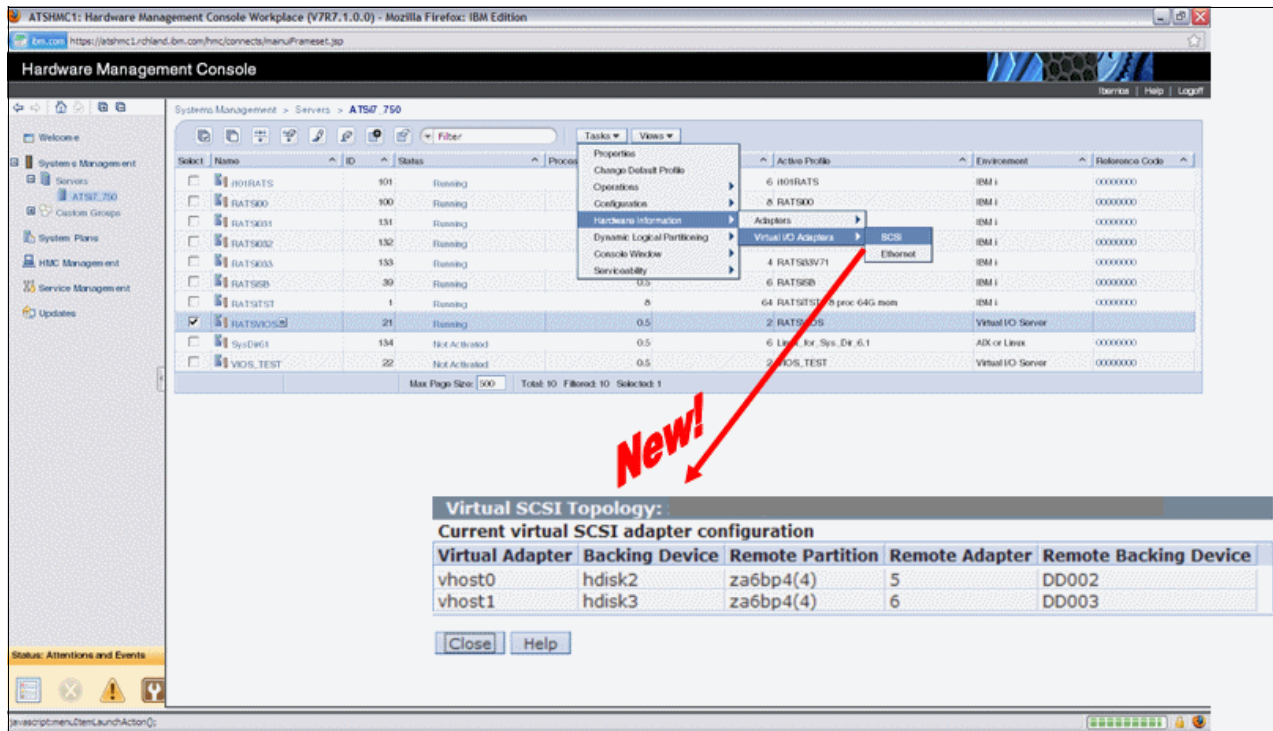


Figure 8-12 Virtual Device Information in the HMC

Tip: To access the Virtual Device Information in the HMC navigate to **Server** → **Partitions** → **VIOS LPAR** → **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 located 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 owned by the logical partition with the virtual server SCSI adapter. This column contains a value only if the virtual server SCSI adapter is already connected to a virtual client SCSI adapter.

Note: You can create virtual server SCSI adapters only for Virtual I/O Server and IBM i logical partitions. This window will always be 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 7.3.5 or later
- VIOS 2.1.2 (FP 22.1) or later
- IBM i 6.1.1 or later (+latest fixes)

Similar information is available by using the existing **lshwres** command on the HMC or Integrated Virtualization Manager (IVM) and 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 --subtype 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 visit the Hardware Information Center at the following web page:

<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/iphcg/lshwres.htm>

8.4 Virtualizing an optical device to IBM i client partitions

An IBM i 6.1 or later server partition on IBM Power Systems Power6 or later can virtualize a natively attached optical device to IBM i 6.1 or later client LPARs as illustrated in Figure 8-13.

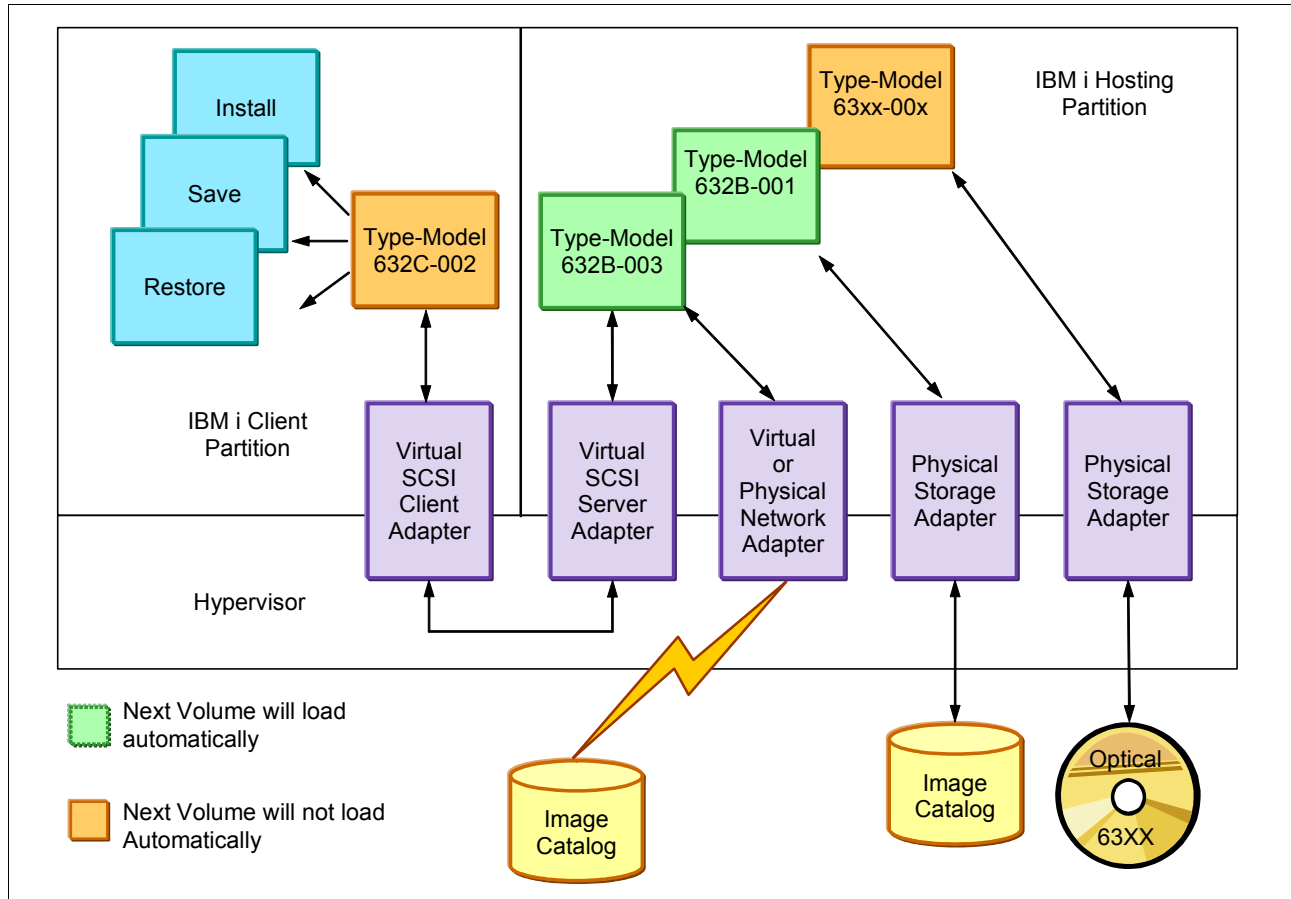


Figure 8-13 Virtualize a natively attached optical device to IBM i 6.1 or later client LPARs

This is achieved by creating a Virtual SCSI server (VSCSI) adapter in the IBM i server partition and a VSCSI client adapter in the IBM i client partition. In addition to the VSCSI connections, a Network Server Description (NSWD) is needed in the IBM i server partition. This is the same NSWD technology used to manage and provide a link of virtual storage resources to AIX and Linux client partitions and System x or BladeCenter x86 servers attached to Power Systems through iSCSI. By default, when a NSWD is enabled (varied on), a virtual SCSI connection exists between the IBM i server partition and the client partition where all optical devices available in the server partition is virtualized to the client partitions. The NSWD parameter RSTDDEVRS (Restricted device resources) can be used to specify which virtualized optical devices the client partition cannot access. An existing VSCSI server adapter and NSWD in the IBM i server partition can be used but creating new ones explicitly for optical I/O traffic is recommended.

Note: Power6 or later and IBM i 6.1 or later is required to support IBM i client partitions. However, prior to IBM i 6.1, an IBM i server partition can virtualize a natively attached optical device to linux client partitions or System x and BladeCenter x86 servers attached through iSCSI.

Note: VIOS server partitions can also virtualize a natively attached optical device to IBM i 6.1 or later client partitions. For more information see IBM Redbooks publication *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590. In this scenario only one client partition can use the virtualized optical device at a time. This is unlike when using IBM i as a server partition where multiple client partitions have shared access to the virtualized optical device.

A virtualized optical device in the server partition can be used for a D-mode Initial Program Load (IPL) and install of the client partition, and for installing program temporary fixes (PTFs) or applications. An image catalog can be used to automate the process by eliminating the need to physically load the next volume. If the image catalog is located in the Integrated File System (IFS) the IBM i server partition will see the device as a 632B-001 or as a 632B-003 if it is located in Network File System (NFS). The IBM i client partition sees the virtualized optical device provided by the IBM i server partition as a device 632C-002 regardless if the virtual device is a natively attached optical device or an image catalog.

Important: When an IBM i partition uses an image catalog natively for a multivolume save, restore, or install, the next volume loads automatically. When an IBM i server partition uses an image catalog to provide a virtualized optical device to an IBM i client partition, the next volume will not load automatically for the client partition.

Follow these steps to virtualize a natively attached optical device from an IBM i 6.1 or later server partition to IBM i 6.1 or later client partitions:

1. Use the managing HMC to create a new VSCSI server adapter on the IBM i server partition (as illustrated in Figure 8-14 on page 260). Perform the following steps to do so:
 - a. In the navigation pane open **Systems Management** → **Servers**, and click the managed system on which the server IBM i server partition resides.
 - b. Select the IBM i server partition, click the **Tasks** button, and choose **Dynamic Logical Partitioning** → **Virtual Adapters**.
 - c. Click **Actions** and choose **Create** → **SCSI Adapter**.

- d. Use the default VSCSI Adapter number or provide your own. Write down the VSCSI Adapter number as you need it in a later step. In Figure 8-14, the number 31 was provided as the Virtual SCSI Adapter number. In the **Type of adapter** field, select **Server**, and click **OK**.
- e. Create the VSCSI adapter within the partition profile for the IBM i server partition so that the VSCSI adapter continues to exist after you restart the partition.

Note: In the Create Virtual SCSI Adapter window, you can select **Any client partition** or **Only selected client partitions** to determine the virtualized optical device as shown in Figure 8-14.

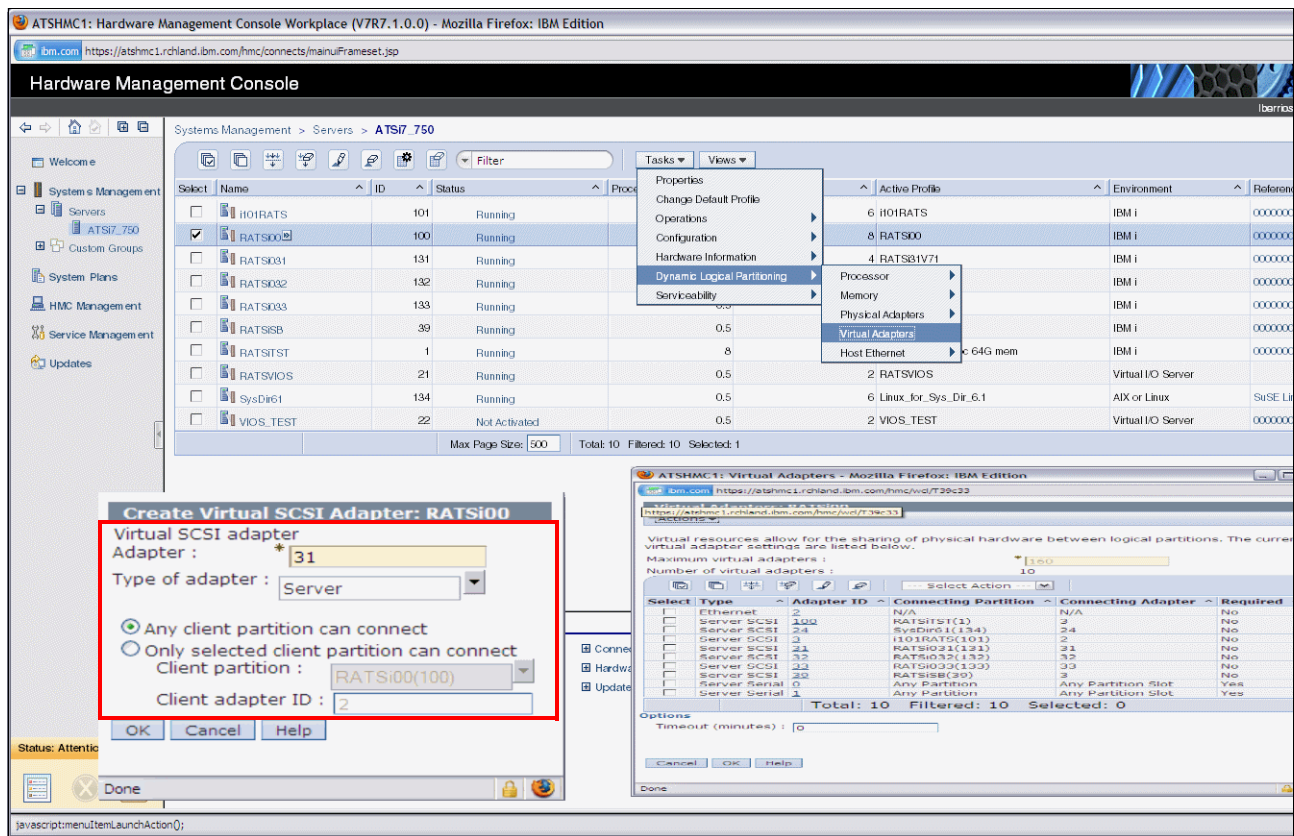


Figure 8-14 Create VSCSI Server Adapter

1. Create a Network Server Description (NWSA) on the IBM i server partition. Perform the following steps to do so.
 - a. Determine the correct VSCSI server resource name by entering the WRKHDWRSC *CMN command on the command line of the IBM i server partition. Look at the controller resources with type 290B. Use **option 7** to display the resource details (the newly created resource from step 1d is located at the bottom of the page).

Look at the last digits of the location code listed as Cxx where x corresponds to the virtual adapter number you wrote down on step 1d. Write the Resource Name (CTLxx) as it be used later to specify the VSCSI Server Resource Name (see Figure 8-15)

Work with Communication Resources System: RATS100

Type options, press Enter.
 5=Work with configuration descriptions 7=Display resource detail

Opt	Resource	Type	Status	Text
	CTL08	290B	Operational	Comm Adapter
	CMB11	290B	Operational	Comm Processor
	CTL03	290B	Operational	Comm Adapter
	CMB07	290B	Operational	Comm Processor
	CTL09	290B	Operational	Comm Adapter
	CMB19	290B	Operational	Comm Processor
	CTL10	290B	Operational	Comm Processor
	CMB03	290B	Operational	Comm Processor
7	CTL05	290B	Operational	Comm Processor

Display Resource Detail

Resource name : CTL05
 Text : Comm Adapter
 Type-model : 290B-001
 Serial number : 00-00000
 Part number :

Location : U0233.C0B.100417P-V100-C31

Logical address:
 SPD bus: 255
 System bus: 128
 System board:

Press Enter to continue.

F3=Exit F5=Refresh F6=Print F12=Cancel

Figure 8-15 Determining the virtual SCSI server resource name

- b. At the IBM i command line on the IBM i server 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 8-16, which corresponds to the name of the IBM i client partition in this example.

Resource Name: Provide the resource name from step 2a. CTL05 was used in this example.

Network server type: *GUEST

Server operating system: *OPSYS

Online at IPL:*YES

Code page: 437

Restricted service resources: Use ***NONE** for a NWSD that is to only provide virtual optical device resources.

Power Control: Use ***NO** if the NWSD created is to only to virtualize an optical device. Use ***Yes** if the NWSD is also used to provide virtual disk storage)

Note: To restrict the optical devices on all of the other NWSDs, use the restricted service resources (*ALLOPT) parameter in the CRTNWSD command for the particular NWSD or you can provide a specify the optical device name.

```

Create Network Server Desc (CRTNWSD)

Type choices, press Enter.

Network server description . . . NWSD          > CLIENT
Resource name . . . . . RSRCNAME             > 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 . . . . . MLTPTHGRP          *NONE
                                + for more values
                                                                More...

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

```

Figure 8-16 Provide NWSD name, resource name, server type and operating system

2. Start the NWSD on the IBM i server partition by entering the WRKCFGSTS *NWS IBM i command.
3. Type **1** next to the NWSD that you want to start and press **Enter**, as shown in Figure 8-17.

```
Work with Configuration Status                                     System:  CLIENT31

Position to . . . . . Starting characters

Type options, press Enter.
  1=Vary on   2=Vary off   5=Work with job   8=Work with description
  9=Display mode status   13=Work with APPN status...

Opt Description      Status      -----Job-----
   B3SI              VARIED OFF
   B3SIPP             VARIED OFF
   B3SIPNET           VARIED OFF
   B3SIPTCP           VARIED OFF
1  CLIENT31          ACTIVE
   EETEST             VARIED OFF
   FANTTEST           ACTIVE
   I101RATS           ACTIVE
   KNORR              VARIED OFF

More...

Parameters or command
====>
F3=Exit   F4=Prompt   F12=Cancel   F23=More options   F24=More keys
```

Figure 8-17 Starting the NWSD

4. Create the VSCSI client adapter in the IBM i client partition as shown in Figure 8-18.
 - a. In the navigation pane, open **Systems Management** → **Servers**, and click the managed system on which the IBM i client logical partition resides.
 - b. Select the IBM i client partition, click **Tasks**, and choose **Dynamic Logical Partitioning** → **Virtual Adapters**.
 - c. Click **Actions** and choose **Create** → **SCSI Adapter**.
 - d. Use the default VSCSI adapter number or provide your own. In the example of Figure 8-18 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, as shown in the example in Figure 8-18 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.

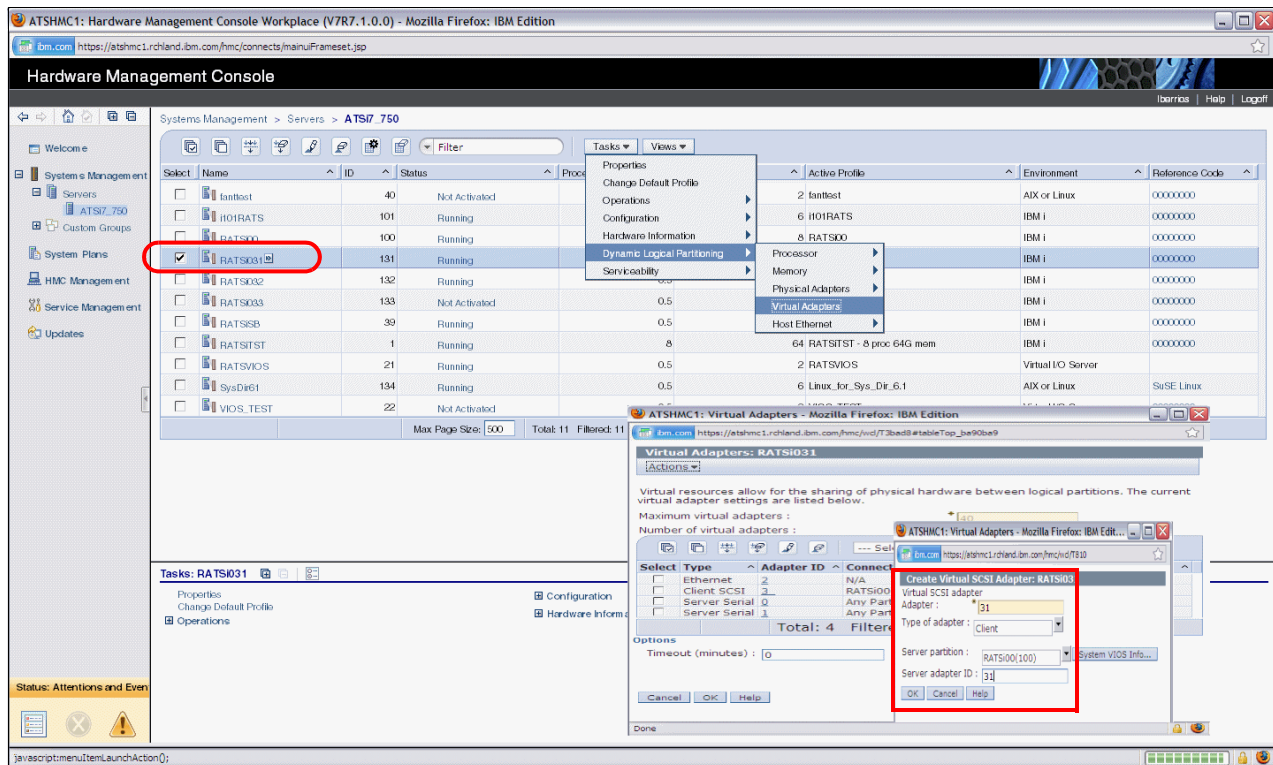


Figure 8-18 Create VSCSI client adapter

5. Enter the WRKHDWRSC *STG IBM i command to locate the virtual optical device in the IBM i client partition.
 - a. Select **option 7** to display resource detail next to each of the CMBxx resources listed, as shown in Figure 8-19.

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 8-19 Use Option 7 to locate Virtual Adapter number

- i. Look at the last digits for the location code Cxx where xx corresponds to the virtual adapter number you wrote down in step 1d, as shown in Figure 8-20.

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
		More...
Press Enter to continue.		
F3=Exit F5=Refresh F6=Print F12=Cancel		

Figure 8-20 Location of the virtual adapter

- b. When you find the correct *CMBxx* resource, look for the *DC0xx* resource and select **option 9** to work with resources, as shown in Figure 8-21.

Work with Storage Resources		System: CLIENT031
Type options, press Enter.		
7=Display resource detail 9=Work with resource		
Opt	Resource	Type-model Status Text
	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 8-21 Use option 9 to Work with resource

The virtualized optical device provided by the IBM i server partition is shown as device type **632C-002** in the IBM i client partition as shown in Figure 8-22.

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
F3=Exit F5=Refresh F6=Print F12=Cancel				Bottom

Figure 8-22 Virtualized optical device shown on IBM i client partition as a 632C-002

For more information about image catalog see the IBM i Information Center. Search to “Virtual optical storage.” The IBM i Information Center is located at the following web page:

<http://publib.boulder.ibm.com/iseriess/>

Tip: For Image Catalog information look in the IBM i Information Center by navigating to: **IBM i and related software → Installing, upgrading, or deleting IBM i and related software → IBM i software reference → Installation devices and media → Image catalog for a virtual device.**

8.4.1 Support for embedded media changers

This 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 use of image catalogs for unattended installs of client partitions. This switching capability also allows users to manually switch media in a client virtual optical device without requiring authority to the serving partition. This is accomplished via the image catalog interface WRKIMGCLGE *DEV command interface on the client partition.

8.4.2 IBM i to IBM i virtual tape support

A simple, cost-effective virtual tape solution is now provided. An IBM i server partition can be used to share a tape drive among multiple client partitions without the use of VIOS. With an IBM i 7.1 server partition and either an IBM i 7.1 client partition with the latest Technology Refresh or an IBM i 6.1 client partition with 6.1.1 machine code, IBM i 7.1 with the latest

Technology Refresh supports virtualizing LTO3, LTO4, LTO5, DAT160, and DAT320 tape drives, including drives in a TS2900, TS3100, and TS3200 tape library when the tape library is in sequential mode. See info APAR II14615 for a complete list of supported devices and required PTFs.

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 are enabled in VPM. Client IBM i partitions, that 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. The user is then be able to use Network Storage Spaces (NWSSTG) and Network Storage Descriptions (NWSD) in the IBM i I/O server partition to define the storage for the client partitions. Tape, disk, and optical are allowed to 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 is similar in functionality to the Shared Ethernet Adapter (SEA) support provided by a Power Systems Virtual I/O Server (VIOS) partition.

This works by putting two Ethernet adapters (one physical, one virtual) into a mode where they can receive traffic that is not specifically destined for their address, and selectively sending those frames onto the other network according to the IEEE 802.1D standard ("bridging" the frames). Because of this, frames 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.

- ▶ Any Ethernet resource that supports line speeds of 1Gbps or greater is supported, except for Host Ethernet Adapter (HEA) resources (Host Ethernet Adapter already supports the ability for multiple partitions to use a single physical port by assigning each partition a logical port).
- ▶ It 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 using a Hardware Management Console, create a virtual Ethernet adapter for the desired VLAN ID, and check the "Access external network" box to indicate that this virtual Ethernet adapter will be used to bridge traffic to the physical network.
- ▶ If using the IBM i Virtual Partition Manager, the virtual Ethernet adapter will automatically be created with the ability to 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 Recommended practices

It is recommended that the selected Ethernet resources only be used for layer-2 bridging (not for IBM i TCP/IP configuration). There is significant processing overhead 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 will be bridged between the two networks, and any other partitions with virtual Ethernet adapters on the same VLAN as the new virtual Ethernet resource will be 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 prior to the latest Technology Refresh. If equivalent behavior is desired:

- ▶ Use the "Copy" option on WRKLIND to make a temporary copy of the line description.
- ▶ Delete the existing line description.
- ▶ Use the "Copy" option again on WRKLIND to replicate the original line description, specifying the desired Bridge identifier.
- ▶ 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 will be returned when trying to vary on any more line descriptions with that Bridge identifier, indicating that the configuration is in error. For a given 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.

As mentioned above, 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 will be returned when trying 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).

Note: In IBM i V7R1, an Ethernet line description's Bridge identifier will not be 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, thereby 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 make use of the resources from the suspended partition.

8.6.1 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
- ▶ AIX Version 6.1 Technology Level 6 Service Pack 3
- ▶ For an IBM i logical partition, the logical partition must be running IBM i Version 7.1 with the latest Technology Refresh.
- ▶ When a logical partition is suspended, the reserved storage device contains the state 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 is capable of suspension without any restrictions. However, when you activate a logical partition with a specific profile, checks are performed for any of the restrictions 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 Restrictions for Suspend / Resume

- ▶ The logical partition must not have physical I/O adapters 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 assigned to the logical partition.
- ▶ The logical partition must not have physical I/O adapters 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.
- ▶ The logical partition must not have a virtual SCSI optical or tape device assigned to the logical partition.

The following additional restrictions apply for IBM i logical partitions that are enabled for suspension:

- ▶ You cannot activate the logical partition with a partition profile which has a virtual SCSI server adapter.
- ▶ You cannot activate the logical partition with a partition profile which 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 visit the IBM Prerequisite at the following web page:

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 the following web page:

<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/p7hat/iphatphibreqs.htm>

8.7 HEA Daughter cards

POWER7 HEA Daughter cards provide integrated I/O connectors for a CEC enclosure. The connections can be virtualized into the system LPARs. All of the connectors are on the rear bulkhead of the CEC enclosure. Choices of features are:

- ▶ #1824 and #1832 for a 1 Gb HEA daughter card with 4 ports
- ▶ #1825 and #1833 for a Fiber 10 Gb HEA daughter card with 2 ports
- ▶ #1826 and #1837 for a Copper 10 Gb HEA daughter card with 2 ports

8.8 10 Gb FCoE PCIe Dual Port Adapter

10 Gb FCoE PCIe Dual Port Adapter (#5708) is a high-performance, 10 Gb, dual port, PCIe Converged Network Adapter (CNA) utilizing SR optics. Each port can provide NIC (Network Interface Card) traffic and Fibre Channel functions simultaneously. IBM i supports use of this adapter through VIOS.



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:

- ▶ 9.1, “General and direct-attached storage management enhancements” on page 274
This section discusses enhancements in the areas of disk management, ASP encryption, performance instrumentation, and tape applicable to IBM i internal direct-attached and possibly external storage.
- ▶ 9.2, “SAN storage management enhancements” on page 290
This section summarizes the enhancements 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.
- ▶ 9.3, “SSD storage management enhancements” on page 302
This section discusses IBM i solid state drive management improvements 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 applicable to IBM i internal direct-attached storage (DAS) and possibly external storage area network (SAN) storage are discussed:

- ▶ 9.1.1, “Concurrent removal of disk units” on page 274
- ▶ 9.1.2, “Hot spare for mirroring” on page 276
- ▶ 9.1.3, “Dual storage I/O adapters” on page 277
- ▶ 9.1.4, “Encrypted ASP enhancements” on page 282
- ▶ 9.1.5, “Disk response time buckets enhancements” on page 284
- ▶ 9.1.6, “CEC node level mirroring” on page 285
- ▶ 9.1.7, “EXP24S SFF Gen2-bay drawer (#5887)” on page 288
- ▶ 9.1.8, “Higher Capacity 10K RPM SFF SAS disk drives” on page 288
- ▶ 9.1.9, “Tape performance instrumentation” on page 289
- ▶ 9.1.10, “Tape library resource name changes for IOP-less IOA attachment” on page 289
- ▶ 9.1.11, “Tape library unreadable barcode changes for IOP-less IOA attachment” on page 289
- ▶ 9.1.12, “DVD/Tape SAS External Storage Unit for Power 795 CEC Rack” on page 290

9.1.1 Concurrent removal of disk units

Concurrent removal of disk units with IBM i 7.1 is now also supported for SYSBAS (for example, system ASP and user ASP disk units), eliminating the previous need for downtime to IPL to 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 navigating to **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 an non-configured after the removal action.

Remove Units from Configuration

Type options, press Enter.

4=Remove unit from configuration

OPT	Unit	Serial	Type	Model	Resource	Status
		ASP Number			Name	
	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 solid state drives (SSDs) are respected by the remove function (for example, DB2 or UDFS media preferences) (See 9.3, “SSD storage management enhancements” on page 302) are honored as much as possible as long as 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 cancelled at any time. The pause operation prevents further data allocations on the disk units selected for removal similar to the *ENDALC option in STRASPBAL.

Note: 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

Usage of hot spare disk units for mirroring is newly supported with 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 and that no manual intervention is required for resuming drive protection. Though 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, the user might still want to manually replace the newly used hot spare with a replacement drive to reestablish the original protection level, e.g. 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 has been 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 **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.

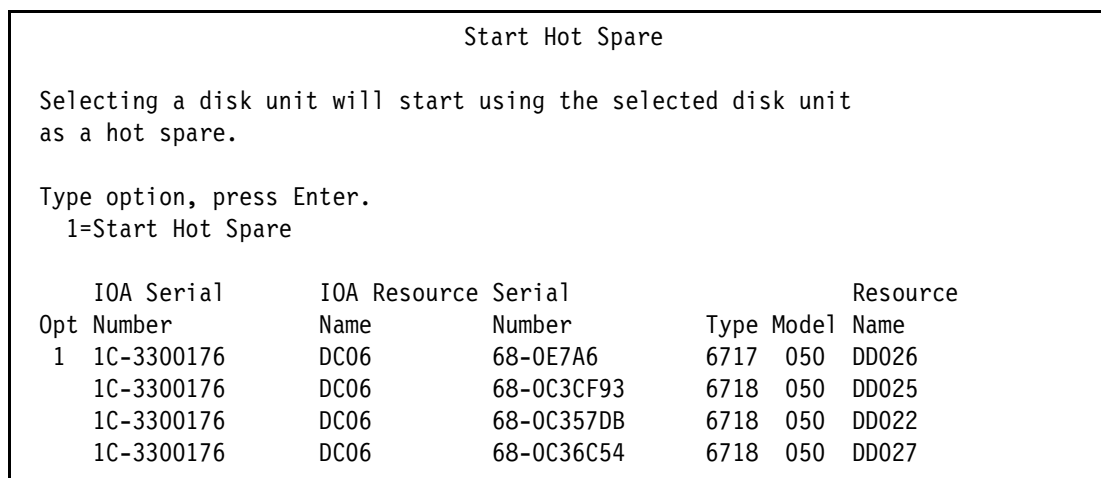


Figure 9-3 IBM i Start Hot Spare panel

The disk IOA does not control mirror protection, so when a mirror protected disk unit fails, the System Licensed Internal Code (SLIC) detects the failed drive and performs the following recovery steps (transparent for the user):

1. SLIC tells IOA to disable the hot spare.
2. Hot spare becomes non-configured.
3. Replace configured unit function is run to replace the failed drive with the now non-configured previous hot spare.
4. 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 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 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 only used after controller failovers (for example, if the active path fails). Dual SAS adapters are redundantly interconnected through a SAS adapter-to-adapter (AA) cable connecting 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 are capable of performing I/O to the attached disk array units. The SAS adapter optimized for a configured RAID set¹ 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.

¹ For IBM i mirroring configurations the disk units attached to a dual SAS adapter are each treated as a one-drive parity set for this purpose

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 (that is, whether it currently operates in a mode as the primary or the secondary adapter) can be seen by navigating to **System Service Tools** → **Start a service tool** → **Hardware service manager** → **Logical hardware resources** when displaying the details for a dual SAS storage IOA and selecting **F14=Dual Storage IOA Configuration**, as shown in Figure 9-5.

Dual Storage IOA Configuration

Type options, press Enter.
 2=Change detail 5=Display detail 6=I/O debug
 8=Associated packaging resource(s) 9=Resources associated with controlling IOP

Resource Opt Name	Type- Model	Status	Serial Number	Operating Mode
DC07	572F-001	Operational	YL3229019FB5	Secondary Storage IOA
DC04	575C-001	Operational	YL3229019FB5	Auxiliary IOA
DC06	572F-001	Operational	YL3229021017	Primary Storage IOA
DC05	575C-001	Operational	YL3229021017	Auxiliary IOA

F3=Exit
F5=Refresh
F6=Print
F12=Cancel

Figure 9-5 IBM i dual storage IOA configuration panel

The disk unit paths for dual SAS adapter connected disk units that are reported as DMPxxx multi-path disk unit resources can be viewed by navigating 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

Best performance without compromising availability in dual a SAS adapter RAID configuration is provided 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, the parity optimization method must be set to **Performance** before creating any RAID5 or RAID6 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 further information about IBM i dual SAS adapter support see the *Dual storage IOA configurations* article in the of the IBM Systems Hardware Information Center at the following web page:

<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/ared5/ared5dualioaconfig.htm>

9.1.4 Encrypted ASP enhancements

The following new functions have been implemented in IBM i 7.1 for ASP encryption supported with 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 has been IPLed past DST and for IASPs when they are varied on. This does not mean an IPL is required, but that the asynchronously run encryption or decryption tasks run only on a system that was IPLed past DST.

Note: For geographic mirroring environments, encryption can only be started or stopped on the production IASP. Data is sent either encrypted or unencrypted to the backup node’s mirror copy IASP. This IASP gets the encryption attribute set accordingly.

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 is securely stored in System Licensed Internal Code such as the ASP master key used for IASP encryption. This is the reason why a SAVSYS is recommended after starting encryption or changing 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

Note: For a clustering environment, an identical ASP master key, which protects the IASP data keys, needs to be manually created using the same paraphrase 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 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 previously existing six buckets is still stored in the QAPMDISK file, the new buckets are stored separately in the new QAPMDISKRIB 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 have been 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 gained 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 further information about the new disk response time buckets in QAPMDISKRB see the IBM i 7.1 Information Center at the following web page:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=%2Frzaxh%2Frzaxhqapmdiskrb.htm>

9.1.6 CEC node level mirroring

A system that has multiple CEC nodes will now mirror disk units in a way that allows it to survive a node outage and that allows a concurrent maintenance repair of a CEC node.

When starting mirroring, the operating system takes into consideration the CEC node under which the disks are located, and attempts to place the two subunits of a mirror protected pair under different CEC nodes. This allows concurrent maintenance of a CEC node since the two subunits of each mirrored disk unit pair are located under a different CEC node, thereby allowing at least one subunit of each mirrored disk unit pair to remain operational during the maintenance operation.

After installing the PTF Group that contains this function, the user may want to consider ending and restarting mirroring to recalculate the mirror protected pairs. There is an Advanced Analysis macro named LEVELOFFPROTECTION, accessible through SST or DST, that allows a user to verify the level of protection for each mirrored pair.

The LEVELOFFPROTECTION macro is accessed from either the Dedicated Service Tools (DST) menu or the System Service Tools (SST) menu.

- To access the LEVELOFPROTECTION macro, go to **System Service Tools** → **Start a service tool** → **Display/Alter/Dump** → **Licensed Internal Code (LIC) data** → **Advanced Analysis**
- On the Select Advanced Analysis Command screen, 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 roll 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
  -      RMTMTIME
  -      SANITIZE
  -      SEIZEINFO
  -      SEMAPHOREINFO
  -      SERVICEDOCS
  -      DUMPTRC
  -      LEVELOFPROTECTION
  -      SMARTCHAIN
  -      SPINLOCKTRACE
  -      SSDSANITIZE
  -      SYNCTOKENINFO

F3=Exit  F12=Cancel

More...

```

Figure 9-11 Selecting the LEVELOFPROTECTION macro

- Pressing enter twice on the Select Advanced Analysis Command screen and the help screen is displayed as shown in Figure 9-12 on page 287.

```

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 executes and displays the following screen in Figure 9-14 on page 288.

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

Note 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 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, refer to IBM Hardware Announcement letter 111-065:

<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 10K RPM SFF SAS disk drives

283 GB and 571 GB 10k 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, refer to IBM Hardware Announcement letter 111-065:

<http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&appname=iSource&supplier=897&letternum=ENUS111-065>

9.1.9 Tape performance instrumentation

A new collection services *RMVSTG category is 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, number of tape marks and blocks spaced, etc. tracked by the IBM i tape code when sending requests to the tape device driver. Currently for reviewing the data 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 needs to be used.

For further information about the structured QAPMTAPE database file performance data, see the IBM i 7.1 Information Center at the following web page:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=%2Frzaxh%2Frzaxhqapmtape.htm>

9.1.10 Tape library resource name changes for IOP-less IOA attachment

When upgrading 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. To ensure the tape library device description points to the correct hardware resource after the upgrade, either the hardware resource name needs to be changed through System Service Tools' Hardware Resource Manager or the device description's resource name field needs to be updated (for example, by using the CHGDEVMLB command) with the new tape library resource name, which was assigned with the upgrade of IBM i 7.1.

9.1.11 Tape library unreadable barcode changes for IOP-less IOA attachment

Prior to IBM i 7.1 if at vary-on of the tape library tape cartridges with unreadable barcodes are found, each of these tape cartridges is loaded in a drive to read the volume ID which 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 required to load and read each cartridge with an unreadable barcode, especially if the library barcode reader itself failed, and also that problems with the barcode label or barcode reader are not made transparent 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 (for example, through the ADDTAPCTG command), a cartridge ID in the form of UNKXXX is fabricated with XXX being a sequential decimal number starting with UNK001. If a cartridge is found by the system in a storage slot with an unreadable barcode, a cartridge ID is fabricated in 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 new handling of unreadable barcodes in IBM i 7.1 reveals barcode problems to the user and still allows the user to read occasionally from tapes without barcode labels (which then are removed from the library again) quicker, without requiring a tape drive for generating cartridge IDs.

Note: With the IBM i 7.1 IOP-less IOA tape library attachment it is recommended not to 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 matching the volume ID needs to be stuck on the cartridge.

9.1.12 DVD/Tape SAS External Storage Unit for Power 795 CEC Rack

The #5274 DVD/Tape SAS external storage unit for the Power 795 CEC rack is a 1U storage unit which can hold HH DAT160 drives, the #5638 1.5TB/3.0TB LTO-5 SAS Tape Drive or slimline DVD drives.

For more information, refer to IBM Hardware Announcement letter 111-065:

<http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&appname=iSource&supplier=897&letternum=ENUS111-065>

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:

- ▶ 9.2.1, “Multipathing for virtual I/O” on page 290
- ▶ 9.2.2, “DS5000 native attachment” on page 292
- ▶ 9.2.3, “Level of protection reporting for multi-path disk units” on page 294
- ▶ 9.2.4, “Library control paths for IOP-less Fibre Channel IOA tape attachment” on page 295
- ▶ 9.2.5, “External disk storage performance instrumentation” on page 296
- ▶ 9.2.6, “Thin Provisioning for DS8700, DS8800 storage servers and for VIOS shared storage pools” on page 301

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 (as shown in Figure 9-15) is supported by IBM i multipathing (that is, the IBM i client can use multipathing across two or more VIOS on the same IBM Power Systems server for protection against VIOS outages due to VIOS updates).

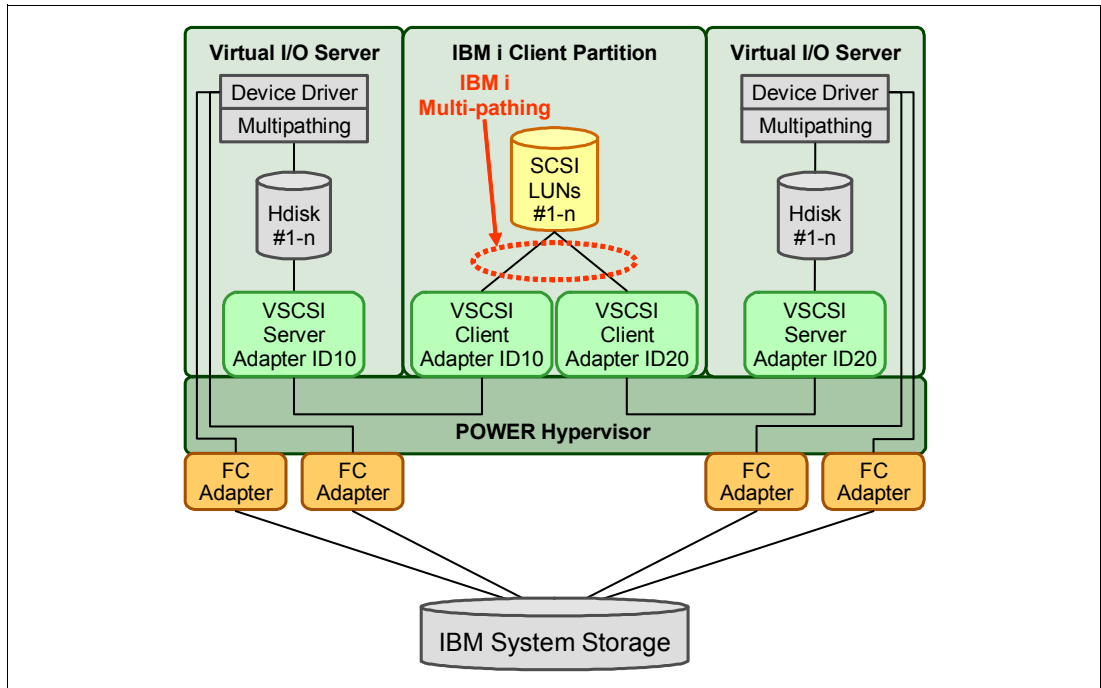


Figure 9-15 IBM i Multipathing with a redundant Virtual I/O Server configuration

This new IBM i multipathing support for virtual I/O eliminates the previous 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 241.

9.2.2 DS5000 native attachment

IBM i SAN storage supported has been 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-attach DS5100/DS5300 storage support provides an easier storage setup (see Figure 9-16) and configuration without needing to deploy the IBM Power VM VIOS. It is only if some other advanced virtualization functions are being used that on IBM i we would need to configure a VIOS.




Figure 9-16 IBM i DS5000 native attachment

Figure 9-17 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
Y2103LQ1HVO0C	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 9-17 Native attached DS5000 LUNs on IBM i

Note: Due to the 4 KB page sector conversion from 8 x 520 bytes sectors to 9 x 512 bytes sectors done 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) honors the DS5000 active/passive controller concept. Under normal working conditions (that is, no controller failover condition and both DS5000 controllers being available), I/O is driven only across the active paths to a disk unit (that is, 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 9-18 shows the active and passive path for disk units from a native attached DS5000 after they were added to an ASP. 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 9-18 IBM i active/passive paths for DS5000 disk units

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, cum 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 respectively 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:

- ▶ Maximum supported LUN size for IBM i is less than 2 TB
- ▶ Usable IBM i net capacity is 8/9 of the configured DS5000 LUN capacity
- ▶ Maximum of 64 LUNs per IBM i Fibre Channel IOA port
- ▶ Unprotected arrays (RAID0) 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
- ▶ Solid state drives in DS5000 are *currently* not supported for IBM i

From an IBM i disk I/O performance perspective the following recommendations need to be followed:

- ▶ To balance workload across both DS5000 controllers, LUNs are to be evenly assigned with regards to preferred controller affinity to controller A and B
- ▶ The generally recommended LUN size for IBM i IOP-less Fibre Channel of 70 GB applies for DS5000 native attachment as well
- ▶ A DS5000 segment size of 128 KB is generally a good comprise for both IBM i transaction and save/restore workload

For further information about the IBM System Storage DS5000 series see the following IBM Redbooks publications:

- ▶ *IBM Midrange System Storage Hardware Guide*, SG24-7676
- ▶ *IBM Midrange System Storage Copy Services Guide*, SG24-7822

For IBM support statements regarding DS5000 Copy Services support with IBM i native attached DS5000 see *IBM i Virtualization and Open Storage Read-me First* available at the following web page:

http://www-03.ibm.com/systems/resources/systems_i_Virtualization_Open_Storage.pdf

IBM STG Lab Services has developed a Copy Services toolkit 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 further information about this Copy Services toolkit offering for DS5000, see IBM STG Lab Services at the following web page:

<http://www-03.ibm.com/systems/services/labservices>

9.2.3 Level of protection reporting for multi-path disk units

With IBM i 7.1 the level of protection for multi-path attached external disk units is now reported for any multi-path disk unit devices from either natively or VIOS-attached disk storage system. The reported levels of multi-path protection reflect the component that can fail without jeopardizing I/O access to the disk units and are exactly the same as those already used for mirrored protection:

- ▶ Remote Bus
- ▶ Ring (HSL / 12X Loop)
- ▶ Tower
- ▶ Bus
- ▶ IOP
- ▶ IOA
- ▶ IOA-Bus

Figure 9-19 shows the new multi-path protection level reporting for the example of DS8000 disk units each attached through three paths from **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 9-19 IBM i protection level reporting for multi-path disk units

9.2.4 Library control paths for IOP-less Fibre Channel IOA tape attachment

Tape library devices 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 has 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 OS pools these together so all the TAPxx resources for the library are in one TAPMLBxx device description

Note: For IBM i 7.1 a second library control path needs to 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.

Prior to 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.

Due to 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. ENDPFCOL FRCCOLEND(*YES)
2. RNMOBJ OBJ(QUSRSYS/QPFCOLDTA) OBJTYPE(*USRSPC)
NEWOBJ(QPFCOLDT2)
3. CALL QSYS/QYPSCOLDTA PARM(*EXTSTG)
4. STRPFCOL

For further information about these DS8000 external storage performance data collection requirements see the *IBM i Memo to Users 7.1* at the following web page:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzaq9/rzaq9.pdf>

Note: 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 when the QAPMXSTGV database file 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 as the SCSI Log Sense command is issued against a LUN.

IBM i Doctor 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 has new functions for visualizing and analyzing this DS8000 or DS6000 storage performance data.

The new iDoctor Collection Services Investigator functions for analyzing the external storage performance log sense data stored in QAPMXSTGV (Log sense page 0x32) and QAPMXSTGD (Log sense pages 0x33 and 0x34) is shown in Figure 9-20.

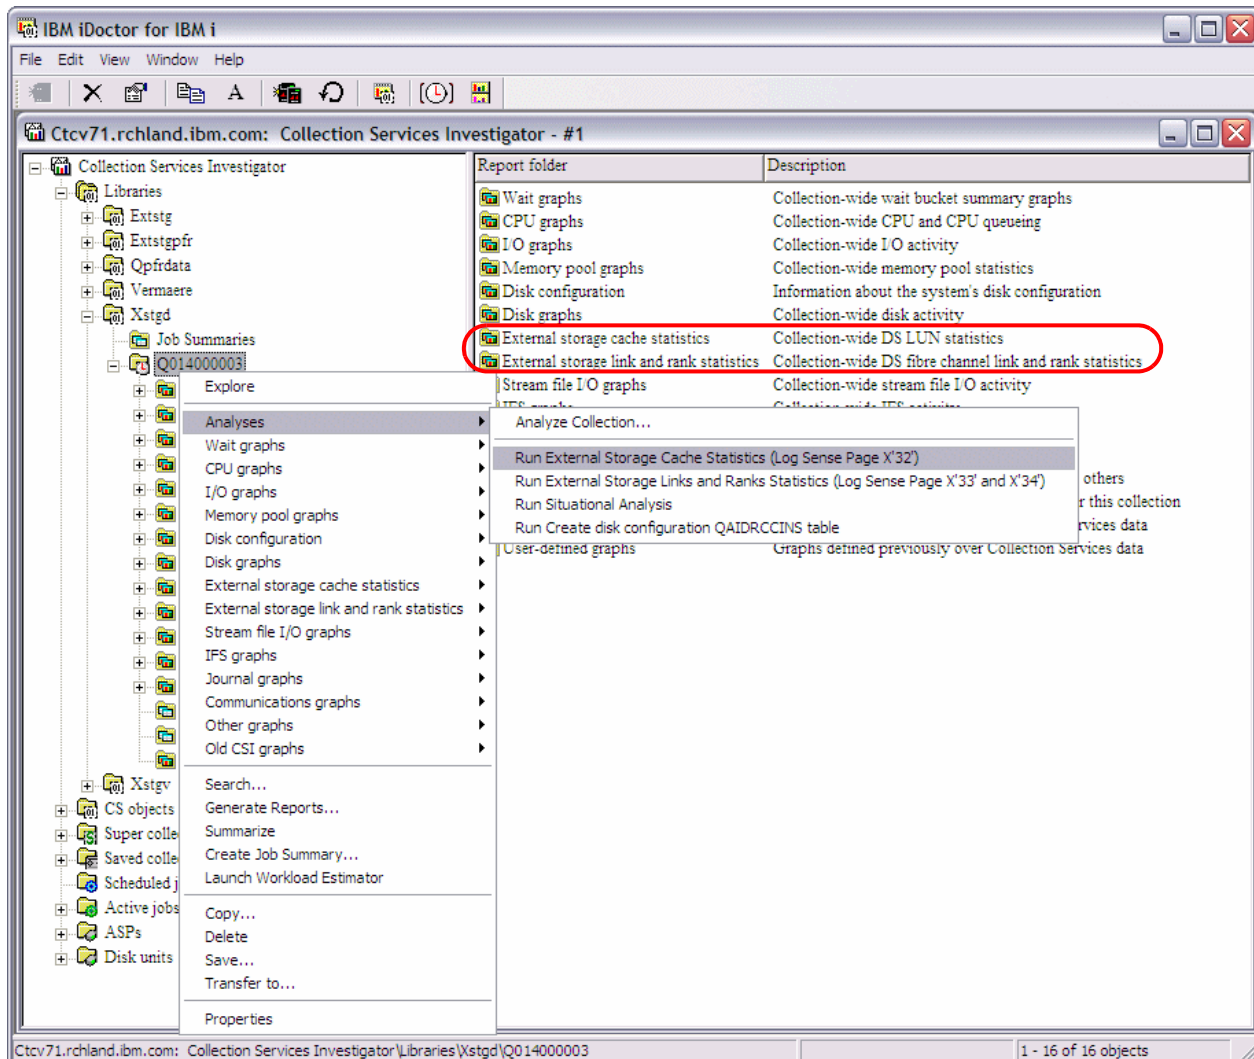


Figure 9-20 iDoctor analysis functions for external storage

Before iDoctor can be used to analyze the external storage performance data, the **Analyses → Run External Storage** functions needs to be run 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 9-21 from the report **External storage cache statistics → by time interval → IO rates totals with cache hits**. The read cache hit% information was available also before from QAPMDISK data but the newly reported write cache hit% from QAPMXSTGV data can be helpful to check for any potential storage subsystem write cache overruns, which is indicated by write cache hits% < 100% and might warrant changes in the workload schedule or a cache size upgrade.

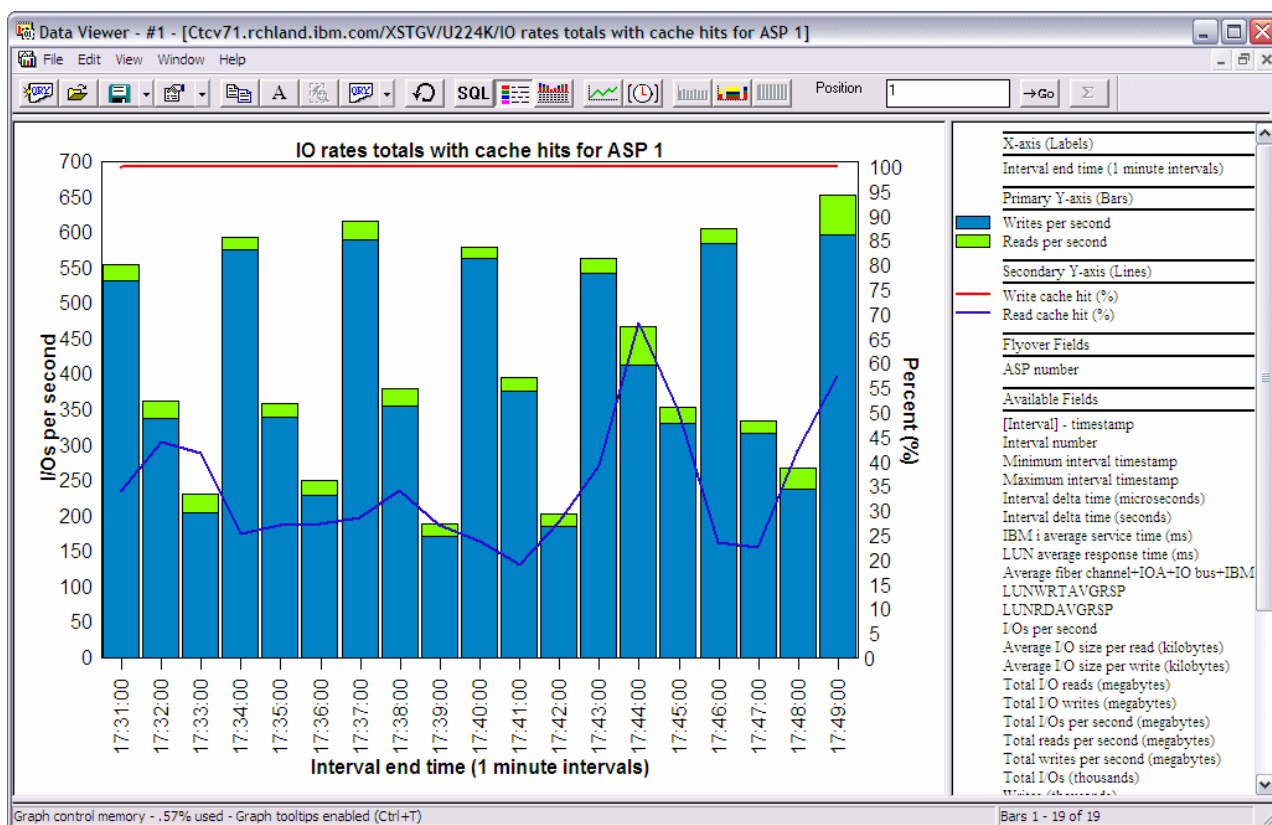


Figure 9-21 iDoctor external storage cache statistics

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 selecting **Rank graphs** → **By rank ID** → **Ranks IO rates totals**. Then, from this view, selecting one or more ranks with a high I/O rate for a more detailed analysis by choosing **Selected Ranks** → **Ranks IO rates** from the right-click context-menu, as shown in Figure 9-22.

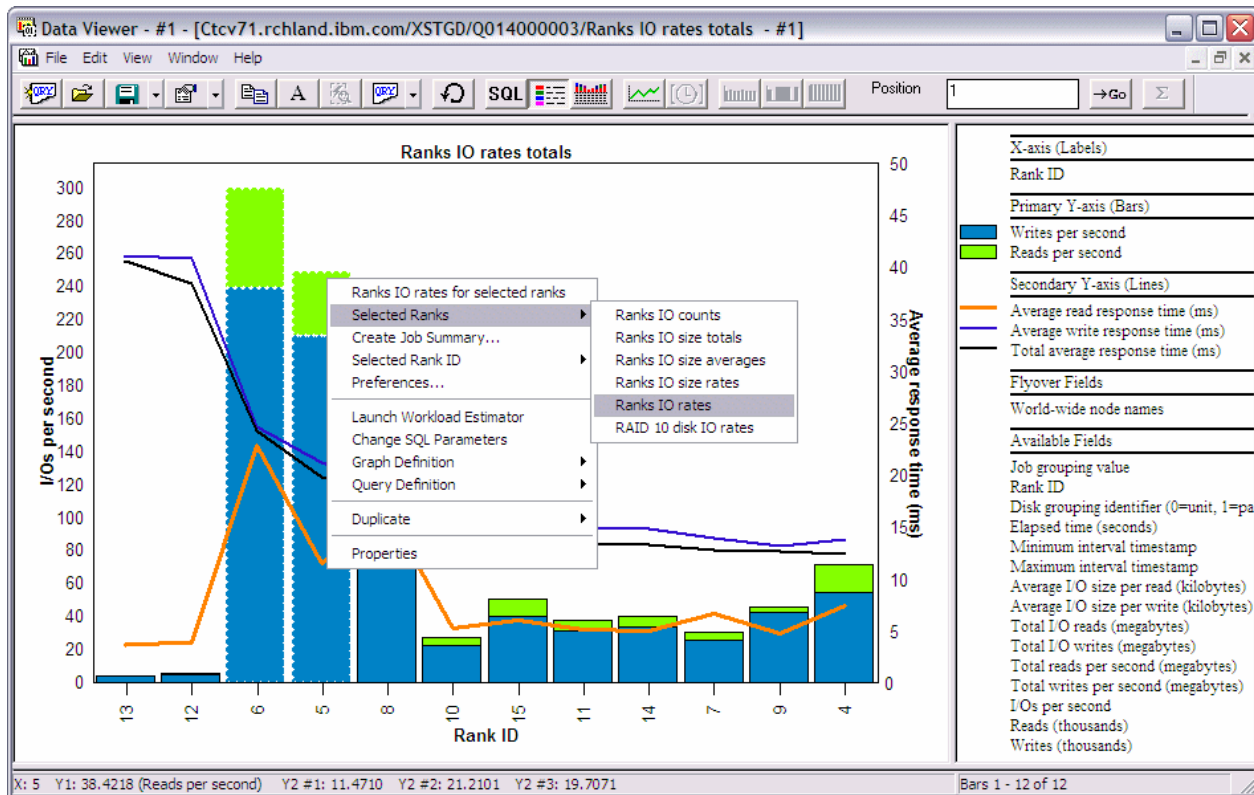


Figure 9-22 iDoctor rank I/O rates ranking

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 9-23. If these were RAID10 instead of RAID5 ranks we can have also chosen the option to display the disk I/O rates. This option is not available for RAID5, because their disk I/O rates cannot be determined accurately from the rank read and write I/O rates.




Figure 9-23 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 have been added to iDoctor, as shown in Figure 9-24.




Figure 9-24 iDoctor link I/O rates

For further information about the IBM iDoctor for IBM i powerful suite of performance tools see the following iDoctor web site, which offers also a 45-day trial version:

https://www-912.ibm.com/i_dir/idoctor.nsf/iDoctor.html

9.2.6 Thin Provisioning for DS8700, 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. Prior to this enhancement, the full amount of configured storage would be allocated at LUN initialization time.

Hardware Requirements:

Thin provisioning enhancement for DS8000 storage servers requires a DS8700 or DS8800 with Release 6.2, available from IBM via FC #1723, or via 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 also for solid state drive (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.

Solid state drives 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 with 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 the IBM White Paper *Performance Value of Solid State Drives using IBM i* which is available at the following web page:

http://www-03.ibm.com/systems/resources/ssd_ibmi.pdf

The SSD Analyzer Tool for IBM i is recommended 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 savefile at the following web page:

<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 article *Requirements for Solid State Drives (SSD)* available at the following web page by searching 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.

- ▶ 9.3.1, “DB2 media preference”
- ▶ 9.3.2, “ASP balancer enhancements for SSDs” on page 307
- ▶ 9.3.3, “User-defined file system media preference” on page 310
- ▶ 9.3.4, “177 GB SFF SSD with eMLC” on page 312
- ▶ 9.3.5, “IBM Disk Sanitizer PRPQ extended to include SSD devices” on page 312

9.3.1 DB2 media preference

DB2 for i has been 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 to be stored, so that DB files that are known to be I/O performance critical can explicitly be placed on high performing SSDs.

Physical and logical DB files, respectively table and indices, were enhanced with a preferred media attribute that can be set through the UNIT parameter for a certain media preference of either SSDs, with UNIT parameter value *SSD (CL commands) or SSD (SQL) for i 6.1 and later, and value 255 (CL commands) for i 5.4.5, or HDDs with UNIT parameter value *ANY

(CL commands) or ANY (SQL) when creating or changing these files through the following CL commands or SQL statements:

CRTPF, CRTLF, CRTSRCPF, CHGPF, CHGLF, CHGSRCPF

CREATE TABLE, CREATE INDEX, ALTER TABLE

Notes:

The UNIT parameter for the aforementioned SQL statements is supported with 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 9-25 shows the new preferred storage unit parameter (UNIT keyword) for the CHGPF command.

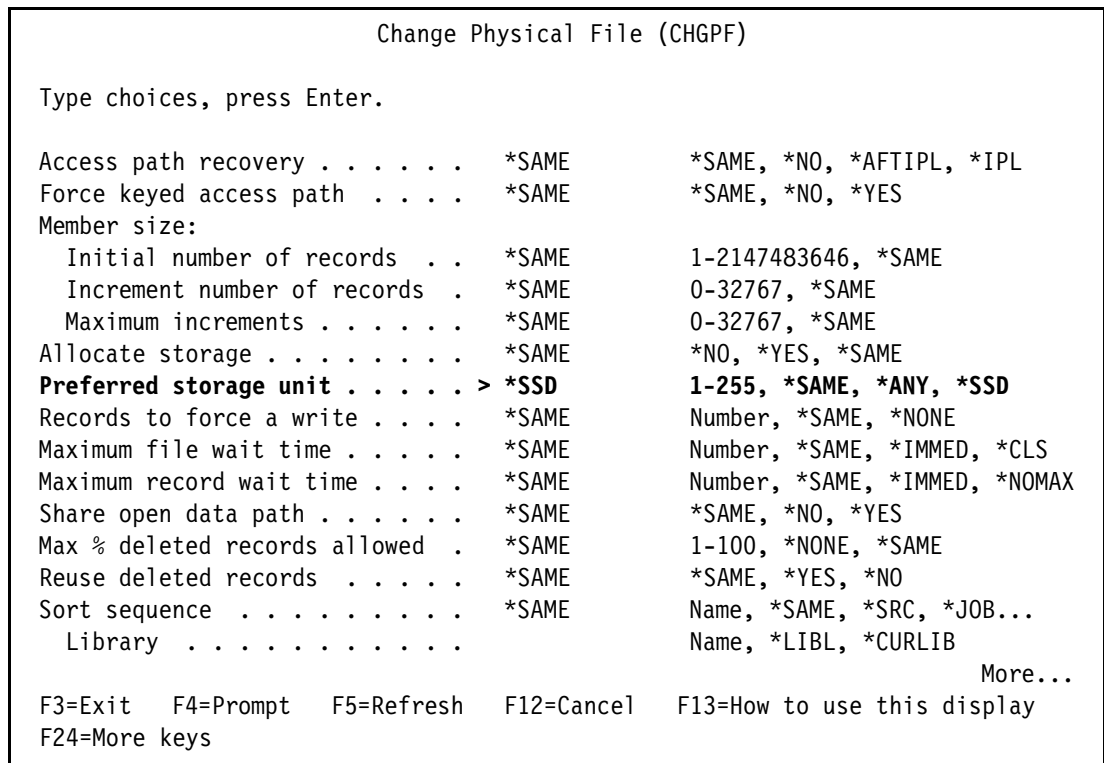


Figure 9-25 IBM i CHGPF command

Dynamic data movement

Enhancements have been implemented for dynamically changing the DB2 media preference, which invokes a dynamic data move done synchronously.²

For releases before IBM i 7.1 the following PTFs are required for a dynamic move of physical or logical database files after changing 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, MF47878

² It is planned to change DB2 media preference to asynchronous data movement in the future.

- ▶ IBM i 6.1.0 PTFs MF47888, MF47892, MF47879
- ▶ IBM i 6.1.1 PTFs MF47889, MF47893, 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) have been 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, are continuously updated by database and reset only at IPL. Because a single logical read I/O can lead to more than one random I/O (for example, due to 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 recommended proceeding (after database performance optimization has been 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 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 as shown in Example 9-1 or the System i Navigator's Health Center activity tab, as shown in Figure 9-26, can be used by saving the query results and using the View History function to compare the results 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 9-26 System i Navigator database health center

Macro for changing storage management SSD allocations

To prevent newly created objects without media preference from being placed on SSDs in a hybrid ASP with HDDs and less capacity used SSDs the macro **smsetstayoffssd** must be executed after which storage management will try to stay off from SSDs for non media preference objects as shown with the **smgetstayoffssd** macro to query the current setting in Example 9-2. The macro **smresetstayoffssd** 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, macro smgetstayoffssd 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

```
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 SYSPARTITIONDISK 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, i 6.1 and i 7.1.

Figure 9-27 shows a query result example from the new SYSPARTITIONDISK view after we used the ASP balancer to move DB files to SSD and manually moved 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).

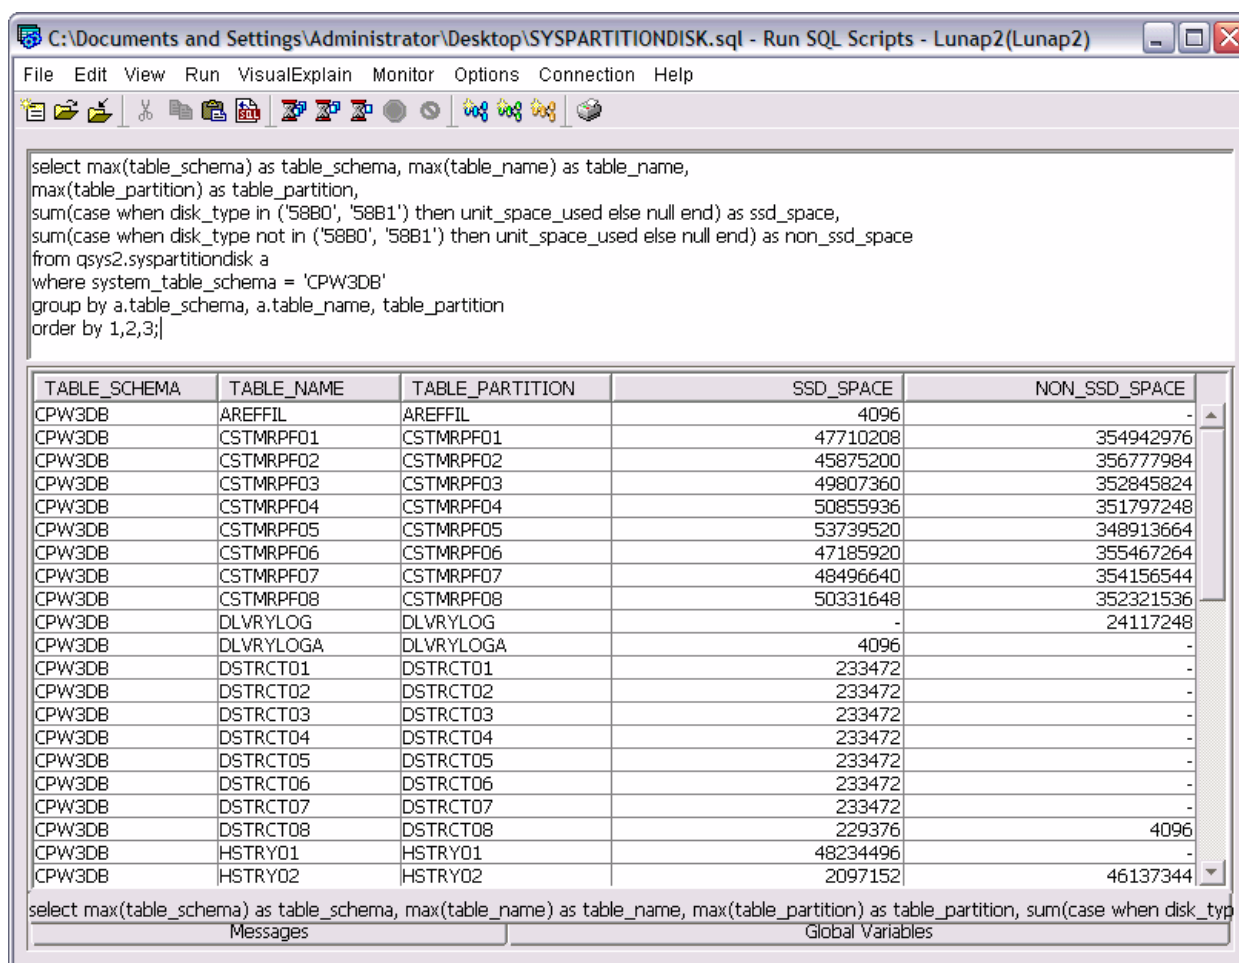


Figure 9-27 SYSPARTITIONDISK view query result

For further information about the new SYSPARTITIONDISK view and function, see the article *IBM DB2 for i Statistical View for Solid State Drive Storage Usage Reporting* at the following web page:

<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) has been extended for support of data migration between SSDs and HDDs as well.

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 now also 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 shall be optimized for performance with using SSDs. Afterwards, the ASP balancer HSM function is started to migrate both cold data off 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 for the extent read I/O counts only has been enhanced with a more efficient migration algorithm in the weighted ASP balancer version and additional functionality regarding SSD media management as described in the following sections.

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 (that is, not only the amount of read I/O accesses to a segment is counted as before but also its read service time is considered for the migration decision).

This weighted ASP balancer enhancement accounting for the read service times provides for a more efficient data media placement because, 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 has been enhanced with an option allowing the user to specify the migration priority for *MOVDTA, *HSM or *MP operations in levels of either *LOW, *MEDIUM, or *HIGH (as shown in Figure 9-28), which influences the number of SLIC internal data moving tasks used for the migration. This is always subject to an inherent trade-off between speed of data migration and its effect on disk use.

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, *HDD, *SSD
*CALC	

Bottom

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

Figure 9-28 IBM i ASP balancer migration priority

Also, the STRASPBAL command syntax slightly changed in IBM i 7.1 with the new subtype parameter that for the *HSM balance type now allows data migration between up to three storage tiers:

- ▶ With subtype *SSD data migration between SSDs and high performance HDDs
- ▶ With subtype *HDD 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 being moved off from SSDs first, then hot data being moved to SSDs

ASP balancer media preference balance type

The *MP balance type is a new ASP balancer function in IBM i 7.1 to help correct any issues with media preference flagged DB objects or UDFS files not located on their preferred media type which is either solid state drives (SSD) or hard disk drives (HDD). This sweeper function

moves objects marked with a media preference attribute to SSDs and non-media preference objects to HDDs when using 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]` default is low
This moves data marked with media preference attribute to the SSDs and non-media preference data off the SSDs.
- ▶ `movemediapreferencetossd asp_num priority [L M H]` default is low
This moves data marked with media preference attribute to the SSDs.
- ▶ `movemediapreferenceoffssd asp_num priority [L M H]` default is low
This moves data not having the media preference attribute off the SSDs.
- ▶ `movemediapreferencestatus asp_num`
This gives the current status of the sweeping.
- ▶ `movemediapreferencestop asp_num`
This ends the sweeping.

Note that the ASP number in the `asp_num` variable needs to be specified in hex format.

A scenario for using the media preference sweeper function is after disk units were added to an ASP, then choosing the add and balance option, which currently does not respect the media preference. It can also be used when disk units have been removed from the configuration due to media type capacity constraints within an ASP. The sweeper function can be used to correct these media preference issues after the capacity constraints got solved.

Script for ASP balancer scheduling

The CL script in Example 9-5 is provided to help set up a repeatedly-running TRCASPBAL and STRASPBAL configuration 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, it is usually recommended to run the TRCASPBAL for the period of critical workload such as a batch window which 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

New support is provided that allows you to specify that storage for objects created in user-defined file systems (UDFS) are to be allocated from solid state drives (SSD), 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 has been added to the CRTUDFS (Create User-Defined FS) command, as displayed in Figure 9-29. 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 to be allocated from SSD storage media, if available. Online help text for the new keyword is not included in the PTF.

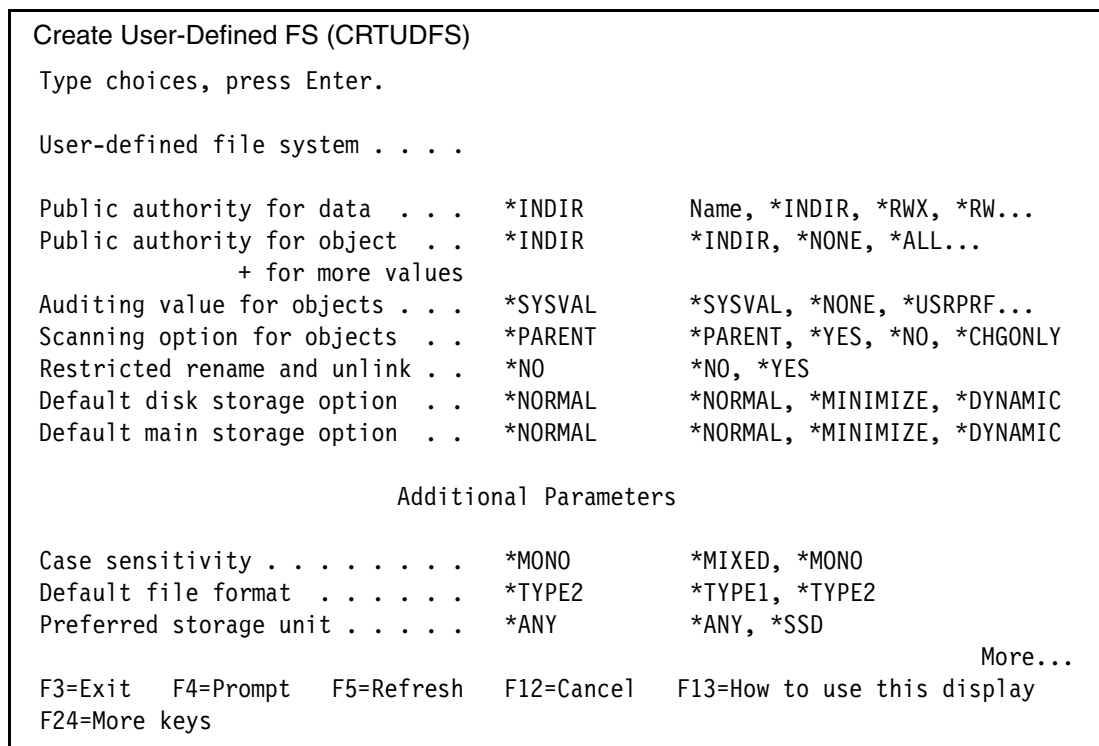


Figure 9-29 CRTUDFS command

Changes have been made to the Qp0IGetAttr()--Get Attributes, QP0LFLOP (Perform File System Operation), and statvfs()--Get File System Information, and related APIs to provide support for determining the preferred storage media for a file system. The RTVDIRINF (Retrieve Directory Information) CL command has also been enhanced to provide this information. See the PTF special instructions for more information about these changes.

The following additional considerations apply when specifying a storage media preference for a UDFS:

- ▶ Specifying a media preference does not guarantee 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 only display or retrieve the storage media preference of a user-defined file system, not the individual objects within a file system.
- ▶ Objects 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 restoring 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, refer to IBM Hardware Announcement letter 111-132:

<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 via a macro interface from either the Dedicated Service Tools (DST) menu or the System Service Tools (SST) menu.

- From DST or SST, select 'Start a service tool'
- Select 'Display/Alter/Dump'
- Select 1 - 'Display/Alter storage'
- Select 2 - 'Licensed Internal Code (LIC) data'
- Select 14- 'Advanced Analysis' (you will need to scroll down to see this option)
- On the Select Advanced Analysis Command screen, 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 9-30. 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 9-30 Selecting the SSDSANITIZE macro

- Press the enter key twice and a help panel is displayed as shown in Figure 9-31 on page 313.

```

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 RESOURCE NAMES> <-IMMED>
More...

F2=Find    F3=Exit    F4=Top    F5=Bottom    F10=Right    F12=Cancel

```

Figure 9-31 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 discusses the following topics relating to networking enhancements enabled in IBM i 7.1:

- ▶ 10.1, “TCP/IP enhancements summary” on page 316
- ▶ 10.2, “HTTP server supports PASE FastCGI PHP Processing” on page 325
- ▶ 10.3, “Telnet client support for SSL encryption” on page 326
- ▶ 10.4, “Sockets programming: New user exits” on page 327
- ▶ 10.5, “IKEv2 support” on page 330
- ▶ 10.6, “IPv6 TCP/IP applications and V6R1 enablement PTFs” on page 331
- ▶ 10.7, “AnyNet is no longer supported on IBM i 7.1” on page 332
- ▶ 10.8, “Ethernet link aggregation” on page 332
- ▶ 10.9, “Sharing physical Ethernet connections across partitions (Ethernet layer-2 bridging)” on page 335
- ▶ 10.10, “IBM Portable Utilities for i (5733-SC1) supported versions” on page 339

10.1 TCP/IP enhancements summary

The following sections summarize the key TCP/IP enhancements in IBM i 7.1. These 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). See 10.6, “IPv6 TCP/IP applications and V6R1 enablement PTFs” on page 331 for additional details.

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, users have an option of using the new DHCP server based on the open source code provided by the Internet Systems Consortium (ISC). Existing customers can continue to use the old DHCP server 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. Currently, there is no GUI support for managing the ISC DHCP server configuration files or for monitoring leases, such as we have with the old DHCP server. Therefore by default, the old DHCP server is the one 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 the IBM i” on page 317. Then, end your DHCP server using the ENDTCPSPVR command (if it is currently 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 made to the old configuration file is 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 help get started with the new server. Certain functions 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', then end and restart the DHCP server. If the ISC DHCP server had assigned any IP addresses when it was running, those leases are not available to the old DHCP server. In other words, there is no backward migration.

Using the ISC DHCP IPv6 server on the IBM i

DHCP has moved from the base OS to 5770-SS1 Option 31 and also requires that 5770-SS1 Option 33 be installed. To use the ISC DHCP IPv6 server on the IBM i, follow these basic steps:

1. Ensure 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 with the following command:

```
ADDENVVAR ENVVAR('QIBM_ISC_DHCP') VALUE('Y') LEVEL(*SYS)
```

3. Run the CHGDHCPA (Change DHCP Attributes) command.

This command migrates any existing DHCP configuration into the configuration files used by the ISC DHCP server. Determine if you want to run an IPv4, IPv6, or both. This 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.

Note: A graphical interface is not provided for managing the ISC DHCP server and monitoring the leases 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 Systems Director 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 is 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 you have at least one line 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 additional 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*, 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.

```
ADDTCPIFC *IP6SAC LIND(line-name)
```

The system only tries to acquire 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, we 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, we 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 configured through CHGTCPA.

To identify ourselves to the DHCPv6 server, the client uses a DHCP Unique Identifier (DUID). This is generated automatically from a MAC address on the system and a time stamp, and is saved by the TCP/IP configuration. This 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 they 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, we support the DNS Server List and Domain Search List options and will add 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 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, we request just the DNS Server List and Domain Search List options from the DHCPv6 server.

Note: 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 have been added.

Enabling the agent to handle SNMPv3 requests

To enable the agent to handle SNMPv3 requests, perform the following steps:

1. If the SNMP server is currently running, end it using the ENDTCPSPVR *SNMP command.
2. Change the SNMP server attributes to allow version 3, using the CHGSNMPA ALWSNMPV3(*YES) command.

Note: The SNMP agent will still be able to receive and handle packets and requests from older versions of SNMP v1 even after the changing the SNMP attributes to specify ALWSNMPV3(*YES).

3. Check the engine identifier 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 using 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 used to specify the engine ID.

Note: An invalid SNMP engine ID can prevent an SNMP manager from communicating with the agent.

Note: Another new SNMPv3 parameter, SNMPENGB, has been added to the CHGSNMPA command is the SNMP engine boots counter. It is recommended that you do not manually change this unless you need to reset it to a value of zero. This indicates the number of times the SNMP engine (agent) has been 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 with 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 will add 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 care must be taken when adding 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 specified are not stored.

The SNMP_USER statements in the configuration file is never edited manually. Instead, the ADDUSRSNMP, CHGUSRSNMP, and RMVUSRSNMP commands are used for maintaining the SNMP users. The CFGTCPSNMP command can be used to display and maintain the list of SNMP users.

5. Start the SNMP server using the STRTCPSVR *SNMP command
6. Start an SNMPv3 manager application and configure it for the SNMP user that was added in step 4.

Note: The IBM i 7.1 SNMP manager APIs snmpGet, snmpSet, and snmpGetnext currently do not support SNMPv3, so a non-native manager such as a PC-based manager must be used. There are a number of these 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 determining the agent's SNMP engine ID and SNMP engine boots values. In addition, a time synchronization occurs. These steps are necessary prior to doing any actual SNMP operations such as Get or GetNext. If you are having problems in these initial steps, the SNMP manager can indicate that a time synchronization error occurred. If this happens, you can try to do the following:

1. End 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 only works for IPv4 addresses. SMTP uses the getaddrinfo() API to look up e-mail DNS records. They are looked up first as IPv6 and then as IPv4, which is different from what Request for Comments (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 has been removed. Starting in IBM i 7.1, Simple Mail Transfer Protocol (SMTP) only supports RFC 2821 and 2822. RFC 2821/2822 deprecate many parts of the 821/822 email standard. 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 will still be honored from RFC 821, other parts of the source route is not contacted. The absolute address is the recommended 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 that allowed the forwarding of email to a single address. FWDMAILHUB always forwards the email and does not attempt a resolve.

MAILROUTER only supports 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.

- ▶ 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 selected 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 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 configuration changes required to enable IPv6 for a connection profile.

Note: PPP configuration enhancements for IPv6 is only available using System Director Navigator for IBM i. It is not available using the PC based client System i Navigator.

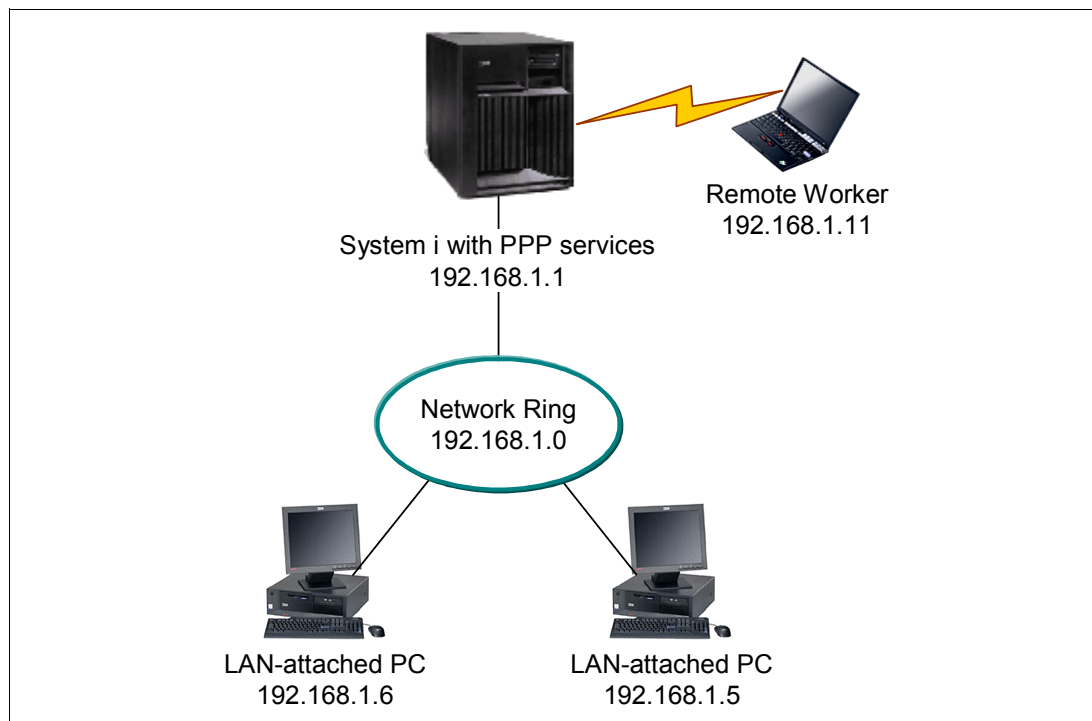


Figure 10-1 Tradition IPv4 PPP remote access configuration

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, if we 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 can be configured in the Receiver Connection profile using System Director Navigator for IBM i as shown in Figure 10-2.

Select **Network** → **Show all Networks Tasks**.

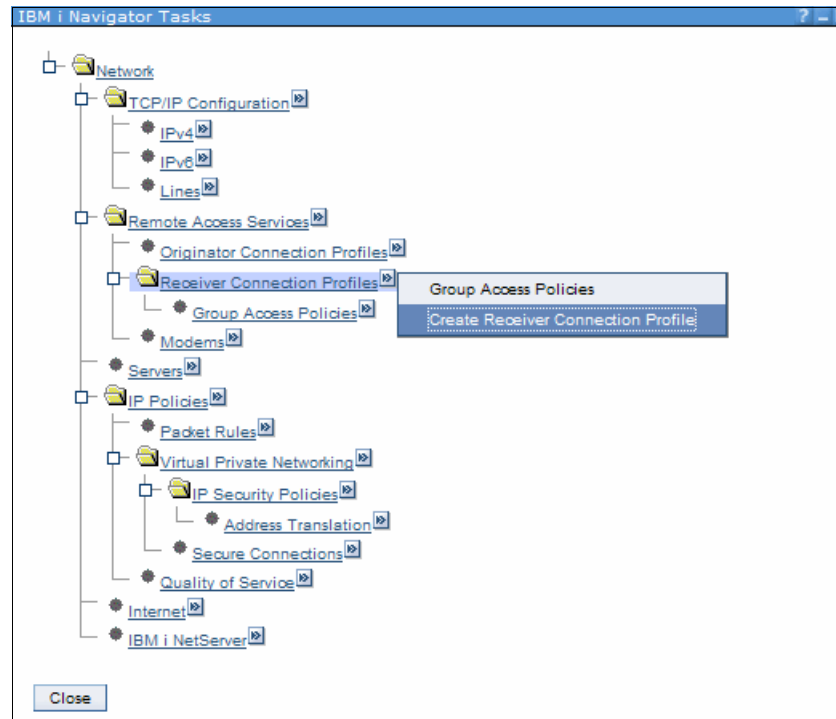


Figure 10-2 From Navigator for IBM i - Network - Show all Tasks - Remote Access Services

Select **Create Receiver Connection Profile**, then **TCP/IP IPv6 Settings**. See Figure 10-3.

The screenshot shows a window titled "New Point-to-Point Profile Properties - localhost". On the left is a vertical sidebar with tabs: "General" (selected), "Connection", "Authentication", "TCP/IP IPv4 Settings", "TCP/IP IPv6 Settings", and "Other". The main content area on the right contains the following fields and options:

- Name:** A text box containing "ITSOPPP".
- Description:** A text box containing "redbook example".
- Protocol type:** A dropdown menu showing "PPP".
- Mode type:** A dropdown menu showing "Switched line-answer".
- Start profile with TCP:** An unchecked checkbox.

At the bottom of the window are two buttons: "OK" and "Cancel".

Figure 10-3 PPP - Create Receiver Connection profile panel with 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. A global IPv6 address is configured in the connection profile as follows:

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, select **None**.
3. Select **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.

Figure 10-4 Configure PPP IPv6 preferences

10.2 HTTP server supports PASE FastCGI PHP Processing

FastCGI is an open standard extending the Common Gateway Interface (CGI) standard supported by many common web servers today. This standard defines how information is exchanged between a web server and FastCGI programs 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 resulting in faster HTTP request processing.

Further implementation details are available from the PDF file at the following web page:

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.

Required components and PTF information for 5.1 and 6.1 follow:

IBM i 5.4 required components

Products

- ▶ 5722SS1 30 Qshell (only needed to create CCSID 819 files as shown above)
- ▶ 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)

IBM i 6.1 required components

Products

- ▶ 5761SS1 30 Qshell (**only needed to create CCSID 819 files as shown above**)
- ▶ 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, new in IBM i 7.1 is the ability to 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.

Notice that the PORT parameter on the TELNET command prompt has moved to a new location in the parameter string, and a new parameter, Secure Connection (SSL), has been added to the command format. If the environment variable has been set up for a secure connection, or SSL(*YES) parameter is selected, the target port number will default to 992.

Requirements are as follows:

- ▶ Must have the Digital Certificate Manager (5770-SS1 Option 34) installed.
- ▶ Must have certificates setup for the client application *QIBM_QTV_TELNET_CLIENT*. Either a trust list or the remote server certificates need to be assigned to this application. The trust list needs to have the certificate authorities of the desired 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 will fail. See Figure 10-5.

Start TCP/IP TELNET (TELNET)			
Remote system	RMTSYS	
Internet address	INTNETADR	
Port	PORT	*DFT
Secure connection	SSL	*ENVVAR

Figure 10-5 New parameter on *STRTCPTELN* command for SSL enablement

Note: 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 the Socket Programming Guide include three sockets-related user exit points. These were added to give a user-defined exit program the ability to control connections based on specific runtime characteristics. User-defined exit programs registered with the exit points defined in the user registry are able to limit incoming and outgoing connections.

10.4.1 Exit points are defined in the user registry

User-defined exit programs 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 Point

User Exit Point	Description
QIBM_QSO_ACCEPT	Enables a custom exit program to allow or deny incoming connections based on the restrictions set by the programs.
QIBM_QSO_CONNECT	Enables a custom exit program to allow or deny outgoing connections based on the restrictions set by the programs.
QIBM_QSO_LISTEN	Enables a custom exit program to allow or deny a socket the ability to listen for connections based on the restrictions set by the programs.

Notes:

1. 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_ACCEPTPERM*. When the socket option is enabled, sockets APIs accepting connections fail with EPERM for each incoming connection rejected by the user exit program registered for *QIBM_QSO_ACCEPT*.
2. Any user trying 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

In this section we illustrate an example of the user exit program fro QIBM_QSO_ACCEPT user exit point.

Example 10-3 rejects all incoming connections to the Telnet server coming from a particular remote IP address between the hours of 12 a.m. and 4 a.m. and determines if the incoming connection is allowed to be accepted by the socket API accepting connections or rejected.

Example 10-3 Socket program example using 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;
}

```

Note: By using the example, you agree to the terms of the code license and disclaimer information which is available at this web page:

<https://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/rzahg/legalsnoticesSW.htm>

10.5 IKEv2 support

You can now use the enhanced IKE version 2 when performing 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 enabling IKE version 2. See Figure 10-6.

- ▶ 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 Director Navigator for i is required to configure an IKEv2 connection.

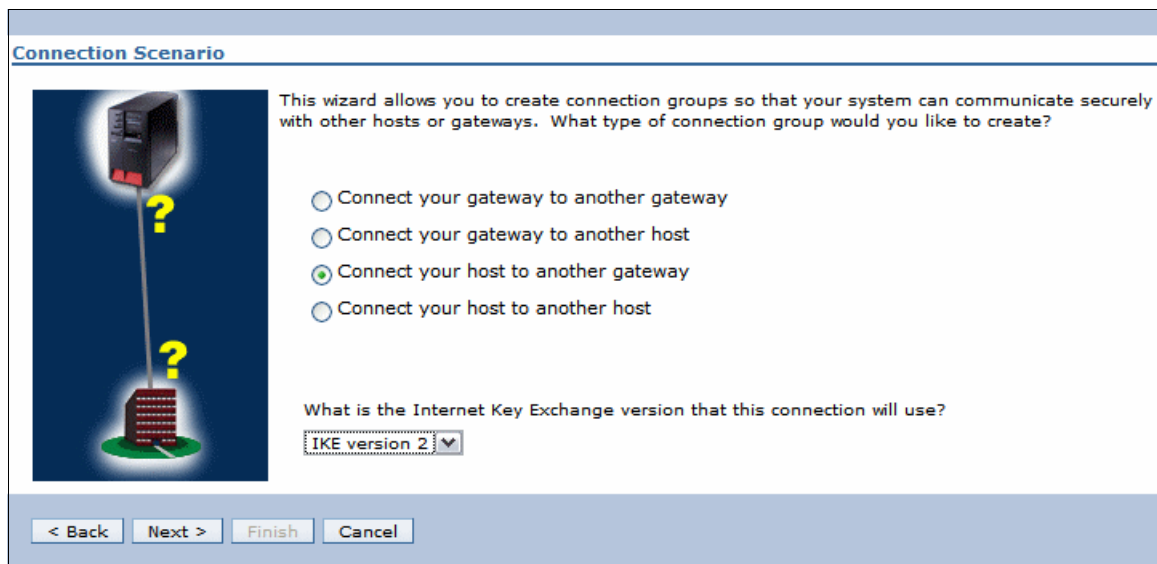


Figure 10-6 Internet Key Exchange V2 enablement for VPN

10.5.1 Enhanced Cryptographic Algorithms have also been added

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

- ▶ Authentication
 - AES-XCBC-MAC
 - HMAC-SHA-256

Diffie-Hellman for PFS

- ▶ Group 14
- ▶ Group 24

For more information and configuration details, see the article *IBM i Security Virtual private networking 7.1*, available at the following web page:

<https://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp?topic=/rzaja/rzajagetstart.htm>

10.6 IPv6 TCP/IP applications and V6R1 enablement PTFs

TCP/IP applications and protocols for IBM i operating system are incrementally adding support for IPv6. The following information center entry provides a complete a 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 web page 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 AnyNet is no longer supported 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 Anynet. Although Anynet has not been removed in IBM i 7.1, IBM has stated that there is no further enhancements beyond IBM i 6.1 and IBM will no longer offer support on 7.1

The IBM i 7.1 Information Center article *Migrating from AnyNet to Enterprise Extender* provides detailed migration considerations and requirements, and is available from the following web page:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/rza/jt/rzajtanytoee.htm>

10.8 Ethernet link aggregation

The following sections discuss the IBM i Ethernet link aggregation function available in IBM 7.1.

- ▶ 10.8.1, "Overview of Ethernet link Aggregation"
- ▶ 10.8.2, "Prerequisites for Ethernet link aggregation"
- ▶ 10.8.3, "Configuring Ethernet link aggregation" on page 333
- ▶ 10.8.4, "Example - Configuring four aggregated network links" on page 334

10.8.1 Overview of Ethernet link Aggregation

Link Aggregation binds several full-duplex Ethernet links running at the same speed together into one logical link with a single Media Access Control (MAC) address. This 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 8 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.

Prior to 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 will be spread across the links as determined by a user-selected policy and incoming traffic by a policy configured at the network switch. This also enables more capacity for a given 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 enabled for static configuration.

Note: 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. Refer to the example command in Figure 10-7 on page 334.

- a. The user creates a line description with the resource name (RSRCNAME) parameter set to a new special value *AGG, which enables new aggregation-specific parameters.
- b. The user chooses which aggregate policy to use for spreading outgoing frames across the aggregated links using the Aggregate Policy (AGGPCY) parameter.
 - i. The first element specifies the standard aggregation technology to be used. Currently only the EtherChannel technology (*ETHCHL) value is supported.
 - ii. The second element specifies the Policy type. These 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 a large number of clients.
 - *SRCPORT
The adapter selection algorithm uses the source TCP/IP or UDP port value.
 - *DESTPORT
The outgoing adapter path is selected via an algorithm using the destination TCP/IP or UDP port value.
 - *SRCDESTP
The outgoing adapter path is selected via 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 host connected back-to-back (i.e. without a intervening switch).
- c. The user chooses a list of the communication resources that should be aggregated and specifies them in the Aggregated resource list parameter.
- d. The user must select a single speed for all of the adapters as specified in the Line speed parameter.
- e. The user must set the DUPLEX parameter to full-duplex (*FULL). The user then creates the aggregated line description.

- f. Last, 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).

```

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 . . . . .                   > *RDRBN
Aggregated resource list . . . . . AGGRSCL > CMN01  ← c)
                                         > CMN02
Local adapter address . . . . . ADPTADR  > *ADPT
Exchange identifier . . . . . EXCHID     > *SYSGEN
Ethernet standard . . . . . ETHSTD       > *ETHV2  ← d)
Line speed . . . . . LINESPEED          > 1G
Duplex . . . . . DUPLEX                 > *FULL   ← e)
Serviceability options . . . . . SRVOPT  > *NONE

More...

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

```

Figure 10-7 Creating an aggregated Ethernet line description.

10.8.4 Example - Configuring four aggregated network links

A logical view of another example follows in Figure 10-8

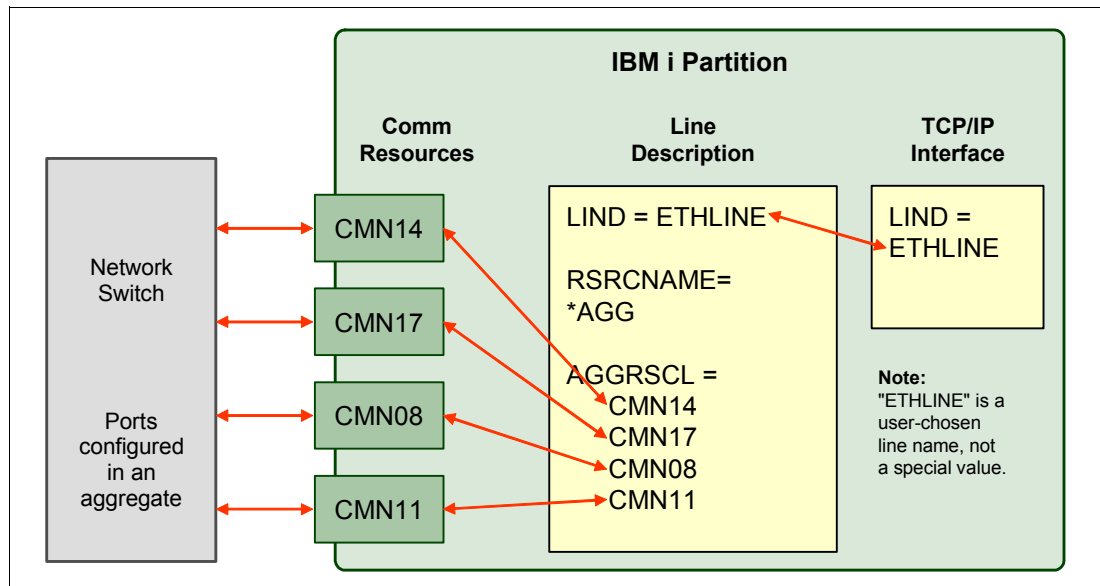


Figure 10-8 Logical view of example with four aggregated links

In the example, four links and IBM i communication resources (CMN14, CMN17, CMN08, CMN11) are aggregated together with one line description named ETHLINE.

The following command in Figure 10-9 creates the line description for the aggregated configuration:

```
CRTLINETH LIND(ETHLINE)
          RSRNAME(*AGG)
          AGGPCY(*ETHCHL *RNDRBN)
          AGGRSCL(CMN14 CMN17 CMN08 CMN11)
          LINESPEED(1G)
          DUPLEX(*FULL)
          TEXT('Four link aggregated line')
```

Figure 10-9 Example CRTLINETH command for four aggregated links

For more information about configuring Ethernet resources and link aggregation, refer to the IBM i Information Center. For Ethernet requirements, refer to the hardware requirements section.

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

10.9 Sharing physical Ethernet connections across partitions (Ethernet layer-2 bridging)

The following sections describe the sharing of physical Ethernet connections via Ethernet level-2 bridging using IBM i 7.1.

- ▶ 10.9.1, “Introduction to Ethernet layer-2 bridging”
- ▶ 10.9.2, “How Ethernet layer-2 bridging works on IBM i”
- ▶ 10.9.3, “IBM i Prerequisites for Ethernet layer-2 bridging” on page 336
- ▶ 10.9.4, “Configuring a shared network connection via Ethernet layer-2 bridging” on page 337

10.9.1 Introduction to Ethernet layer-2 bridging

Logical partitions in a Power 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 System.

One answer to this dilemma is the new Ethernet layer-2 bridging function in IBM i 7.1. While 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 the hassle, expense and complication of 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 may be reduced administration costs since there are fewer physical resources to manage,
- Complexity may 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

- A single partition is defined to have a physical Ethernet adapter.
- A bridge using the IEEE 802.1D standard, is configured to link a virtual Ethernet LAN to the physical Ethernet adapter.
- Frames transmitted by virtual Ethernet adapters on the same VLAN as the bridging virtual Ethernet adapter can be sent to the physical network.
- Frames sent from the physical network can be received by adapters on the virtual network.
- Once 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 follows in Figure 10-10 on page 336.

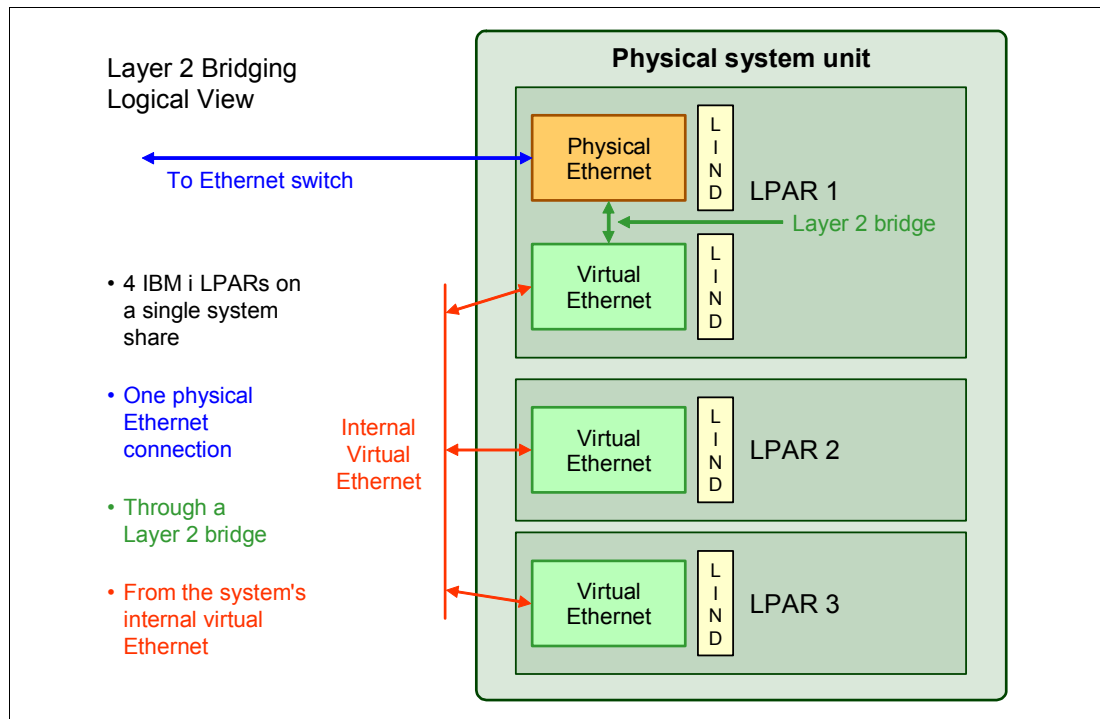


Figure 10-10 Layer-2 Bridging Logical View on IBM i

10.9.3 IBM i Prerequisites for Ethernet layer-2 bridging

To use Ethernet layer-2 bridging, users must have the following prerequisites:

- 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 of:
 - the Hardware Management Console
 - the IBM i Virtual Partition Manager

Note: IBM recommends that the selected Ethernet resources be used for only layer-2 bridging and not for IBM i TCP/IP configuration. There is a significant increase in processor usage for any host traffic that uses bridged resources.

10.9.4 Configuring a shared network connection via Ethernet layer-2 bridging

The following steps are used to configure Ethernet layer-2 bridging to share a network card:

- a. The administrator uses the management console to create a virtual Ethernet adapter in the IBM i partition, indicating that the adapter will be used for external access.

Configuring a virtual adapter using an Hardware Management Console

If using a Hardware Management Console, navigate to **Systems Management** → **Servers** → **Configuration** → **Manage Profiles** → **Edit profile** → **Virtual Adapters** to reach the screen shown in Figure 10-11 on page 337.

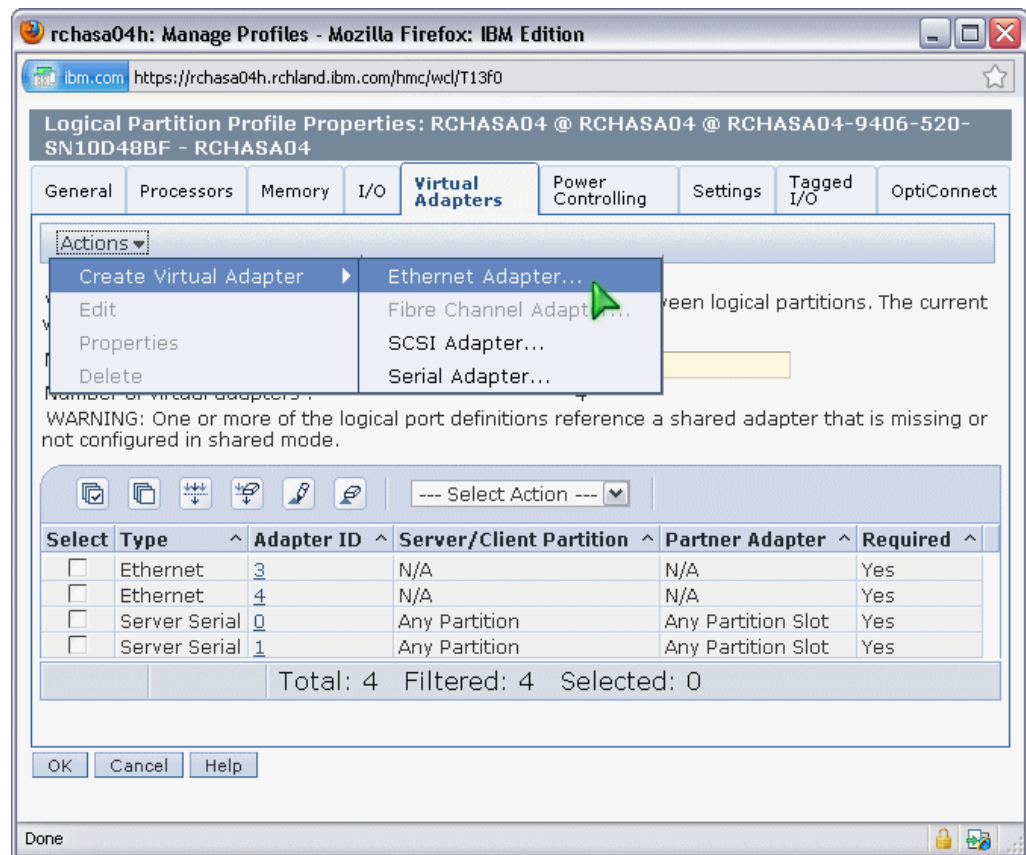


Figure 10-11 Select Create Virtual Adapter - Ethernet Adapter on an HMC

When creating the virtual Ethernet adapter, check the Access external network box to indicate that this virtual Ethernet adapter is used to bridge traffic to the physical network as shown in Figure 10-12 on page 338.

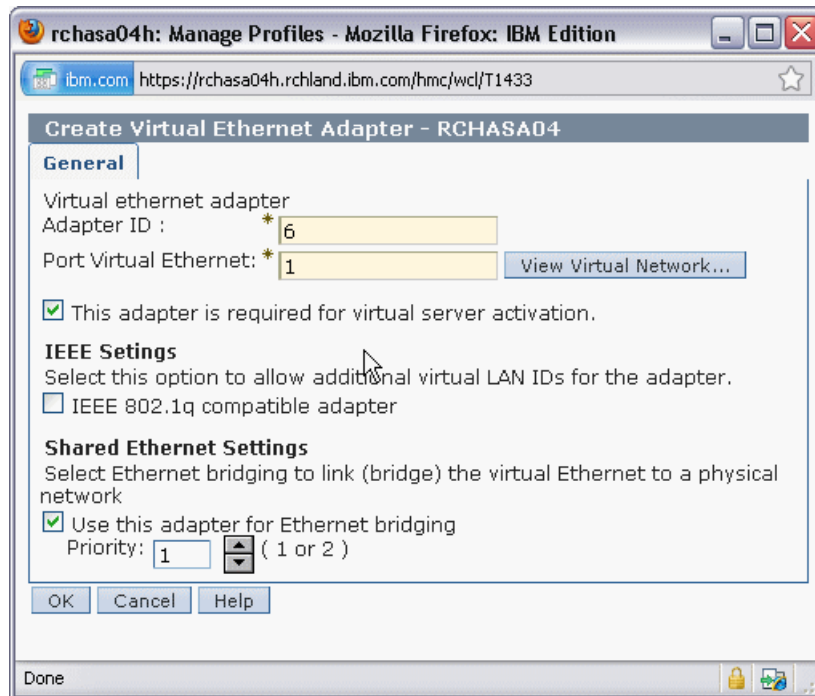


Figure 10-12 Creating the virtual Ethernet adapter on an HMC

Configuring a virtual adapter using IBM i Virtual Partition Manager

If 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 **STRSST** → **Work with system partitions** → **Work with partition configuration** → **Change**, then create the virtual adapter by changing one or more of the virtual ethernet identifiers to 1 (Yes) as shown in Figure 10-13 on page 338.

Change Partition Configuration		System: RC
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	
	1 2 2 2	

Figure 10-13 Virtual Partition Manager with virtual Ethernet ID1 activated

- b. On the IBM i partition with the physical adapter, create two Ethernet line descriptions:
 - i. One line description for the Ethernet link (physical communications resource CMN09) connected to the physical network as shown in Figure 10-14.

```
CRTLINETH LIND(ETHLINEP) RSRNAME(CMN09) BRIDGE(COVERED) LINESPEED(1G)
DUPLEX(*FULL) TEXT('Line for physical Ethernet link')
```

Figure 10-14 Command to create line description for the physical link

- ii. One line description for the new virtual Ethernet adapter (virtual resource CMN14) as shown in Figure 10-15.

```
CRTLINETH LIND(ETHLINEB) RSRNAME(CMN14) BRIDGE(COVERED) LINESPEED(1G)
DUPLEX(*FULL) MAXFRAME(8996) TEXT('Line for virtual Ethernet bridge')
```

Figure 10-15 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 via the Work with Hardware Resources (WRKHDWRSC) command specifying the TYPE(*CMN) parameter.

- c. 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”.
 - d. When both line descriptions are active, traffic will be bridged between the physical network and the virtual networks.
 - e. On any IBM i partition which is to use the bridged connection, a line description must be created specifying a virtual communications resource and the bridge name of “COVERED”. The command to do this is shown in Figure 10-16 on page 339.

```
CRTLINETH LIND(ETHLINVRT) RSRNAME(CMNxx) BRIDGE(COVERED) LINESPEED(1G)
DUPLEX(*FULL) MAXFRAME(8996) TEXT('Line for virtual Ethernet')
```

Figure 10-16 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, refer to the IBM i Information Center. For Ethernet requirements, refer to 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 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	7.1 PTFs	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 discusses the following topics related to printing and output that were enhanced in IBM i 7.1:

- ▶ 11.1, “Printing overview” on page 342
- ▶ 11.2, “Overview of IBM i 6.1 print enhancements” on page 343
- ▶ 11.3, “Enhancements to base print” on page 344
- ▶ 11.4, “Print Services Facility for IBM i (PSF) enhancements” on page 350
- ▶ 11.5, “Transform services enhancements” on page 353
- ▶ 11.6, “IBM Systems Director Navigator for i print enhancements” on page 354
- ▶ 11.7, “IBM i Access for web print enhancements” on page 362
- ▶ 11.8, “Host Print Transform enhancements” on page 362

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 a variety of 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 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:

- ▶ 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 contained in the printer file.
- ▶ 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.
- ▶ The destination of the output data is based on values 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.
- ▶ 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
- ▶ The printer writer program interacts between the output queue and the printer and can be used to convert the printer data stream.
- ▶ The printer writer program included in IBM i supports a variety of printer data streams. Print Services Facility™ for IBM i provides 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 a variety of attachment methods.

- ▶ A remote writer allows you to route spooled files from an output queue on your system to another system.

Figure 11-1 shows the IBM i printing process.

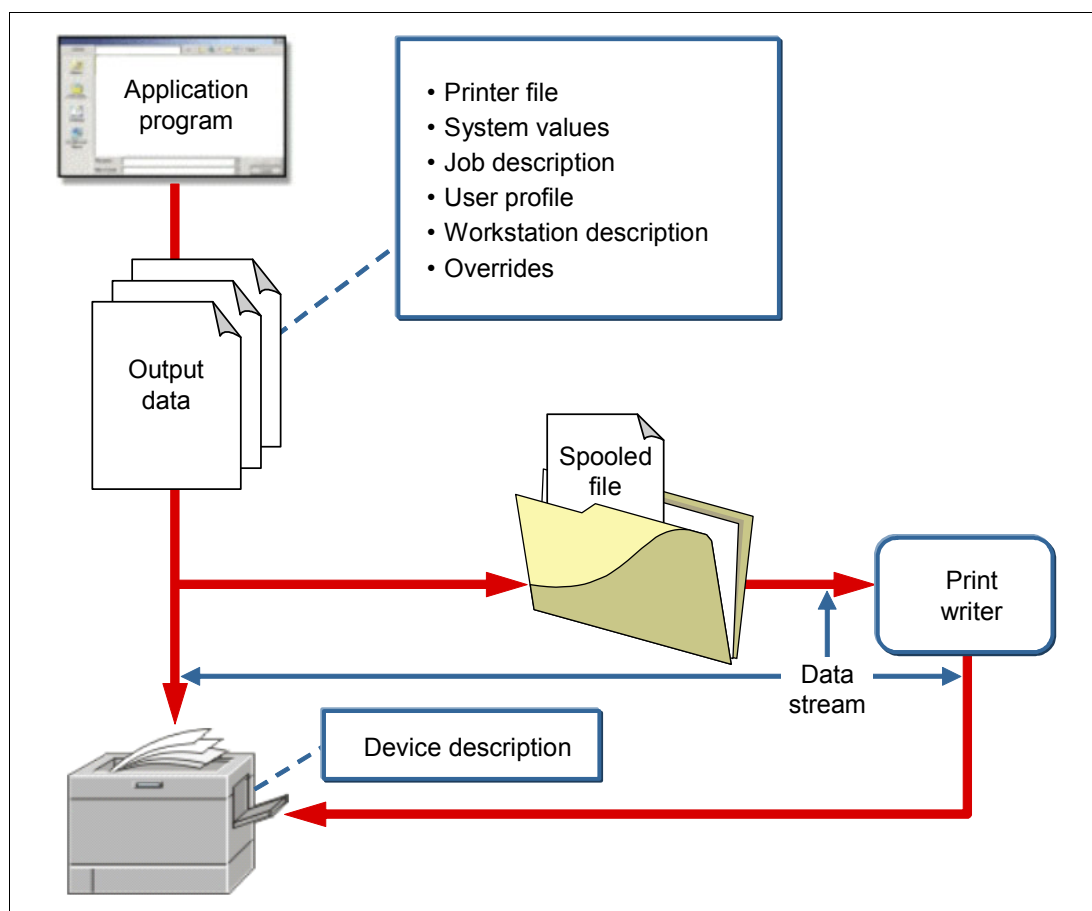


Figure 11-1 IBM i printing process

11.2 Overview of IBM i 6.1 print enhancements

Major enhancements to print were made in IBM i 6.1:

- ▶ Date and timestamps in the spooling function were changed to Universal Coordinated Time (UTC). This also required enhancements in data queue formats and addition of a third format that supports the UTC time/date format.
- ▶ System i Navigator and IBM Systems Director 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 included 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 data be transformed.

For more information about the IBM i 6.1 print enhancements, see the IBM i 6.1 Technical Overview available at the following web page:

<http://www.redbooks.ibm.com/redbooks/pdfs/sg247713.pdf>

11.3 Enhancements to base print

The sections that follow discuss the following enhancements, which have been made to the system base print functions:

- ▶ 11.3.1, “CPYSPLF command enhanced to copy to PDF or TIFF stream files”
- ▶ 11.3.2, “New QIBM_QSP_SECURITY exit point and formats” on page 347
- ▶ 11.3.3, “Encryption of PDF output (Infoprint Server LP)” on page 349

11.3.1 CPYSPLF command enhanced to copy to PDF or TIFF stream files

The Copy Spoolfile (CPYSPLF) command has been 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 data base 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 to be 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 data base file (TOFILE) parameter is *TOSTMF.

*NONE

The output is written to a user-defined physical file. This value is only valid if the To data base 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 data base file (TOFILE) parameter specifies *TOSTMF.

Note: If the stream file exists, the CCSID associated with the stream file will not be changed.

Workstation customizing object (WSCST)

Specifies the workstation customizing object to use to transform the spooled file output to final form before writing it to a stream file. If the To data base file (TOFILE) parameter specifies a physical database file, the WSCST parameter is ignored. In order 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 to be used.

If the To data base file (TOFILE) parameter specifies *TOSTMF and the device type of the spooled file is *AFPDS or *USERASCII, the spooled file data will be copied directly

to the stream file. If the stream file does not exist, the associated CCSID of the stream file will be set to 65535.

For other types of spooled files, the spooled file data will be copied to the stream file using the Control character (CTLCHAR) parameter to format the data. Lines will be ended 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 will be associated with the stream file.

If the spooled file has a CHRID attribute other than *DEV D, the CHRID attribute will be 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 which created the spooled file will be 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 will be 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 will be 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 already exists. If the stream file does not exist, it is created.

*NONE

No records are copied and the operation will fail.

*REPLACE

The records replace the existing stream file records.

Since 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 on page 346, the CPYSPLF command prompter is displayed and F11 has been pressed to display the parameter names. The *TOSTMF value is specified for the TOFILE parameter as shown in the red box. This 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 on page 347, the next page of the CPYSPLF command is shown. In the red box, specification of the new TOSTMF, WSCST and STMFOPT parameters is shown.

- the TOSTMF parameter directs to copy output to file /pdfoutput,
- the WSCST parameter specifies to convert to a PDF file and
- the STMFOPT parameter specifies to replace 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

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

```

Figure 11-3 Copy spooled file command prompt - new TOSTMF, WSCST and STMFOPT parameters

Note: 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 product.

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 created by that user.

Spooled file security has been 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 previously mentioned 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 mentioned in the previous list can be overridden by the spooled file security exit program.

This is another way to access spool files for the CPYSPLF (Copy Spooled File), DSPSPLF (Display Spooled File), and SNDNETSPLF (Send Network Spooled File) commands, if DSPDTA(*YES) was specified when the output queue was 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 already listed previously, using DSPDTA(*NO) when creating the output queue does not restrict the user from displaying, copying, or sending the file. DSPDTA authority is only checked 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 resides 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 executed in a system job (including SCPF), a subsystem monitor job, or any job that is running under one of the system user profiles 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
- ▶ DLTSPFL: Delete Spooled File
- ▶ DSPSPLF: Display Spooled File
- ▶ EXPORT: iSeries Navigator API to export an EBCDIC spooled file to an ASCII desktop
- ▶ HLDSPLF: Hold Spooled File
- ▶ QGSLRSC: 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
- ▶ SNTCPSPFL: Send TCP/IP Spooled File
- ▶ WRKPRTSTS: Work with Printing Status
- ▶ WRKSPLFA: Work with Spooled File Attributes

More details regarding the exit program format names, formats and parameters are available in the IBM i 7.1 Information Center:

<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":

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.

Note: 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 CHGSPLFA (Change Spooled File Attributes) command.
- ▶ It can be specified in a printer file using the CRTPTF (Create Printer File) command.
- ▶ It can be specified by using the OVRPTF (Override Printer File) 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 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 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, perform 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 parameter PDFMAP in the Print Services Facility (PSF) configuration object.
4. Specify the default IBM-supplied user program with parameter PDFMAPPGM in the PS configuration object.
5. Specify the PDFENCRYPT parameter on the printer file's or spooled file's USRDFNDA 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:
 - 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:
USRDFNDA('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:
 - 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:
USRDFNDA('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:
 - 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:
USRDFNDA('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:
 - 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:
USRDFNDA('PDFENCRYPT(*STMF *NOMAIL)')

11.4 Print Services Facility for IBM i (PSF) enhancements

This section overviews the Print Services Facility for IBM i (PSF) enhancements in IBM i 7.1. They include:

- ▶ 11.4.1, "Disabling offset stacking and edge mark printing" on page 351.
- ▶ 11.4.2, "Specify public data authority for directories created by PSF" on page 351.
- ▶ 11.4.3, "Improved PSF debugging capabilities" on page 352.

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 (PSFDFNOPT) parameter named OFFSTACK:

- ▶ PSFDFNOPT (OFFSTACK(*YES))

This is the default. This command specifies that there is no change to offset stacking and edge marking. That is, offset stacking and edge marking occurs between each spooled file. All offset stacking and edge mark changes in the form definition are honored.

- ▶ PSFDFNOPT (OFFSTACK(*NO))

This command specifies that no offset stacking or edge marking is done. This 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 CRTPSFCFG (Create PSF Configuration) or specified by using the CHGPSFCFG (Change PSF Configuration) command.

11.4.2 Specify public data authority for directories created by PSF

PSF has been 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 PSFDFNOPT (PSF Defined Option) parameter named PDFDTAAUT (PDF data authority).

The various PDFDTAAUT values and their functions are as follows:

- ▶ PSFDFNOPT (PDFDTAAUT(*INDIR))

The authority for the directory to be created is determined by the directory in which it is to be created. The directory immediately preceding the new directory determines the authority.

A directory 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 to be created.

A directory 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))

The user can change the object and perform basic functions on the object except those 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))

The user 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.

The data area must be created before starting the printer writer, must be created in library QGPL, and the name must match the printer device description name. When the PSFTRACE file is no longer needed, delete the data area with the DLTDTAARA (Delete Data Area) command.

For more information about PSFTRACE and interpreting the data within it, see the Troubleshooting Mapping Problems section of the Advanced Function Presentation PDF of the IBM i 7.1 Information Center at the following web page:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzau6/rzau6.pdf>

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'.

11.5 Transform services enhancements

Transform Services functional enhancements are discussed in this section:

- ▶ 11.5.1, “Generate PDFs from existing spooled files” on page 353
- ▶ 11.5.2, “PDF transform enhancements” on page 353

11.5.1 Generate PDFs from existing spooled files

Prior to IBM i 7.1, data was sent directly to transform services. If the job generating the spool file ended abnormally without closing the spool file, no PDF output was generated.

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.

This capability is covered in more detail in 11.8, “Host Print Transform enhancements” on page 362.

11.5.2 PDF transform enhancements

The following sections discuss enhancements to the PDF transform.

Additional barcode support

Barcode support of the PDF transform has been enhanced to render all barcode types 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. Text is limited to ANSI characters.

Improved national language support

Eastern European languages require embedded fonts to display all 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 forth 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 Systems Director Navigator for i print enhancements

IBM Systems Director Navigator for i has the following enhancements related to printing:

- ▶ The View as PDF task was added to enable viewing spooled files as PDFs. This includes AFPDS and SCS printer output files. This function allows users to open and view the contents of a printer output file using Acrobat Reader.
- ▶ The Export task was replaced with an “Export as” pop-up menu. The pop-up menu adds the ability to export spooled files as PDF files. The new export of PDF files function allows the user 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.

Note: The “Export as PDF” to IFS, to an output queue and to an email require the 5722IP1 Infoprint Server for iSeries licensed program product.

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 is displayed as shown.

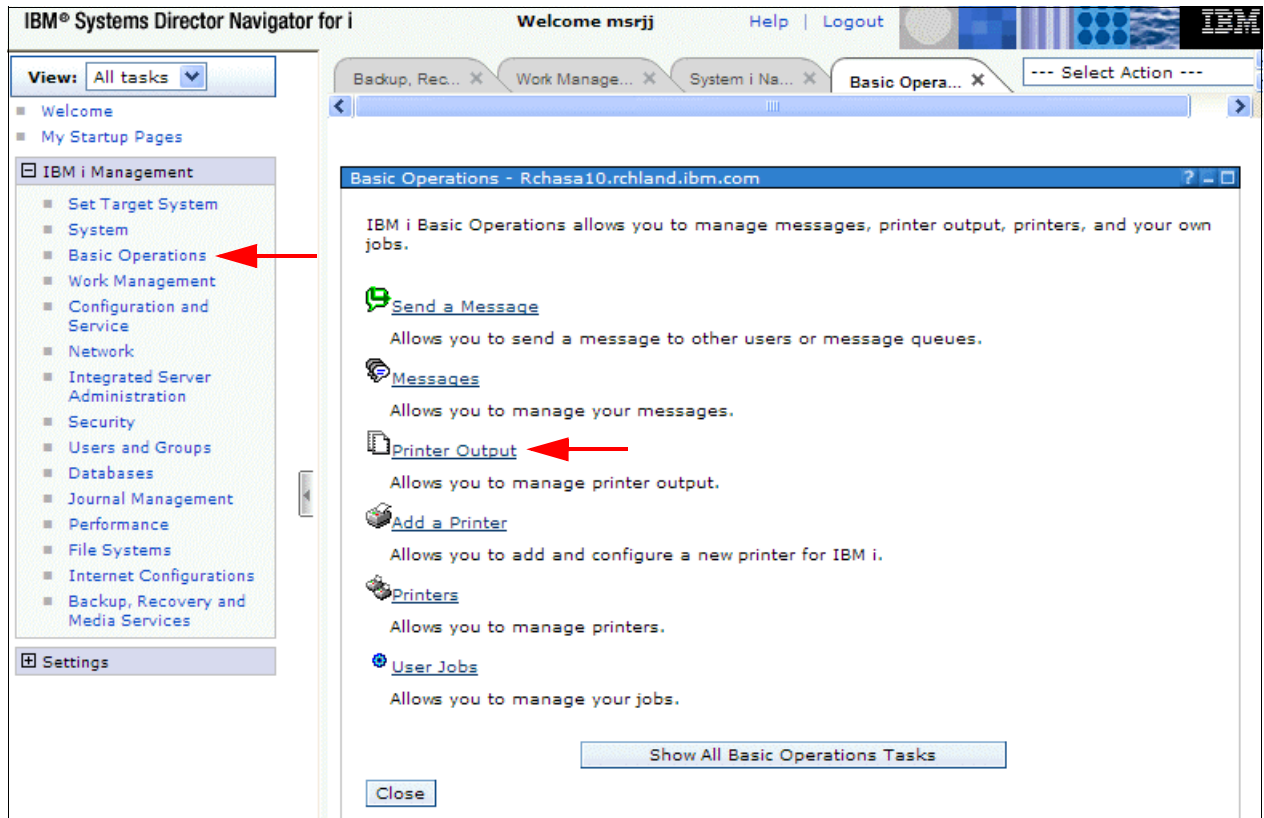


Figure 11-4 Navigating to the printer output list in IBM Systems Director Navigator for i

The second arrow in Figure 11-4 on page 355 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.

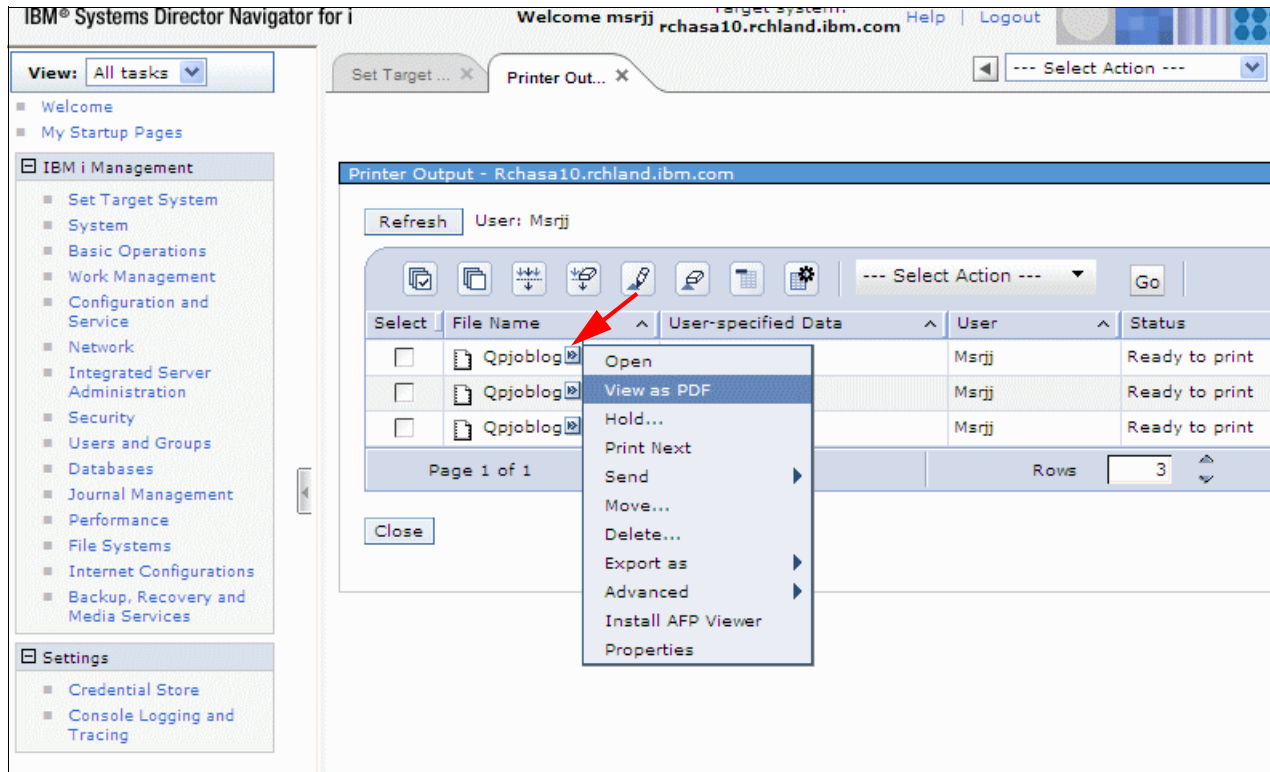


Figure 11-5 Printer Output list with pop-up menu in IBM Systems Director Navigator for i

The arrow points to the pop-up menu icon after the print output file. When this is selected, the pop-up menu is shown. Highlighted in the pop-up menu is the View as PDF task.

When the View PDF task is selected, you see the output as PDF, as shown in Figure 11-6.

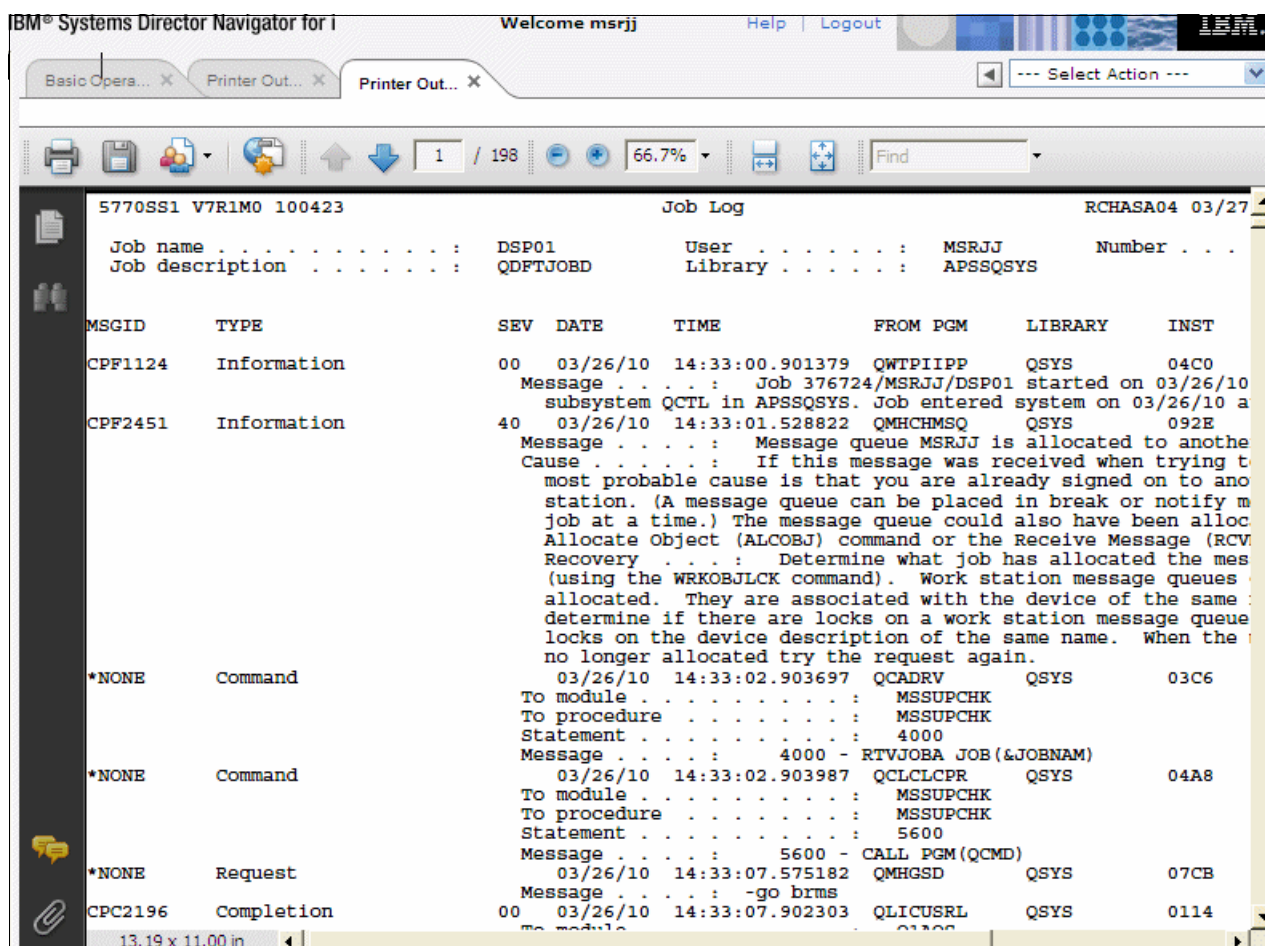


Figure 11-6 PDF displayed from View PDF in IBM Systems Director Navigator for i

11.6.2 Export printer output in PDF format using InfoPrint Server

Print functions of IBM Systems Director Navigator for i now uses transform services to export spooled files in PDF format. The output can be placed in 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

Note: For the latter three options, the InfoPrint Server Licensed Program Product (5722-IP1) is required. User's can use the native IBM Transform Services for i (5770-TS1) Licensed Program Product 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 product to convert spooled files to PDF remains.

The navigation to a specific printer output file is identical to Figure 11-4 on page 355. Perform the following steps to export printer output in PDF format using transform services:

1. In the Printer Output list panel, select the Printer Output menu item on the Basic Operations menu. the pop-up menu is displayed, as shown in Figure 11-7.

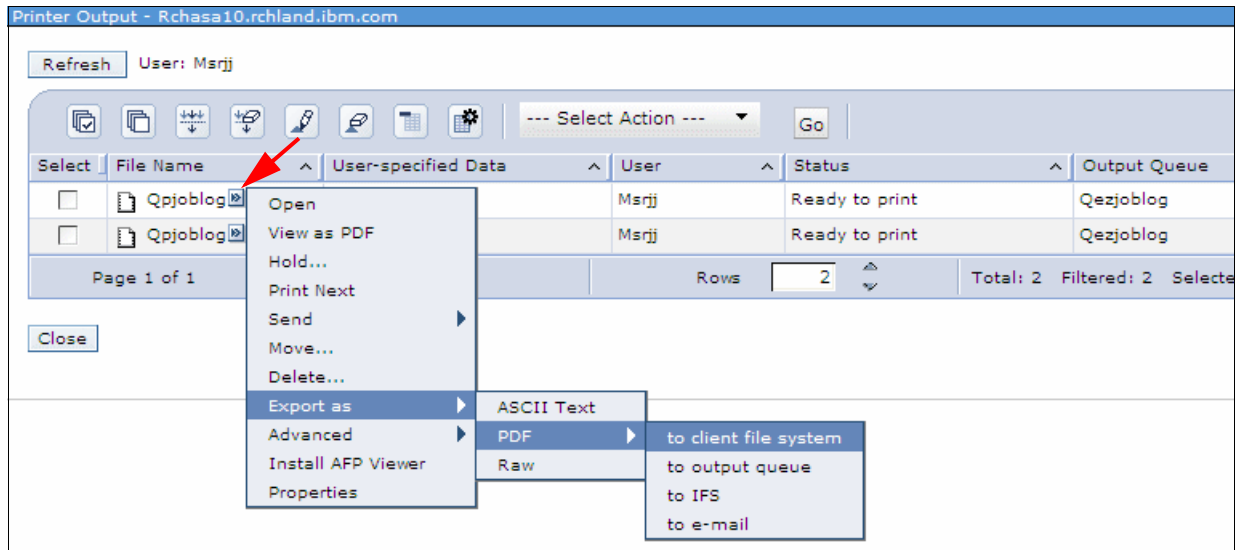


Figure 11-7 Printer output list with Export PDF options shown

2. Select the **Export as** task. A pop-up menu with PDF options is displayed, as in Figure 11-7. Select **PDF**. In the next pop-up menu that is displayed, select **to IFS**. The Convert Printer Output to PDF Wizard (Figure 11-8) is displayed.

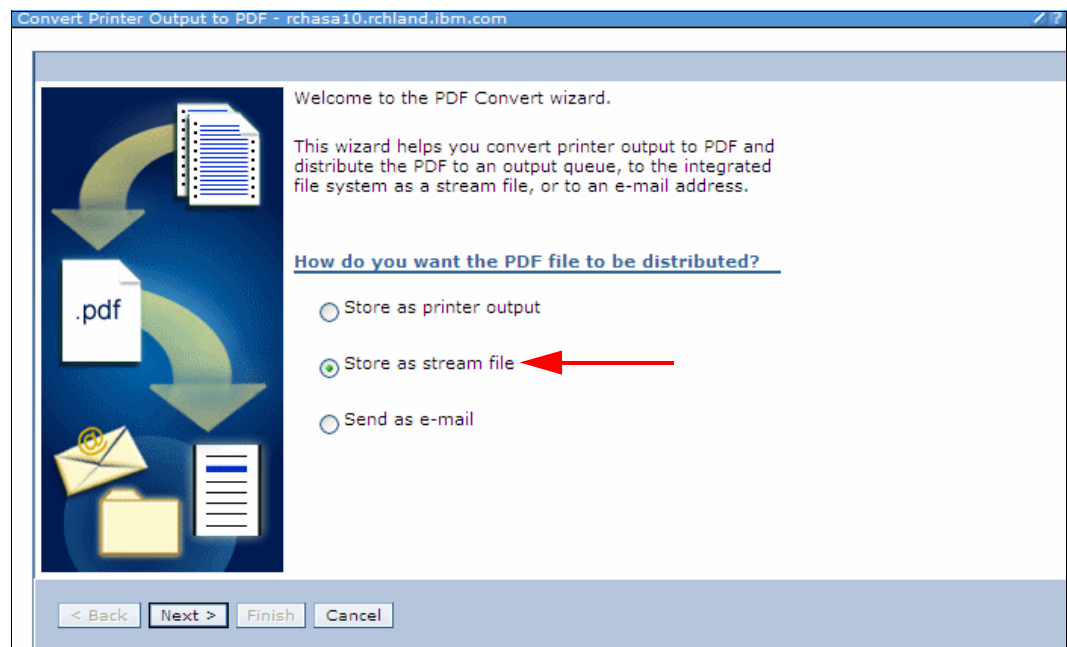


Figure 11-8 Convert printer output to PDF wizard

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 11-9.

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.

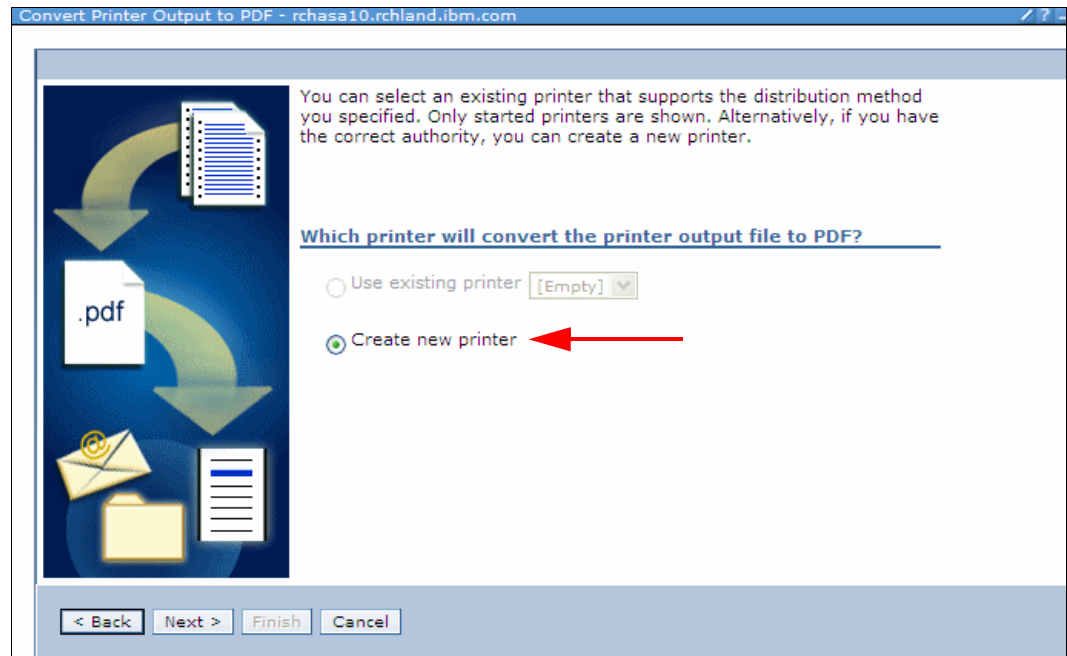


Figure 11-9 Convert Printer output to PDF - select printer

4. Click **Next**. You are prompted for a printer name and port, as shown in Figure 11-10. The Port defaults as shown.

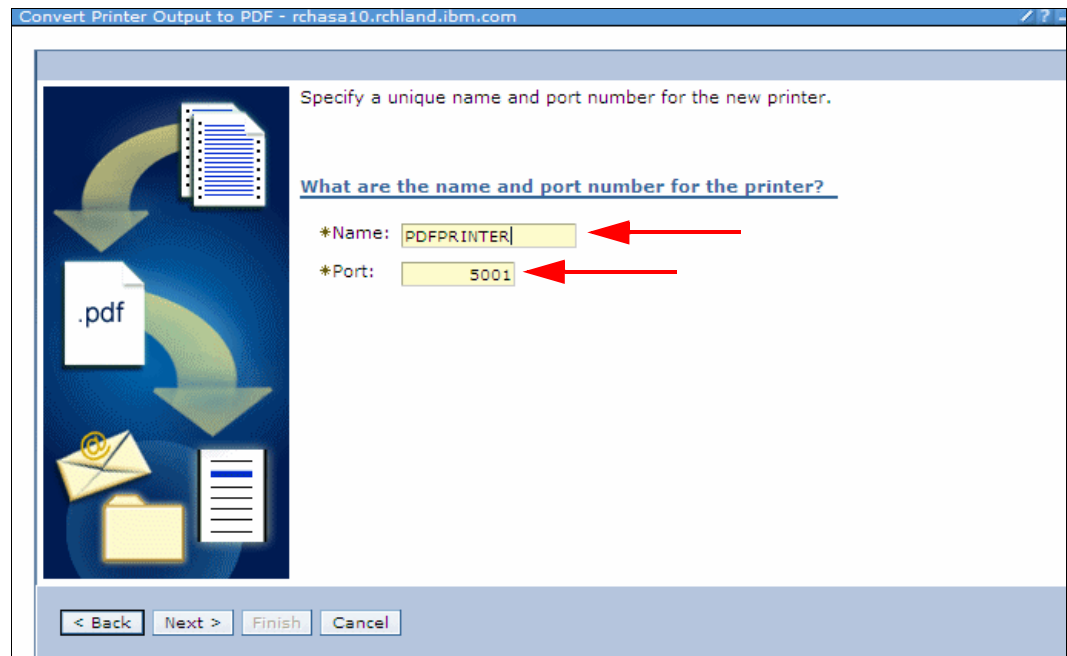


Figure 11-10 Creating a PDF printer

5. Set the printer name to PDFPRINTER, and click **Next**.

- Specify Advanced configuration parameters in the printer configuration panel shown in Figure 11-11. Click **Next**.

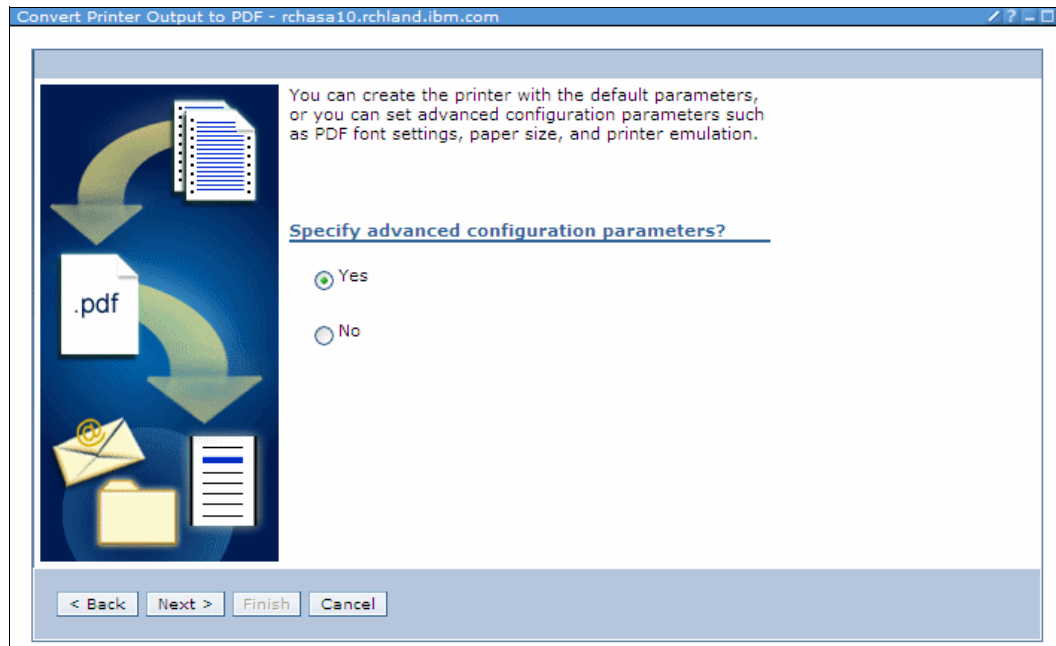


Figure 11-11 Specify Advanced configuration parameters

An advanced parameters window (Figure 11-12) is displayed.

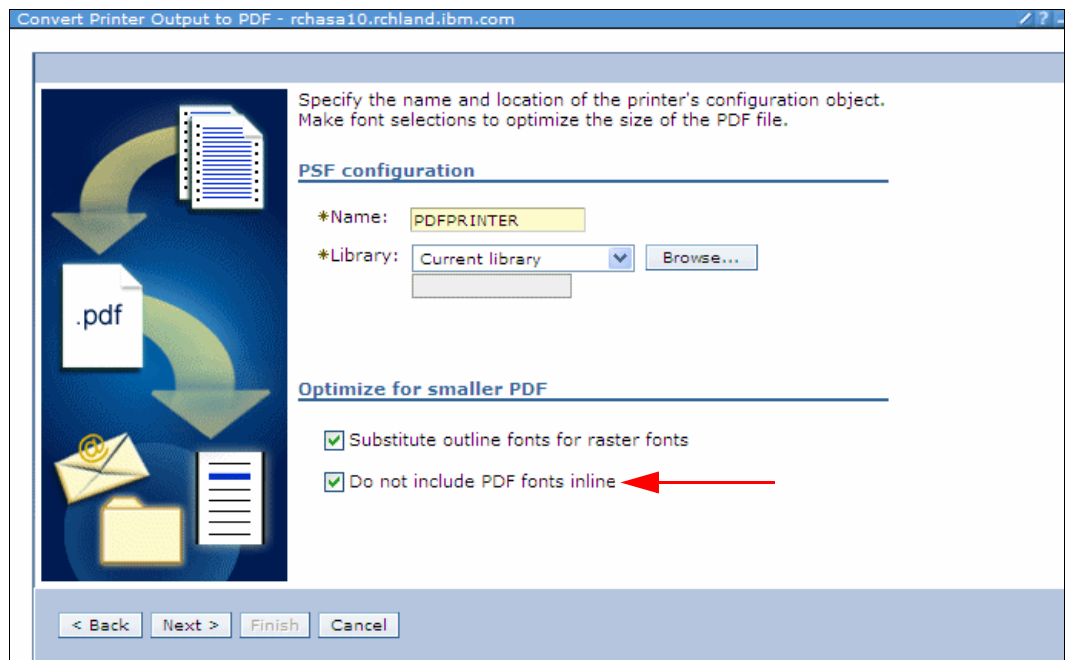


Figure 11-12 Advanced PDF printer configuration parameters

The arrow in Figure 11-12 points to an important function to minimize the PDF size. Transform services embeds the PDF fonts in the PDF to preserve text appearance and text content. This 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 11-13) is displayed. Accept the defaults and click **Next**.

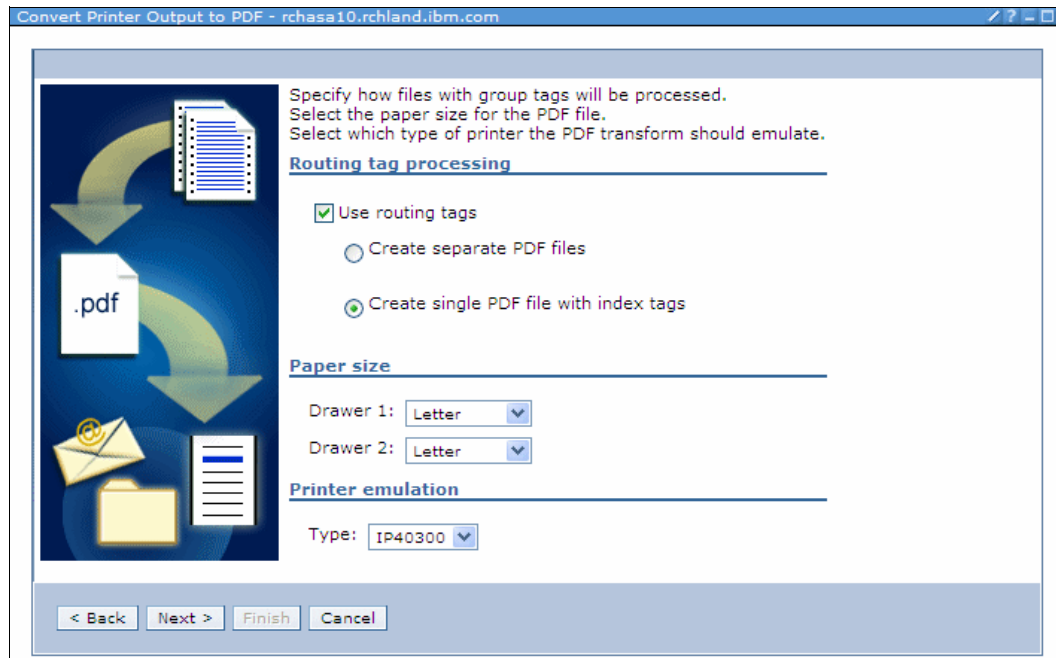


Figure 11-13 Advanced parameters

We now see the IFS path where the PDF is to be stored in Figure 11-14.

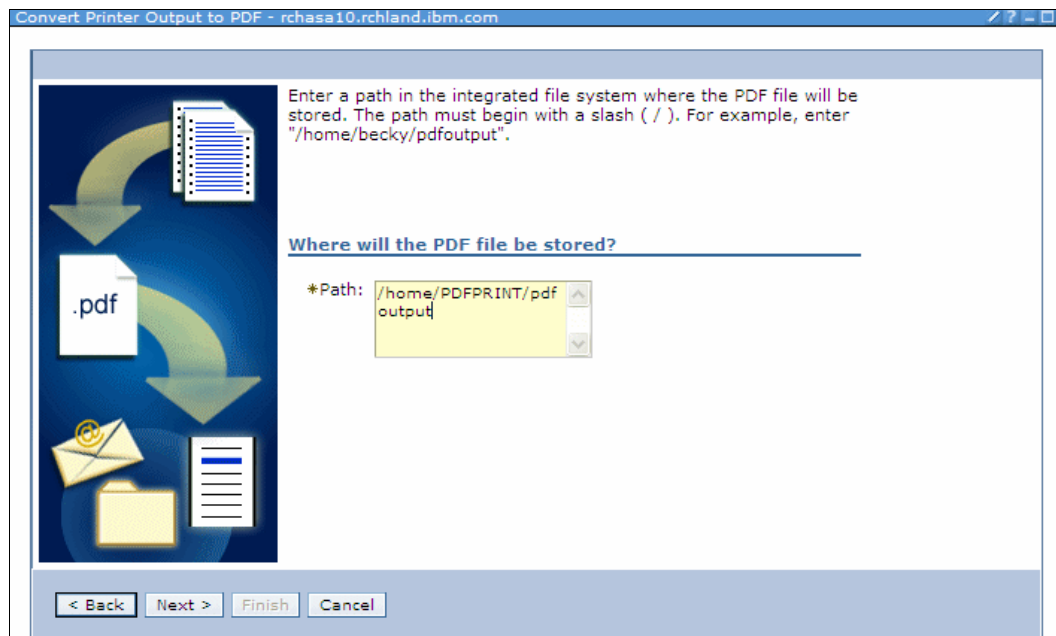


Figure 11-14 Setting the destination in the IFS

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 panel 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 has been added for the AFP to PDF Transform (option 1) of the IBM Transform Services for i (5770-TS1) licensed program product when viewing printer spool output as a PDF document.

Output can be viewed in a browser or placed in the IBM i integrated file system (IFS).

More 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 the following web page:

<http://publib.boulder.ibm.com/infocenter/iseres/v7r1m0/topic/rzamm/rzamm.pdf>

Note: 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 Transforms now uses Transform Services. The following sections explains the changes with examples.

11.8.1 The Host Print Transform API now uses Transform Services

The QwpzHostPrintTransform (Host Print Transform) 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 attribute CTXFORM. Three such WSCST objects with attribute CTXFORM are shipped with the system. One of these can be used, or you can create your own using the CRTWSCST (Create WSCST) 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 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 is 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 generally results in the smallest possible PDF file size, but only reliably presents Latin and CJK languages.

► QSYS/QCTXPDMBD

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 invoked with a WSCST object with attribute CTXFORM, the job reads the input data stream from the spooled file 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 provided on the API. See the API documentation in the IBM i 7.1 Information Center for details.

11.8.2 Example: Host Print Transform API with Transform Services

This sample flow illustrates the use 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 if it will transform the spooled file. It returns to the caller an indication whether it will transform the spooled file. It also passes back an indication that it will read 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 will not hold all of the transformed data, the Host Print Transform API will indicate 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 has been passed back to the caller at which point the Host Print Transform API will indicate 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 starting with the process file (20) process option are repeated.
6. All transforming is completed.
The Host Print Transform API is called with the terminate (50) process option. The Host Print Transform API cleans up any work spaces that it has created and returns to the caller.

11.9 References

The following references are for additional information regarding IBM i printing.

IBM i Printing Basic Printing PDF

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzalu/rzalu.pdf>

IBM i Printing Advanced Function Presentation PDF

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzau6/rzau6.pdf>

IBM Advanced Function Printing Utilities for iSeries: User's Guide PDF

http://publib.boulder.ibm.com/infocenter/iseriess/v6r1m0/topic/books_web/s5445349.pdf

IBM i Files and File Systems Spooled files PDF

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzata/rzata.pdf>

System i Programming - DDS for Printer Files PDF

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzakd/rzakd.pdf>

iSeries Guide to Output

<http://publib.boulder.ibm.com/infocenter/iseriess/v6r1m0/topic/rzalu/s5445319.pdf>

InfoPrint AFP Font Collection web page

http://www-03.ibm.com/systems/i/software/print/afpfonthome_m_ww.html

InfoPrint AFP Print Utilities for System i

http://www-03.ibm.com/systems/i/software/print/afputilhome_i_ww.html

Print Services Facility for IBM i for i5/OS

http://www-03.ibm.com/systems/i/software/print/ipfontscmp_m_ww.html



Integration with BladeCenter and System x

This chapter describes the new enhancements to the integration with BladeCenter and System x with IBM i 7.1

The following topics are discussed:

- ▶ 12.1, “iSCSI software targets” on page 366
- ▶ 12.2, “Defining the iSCSI software target support” on page 367
- ▶ 12.3, “Service Processor Manager function” on page 370
- ▶ 12.4, “VMware support changes” on page 370
- ▶ 12.5, “New planning work sheets” on page 372
- ▶ 12.6, “IBM Systems Director Navigator for i” on page 372
- ▶ 12.7, “New IBM i CL commands” on page 378
- ▶ 12.8, “IBM i changed CL commands” on page 380
 - ▶ 12.8.2, “CRTNWSCFG (Create NWS Configuration) and CHGNWSCFG (Change NWS Configuration) CL commands” on page 380
 - ▶ 12.8.3, “INSLNXSVR (Install Linux Server) CL command” on page 381
 - ▶ 12.8.4, “No new integrated Linux servers” on page 381
- ▶ 12.9, “Fewer IBM i licensed programs required” on page 381
- ▶ 12.10, “Changes to IBM i integration with BladeCenter and System x documentation” on page 381
- ▶ 12.11, “Additional information” on page 383

12.1 iSCSI software targets

With i 7.1, IBM i is now supporting iSCSI software targets using standard Ethernet Network Interface Cards (NICs) as shown in Figure 12-1. Software targets are providing 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 the customer is limited to 1 Gb connectivity if using 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 using iSCSI software initiator in combination with iSCSI software target, you have full 10 Gb connectivity.

12.1.1 IBM i Integrated server object model with a Target iSCSI

With a physical iSCSI Target HBA, the NWSH object identified the hardware resource and configured the IP address information (and other attributes) as shown in Figure 12-2.

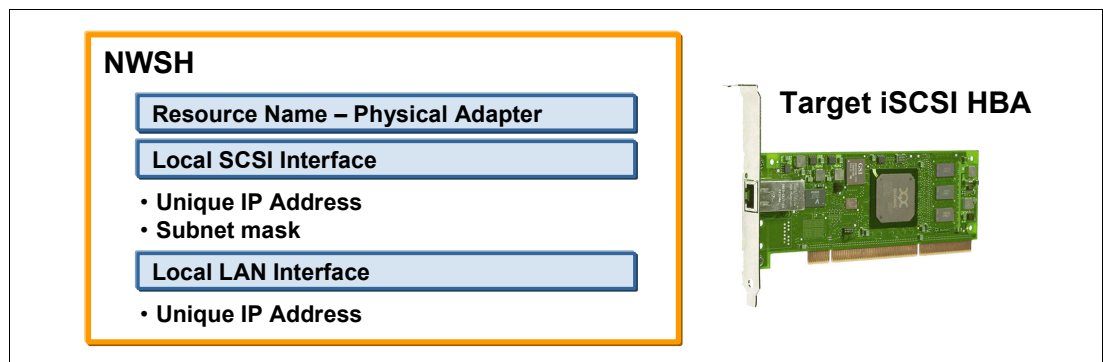


Figure 12-2 NWSH object for a physical iSCSI target HBA

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.

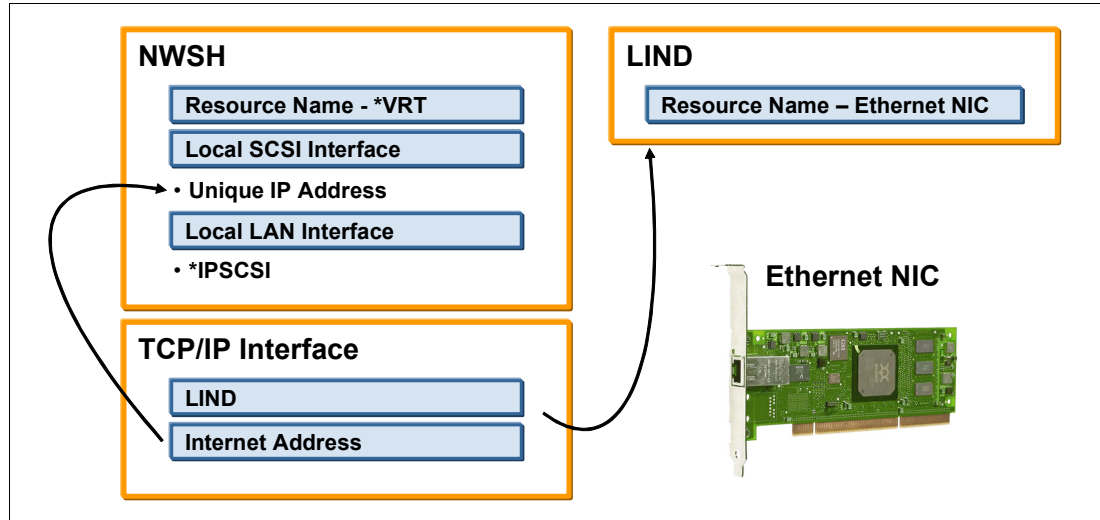


Figure 12-3 NWSH object for a software target

The physical iSCSI HBA has been replaced by an Ethernet NIC, which is defined in the Ethernet line description.

Within the TCP/IP interface, we 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.

12.2 Defining the iSCSI software target support

In IBM i 7.1, it is possible to define the iSCSI software target support in the green-screen interface and by using IBM Systems Director Navigator for i.

12.2.1 CRTDEVNWSH CL command interface

To define a iSCSI software target (Figure 12-4), we must specify *VRT for the resource parameter on the CRTDEVNWH (Create Device Description for a Network Server Host) 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

On the LCLIFC parameter, we specify the *IPSCSI option, which indicates the local LAN interface internet address is the same as the local SCSI interface internet address.

12.2.2 IBM Systems Director Navigator for i changes for iSCSI software target support

It is possible to define the iSCSI software target support by using IBM Systems Director Navigator for i. You can use the *New Network Server Host Adapter* option on the *Integrated Server Administration* page, as shown in Figure 12-5.

The screenshot shows a window titled "New Network Server Host Adapter - Localhost". On the left is a tree view with "General" selected, and sub-items "Communications" and "Local (Target) Interface". The main area is titled "Network server host adapter (NWSH) device description:". It contains the following fields and controls:

- Name:** A text box containing "SWTARGET".
- Description:** A text box containing "SWTARGET example for NWSH".
- Hardware resource:** A dropdown menu with "Virtual" selected.
- Online at IPL:** An unchecked checkbox.
- Object Authority:** A button.
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

Figure 12-5 Network server host adapter - Virtual resource

Here you can specify Virtual for the Hardware resource parameter to create the Network server host adapter device description for the iSCSI software target.

Within *IBM Systems Director Navigator for i*, it is possible to create a TCP/IP interface and a corresponding line description when creating a NWSH configuration object. This can be done by clicking **New**, as shown in Figure 12-6.

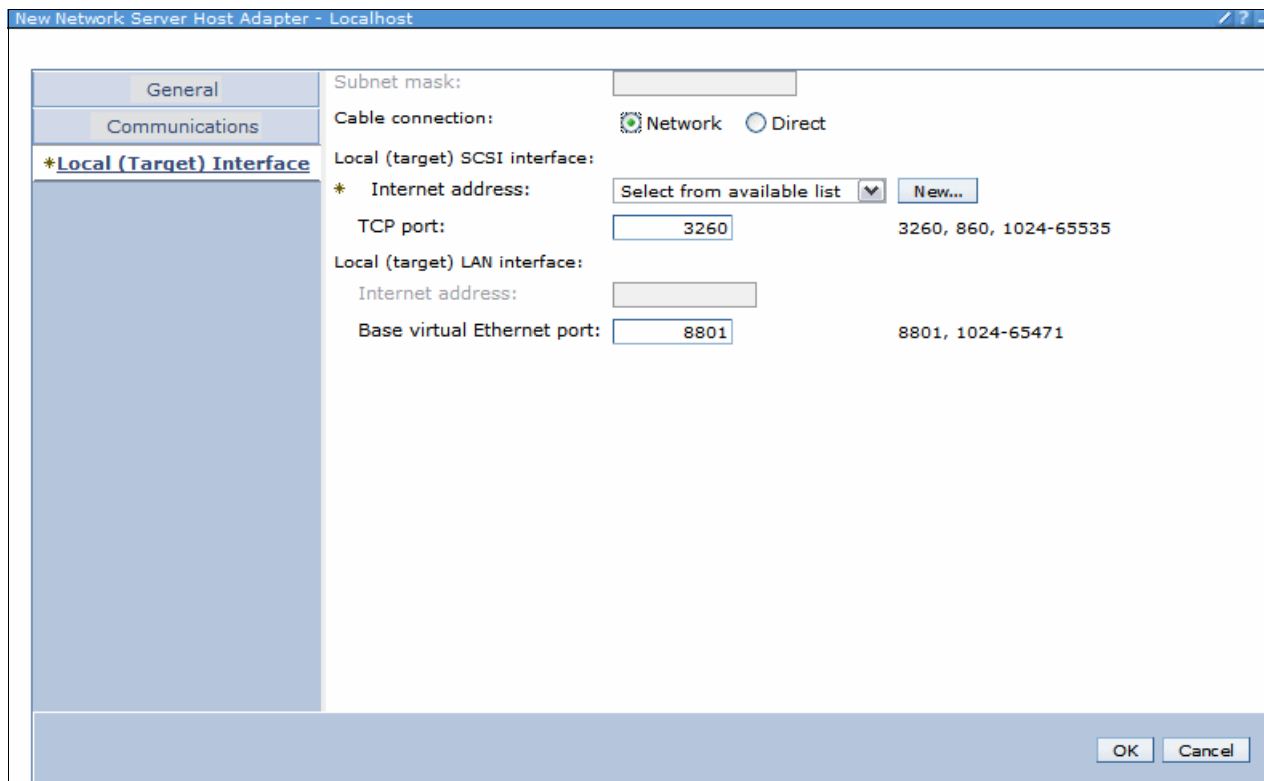


Figure 12-6 Create new TCP/IP interface for NWSH configuration object

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.

See the *Integrated Server Operating System (Server OS) Versions* section of the *IBM i iSCSI Solution Guide* for the specific OS versions that are currently supported with each IBM i release:

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 on 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 *ESX Server operating system

For VMware ESXi embedded servers, the system drive (the first drive) is no longer required as well.

Note: VMware ESX servers that were installed on prior IBM i releases must be migrated to the new NWSD type after installing 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 allows a concurrent save of ESX data without requiring the ESX server to be shut down or applications ended.

See the *IBM i iSCSI Solution Guide* for additional information:

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

12.5 New planning work sheets

New planning work sheets have been added to the *IBM i iSCSI Solution Work Sheets* PDF.

- ▶ IBM i TCP/IP interface work sheet
This worksheet is used when planning for iSCSI software targets.
- ▶ IBM i line description work sheet
This worksheet is used when planning for iSCSI software targets.
- ▶ Integrated server installation work sheet
This worksheet replaces the INSWNTSVR (Install Windows Server) command work sheet.
- ▶ VMware ESX post-installation work sheet
This worksheet is used for VMware ESX server post-installation tasks.

In addition, these work sheets have been enhanced to allow them to be filled in and saved softcopy.

The instructions for filling out these work sheets are in the *IBM i iSCSI Solution Guide* PDF.

Both PDFs are available on the following Web page:

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

Note: The instructions and work sheets were previously in the *iSCSI Network Planning Guide* topic in the Information Center.

12.6 IBM Systems Director Navigator for i

The *IBM Systems Director 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.

Note: The *System i Navigator* GUI is still available in IBM i 7.1 and works adequately for many tasks. However, the new GUI tasks 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 Systems Director Navigator for i* web GUI, and are discussed in the following sections:

- ▶ 12.6.1, “Create Server task” on page 373

- ▶ 11.6.2, “Clone Integrated Server task” on page 281
- ▶ 12.6.3, “Delete Server task” on page 377
- ▶ 12.6.4, “Launch web console” on page 378

12.6.1 Create Server task

This task, as shown in Figure 12-8, creates an iSCSI-attached integrated server.

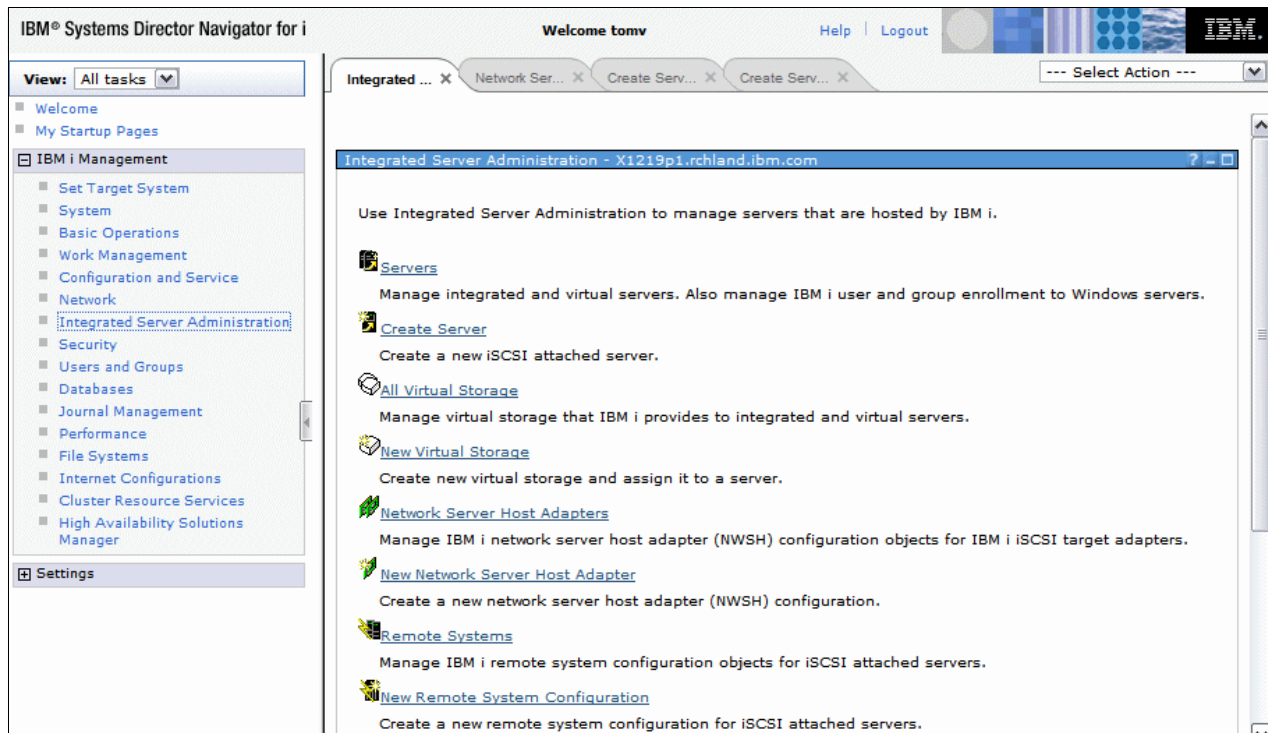


Figure 12-8 Create Server option in the IBM Systems Director Navigator for i web GUI

Creating a server in an IBM i Integrated Server environment has never been easier. It is recommended to use the *IBM Systems Director Navigator for i* to walk you through a server installation.

Be aware of the fact that upfront, you need to configure the following objects on 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 release IBM i 7.1, the following operating systems are supported:

- ▶ Windows Server 2008 and Windows Server 2003 Server editions.
- ▶ VMware ESX 4.0
- ▶ VMware ESXi Embedded 4.0, 4.1 and 5.0
- ▶ VMware ESXi Installable 4.1 and 5.0

This will startup the *Create Server Wizard* as shown in Figure 12-9. This is helpful to create a server that uses IBM i virtual storage.

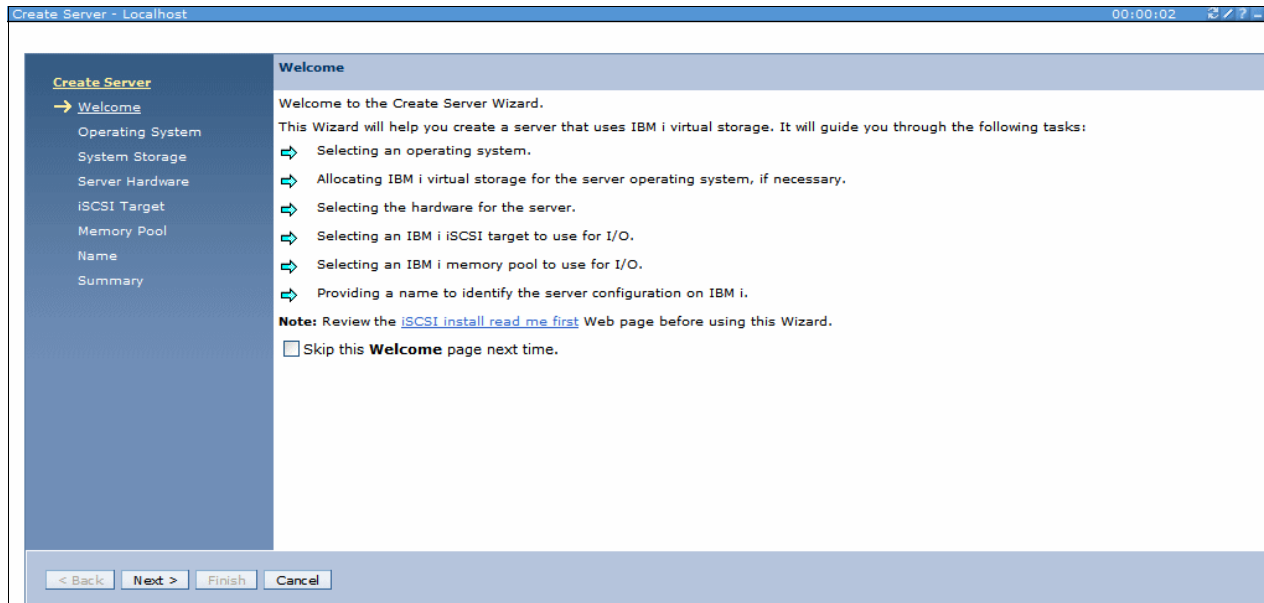


Figure 12-9 Create Server wizard

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

Note: Review the *Server installation road map and checklist* chapter of the *IBM i iSCSI Solution Guide* before using this Wizard

For more information, see the following web page:

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, as shown in figure Figure 12-10 creates an iSCSI-attached integrated Windows server based on one that was previously installed.

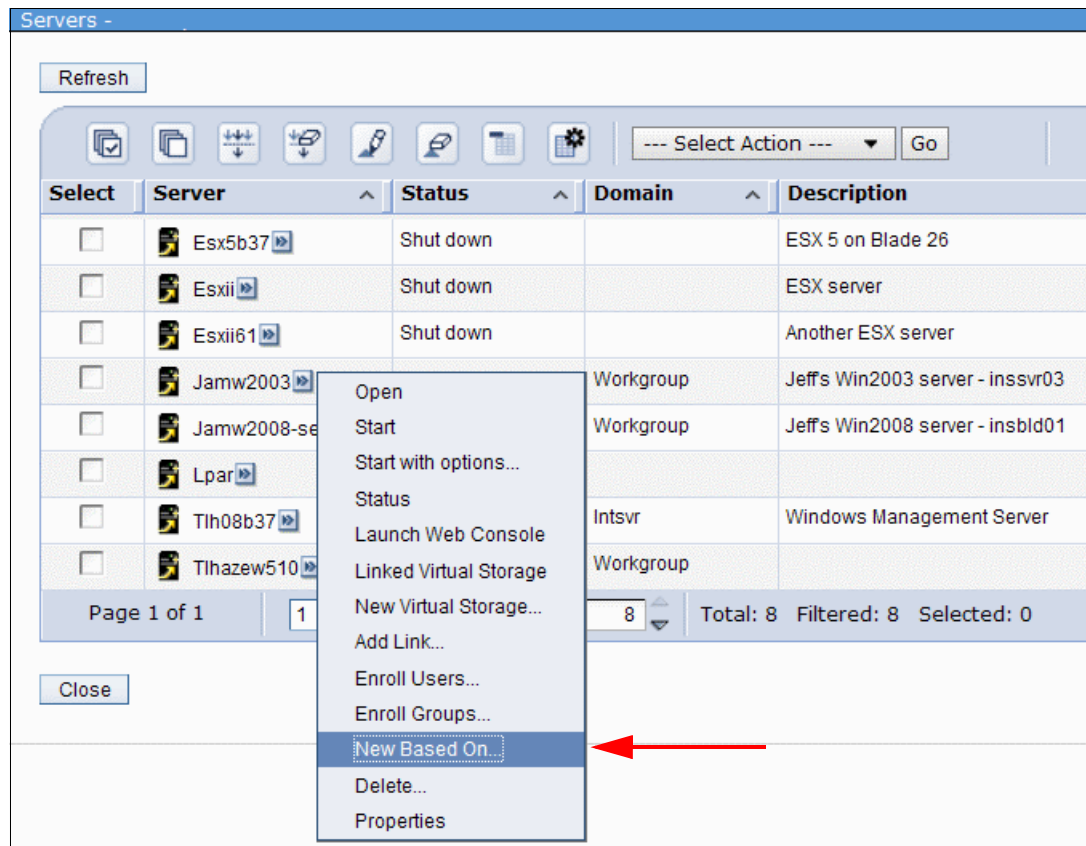


Figure 12-10 *New Based On...*(cloning) option in the IBM Systems Director Navigator for i Web GUI

This starts the *Create Server Based On* (cloning) Wizard as shown on Figure 12-11 on page 376.

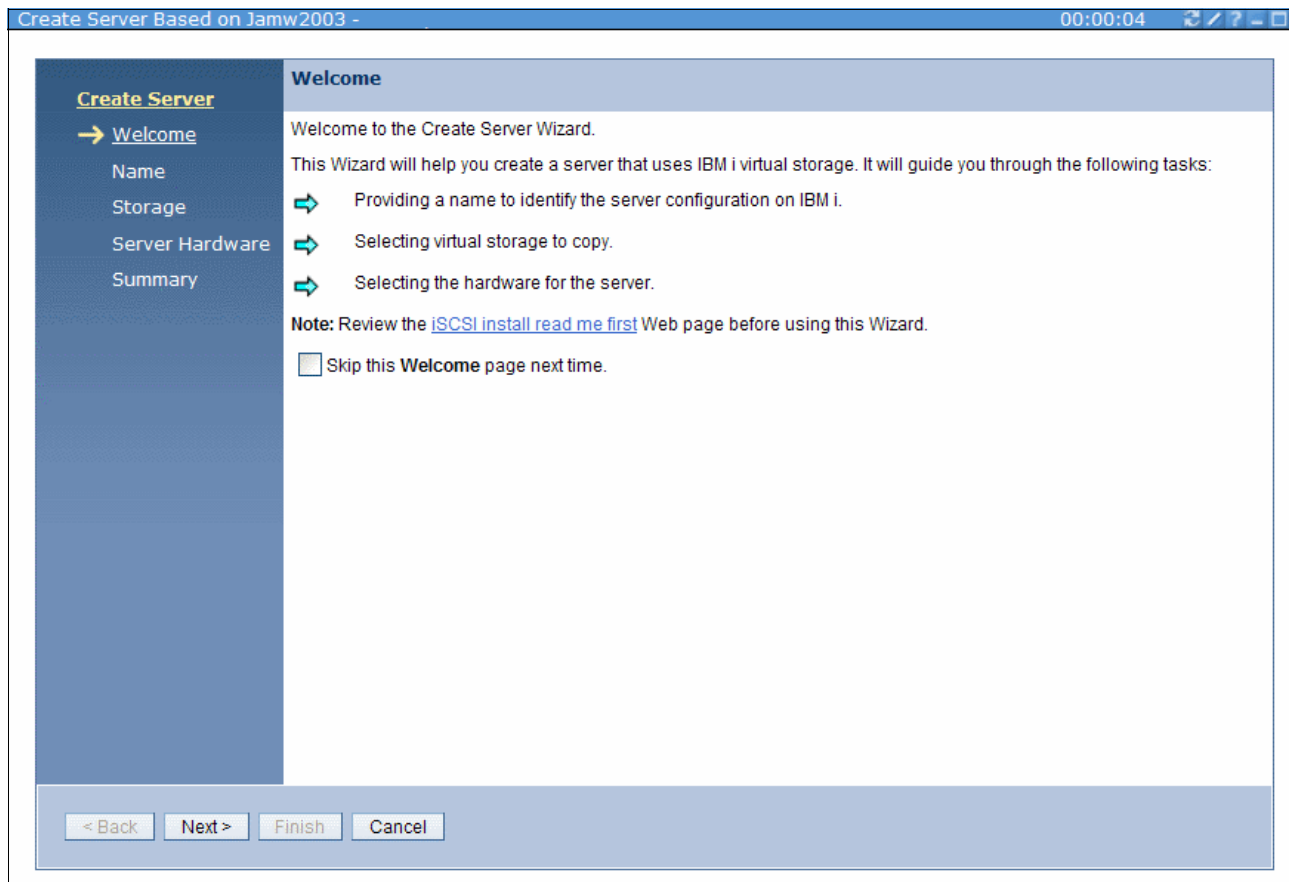


Figure 12-11 Create Server Based On (Cloning) Wizard

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 2008 and Windows Server 2003 Server editions. The cloning process requires that you prepare the base server for cloning before using the cloning task. Additional configuration is required after the server is cloned.

Note: Review the *Server cloning road map and checklist* chapter of the *IBM i iSCSI Solution Guide PDF* before using this Wizard:

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.

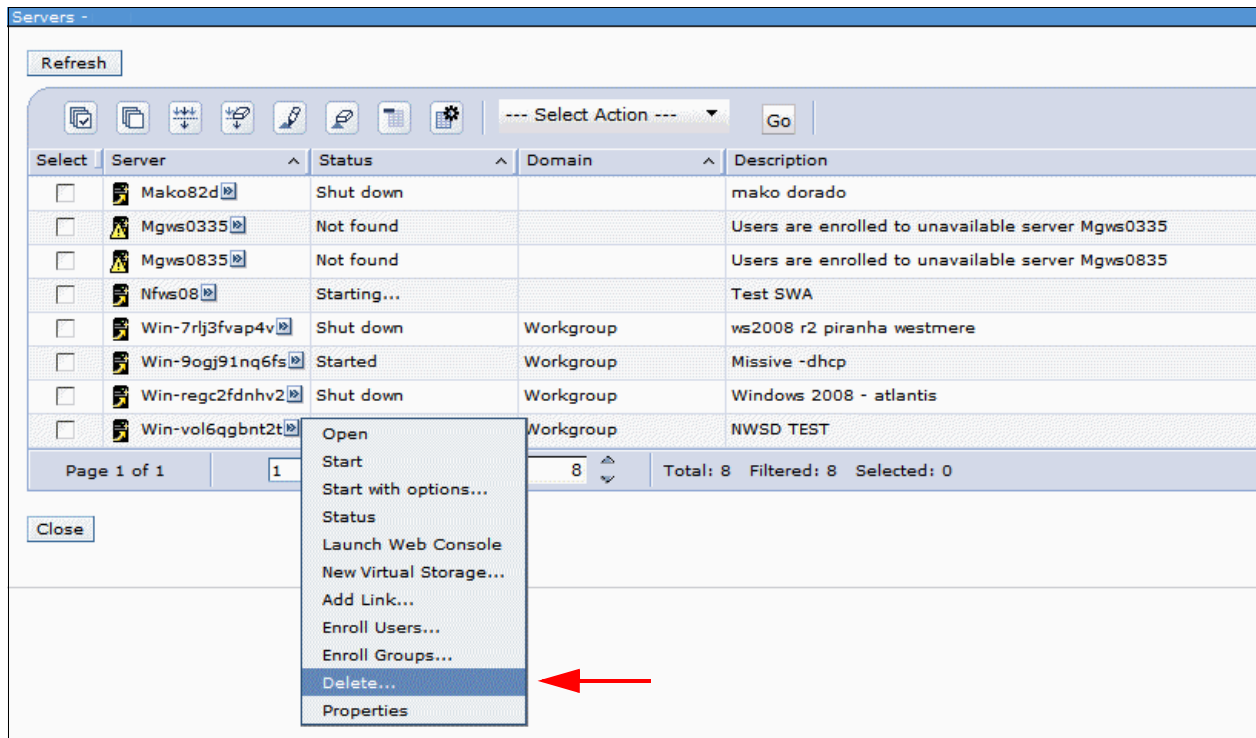


Figure 12-12 Delete Server option in the IBM Systems Director Navigator for i web GUI

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. In case of an IBM BladeCenter, it can launch the Advanced Management Module web interface, as shown in Figure 12-13.

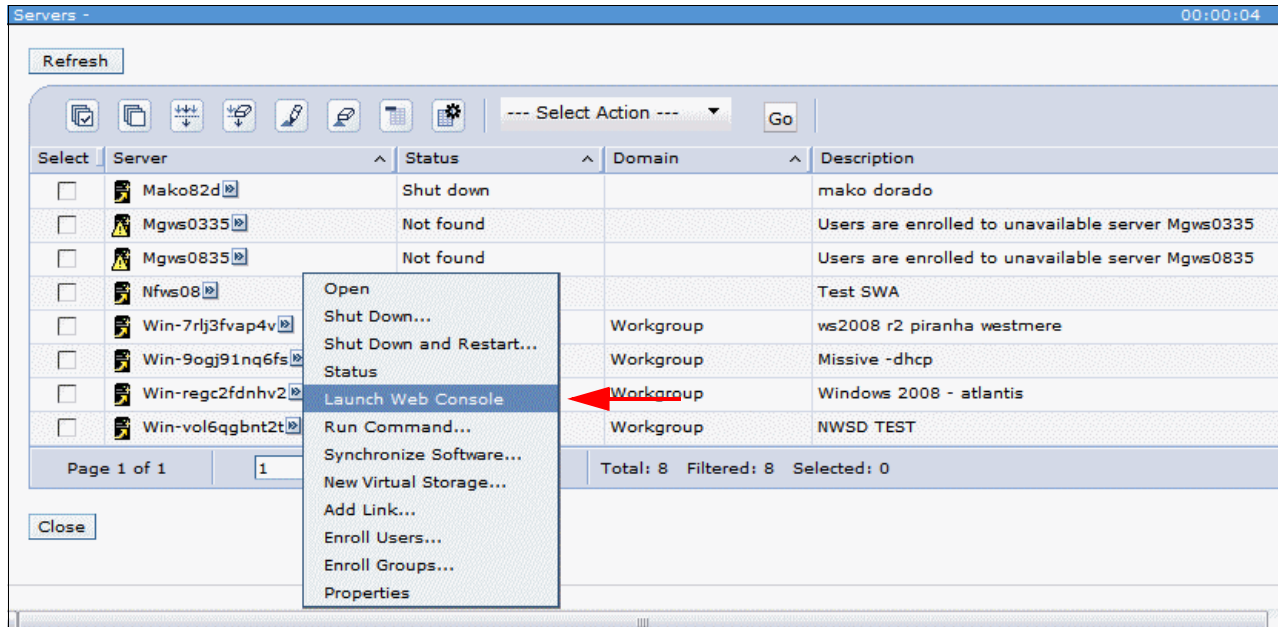


Figure 12-13 Launching the web Console from the IBM Systems Director web GUI

12.7 New IBM i CL commands

The following new IBM i control language (CL) commands are available for integrated servers:

- ▶ INSINTSVR (Install Integrated Server)
- ▶ DLTINTSVR (Delete Integrated Server)

12.7.1 INSINTSVR (Install Integrated Server)

The INSINTSVR command installs an iSCSI-attached integrated Windows Server 2008 or VMware ESX server as shown in Figure 12-14.

```

                                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 12-14 INSINTSVR command

For more information, see Instal Integrated Server at the following web page:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/cl/insintsvr.htm>

12.7.2 DLTINTSVR (Delete Integrated Server) command

The DLTINTSVR command, shown in Figure 12-15, 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 12-15 DLTINTSVR command

12.8 IBM i changed CL commands

The following new IBM i control language (CL) commands are changed for integrated servers:

- ▶ INSWNTSVR (Install Windows Server) CL command
- ▶ CRTNWSCFG (Create NWS Configuration) and CHGNWSCFG (Change NWS Configuration) CL commands

12.8.1 INSWNTSVR (Install Windows Server) CL command

The INSWNTSVR (Install Windows Server) command has a number of parameter changes that might now require you to recompile any existing CL programs that use this command:

- ▶ The WNTVER (Windows server version) 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 LICMODE (License mode) parameter.
- ▶ The WNTVER (Windows server version) parameter no longer supports the installation of Windows Server 2008. To install Windows Server 2008, use the new INSINTSVR (Install Integrated Server) command.
- ▶ The INSTYPE (Installation type) parameter is obsolete and has been removed. ServerGuide assisted installations are no longer supported.
- ▶ The OPTION (Install option) parameter is obsolete and has been removed. There are no supported upgrades that can be processed using the INSWNTSVR (Install Windows Server) command.
- ▶ The ENBUNICAST (Enable unicast) parameter is obsolete and has been removed. Communication with the service processor is configured by specifying a value for SPNAME (Service processor name) or SPINTNETA (SP internet address) parameter.
- ▶ The BOOTDEVID (Boot device ID) parameter is obsolete and has been removed.
- ▶ The RMTIFC Remote (initiator) interfaces) parameter, SCSI and LAN Gateway addresses are not supported. The iSCSI attached servers do not support bridged networks requiring the gateway address. These elements have been removed.

12.8.2 CRTNWSCFG (Create NWS Configuration) and CHGNWSCFG (Change NWS Configuration) CL commands

The ENBUNICAST (Enable Unicast) parameter has been removed from the CRTNWSCFG (Create NWS Configuration) and CHGNWSCFG (Change NWS Configuration) 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 ENBUNICAST (Enable Unicast) configured to *NO must use the CHGNWSCFG (Change NWS Configuration) command to specify either the SPNAME (Service Processor Name) or SPINTNETA (Service Processor IP Address) parameter.

iSCSI-attached network server descriptions cannot vary on until the network server configurations of type *SRVPRC with *Enable unicast* configured to *NO have been changed.

12.8.3 INSLNXSVR (Install Linux Server) CL command

The INSLNXSVR (Install Linux Server) CL command is no longer supported in IBM i 7.1. The INSLNXSVR command was used for Linux and VMware ESX server installs on prior IBM i releases.

Note: 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 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 have been implemented to the 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:

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

This guide provides the information 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 matrices: 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 road map and checklist: Required information to install a server that is integrated with IBM i.
- ▶ Server cloning road map and checklist: Required information to clone a Windows server that is integrated with IBM i.

- ▶ BladeCenter and System x configuration: See iSCSI configuration tasks for BladeCenter blade and System x servers.
- ▶ Additional topics: Other topics related to the iSCSI solution.

Note that this guide consolidates and replaces the following information:

- ▶ iSCSI installation road map PDFs for i 7.1, i 6.1, and i 5.4
- ▶ iSCSI Initiator Hardware Configuration PDF
- ▶ iSCSI install read me first Web page
- ▶ BladeCenter and System x models supported with iSCSI Web page
- ▶ iSCSI target and iSCSI initiator tables that were formerly on the iSCSI solution (iSCSI) Web page
- ▶ Ethernet switches for iSCSI Web page
- ▶ Ordering channels for iSCSI Web page
- ▶ Various Web pages related to tape and optical device support
- ▶ Various Web pages 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 has been updated:

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp?topic=/rzahq/rzahqntspo.htm>

Here are the most significant updates to this topic:

- ▶ Information related to various i 7.1 enhancements has been added
- ▶ The following documentation has been migrated to the *IBM i iSCSI Solution Guide* PDF:
 - The Integrated server installation road map 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 work sheets have been migrated to the *IBM i iSCSI Solution Work Sheets* PDF
- ▶ Integrated Windows servers that run on an Integrated xSeries® Server (IXS) or a server attached using an Integrated xSeries Adapter (IXA) are no longer documented
- ▶ Since the new *Create Server* task is now available within *IBM Systems Director 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 developerWorks

This *developerWorks* group is the place where developers and IT professionals who work with the IBM i iSCSI solution congregate to connect, share, and collaborate:

<http://www.ibm.com/developerworks/groups/IBMiIntegratedServer>

This group provides:

- ▶ A message board.
- ▶ Bookmarks for important Web sites and reference material.

- ▶ A set of Wiki pages, including:
 - Documentation
 - Learning Resources
 - Service & 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 BladeCenter and System x* page on the *IBM i Technology Updates* Wiki lists the latest integrated server enhancements including the required group PTF levels:

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

12.10.5 IBM i integration with BladeCenter and System x Marketing Web site

This Marketing Web site has been streamlined:

<http://www.ibm.com/systems/i/advantages/integratedserver>

Most of the technical content that was previously on this Web site (for example, the iSCSI instal read me first Web page) has been migrated to the *IBM i iSCSI Solution Guide* or to *developerWorks*. See the previous sections for more information

12.11 Additional information

IBM i integration with BladeCenter and System x Marketing Web site:

<http://www.ibm.com/systems/i/advantages/integratedserver>

IBM i iSCSI Solution Guide and associated planning work sheets:

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

IBM i integration with BladeCenter and System x topic in the Information Center:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/rzahq/rzahqntspo.htm>

IBM i integration with BladeCenter and System x group on *developerWorks*:

<http://www.ibm.com/developerworks/groups/IBMiIntegratedServer>

Integration with BladeCenter and System x page on the *IBM i Technology Updates* Wiki:

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



IBM Advanced Job Scheduler for i enhancements

This chapter describes new functions available in IBM i 7.1 related to job scheduling, as available using the 5770-JS1 IBM Advanced Job Scheduler for i licensed program product.

The chapter covers the following topics:

- ▶ 13.1, “Advanced Job Scheduler capabilities” on page 386
- ▶ 13.2, “IBM Systems Director Navigator for i AJS support” on page 387
- ▶ 13.3, “Other AJS enhancements” on page 425
- ▶ 13.4, “References” on page 430

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 through an e-mail 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 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 grouped together to run consecutively in the order specified in the group sequence field. This 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

The notification feature in the AJS allows you 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 will pass before the next person in the list is notified. Escalation continues until someone stops escalation. Notification, coupled with report distribution, allows you 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 RUNRMTCMD (Run Remote) command is used to issue the commands on the remote systems. RUNRMTCMD allows server users 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 REXECD (Remote Executing Daemon). 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 has the ability to manage multiple scheduling environments. A scheduling environment is all the objects in the QUSRIJS data library duplicated into another library. QUSRIJS is the main data library 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 defining 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, nor could it 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 function as 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 discusses the new functions during the walk-through.

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 logging in to the IBM Systems Director Navigator for i web page, the Welcome page is presented as shown by 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.

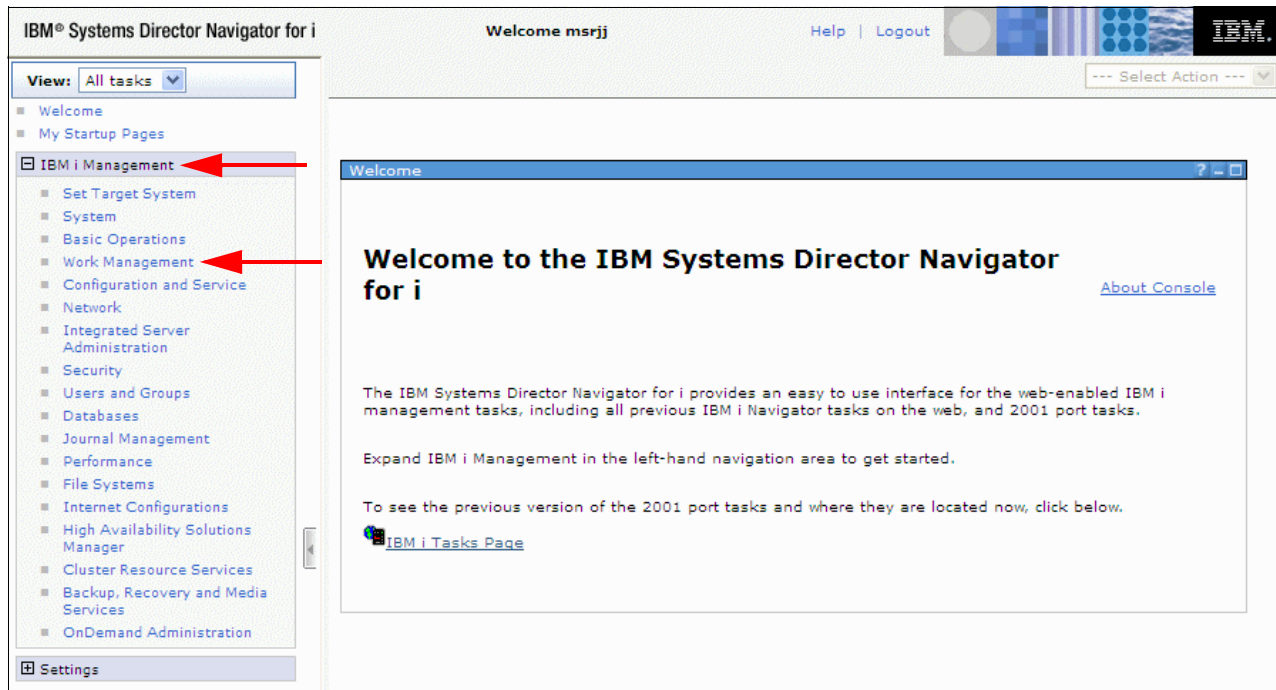


Figure 13-1 Welcome page with IBM i Management item expanded

When the Work Management function is selected, the Work Management main menu in Figure 13-2 is displayed.

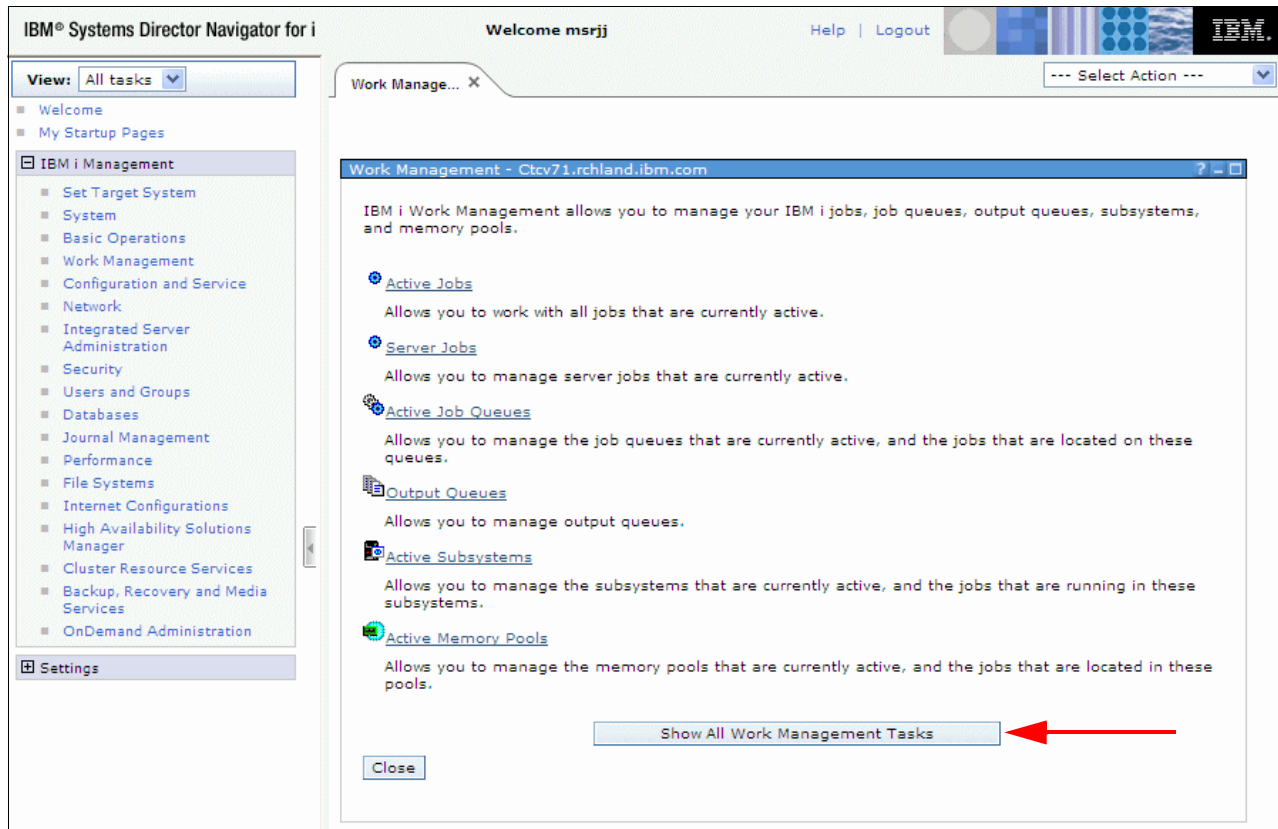


Figure 13-2 Work Management main menu

The Work Management menu includes many of the operations level actions of IBM i work management. None of the job scheduler actions appear on this page. To access the job scheduler functions, the **Show All Work Management Tasks** button to which the arrow points, must be selected.

To make subsequent pages larger, the icon was selected to hide the navigation area, then the **Show All Work Management Tasks** button was selected.

The Work Management Advanced menu is shown in Figure 13-3.

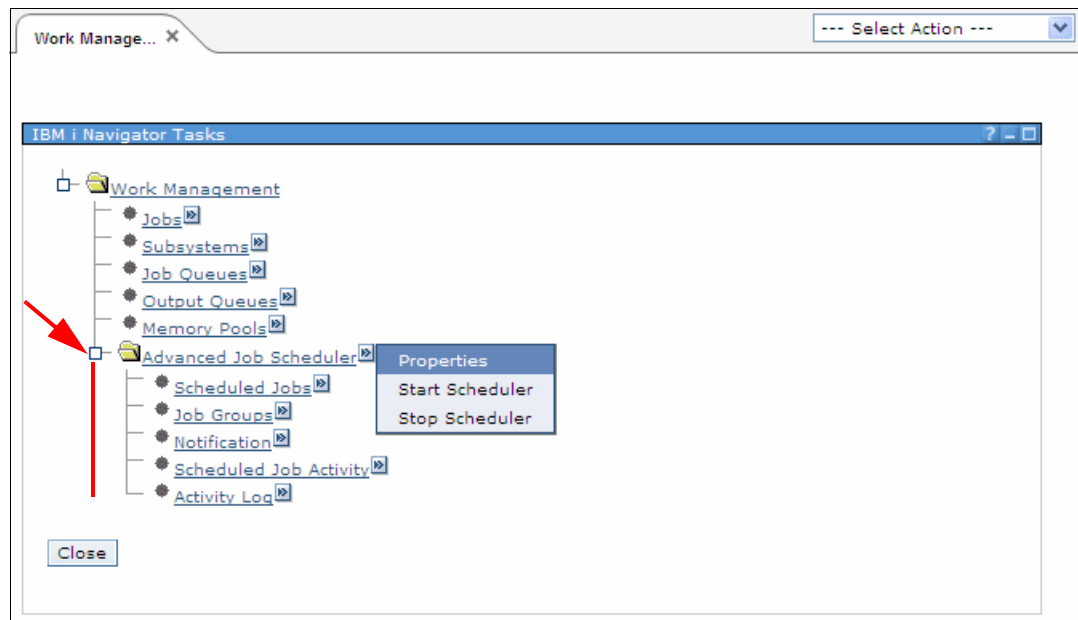


Figure 13-3 Work Management all tasks menu

The first time the Work Management advanced menu is displayed, 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 pop-up menu of actions. In Figure 13-3, the pop-up menu box for the AJS menu item is shown.

Following each of the menu items are icons indicating that pop-up menus are available for that item. The AJS item's pop-up menu has vital actions that are discussed in 13.2.2, "AJS pop-up menu" on page 394.

IBM Systems Director navigation

After logging in to the IBM Systems Director, you must navigate to an IBM i system. One way to do this is shown in the following procedure:

1. Select the **Navigate Resources** task in the Navigation panel. (See the red arrow on the left)
2. Select the **All Systems** systems group from the Navigate Resources panel as shown in Figure 13-4. (See the red arrow in the middle.)

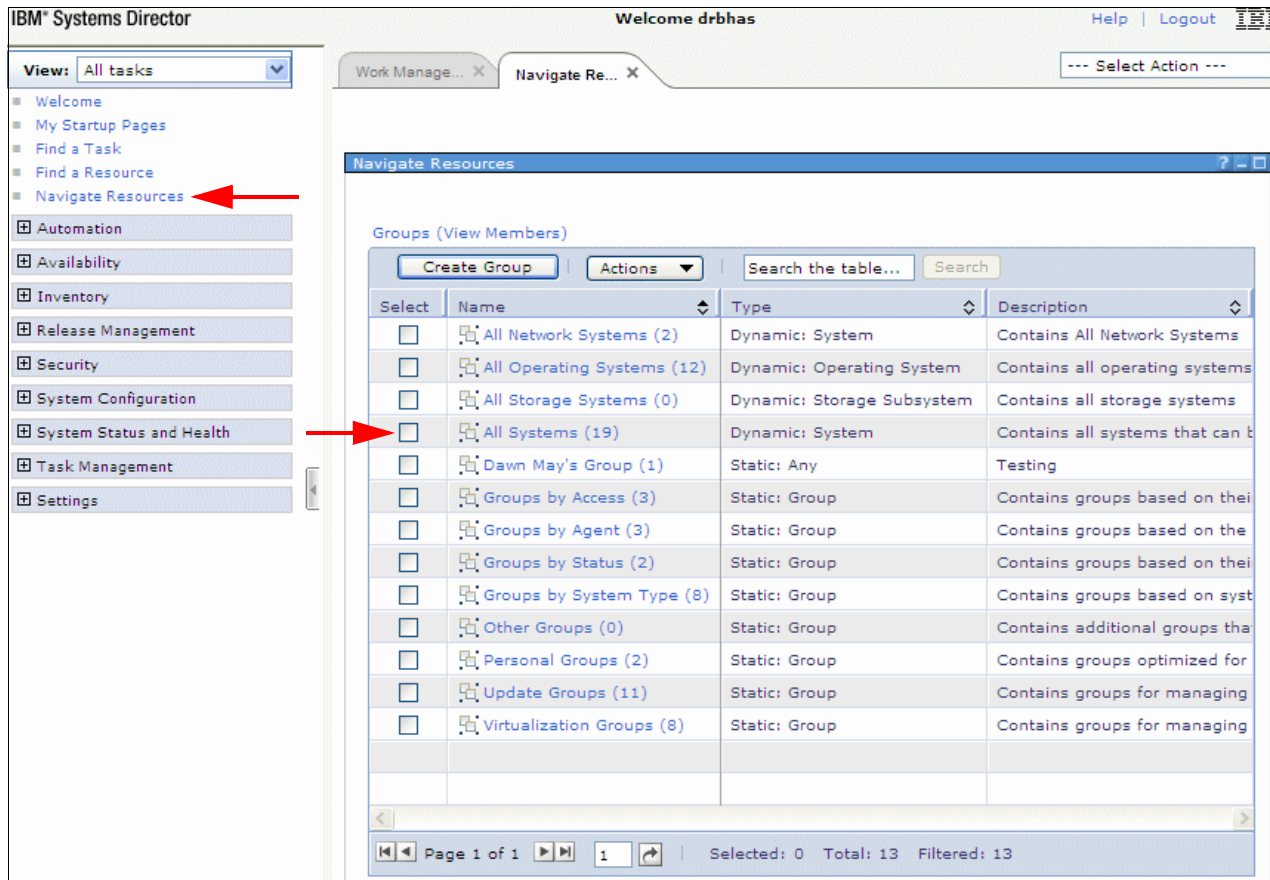


Figure 13-4 Selecting Navigate Resources, then selecting the All Systems resource group

- The **All Systems** group was selected on the previous page and the All Systems group systems were displayed. The window was then scrolled to the second page in search of a specific IBM i system.

[illegible]

4. Click the **Actions** button after the system is selected. A drop-down menu is displayed.

5. Select the **IBM i Management** item from the Actions drop-down menu. A second drop-down menu is displayed. On this second menu, select **Work Management**, as shown in Figure 13-6.

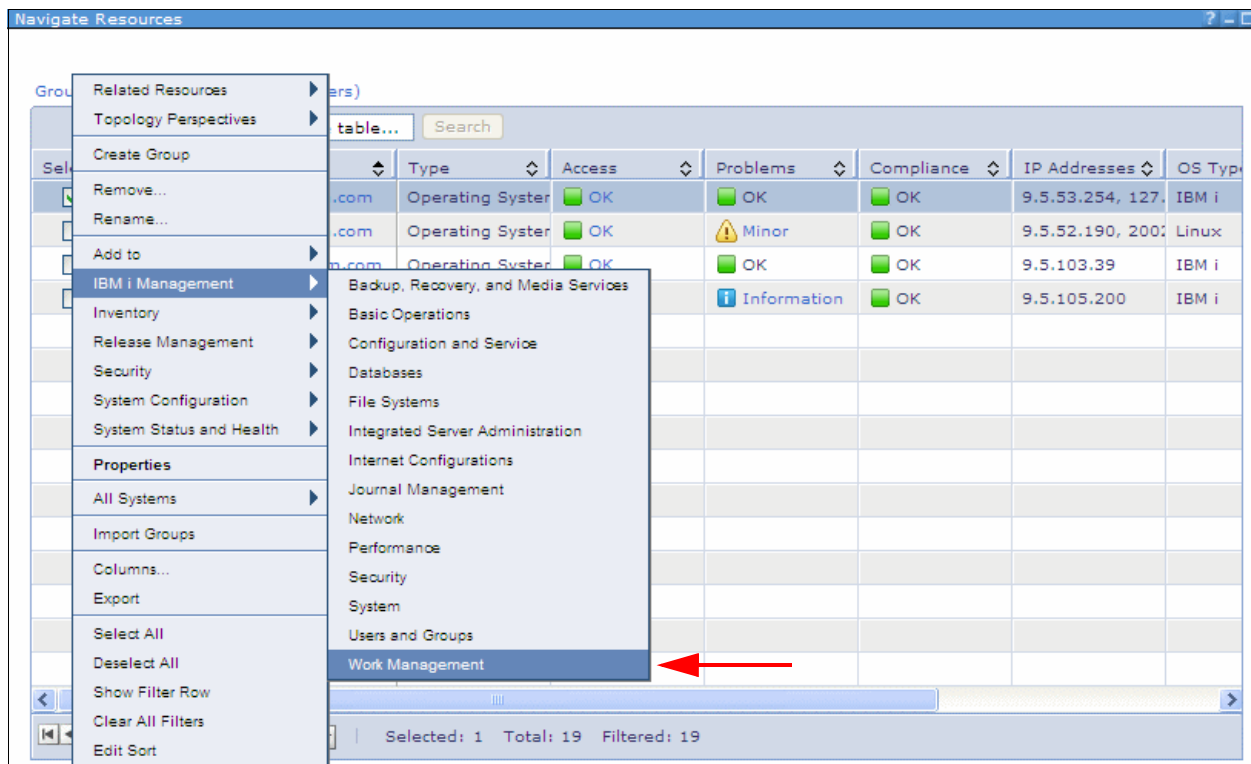


Figure 13-6 Selecting Work Management from the Actions drop down menu

When you select Work Management, a menu for basic Work Management tasks is displayed, as shown in Figure 13-7.

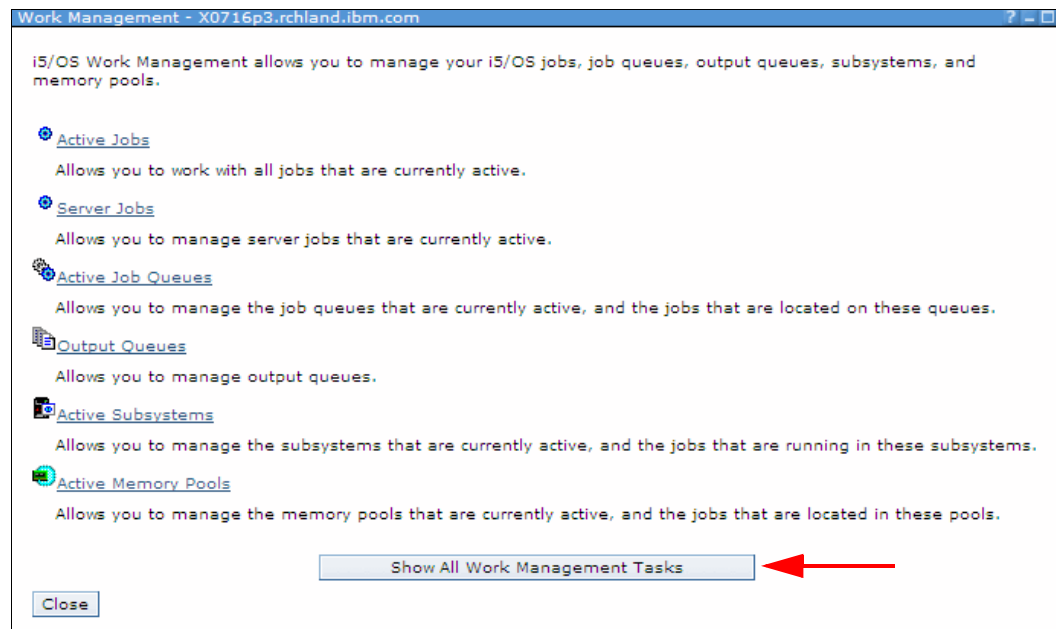


Figure 13-7 IBM i Work Management basic menu accessed from Systems Director

The Work Management menu that is displayed 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 the **Show All Work Management Tasks** button. An advanced Work Management page is displayed.

When first displayed, the AJS task menu is minimized. Upon expanding the menu by selecting the icon to which the arrow points, the panel in Figure 13-8 is displayed.

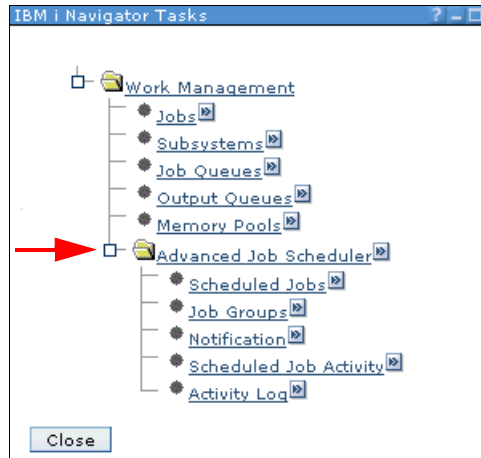


Figure 13-8 Systems Director - IBM i Work Management all tasks menu

13.2.2 AJS pop-up menu

In the previous section, the Advanced Jobs Scheduler pop-up menu was shown in Figure 13-3 on page 390.

The AJS pop-up menu contains items of system-wide scope, rather than of job scope. This is the menu where the job scheduler functions themselves are configured and maintained. In the following sections, the Advanced Jobs Scheduler pop-up menu actions are reviewed.

13.2.3 Start Scheduler action

This action 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 390).
2. Select the Start Scheduler option. The confirmation panel in Figure 13-9 is shown.




Figure 13-9 Start AJS monitor confirmation panel

Clicking **Yes** starts the AJS monitor.

Clicking **No** returns you to the panel shown in Figure 13-3 on page 390.

13.2.4 Stop Scheduler action

This action 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 390).
2. Select the Stop Scheduler option. The confirmation panel in Figure 13-10 is shown.

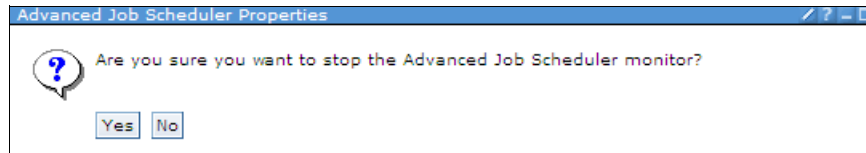


Figure 13-10 Stop AJS monitor confirmation panel

Clicking Yes stops the AJS monitor.

Clicking No returns you to the previous panel.

13.2.5 AJS properties menu

The AJS properties are divided into six tabs on the left side of Figure 13-11. The following sections discuss each tab.

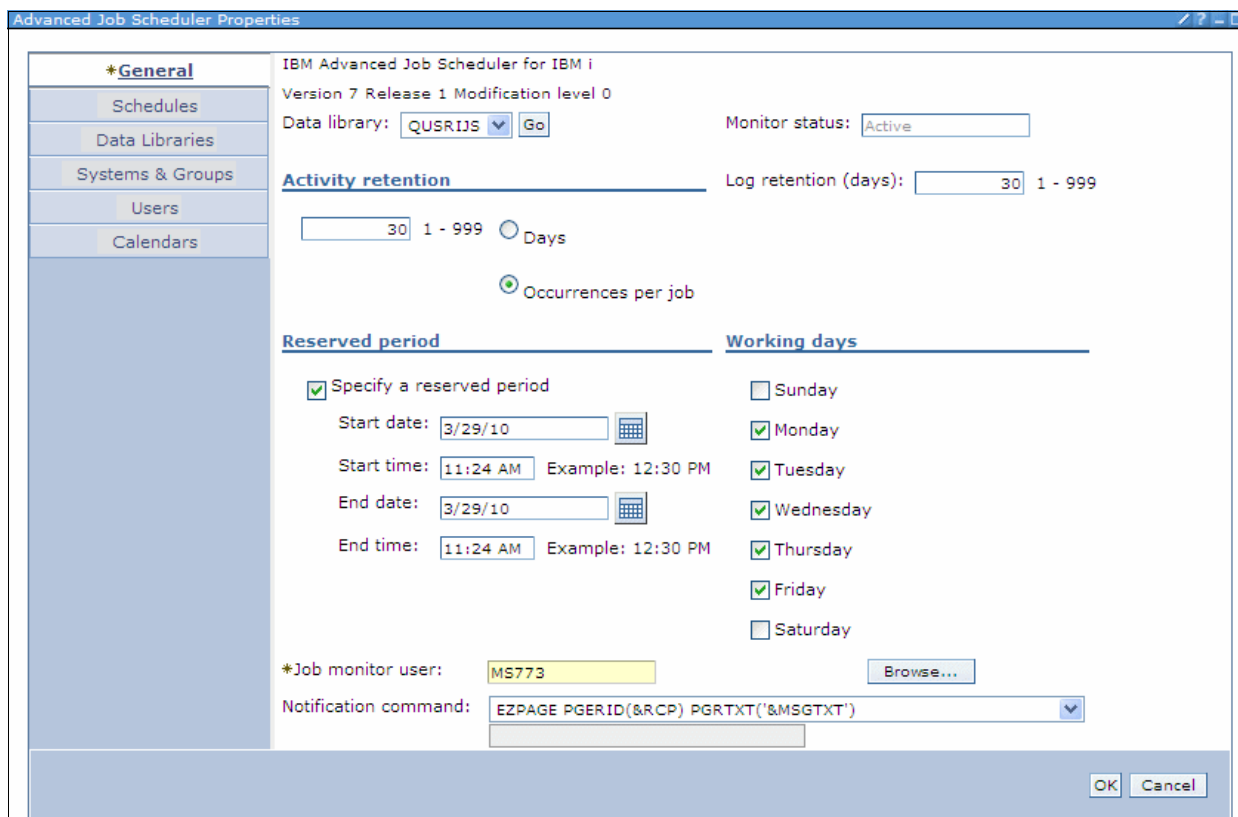


Figure 13-11 AJS properties - General tab

General tab

The General tab shown in Figure 13-11 on page 395 is used to view and set the general properties of the job scheduler. The user 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 will send a notification whenever a job completes successfully or fails

Schedules tab

The Schedules tab shown in Figure 13-12 provides you with a display of the existing schedules on your system. The buttons at the right of the page provide the following functions

- ▶ **New:** Enables the user to create a new schedule.
- ▶ **New Based On:** Enables creation of a new schedule based on an existing one.
- ▶ **Remove:** Removes an existing schedule.
- ▶ **Properties:** Enables one to modify an existing schedule.

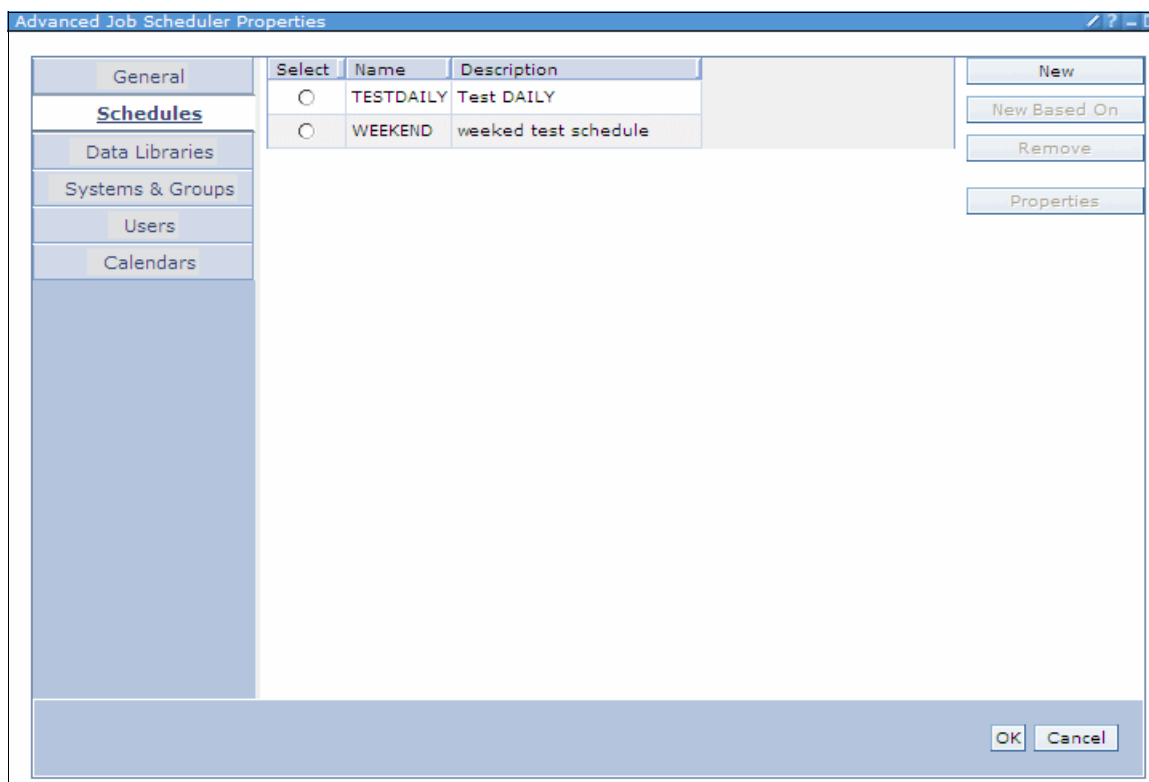


Figure 13-12 AJS Properties Schedules tab

Selecting NEW button in Figure 13-12 on page 396 displays the New Schedules panel, as 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

--

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 page

A schedule is a specification of days on which a job can run. Notice that there are no times listed. Scheduled jobs and group jobs can select a schedule rather than make their own when they are created.

Data Libraries tab

The Data Libraries tab displays all job scheduler data libraries. As shown in Figure 13-14, the user can add, remove, and modify a data library, and can start and end the job scheduler monitor job for a specific data library.

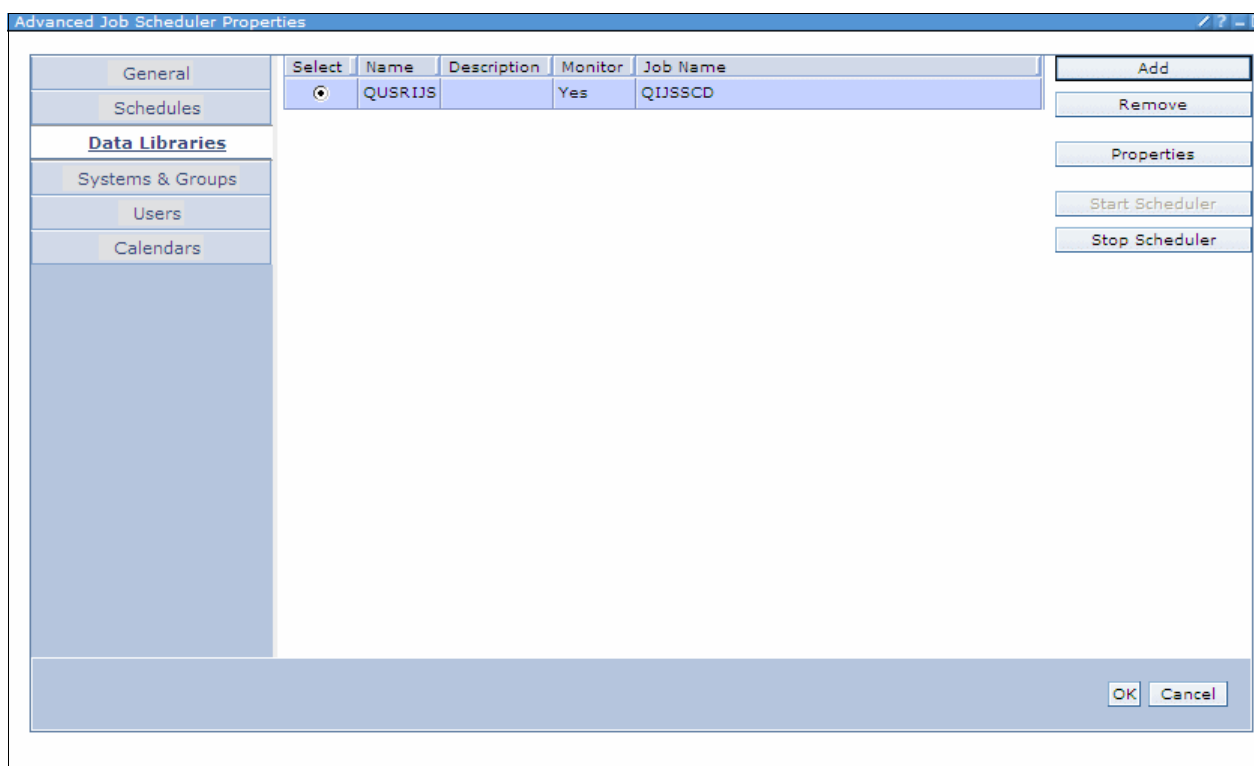


Figure 13-14 AJS Properties - Data Libraries page

The **Add** button causes the New Data Library page to be displayed, as shown in Figure 13-15.

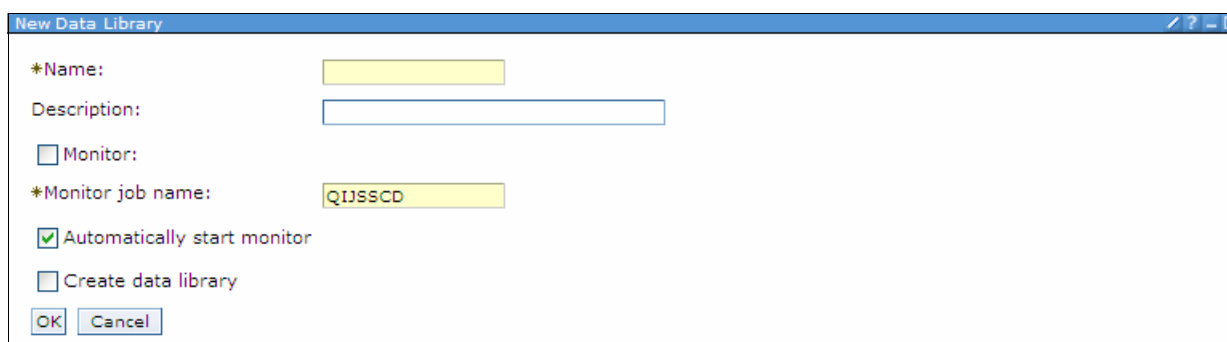


Figure 13-15 AJS New Data Library page

From this panel, the user can accomplish the following tasks:

- ▶ Create a new jobs scheduler data library
- ▶ Specify a job scheduler monitor
- ▶ Specify 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 running simultaneously. The user can switch from one job scheduler running a production environment to the other job scheduler library running a test environment. This capability is covered in more detail in 13.3.2, “Multiple Scheduling Environments function” on page 426.

Systems & Groups tab

The Systems & Groups tab enables a user to add, remove, and change existing IP addresses and other information for IBM i and non-IBM i systems used by the AJS.

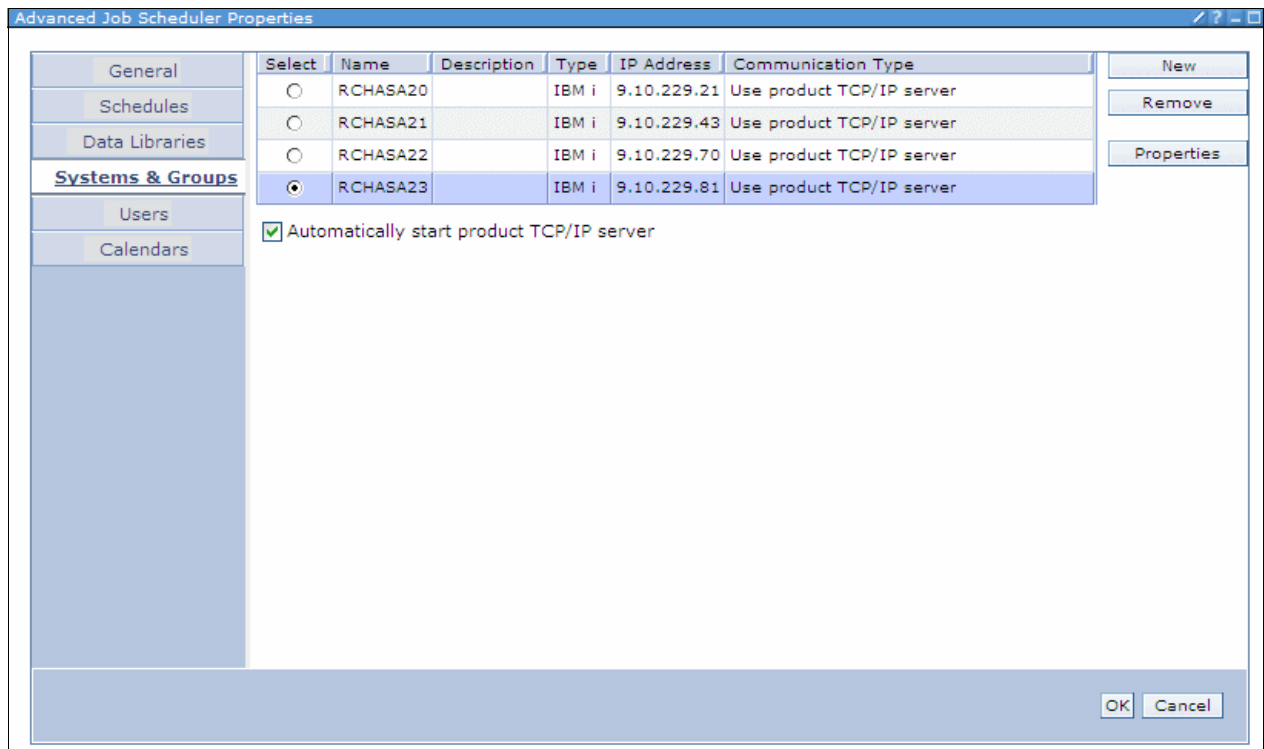


Figure 13-16 Job Scheduler properties - Systems and Groups tab

Users tab

The Users tab (Figure 13-17) enables maintenance of a list of job scheduler users 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.

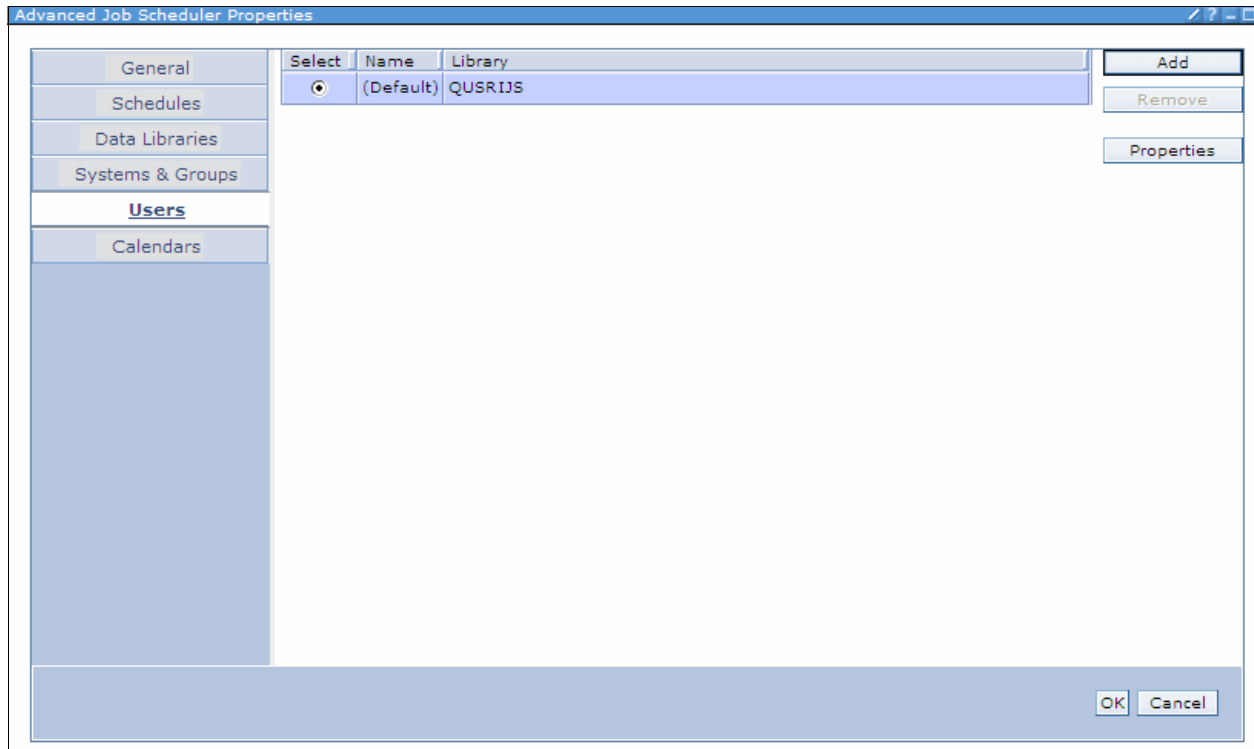


Figure 13-17 AJS Properties - Users tab

Suppose 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 only access one job scheduler data library, the administrator effectively secures the users to their own application's job scheduler when locking access to the others.

Calendars tab

The buttons at the right of the page (Figure 13-18) provide the following functions:

- ▶ **New:** Create new calendar
- ▶ **New Based On:** Create a new calendar based on an existing one
- ▶ **Remove:** Remove calendars
- ▶ **Properties:** Maintain calendars

Scheduling calendars and holiday calendars have different pages because they have significantly different parameters.

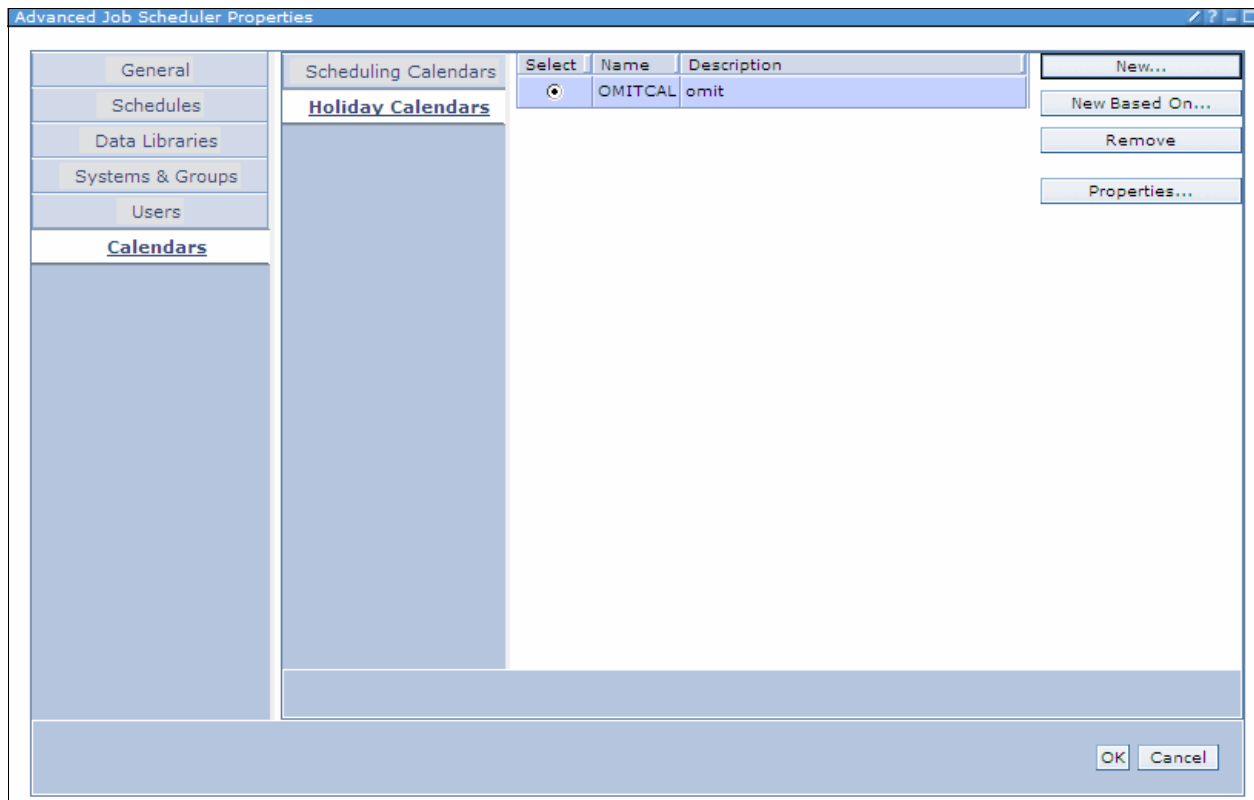


Figure 13-18 AJS properties - Calendars tab

The Holiday Calendar Properties panel is shown in Figure 13-19.

Holiday Calendar Properties

*Name: OMITCAL

Description:

Reference calendar: Use entry from below

Date: 3/30/10

☐ Every year ☒ Special: Do not run

☐ Specific:

Details:

Select	Date	Alternate day to run
<input type="radio"/>	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 page

Next, we examine job-specific functions and actions.

13.2.6 Scheduled Jobs menu

In this section, the various actions of the Scheduled Jobs pop-up menu are reviewed. This menu is accessed through the Work Management Advanced menu (Figure 13-3 on page 390), by expanding the Advanced Job Scheduler option.

The Scheduled Jobs Menu is shown in Figure 13-20 on page 403. The menu offers the following options:

- Scheduled Jobs

This action displays a table of scheduled jobs and enables the user to perform a variety of actions to them. The user can customize which columns appear in the table, sort on any column, and filter the table. Using the scheduled job table is discussed in 13.2.8, “Using the Scheduled Jobs Table” on page 406.

- New Scheduled Job

This action enables the user to create a new scheduled job.

► Properties

This action enables a user to change scheduled jobs properties for all scheduled jobs. The properties include:

- Application for Scheduled Job
- Base Periodic frequency on start time
- Reset held jobs
- Reset after change
- Start time of day

► Reset Scheduled Jobs

This action enables the user to reset all scheduled jobs.

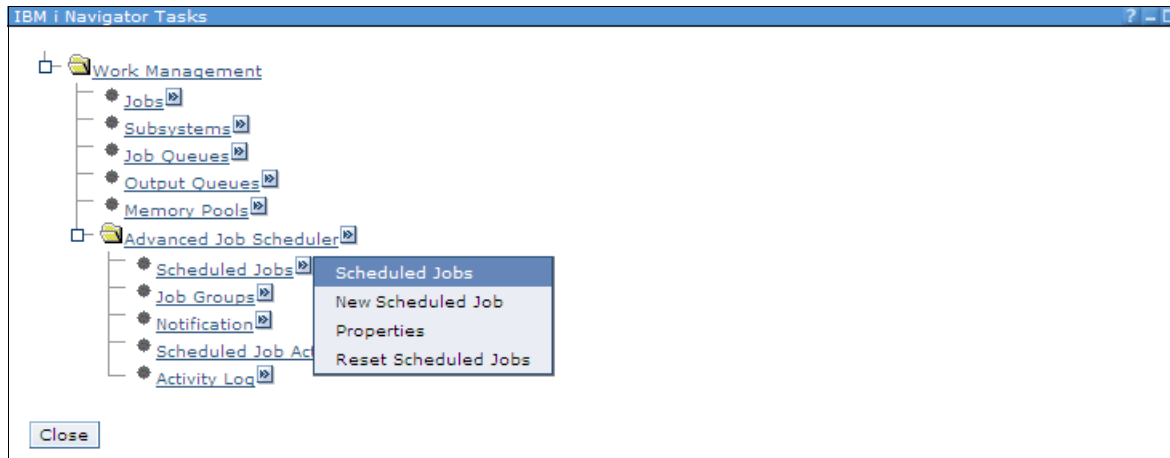
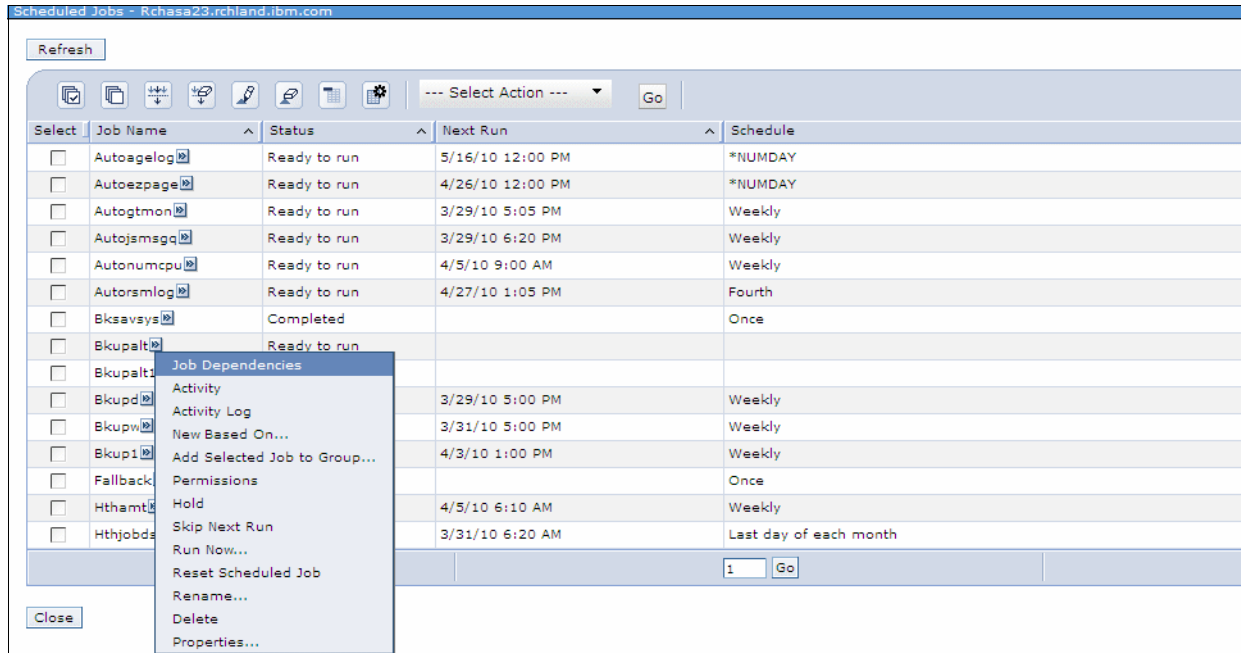


Figure 13-20 Scheduled Jobs pop-up menu

13.2.7 Using the Scheduled Jobs action

In Figure 13-20, the Scheduled Jobs pop-up menu is highlighted. This section provides an overview of what this action can do.

When the user selects the Scheduled Jobs action on the Scheduled Job pop-up menu, a table of scheduled jobs on the system is displayed, as shown in Figure 13-21.



The screenshot shows a web-based interface 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 "BKUPD" 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
<input type="checkbox"/>	Autoagelog	Ready to run	5/16/10 12:00 PM	*NUMDAY
<input type="checkbox"/>	Autoezpage	Ready to run	4/26/10 12:00 PM	*NUMDAY
<input type="checkbox"/>	Autogtmon	Ready to run	3/29/10 5:05 PM	Weekly
<input type="checkbox"/>	Autojmsgq	Ready to run	3/29/10 6:20 PM	Weekly
<input type="checkbox"/>	Autonumcpu	Ready to run	4/5/10 9:00 AM	Weekly
<input type="checkbox"/>	Autorsmlog	Ready to run	4/27/10 1:05 PM	Fourth
<input type="checkbox"/>	Bksavsys	Completed		Once
<input type="checkbox"/>	Bkupalt	Ready to run		
<input type="checkbox"/>	Bkupalt			
<input type="checkbox"/>	Bkupd		3/29/10 5:00 PM	Weekly
<input type="checkbox"/>	Bkupw		3/31/10 5:00 PM	Weekly
<input type="checkbox"/>	Bkup1		4/3/10 1:00 PM	Weekly
<input type="checkbox"/>	Fallback			Once
<input type="checkbox"/>	Hthamt		4/5/10 6:10 AM	Weekly
<input type="checkbox"/>	Hthjobds		3/31/10 6:20 AM	Last day of each month

Figure 13-21 Scheduled Jobs list table

One can select multiple job using the check boxes. There are options that enable selection of all jobs. For now, we review what the user can do to a specific job

To illustrate the many tasks a user can accomplish in a specific job, the pop-up menu for the BKUPD job is displayed in Figure 13-21 on page 404. Although it is beyond the scope of this book to discuss each of these actions in detail, a brief description of each action follows:

- Job Dependencies

This action enables display and update of job dependencies including predecessors and successors, and whether all dependences must be met or just one.

- 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 based on the selected job. The user can override any of the based on job's parameters.

- Add Selected Job to Group

This action allows the user to add the selected job to a job group.

- ▶ Permissions

This action allows the user 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 which are not already held.

- ▶ Release

This action allows a held scheduled job to be released. This action is shown only for jobs which are held.

- ▶ Skip Next Run

This action allows a user to skip (omit) the next scheduled run for the job.

- ▶ Run Now

This action allows a user to run a scheduled job now. It can also run a job at a specified date and time and at a certain number of minutes from now. It 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

This action enables a user to view and change virtually every parameter and value for a scheduled job. The parameters are divided into nine pages of parameters:

- General

Specifies type of scheduled job, job name, job (text) description, the commands and sequence to be run.

- Schedule

Specifies when the job is to be run.

- Batch Information

Specifies information required to submit the job such as run on 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, command logging, etc.

- Notification

Specifies options and parameters for notification of job status

- Documentation

Provides a text field for job documentation

- Problem Recovery

Specifies maximum run time, alternate job in case of failure, action to take when 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 - 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
Displays the last run details.

The Scheduled Jobs properties and display panels are covered more completely in 13.2.9, “Adding a new scheduled job” on page 412.

13.2.8 Using the Scheduled Jobs Table

This section reviews the Scheduled Jobs table (including the following tasks: sorting, changing columns and using filtering).

The Select Action drop down menu at the top of the table enables options not previously discussed. With no scheduled jobs selected, the menu has the following options, as shown in Figure 13-22 on page 406:

- ▶ New: The user can select this to create a new scheduled job.
- ▶ Reset Scheduled Jobs: The user can reset all scheduled jobs.
- ▶ Show find toolbar: Displays a small toolbar used for searching the table.
- ▶ Table Actions: Displays another menu with more actions for manipulating the table.

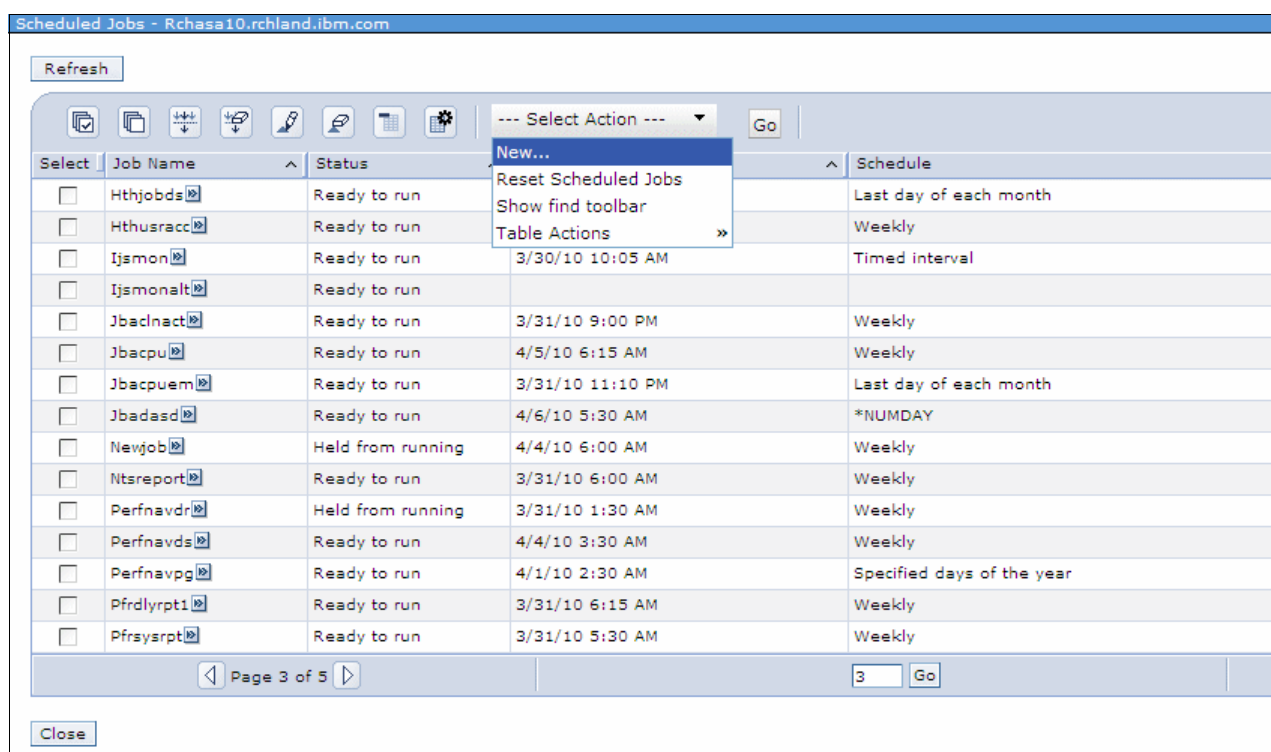


Figure 13-22 Scheduled Jobs Table Action drop down menu

If the user selects one or more jobs in the table, the Select Action menu adds job-specific actions to the list options. These are equivalent to the actions listed in the scheduled job’s pop-up menu shown in Figure 13-21 on page 404.

If the Table Actions menu item is selected, a second drop-down menu is displayed, as shown in Figure 13-23 on page 407. The actions of this menu are:

- ▶ **Select All**
Selects all scheduled jobs.
- ▶ **Deselect All**
Deselects all scheduled jobs.
- ▶ **Show Filter Row**
Shows a row at the top of the table 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 can be helpful in speeding changes to the table as the data does not have to be refreshed on the panel between multiple changes.
- ▶ **Configure Columns**
This action is used to change the columns in the table.
- ▶ **Restore Defaults**
Removes the changes made to the table by sorting, filtering and column changes.

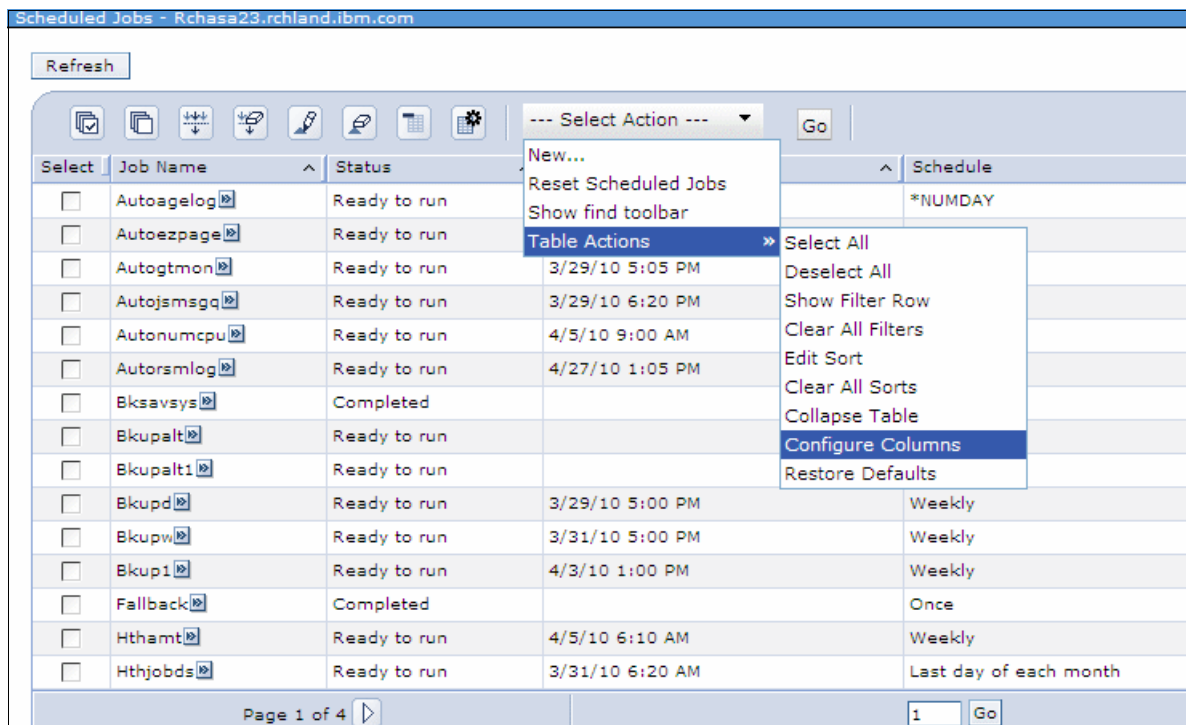


Figure 13-23 Scheduled Jobs - Table Actions

Resizing Columns

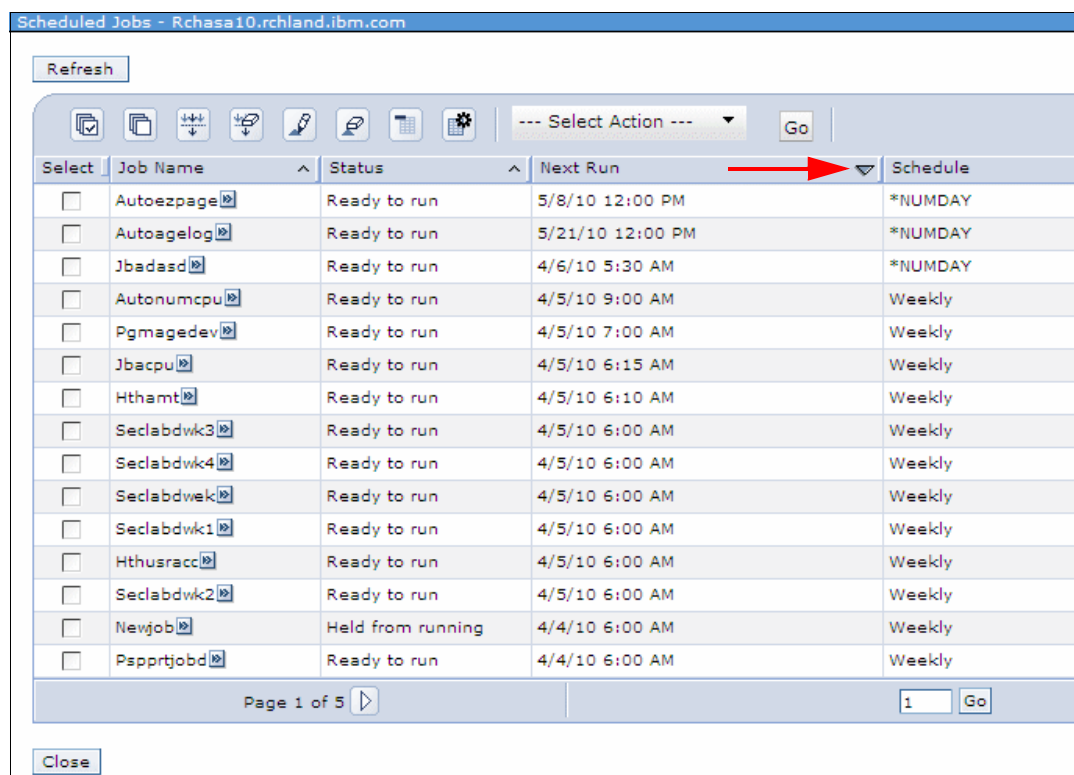
Columns can be resized by dragging the column separator in the table column description (header) line.

Sorting the table

Any of the columns which 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 sorting, the ^ (circumflex) symbol is replaced by a triangle symbol indicating the sort direction.

In Figure 13-24 on page 408, 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, use the **Clear All Sorts** action from the Table Actions menu.



Select	Job Name ^	Status ^	Next Run ▼	Schedule
<input type="checkbox"/>	Autoezpage	Ready to run	5/8/10 12:00 PM	*NUMDAY
<input type="checkbox"/>	Autoagelog	Ready to run	5/21/10 12:00 PM	*NUMDAY
<input type="checkbox"/>	Jbadasd	Ready to run	4/6/10 5:30 AM	*NUMDAY
<input type="checkbox"/>	Autonumcpu	Ready to run	4/5/10 9:00 AM	Weekly
<input type="checkbox"/>	Pgmagedev	Ready to run	4/5/10 7:00 AM	Weekly
<input type="checkbox"/>	Jbacpu	Ready to run	4/5/10 6:15 AM	Weekly
<input type="checkbox"/>	Hthamt	Ready to run	4/5/10 6:10 AM	Weekly
<input type="checkbox"/>	Seclabdwk3	Ready to run	4/5/10 6:00 AM	Weekly
<input type="checkbox"/>	Seclabdwk4	Ready to run	4/5/10 6:00 AM	Weekly
<input type="checkbox"/>	Seclabdwk	Ready to run	4/5/10 6:00 AM	Weekly
<input type="checkbox"/>	Seclabdwk1	Ready to run	4/5/10 6:00 AM	Weekly
<input type="checkbox"/>	Hthusracc	Ready to run	4/5/10 6:00 AM	Weekly
<input type="checkbox"/>	Seclabdwk2	Ready to run	4/5/10 6:00 AM	Weekly
<input type="checkbox"/>	Newjob	Held from running	4/4/10 6:00 AM	Weekly
<input type="checkbox"/>	Pspprtjobd	Ready to run	4/4/10 6:00 AM	Weekly

Page 1 of 5

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 the user can remove columns and change the order of the columns using the **Configure Columns** option of the Table actions menu. When **Configure Columns** is selected, a light shaded column section area is displayed, as shown in Figure 13-25.

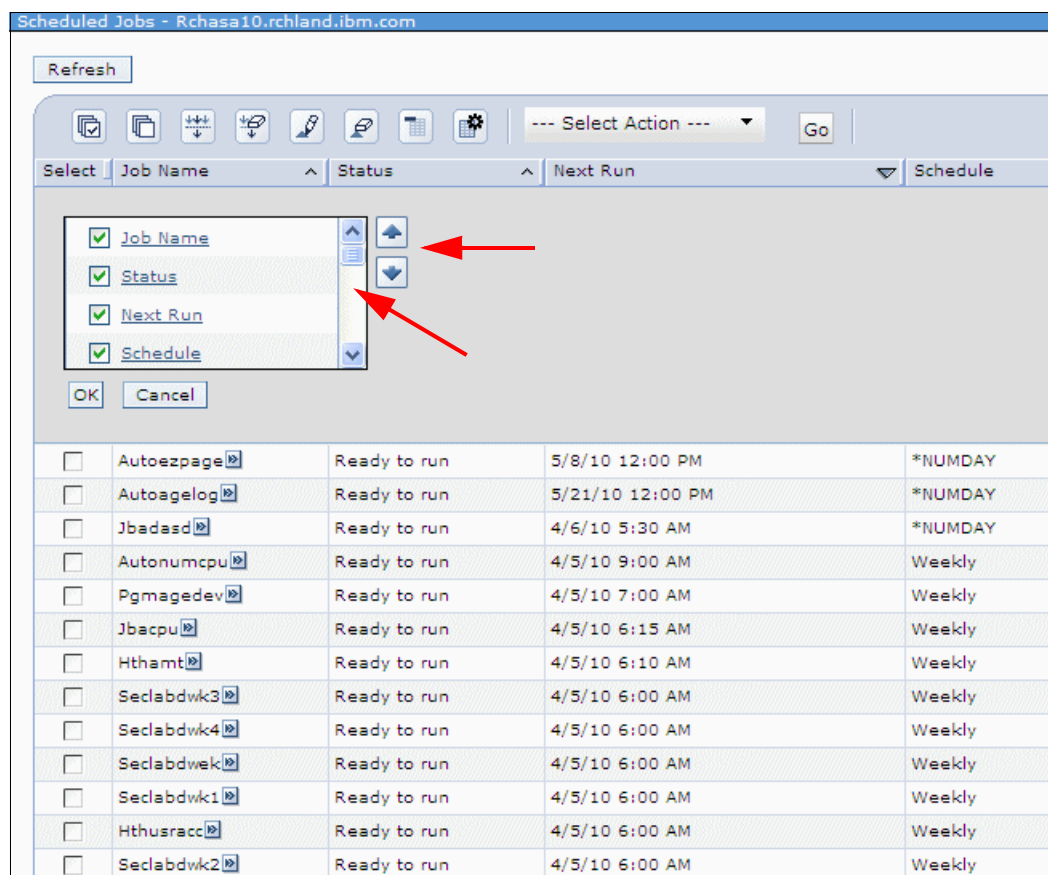


Figure 13-25 Table Column Selection and Reordering

The vertical bar points to a window containing a list of columns with a check box in front of each 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 not to 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.

The **OK** button 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, the user selects **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.

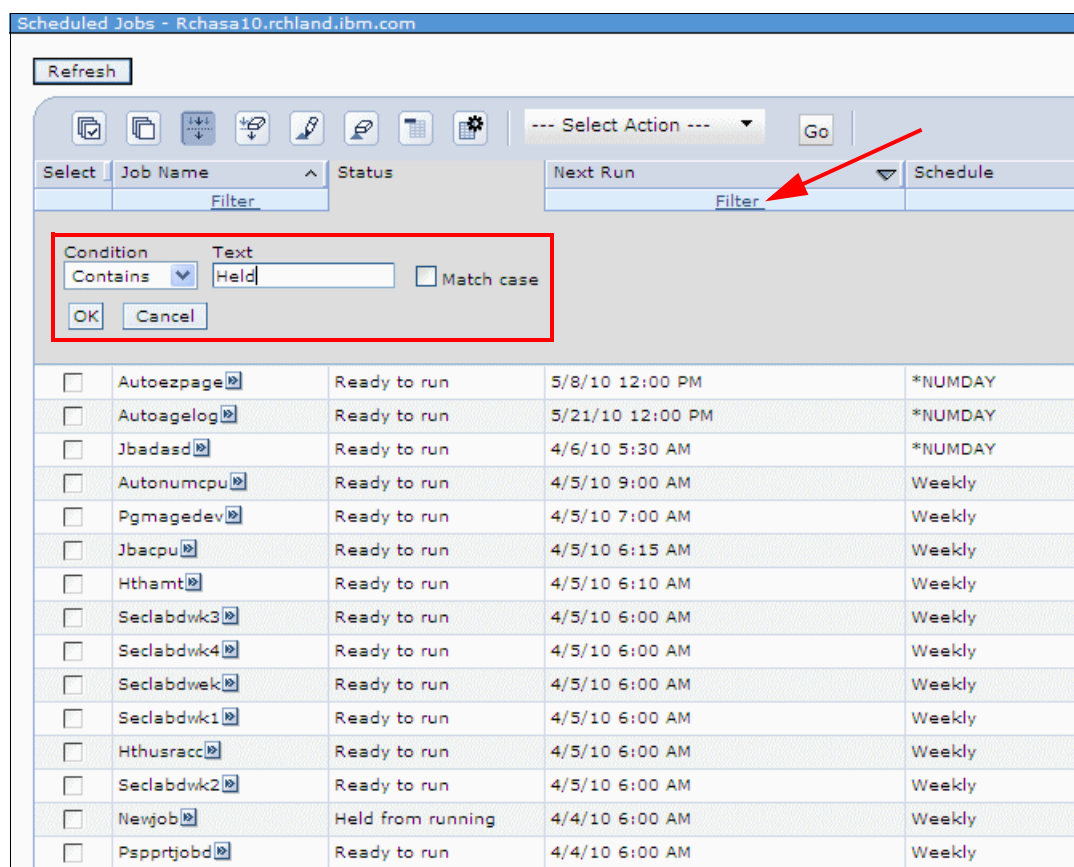


Figure 13-26 Scheduled Job List Filtering

To add a filter for a column, select **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 causes scheduled jobs with status values containing Held or more briefly, all held scheduled jobs.

To disable the filter, the user can select the filter's check box as shown by the top arrow in Figure 13-27. The filter remains, but it does not filter. Is disabled. This is used when a user wants to retain a filter for future use.

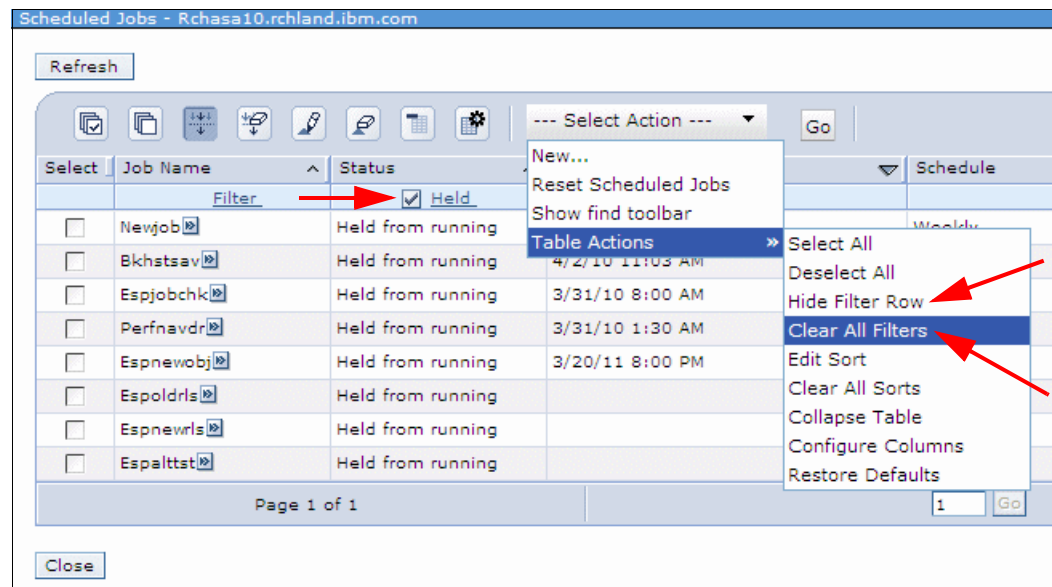


Figure 13-27 Scheduled Jobs Table Filter options

There are two other actions selectable from the Table Actions Menu that impact filters shown in Figure 13-27. They are:

- 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.

This completes our discussion of scheduled job table operations. The other tables in the job scheduler junctions such as Job Groups, Scheduled Job Activity, and Activity Log use similar operations.

13.2.9 Adding a new scheduled job

In this section, the functions and panels to add a new scheduled job are discussed.

To add a new scheduled job, perform the following steps:

1. Select the New Scheduled Job action from the Scheduled Job pop-up menu, as shown in Figure 13-28.

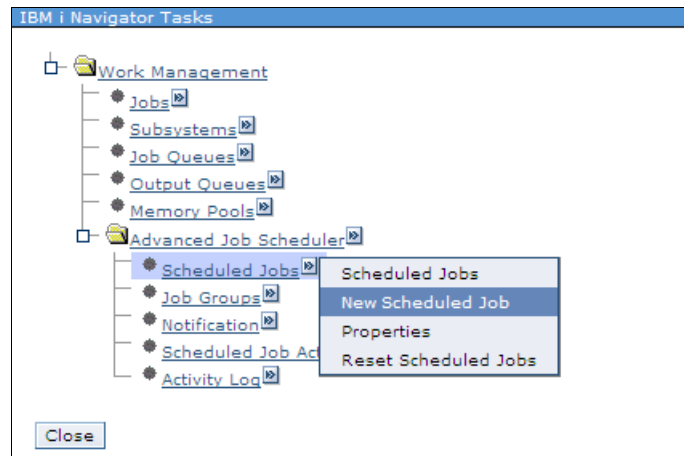


Figure 13-28 Selecting New Scheduled Job action

The General pages of the New Scheduled Job menu (Figure 13-29) is displayed.

Figure 13-29 New Scheduled Job General tab

In Figure 13-29 on page 413, the job name has been set to NEWJOB, the job is a scheduled job (not an alternate job or non-scheduled job) and the job (text) description has been entered.

2. Add one or more commands. To do this, click **Add** in the Commands section. The Command Properties page (Figure 13-30) is displayed.

Figure 13-30 Command Properties panel

3. Specify the command. In this example, RTVDSKINF (Retrieve Disk Information) is entered into the Command box in Figure 13-30.

In this 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 panel shows only the basic parameters. See Figure 13-31.




Figure 13-31 Basic parameters panel

5. Check the advanced parameters by clicking **Advanced**. The panel refreshes with another RTVDSKINF command parameter displayed with a default value, as shown in Figure 13-32.




Figure 13-32 Advanced parameters panel

6. Select **View**. The View menu is displayed, as shown in Figure 13-33.




Figure 13-33 View menu

7. Selects 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 panel refreshes and Figure 13-34 is shown.

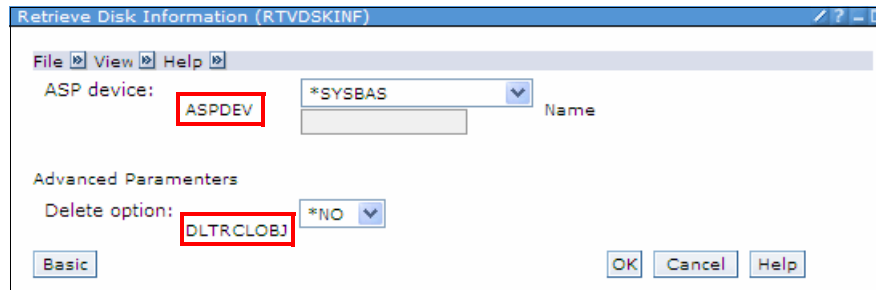


Figure 13-34 Keywords option panel

The last thing of note regarding command prompting is the command help. Figure 13-35 shows the Help pop-up menu.

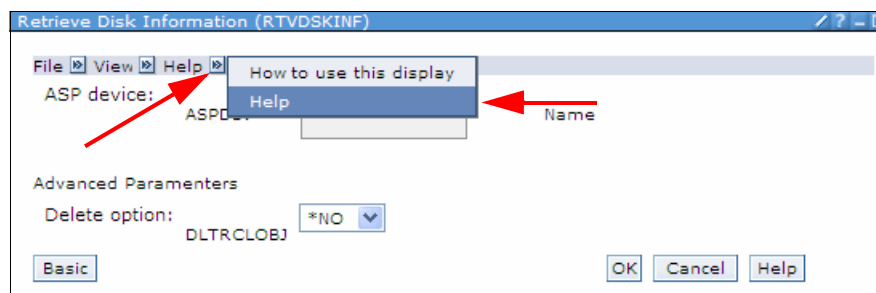


Figure 13-35 Help pop-up menu

Help is selected and the command help is displayed as shown in Figure 13-36 To return from the Help panel, at the bottom of the help page (not shown) is an **OK** button.

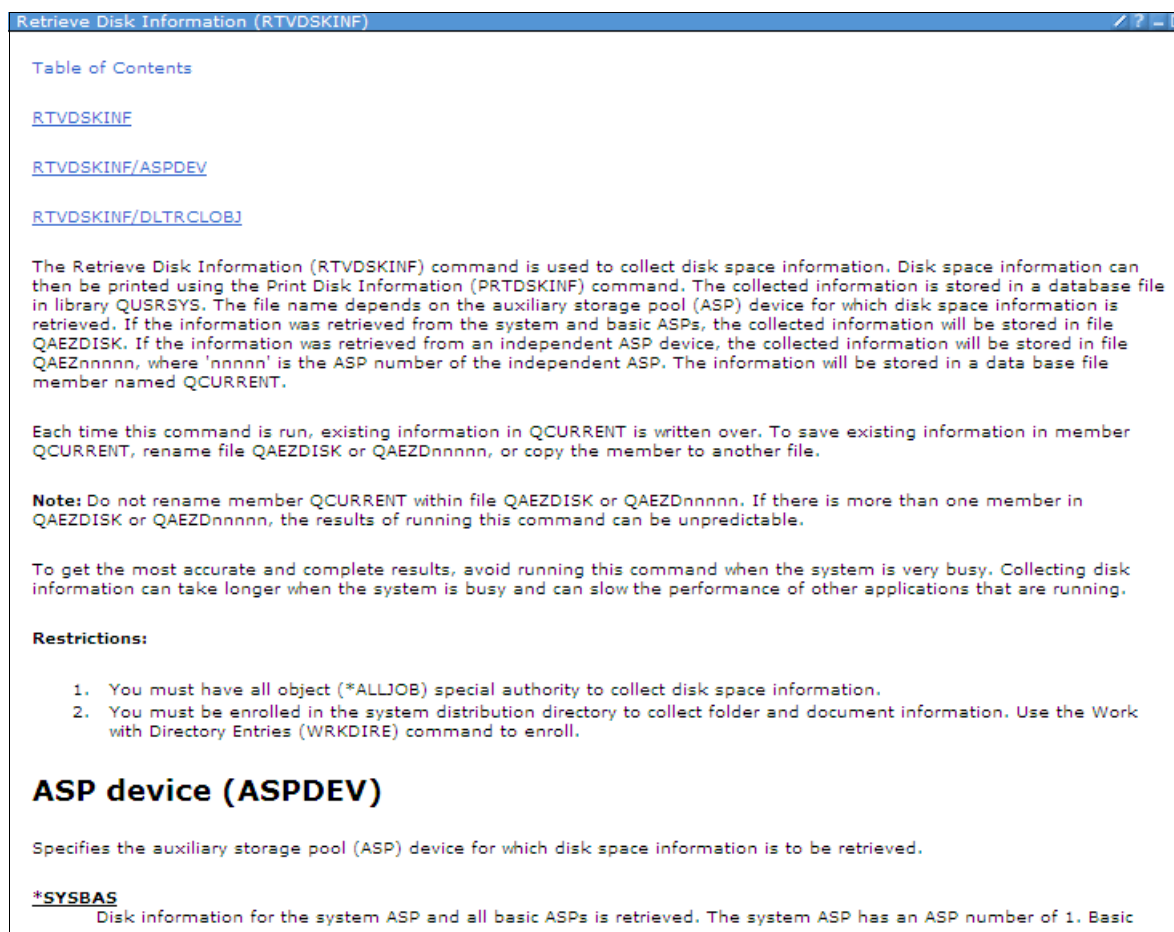


Figure 13-36 RTVDSKINF Command help

Now that the correct command specified, the Help panel is closed by clicking **OK**. The RTVDSKINF command is accepted by clicking **OK** on the RTVDSKINF command prompt display. The command properties for this RTVDSKINF use are accepted by selecting **OK** on the Command Properties panel. The General pages of the New Scheduled Job function is now shown.

In this example, two more commands have been added to the scheduled job (not shown), but the second command is out-of-order, it needs to be the first command to be run. In Figure 13-37, the out-of-order second command is selected, then the **Move Up** button is clicked to move the second command to the first run position.

New Scheduled Job

***General**

*Job name: **NEWJOB**

Type of job

☒ Scheduled

☐ Non-scheduled

☐ Alternate

Description: **This is a new sample job for the IBM i Redbook**

Commands

Select	Enabled	Commands
<input type="radio"/>	<input checked="" type="checkbox"/>	RTVDSKINF
<input checked="" type="radio"/>	<input checked="" type="checkbox"/>	CHGJOB OUTQ(QUSRSYS/DISKINF)
<input type="radio"/>	<input checked="" type="checkbox"/>	PRTDSKINF RPTTYPE(*OBJ) OBJ(*ALL) MINSIZE(20000)

☐ Run as remote command

Application:

Report distribution list:

☐ Use start time of day

Buttons: Add, Add Based On, Remove, Move Up, Move Down, Properties, OK, Cancel

Figure 13-37 Moving a command to change the run order

When the command order is correct, the **OK** button is selected. The next step is to select the Schedule page.

New Scheduled Job: Schedule page

The Schedule page is shown in Figure 13-38. It is used to specify when and how often the scheduled job runs. There are a plethora of scheduling options that are beyond the scope of this book.

The screenshot shows the 'Scheduled Job Properties' dialog box with the 'Schedule' tab selected. The left sidebar contains tabs for General, Schedule, Batch Information, Notification, Documentation, Problem Recovery, Communication, and Local Data Area. The main area displays a calendar for March 2010, with the date 30 highlighted. Below the calendar is a date range selector showing '3/29/2010' to '12/31/2039'. The 'Dates to run' section has two sub-sections: 'Frequency' and 'Details'. In the 'Frequency' section, 'Weekly' is selected. In the 'Details' section, 'Sunday' is checked. The 'On specific times' section is also visible, showing a time of 6:00 AM.

Figure 13-38 New Scheduled Job Schedule page

In the upper left corner, there is a calendar, with which you can select a date range, multiple dates, or one date, as shown in Figure 13-39.

This screenshot is a closer view of the calendar section of the 'Scheduled Job Properties' dialog box. It shows the month of March 2010, with the date 30 highlighted. Below the calendar is a date range selector showing '3/29/2010' to '12/31/2039'. The 'Additional Calendars...' button is also visible.

Figure 13-39 Calendar section of Schedule Job Schedule page

In the upper right corner, shown in Figure 13-40, the user can specify to repeat 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 Schedule Job Schedule page

The “Dates to run” section, shown in Figure 13-41 has a set of radio buttons (circled) to specify the frequency the job is to run. The Details section varies depending on the selected **Frequency** button. In this case, the user has 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 Monthly Yearly Skip count: 0 Select as working days

Details

☒ Sunday ☐ Monday ☐ Tuesday ☐ Wednesday ☐ Thursday ☐ Friday ☐ Saturday

Figure 13-41 Dates to run section of Schedule Job Schedule page

Now that scheduling is completed, the user must move to the Batch Information page.

New Scheduled Job: Batch Information Page

Typically, when a job is submitted there are a host of parameters which must be specified for the job to run. The Batch Information page shown in Figure 13-42 is used 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 Browse...

☐ System group

Submission Information

Job description: Use user's job description Browse... Output queue: Use job description value Browse...

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 Browse... User: MS773 Browse...

Library: Use entry from below Current library: Use job control value Browse...

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

OK Cancel

Figure 13-42 New Scheduled Job Batch Information page

On the Batch Information page, the fields have drop-down options and, in most cases, include a blank field to enter the proper values manually. Parameters such as “Job description” and “User” have **Browse** buttons that display a list panel from which a selection can be made.

After setting these parameters and selecting the **OK** button, the user can choose the Notification page. In many situations, the scheduled job as created thus far executes.

New Scheduled Job: Notification Page

The Notification Page as shown in Figure 13-43 allows the user to set up various methods of notification of a variety of job conditions.

Scheduled Job Properties

Notification

- ☒ **Send completion message**
Message queue: Use entry from below (dropdown) **Browse...** (button) **QSYSOPR** (text field)
Library: Library list (dropdown)
- ☒ **Notify if job completes successfully**
Recipient: Use job control value (dropdown)
Message: Use completion message (dropdown)
- ☒ **Notify if job fails**
Recipient: Use job control value (dropdown)
Message: Use entry from below (dropdown) **Job NEWJOB failed** (text field)
- ☒ **Notify if job does not start within limit**
Recipient: Use job control value (dropdown)
Limit: Use job control value (dropdown)
- ☒ **Notify with error messages**
Recipient: Use job control value (dropdown)
Message: Send job log (dropdown)

OK Cancel

Figure 13-43 New Scheduled Job Notification page

There are five notification options on the Notification Page. Each option has a check box for enabling or disabling the notification.

- **Send completion message**

For this function, the QSYSOPR value was added by selecting the **Use entry from below** drop down option and typing it in the blank field. The user can have selected the message queue by using the **Browse** button and selecting a queue from a list of queues.

- **Notify if job completes successfully**

In this example, the “Use job control value” was selected from the drop-down menu and “Use completion message” was selected from the Message drop-down menu.

- **Notify if Job Fails**

In this example, the Recipient value “Use job control value” was selected from the drop-down menu, and the Message was manually entered.

- **Notify if job does not start within limit**

In this example, the values were allowed to default.

- **Notify with error messages**

In this example the user chose to send the job log.

The drop-down menus prevent the user from entering incorrect options.

A review of the remaining new scheduled job pages follows.

New Scheduled Job: Documentation Page

The Documentation Page, 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, what an operator does when it fails, and so forth).

Scheduled Job Properties

General

Schedule

Batch Information

Notification

Documentation

Problem Recovery

Communication

Local Data Area

Text:

This job is a demonstration job for the IBM i
Technical Overview Redbook.

OK Cancel

Figure 13-44 New Scheduled Job Documentation page

New Scheduled Job: Problem Recovery

The Problem Recovery page is shown in Figure 13-45. On this page, the user specifies problem handling parameters and the message logging parameters for the scheduled job.

The screenshot shows the 'Scheduled Job Properties' dialog box with the 'Problem Recovery' tab selected. The left sidebar contains tabs: General, Schedule, Batch Information, Notification, Documentation, **Problem Recovery**, Communication, and Local Data Area. The main area contains the following fields:

- Maximum run time: No maximum (dropdown), Go (button)
- Maximum run time action: End the job (dropdown)
- ☐ Use an alternate job if job fails
- Job name: (text field), Browse... (button)
- Group: (text field)
- Sequence: 0 (text field), 0 - 99 (text field)
- Action if job cannot start when scheduled: Do not start (dropdown)
- Reply to inquiry message: Use system reply list (dropdown)
- Message logging** (section header)
- Level: Use job description value (dropdown)
- Severity: Use job description value (dropdown)
- Text: Use job description value (dropdown menu is open, showing options: Use job control value, Use job description value, Use message text, Use message text and message help, None, if completes normally)

OK and Cancel buttons are at the bottom right.

Figure 13-45 New Scheduled Job Problem Recovery page

New Scheduled Job: Communication page

Figure 13-46 shows the new scheduled job's Communications page.

The screenshot shows the 'Scheduled Job Properties' dialog box with the 'Communication' tab selected. The left sidebar contains tabs: General, Schedule, Batch Information, Notification, Documentation, Problem Recovery, **Communication**, and Local Data Area. The main area contains the following fields:

- 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 page

New Scheduled Job: Local Data Area Page

Figure 13-47 shows the new scheduled job's Local Data Area page.

The screenshot shows the 'Scheduled Job Properties' dialog box with the 'Local Data Area' tab selected. The left sidebar contains tabs for General, Schedule, Batch Information, Notification, Documentation, Problem Recovery, Communication, and Local Data Area. The main area is divided into a 'Text' section with a large text box and a 'Command Variables' section. The 'Command Variables' section includes an 'Available' list with variables QMM, QDD, QYY, QMDY, and QDMY. An 'Add -->' button is next to this list. Below the list is a 'Position' field with a value of '1' and a range '1 - 1,024'. To the right is a 'Selected' section with a table header 'Select Variable Position' and a single row with 'None'. Below the table are 'Move Up', 'Move Down', and '<-- Remove' buttons. At the bottom right are 'OK' and 'Cancel' buttons.

Select	Variable	Position
	None	

Figure 13-47 New Scheduled Job Local Data Area page

Similar panels are shown when changing an existing scheduled job.

13.3 Other AJS enhancements

In this section we cover other IBM Advanced Job Scheduler for i enhancements.

13.3.1 Notification enhancements

Notification capabilities were enhanced by adding the ability to send a spooled file to an IFS path name with the option whether to replace the file if it already exists. Replacement variables can be used to make the file and the path unique. The path name is specified in a Recipient as shown in Figure 13-48.

The panel in the background shows the New Recipient menu and its General page, as accessed through IBM Systems Director Navigator for i. The panel in the foreground shows the New Recipient Path panel 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.

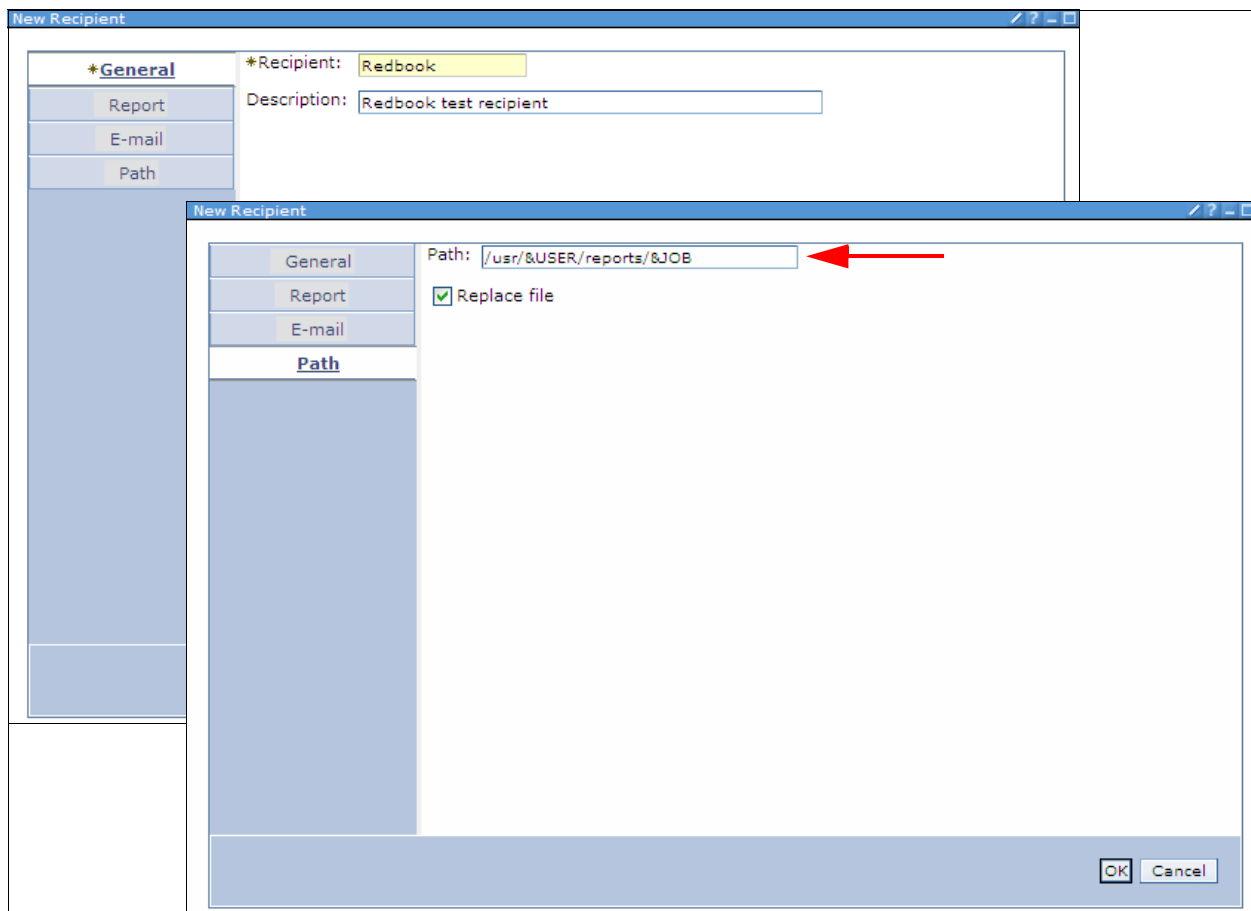


Figure 13-48 The Path page of the Create Recipient menu

13.3.2 Multiple Scheduling Environments function

With the Multiple Scheduling Environments (MSE) function you can set up any number of scheduling environments running 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 replicating live data from a production system to a backup system.

In the character-based interface, the CHGDLJS (Change Data Library using JS) command is used to start additional 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 CHGDLJS command

The user can also specify which users can use which job scheduler environment using the SETDLJS (Set Data Library) command, shown in Figure 13-50.

```

                                Set Data Library using JS (SETDLJS)

Type choices, press Enter.

User . . . . . > MSRJJ           Name, *ALL, *CURRENT...
Data library . . . . . > AJSTEST   Name, *NOCHG

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

```

Figure 13-50 Setting user MSRJJ's access to a job scheduler data library

Restriction: A user profile can be associated with only one job scheduler data library at a time. If not careful, a user 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

In this section we describe the enhancements on 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 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 needs to 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 functions adds flexibility to configure which jobs in a group run on different days and dates without breaking up the group.

Start Group job enhancements

Starting group jobs has been enhanced to add new “Based on” parameters to enable the user to specify an override capability of which jobs in a group run.

The STRGRPJS (Start Group using JS) 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 the user 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 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 Submit Scheduled Job function

When submitting a scheduled job manually, the user can now specify a submit time offset. Suppose it is 6 p.m., and a user wants to submit a scheduled job to run at 9 p.m., but does not want to change the job's schedule, as this is a one-time need. The user can specify a time offset of 180 minutes, causing the job to run at 9 p.m.

The SBMJOBJS (Submit Job using Job Scheduler) 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 from 1 to 720 to determine the submit time.

Similar function is available as the “Minutes from new” parameter through IBM Systems Director Navigator for i, as shown in Figure 13-53.

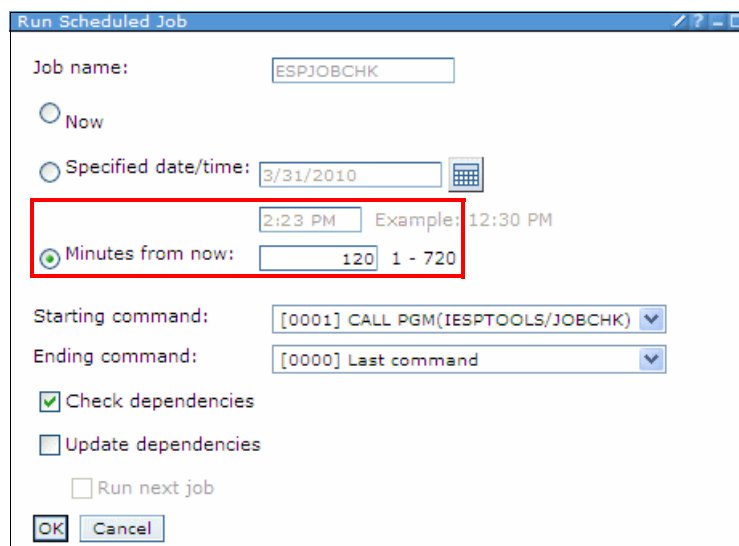


Figure 13-53 Run scheduled job minutes from now

This new function is also available through the System i Access (GUI).

Job dependency enhancements

Jobs can now be dependent on the contents of a data area or a value returned from a program. When specifying a resource dependency for 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

See the following resources for more information:

IBM AJS for i What's New page

http://www-03.ibm.com/systems/i/software/jscheduler/v7r1_feat.html

IBM AJS for i PDF

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzaks/rzaksads.pdf>

Connecting to IBM i IBM Systems Director Navigator for i PDF

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzatg/rzatgdirector.pdf>



Web serving

This chapter provides an overview of the changes made to the IBM HTTP server for i (5770-DG1).

- ▶ 14.1, “Triggered cache manager removed” on page 432
- ▶ 14.2, “Web server search engine and web crawler support removed” on page 432
- ▶ 14.3, “Plug-ins and LoadModule directives” on page 432
- ▶ 14.4, “HTTP Server for i with HA IPv6 Support” on page 438

14.1 Triggered cache manager removed

The Triggered cache manager support (option 1 of DG1) has been 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 have been removed in IBM i 7.1.

The following commands have been 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 made use of the search support will have to be modified to no longer use this removed option.

14.3 Plug-ins and LoadModule directives

This applies to all IBM HTTP servers that have been associated with a WebSphere Application Server Version 6.1 (5733-W61) or Version 7.0 (5733-W70) when upgrading to IBM i 7.1.

All Application Server service programs implementing 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 Servers in IBM i 7.1, see 15.1, “IBM Integrated Web Services for i” on page 442.

In IBM i 7.1, the LoadModule directives used by external HTTP servers associated with Application Server Version 6.1 and Application Server Version 7.0 have changed.

For HTTP servers that have been associated with a WebSphere Application Server Version 6.1 or Version 7.0, 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 Application Server installation on your system contains the program and service program objects for the installed product.

- ▶ For WebSphere Application Server 6.1
 - 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 Version 6.1 product installed on the system is QWAS61A.

- For WebSphere Application Server 7.0
 - The product library name for Version 7.0 is QWAS7x (where x is A, B, C, and so on).
 - The product library for the first WebSphere Application Server Version 7.0 product installed on the system is QWAS7A.

The product library can be determined by examining file `<profile_root>/properties/.instance.properties` 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 Web page.

We now discuss an example, where we assume we created a WebSphere Application Server WAS 7.0 server instance WAS70TOMV on an i 6.1 IBM i.

We start the IBM Systems Director Navigator for i and click the **IBM i Tasks Page** link on the Welcome Page as shown in Figure 14-1.

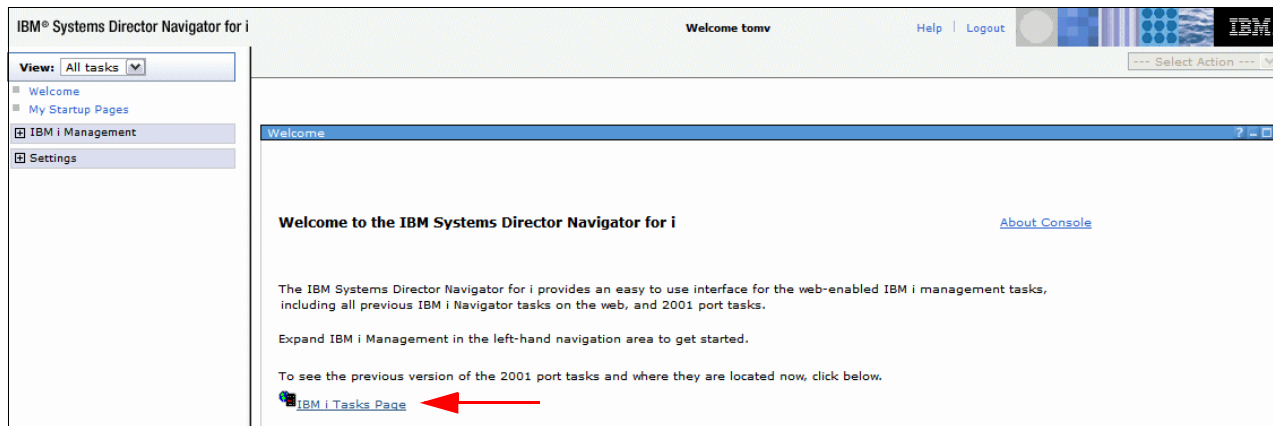


Figure 14-1 IBM Systems Director Navigator for i - Welcome page

This brings us to the panel as shown in Figure 14-2 where we click **IBM Web Administration for i**.

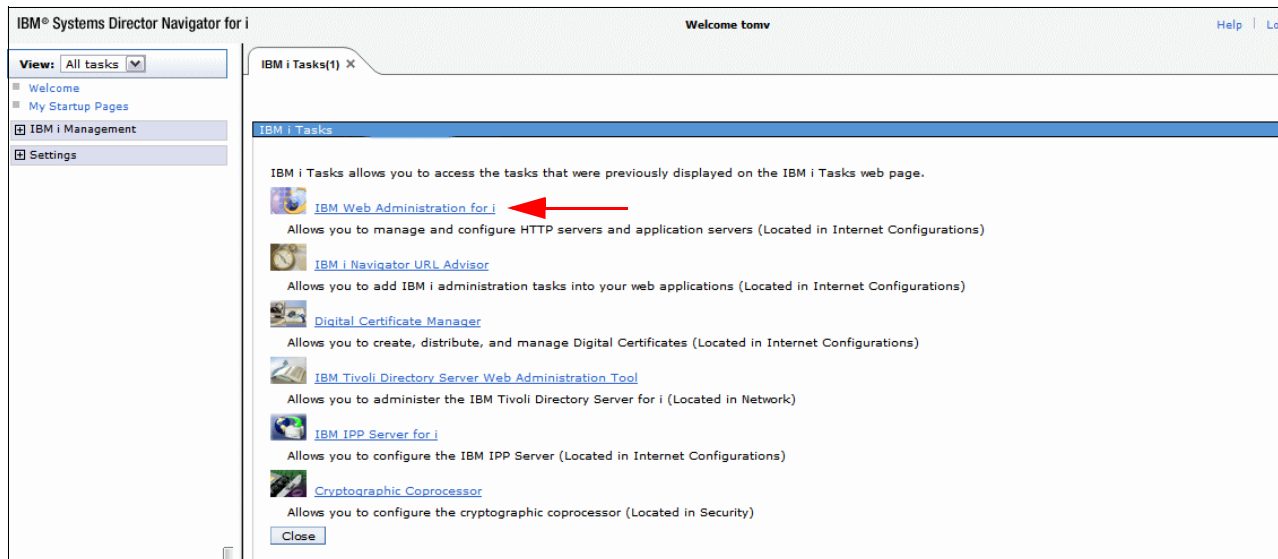


Figure 14-2 IBM Web Administration for i

We now get to the IBM Web Administration for i Web page as shown in Figure 14-3.

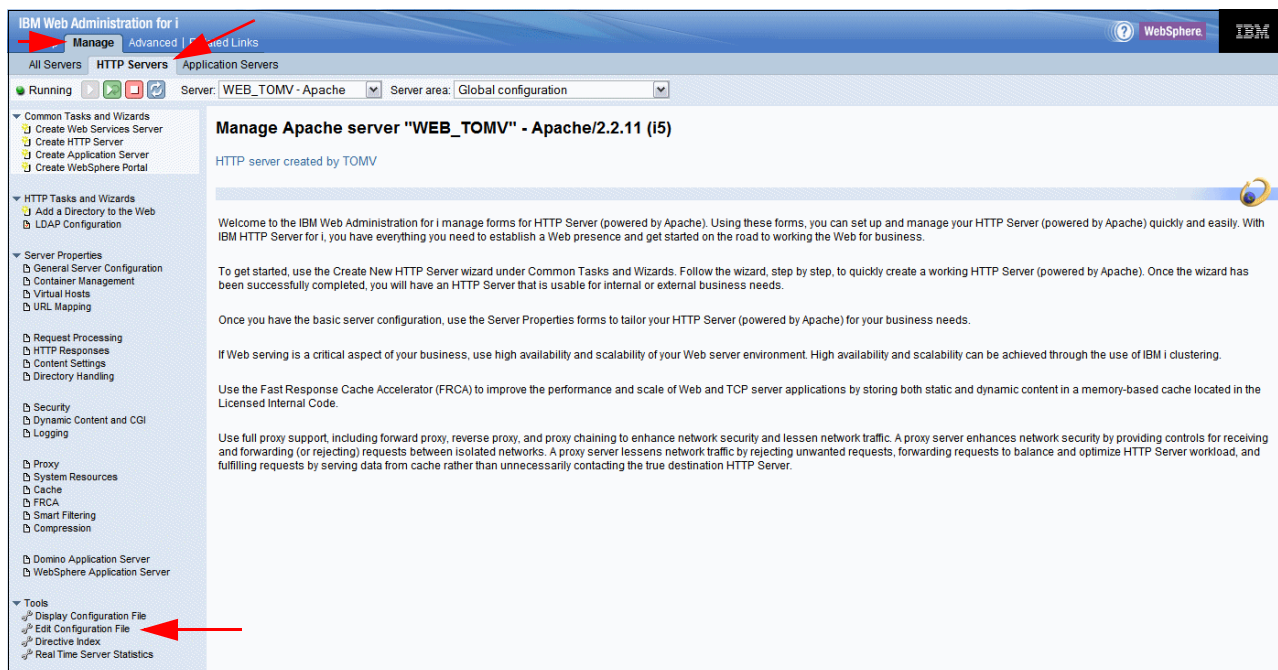


Figure 14-3 IBM Web Administration for i - Manage HTTP Servers

We perform the following steps:

From the Manage tab, click **HTTP Servers**.

- ▶ Select your server from the Server pull-down menu.
- ▶ Click **Edit Configuration file** in the navigation frame.

In Figure 14-4, we have the configuration file of the corresponding HTTP server instance WEB_TOMV that got 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/MERC
URE_WAS70TOMV/nodes/MERCURE.BE.IBM.COM-node/servers/IHS_WEB_TOMV/plugin-cfg.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

We now look for the library name that is associated with our 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

According to the aforementioned explanation, we 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/MERC
URE_WAS70TOMV/nodes/MERCURE.BE.IBM.COM-node/servers/IHS_WEB_TOMV/plugin-cfg.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

On the IBM Web Administration for i Web page as shown in Figure 14-7, you click **Apply** to change the configuration file.

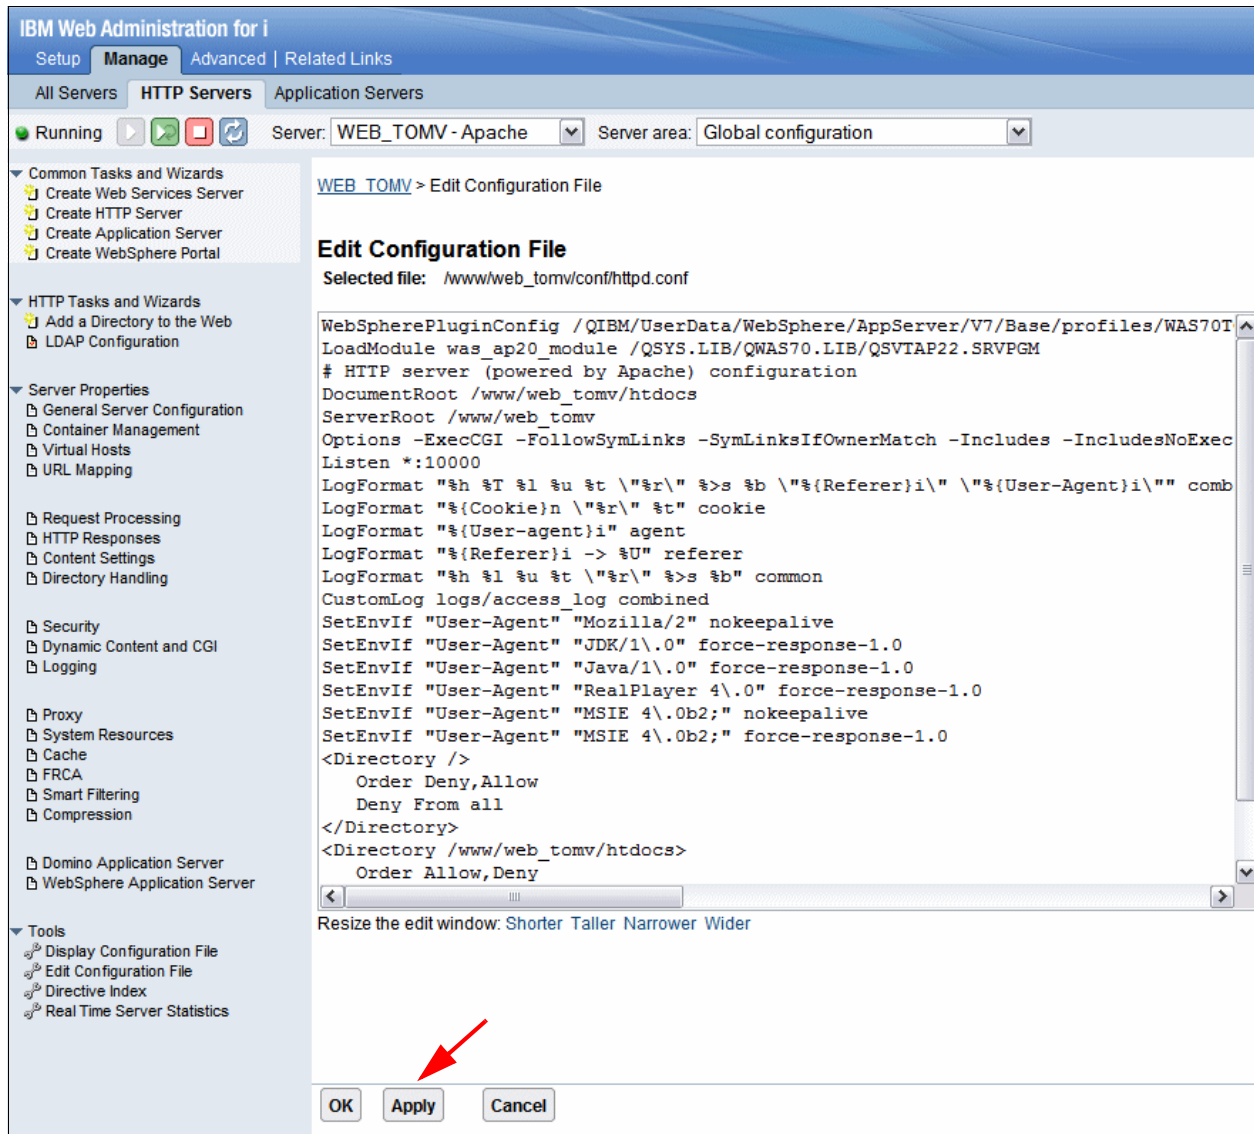


Figure 14-7 IBM Web Administration for i - Applying changes to configuration file

You now get a message that the configuration was successfully changed. Do not restart the server at this time however. Click **OK** and from now on you can start the upgrade from IBM i 6.1 towards IBM i 7.1. Afterwards, it will possible to successfully start the HTTP server 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 Web site. This improves the availability of business-critical Web applications built with static Hypertext Markup Language (HTML) pages or Common Gateway Interface (CGI) programs.

Prior to IBM i 7.1 release, only IPv4 address is supported by PowerHA. Now PowerHA for i on IBM i 7.1 already fully supports IPv6 address (including all HA-related APIs, commands, and GUIs), so HA IPv6 support has also been added to HTTP Server for i on IBM i 7.1. Now users can use IPv6 address to configure all web servers in the cluster and access their web applications running in the highly available Web server environment.

Note: The IPv6 network among the clients and the cluster must have already been setup and available to access. That means every client should be able to PING through the clustered IPv6 address.

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)



Web enablers

This chapter discusses the following topics:

- ▶ 15.1, “IBM Integrated Web Services for i” on page 442
- ▶ 15.2, “Changes in IBM WebSphere Application Server” on page 442
- ▶ 15.3, “IBM Web Administration for i” on page 447

15.1 IBM Integrated Web Services for i

Integrated web services for i enables integrated language environment (ILE) applications to play in the web services and Service Oriented Architecture (SOA) arena with very little effort, knowledge and resources. The convergence of web service 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.

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 discuss enhancements that have been made to the integrated web services support on IBM i.

Note for the latest news on the Integrated Web Services for i support go to the web site at:

<http://www.ibm.com/Systems/i/software/iws/>

15.1.1 Integrated Web Services Client for ILE

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. A user would use the tools to generate code that can be called from ILE applications to consume a web service.

In the past, the tools only supported the generation of C and C++ code and users had to manually generate the program or service program that contained the generated code. The tools have been enhanced to generate RPG code in addition to automatically creating the service program containing the generated code.

To use this support, ensure 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 discuss enhancements that have been made to IBM 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 have been replaced by IBM Technology for Java. If your applications are still using classic Java you need to upgrade to IBM Technology for Java, but before you do that you need to be aware of the following information (see also Table 15-1 on page 443):

- ▶ Classic 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 gigabytes, which is approximately 1000 threads. If you require more than 1000 threads or a Java object heap larger than 3 gigabytes use the 64-bit version of the IBM Technology for Java
- ▶ If you have integrate language environment (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.
- ▶ 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 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. IBM recommends Java SE 6 when migrating from Java Developer Kit 1.4 or 5.0

For detailed information about what's new in Java for IBM i 7.1 see the information center.

Note: 5761JV1 options 6, 7, and 10 are the only options 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, visit the following web site

<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 completing the installation.

15.2.4 Upgrading to IBM i 7.1

If you are currently using IBM WebSphere Application Server V6.0 and upgrading to IBM i 7.1, you need to migrate to the IBM WebSphere Application Server V6.1.0.29 or V7.0.0.7 products. Migration can occur before or after upgrading to IBM i 7.1.

Before upgrading to IBM i 7.1, ensure that all IBM WebSphere Application Server installations meets the minimum required fix levels. The version identifier is contained in the file `<app_server_root>/properties/version/WAS.product`. 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 IBM WebSphere Application Server V6.1, apply fix pack 29 (6.1.0.29) or later if needed. For IBM WebSphere Application Server V7.0, apply fix pack 7 (7.0.0.7) or later if needed.¹

When upgrading to IBM i 7.1, enable IBM 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 upgrading to IBM i 7.1, if you upgraded from V5R4 or earlier, update the IBM WebSphere Application Server service programs for IBM i 7.1. To update the programs, perform the following steps:

1. Start the Qshell interpreter.
2. Change directories to `<app_server_root>/bin`.
3. Invoke **export OSVER=V7R1M0** to export the OSVER environment variable to the Qshell environment.
4. Invoke the `_postfpexit` script as follows:
`_postfpexit<app_server_root><product_library>`. Where `<product_library>` is the product library for the WebSphere Application Server installation as listed in `<app_server_root>/properties/product.properties`.
5. Upload the LoadModule directives. LoadModule directives are used by external HTTP servers associated with Application Server Version 6.1 and Application Server Version 7.0. For information about updating the LoadModule directives see "Plug-ins and LoadModule directives" on page 432.

¹ It is recommend switching any WebSphere servers running classic to J9 prior to the OS upgrade. This would be done with the `'/qibm/proddata/websphere/appserver/<version>/<edition>/bin/enablejvm -jvm std32'` command

15.2.5 IBM Installation Manager for WebSphere Application Server Version 8.0

WebSphere Application Server Version 8.0 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.

Restrictions:

- ▶ Do not use the same response files that are used with WebSphere Application Server Version 7.0 or earlier to install or uninstall Version 8.0 and later; use response files that are based on Installation Manager to install, update, or uninstall Version 8.0 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 IBM Installation Manager:

IBM Installation Manager is a general-purpose software installation and update tool that runs on a range of computer systems. Installation Manager can be invoked 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.

More information about the IBM installation Manager can be found at the IBM Installation Manager Version 1.4 Information Center at the following web site:

<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 only need to run Installation Manager 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 keep track of 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 binaries that are copied from the installation kit and a set of runtime data that describe the products that have been installed by this particular Installation Manager. Before creating an Installation Manager, you must decide in which mode the Installation Manager will run as well as where the binaries and runtime data—called "agent data" or "appdata"—will reside. Then, you issue the Installation Manager installation command from the appropriate user ID to create the Installation Manager.

Accessing product repositories:

All software materials that will be installed with IBM 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 have created 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 Version 8 products by issuing the versionInfo command 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.

15.2.6 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, perform the following steps:

1. Enter the GO LICPGM (Go Licensed Program) command and select Option 10.
2. If you do not see the program listed, then perform 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 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 prior to moving 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 have been made to IBM Web Administration for i to include 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 exposes you to missing important information of which you might need to be aware. This support has been included as part of the IBM Web Administration for i interface.

The Web log monitor gives the user the ability 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 the 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

The web log monitor ensures that important messages are not lost or missed.

Note: Web log monitor inspects the log files as long as Web Administration for i is started. Minimum OS supported is IBM i 6.1 PTF SF99115 level 12 or higher.

Configuring Web Log Monitor

To activate Web Log Monitor follow these steps:

1. Bring up 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 page (Figure 15-1) click the **IBM i Tasks Page** link in the right navigation panel.

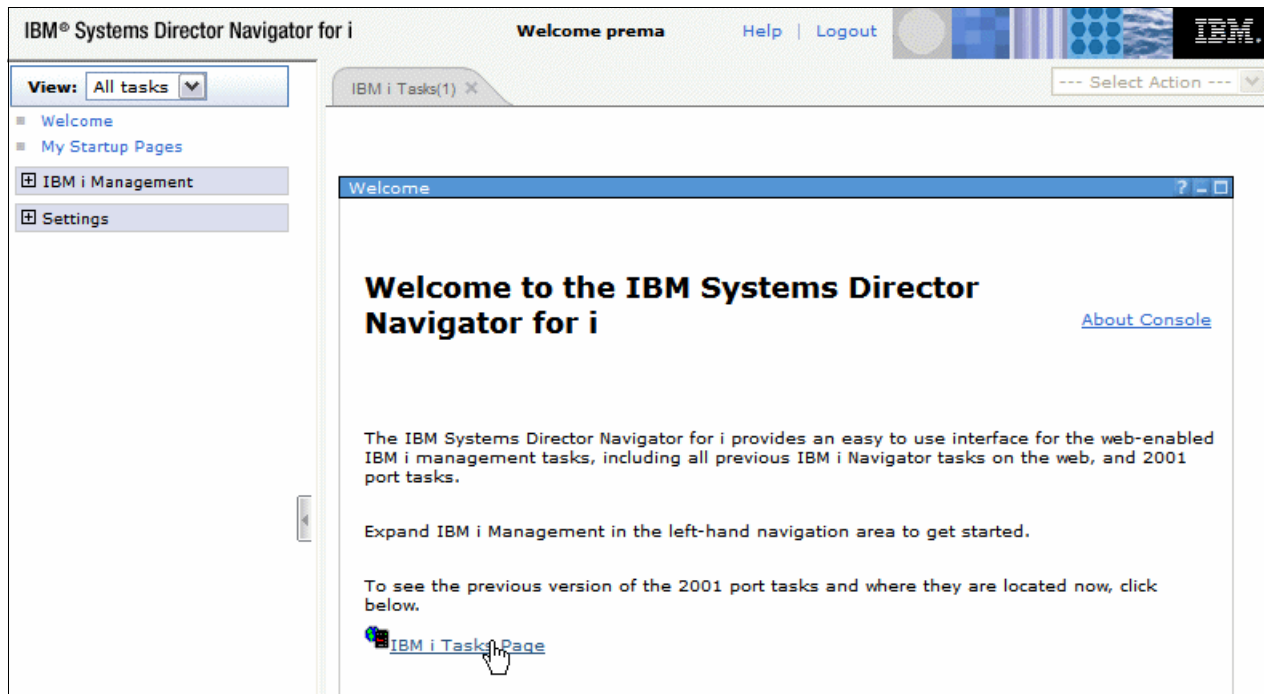


Figure 15-1 IBM Systems Director Navigator for i

3. Click the **IBM Web Administration for i** link (Figure 15-2).

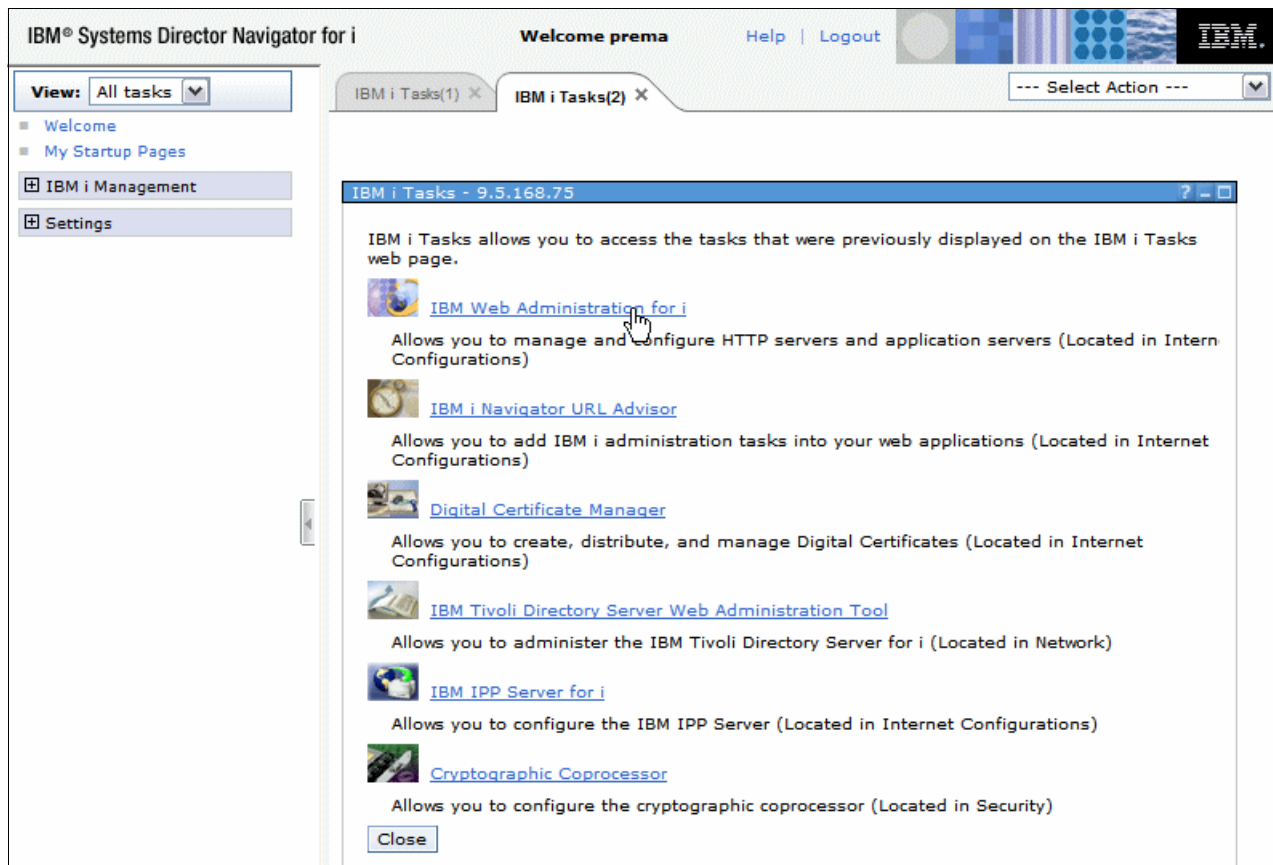


Figure 15-2 IBM Web Administration for i

4. Click **Manage** (Figure 15-3).

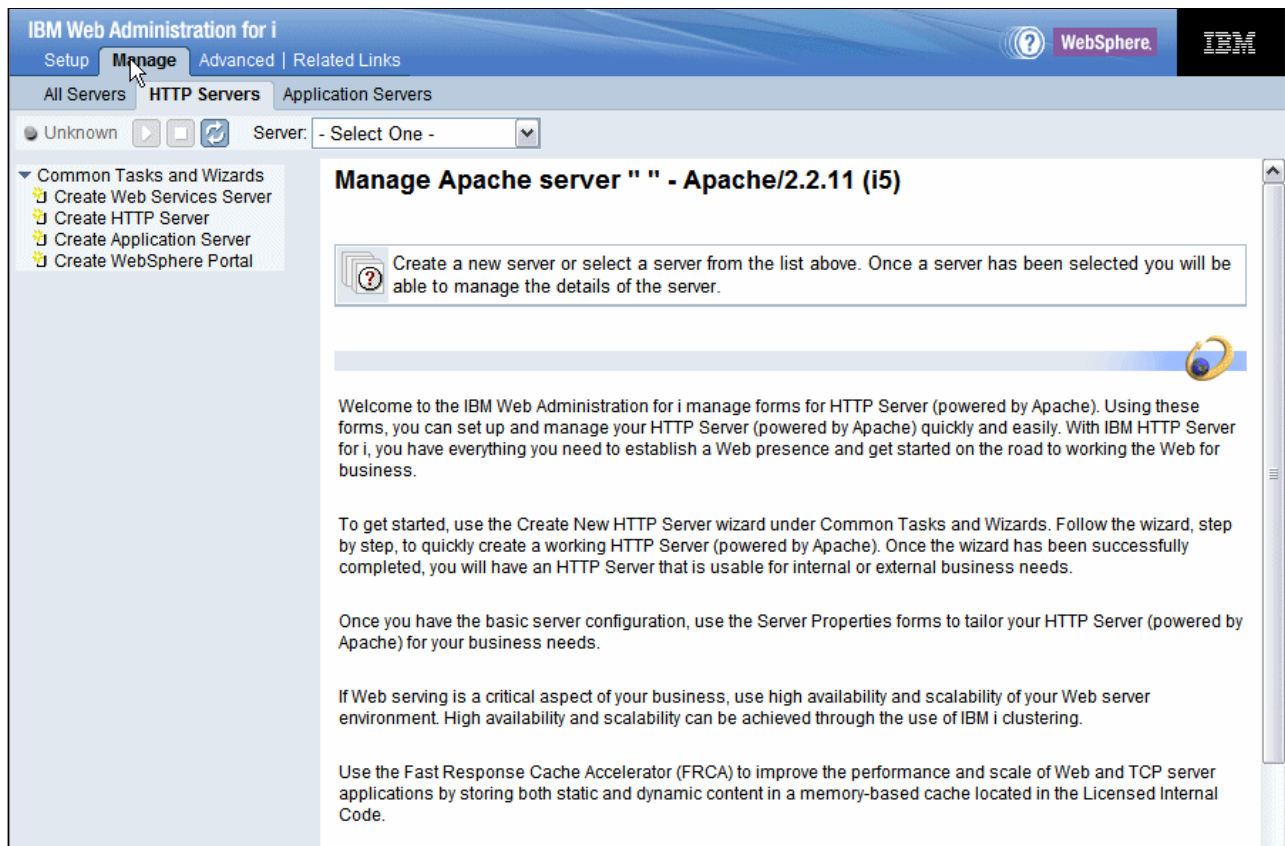


Figure 15-3 Manage tab

5. Click **Application Servers** (Figure 15-4).

The screenshot displays the IBM Web Administration for i interface. The top navigation bar includes 'Setup', 'Manage' (selected), 'Advanced', and 'Related Links'. The 'Manage' tab is active, and the 'Application Servers' sub-tab is selected. The main content area is titled 'Manage Web Services Server' for 'WSERVICE'. It includes a description of the Web services server and a 'Manage Deployed Services' table. The table is currently empty, showing 'There are no entries for this table.' A 'Refresh' button is located below the table. The left sidebar contains a tree view with categories like 'Common Tasks and Wizards', 'Web Services Wizards', 'Server Properties', 'Services', 'Web Performance', and 'Problem Determination'.

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)

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 15-4 Application Servers

6. Click **Web Log Monitor** under the “Problem Determination” link in the left navigation panel (Figure 15-5).



Figure 15-5 Web Log Monitor

7. Click **Enable Web Log Monitor** (Figure 15-6).

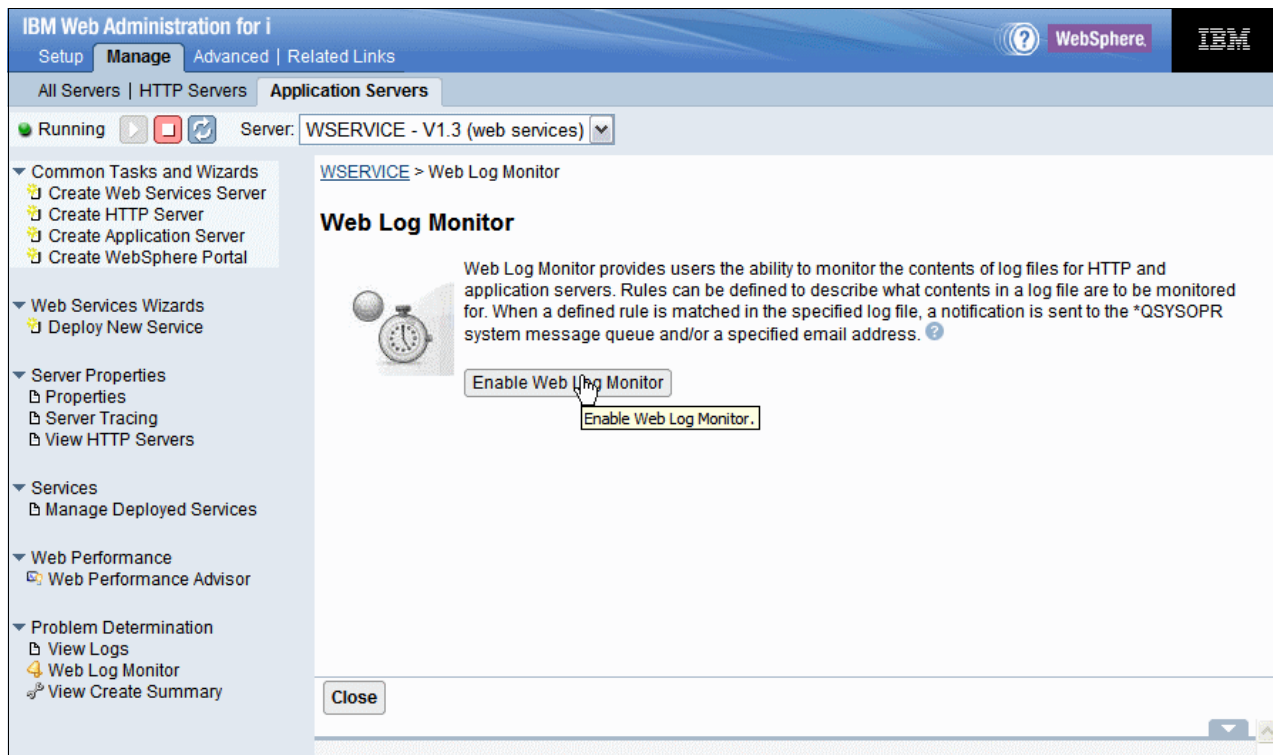


Figure 15-6 Enable Web Log Monitor

8. Specify the log file to monitor. Click **Browse** to select the log file (Figure 15-7). Only log files that you are authorized to use are shown in the browser window. Select the log file and click **Next**.



Figure 15-7 Specify Log file to monitor

When the log file is selected, the Basic Tab (Figure 15-8) or the Advanced Tab can be used to configure the rule.

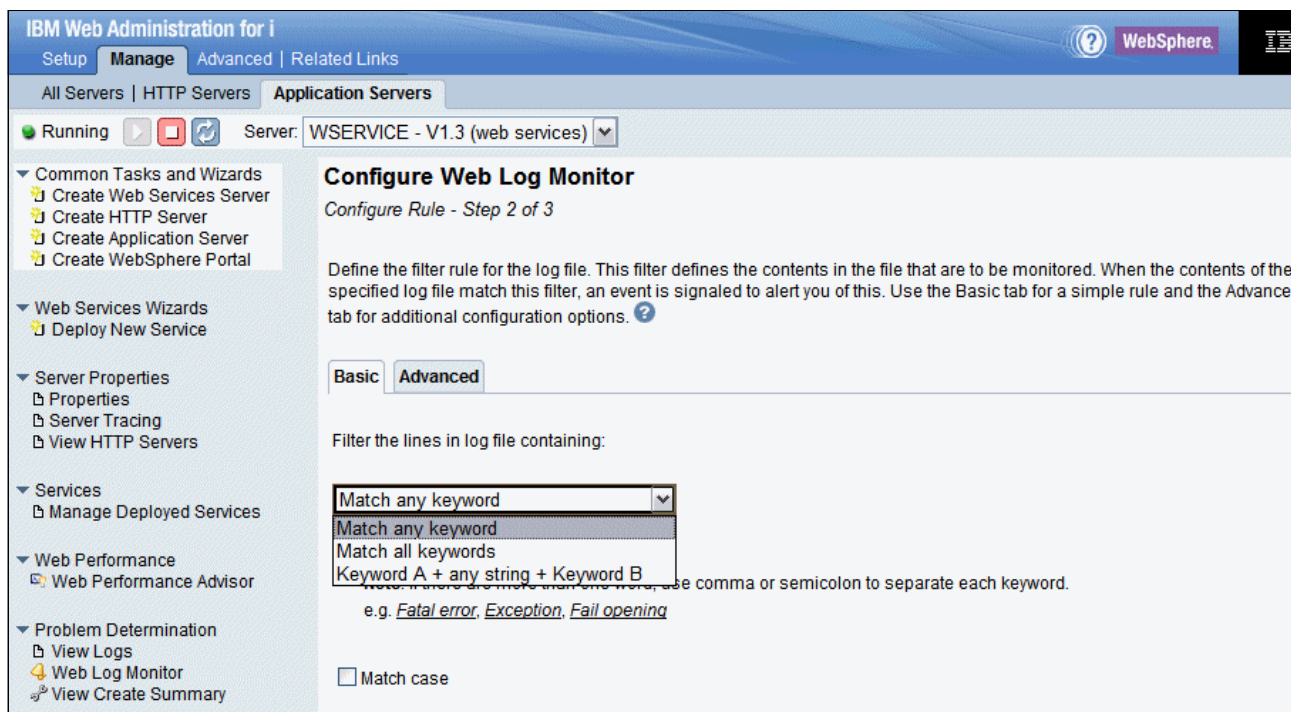


Figure 15-8 Basic tab

Keywords can be used to filter the content of specified log file. To specify more than one keyword use 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 keyword specified here are *Dump*, *Failed*, *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*, *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*, *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 keyword specified here are *Dump* and *detail* then this line is considered a match.

In the Advanced Tab (Figure 15-9) specifies complex patterns to be used in the rule.

- Exclude the following keywords
This option specifies whether or not to ignore a line which contains any of the following keywords.
- Configure regular expressions
This option gives users the ability to customize the rule by 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 lists various tasks and wizards, including 'Common Tasks and Wizards', 'Application Server Wizards', 'Server Properties', 'Applications', 'Resource Configuration', 'Web Performance', and 'Problem Determination'. The main content area is titled 'Configure Web Log Monitor' and 'Configure Rule - Step 2 of 3'. It contains a text box for defining the filter rule, with instructions: 'Define the filter rule for the log file. This filter defines the contents in the file that are to be monitored. When the contents of the file match this filter, an event is signaled to alert you of this. Use the Basic tab for a simple rule and the Advanced tab for additional configuration options.' Below this, there are two tabs: 'Basic' and 'Advanced'. The 'Advanced' tab is selected. It contains two checkboxes: 'Exclude the following keywords' and 'Configure regular expression'. The 'Exclude the following keywords' checkbox is checked, and there is a text box below it with the instruction: 'If any of the following keywords appears in a line, ignore the line. Use semicolon to separate multiple keywords.' The 'Configure regular expression' checkbox is also checked, and there is a text box below it with the instruction: 'This option overrides other rule configurations.' The text box contains the regular expression '(?!).*\\(Open\\E).*' and an example: 'e.g. *Failed opening*on server*, ^[0-9]*[1-9][0-9]*\$', ^s*\\s*\$'. At the bottom, there are three buttons: 'Back', 'Next', and 'Cancel'.

Figure 15-9 Advanced tab

- If you have multiple log files to monitor, click **Add** in the File and Rule Tab as indicated in Figure 15-10, and repeat steps 8 to 9.

IBM Web Administration for i

Setup **Manage** Advanced | Related Links

All Servers | HTTP Servers **Application Servers**

Running 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. Describe what contents in a log file are to be monitored for. When a defined rule is matched in the specified *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				
<input type="radio"/> /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				
<input type="radio"/> /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 15-10 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 15-11). The choices are as follows:

- 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
Notification is send to the specified email address. More than one email address can be specified, use comma to separate multiple email addresses.
- Sender’s email address
Specifies the email address used to send the notification
- SMTP server
Specifies the SMTP (Simple Mail Transfer Protocol) server used to send the mail notification
- Account
Specifies the account used to log on to the SMTP (Simple Mail Transfer Protocol) server
- Password
Specifies the password of the account used to log on to the SMTP (Simple Mail Transfer Protocol) server

The screenshot shows the IBM Web Administration for i console. The top navigation bar includes 'Setup', 'Manage' (selected), 'Advanced', and 'Related Links'. Below this, there are tabs for 'All Servers', 'HTTP Servers' (selected), and 'Application Servers'. The main content area is titled 'Web Log Monitor' and includes a description of its functionality. The 'Notification' tab is active, showing a 'Notification:' section with two checked options: 'Send message to *QSYSOPR message queue.' and 'Send email to [text box] e.g. joe@us.ibm.com. Use comma to separate multiple email addresses.' Below these are input fields for 'Sender's email address:', 'SMTP server:', 'Account:', and 'Password:'. A 'Send Test Notification' button is located below the input fields. At the bottom of the form are 'OK', 'Apply', and 'Cancel' buttons. The left sidebar contains a tree view with categories like 'Common Tasks and Wizards', 'HTTP Tasks and Wizards', 'Server Properties', 'Tools', and 'Web Log Monitor' (highlighted).

Figure 15-11 Notification tab

11. Specify interval in minutes and maximum number of notifications to be sent per hour to prevent notification flood on the Miscellaneous tab (Figure 15-12).

- 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.

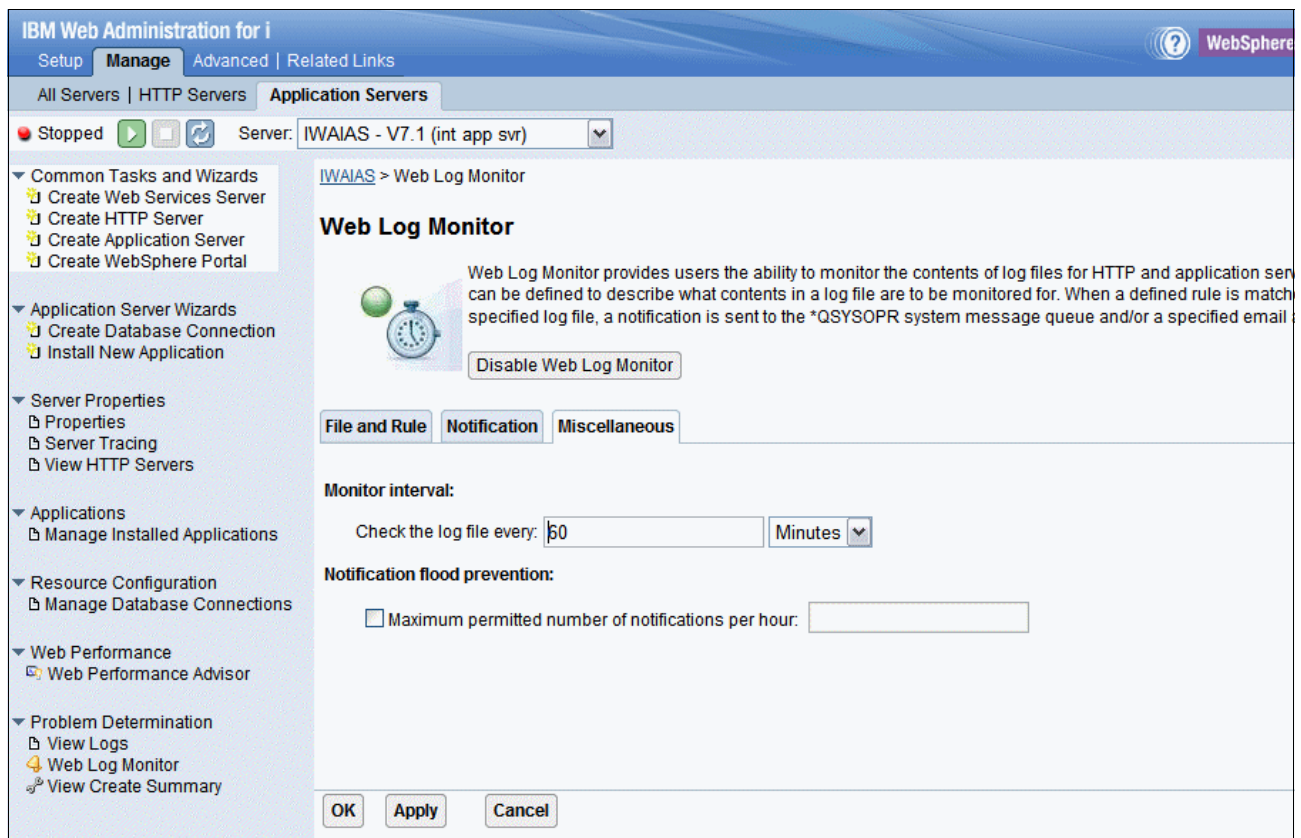


Figure 15-12 Miscellaneous tab

To disable Web Log Monitor, click the Disable Web Log Monitor on the Web Log Monitor introduction page.

Note: Only users who are 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 navigate to **IBM i Task page** → **IBM Web Administration for i** → **Advanced** → **Permissions**.

The Manage Permissions form allows administrators to give permission to users without *IOSYSCFG and *ALLOBJ. There are two roles that can be given:

- 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 have been 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 deployed with integrated Web services servers

b. An administrator granting permissions to user profile needs to explicitly grant the create-server permission

A new feature group profile support adds the convenience to grant or revoke permissions to a group of users all at once, otherwise, these users need to be granted or revoked permissions separately, which is time-consuming and error prone. The algorithm is like this: When a user has one or more supplemental groups, we combine the permissions 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 Web Administration interface.

The group profile support only applies to IBM i 6.1 and above.

Adding permissions

The Add Permissions wizard takes into consideration all aspects of a selected server. If a selected application server is associated with an HTTP Server, the Add Permissions wizard

takes this into account and recommends that permissions are specified correctly for that web environment. This ensures that the specified user can successfully manage the server based on granted permissions.

Users who are being granted permission to servers can be given a different role for each server. When a user is given permission to create new servers, any server that they do create will automatically be updated to give them developer permission to that newly created server.

To add permissions for a user ID, perform the following steps:

1. Click **Add Permissions** (Figure 15-13) Select the user ID to which you want to give permission. Click >> and click **Next**.

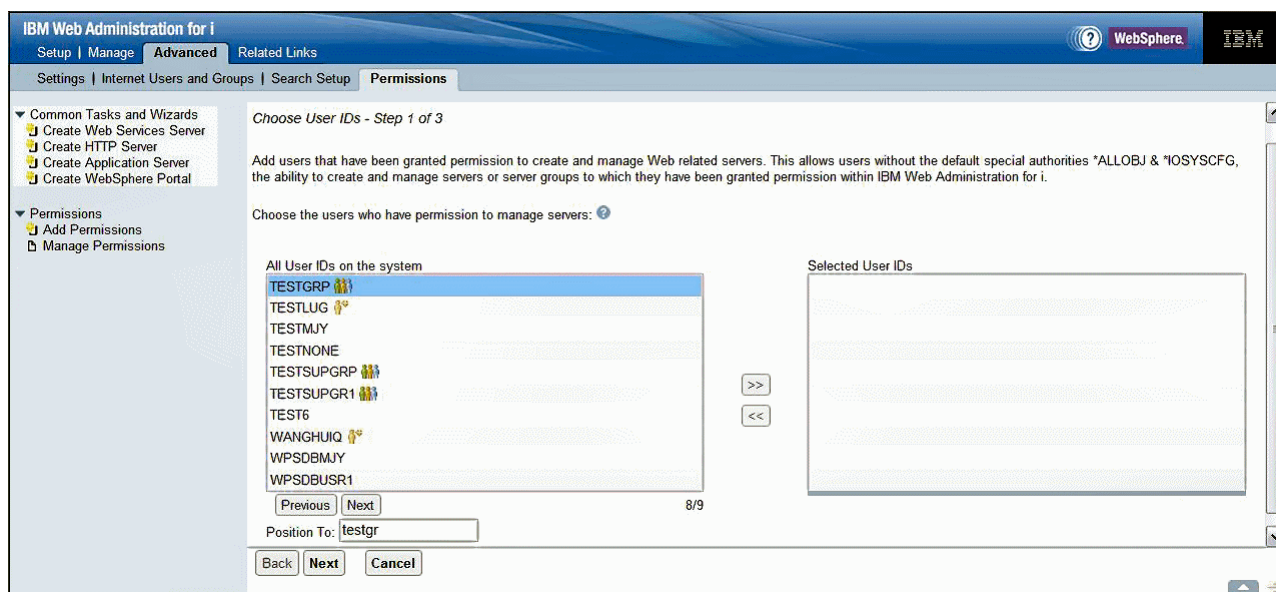


Figure 15-13 Add Permissions Step 1

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. The options are explained as follows:
 - Server
 - This specifies how the servers category are to be handled. There are three values:
 - ***ALL**: The users specified in step 1 on page 460 give permission to all the servers for this category, including all existing server and any other server created in the future.
 - ***NONE**: The users specified in step 1 on page 460 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: List the permissions that is granted to a user for a server
 - **Operator**: Basic permission to the server. See Table 15-2 on page 459 for details
 - **Developer**: All permissions to the server. See Table 15-2 on page 459 for details.

- **Create Server:** The specified users can be granted permission to create new 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	<input checked="" type="checkbox"/>
Integrated Web Services Server	*ALL	Developer	<input checked="" type="checkbox"/>
WebSphere Application Server	*ALL	Developer	<input checked="" type="checkbox"/>
IBM HTTP Server	*ALL	Operator	<input checked="" type="checkbox"/>

Back Next Cancel

Figure 15-14 Add Permissions Step 2

- When you have defined the proper authority and role, click **Next**. Check what is displayed (Figure 15-15) if you need to change click **Back**, or 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 15-15 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 takes into consideration all aspects of a selected server. If a selected application server is associated with an HTTP Server, the Modify Permissions wizard takes this into account and recommends that permissions are specified correctly for that entire web environment. This means either to add or remove the permissions for all servers within that specified web environment. This 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 15-16). This is launched from the By Server view on the Manage Permissions panel when choosing 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 15-16 Modify permissions by Server view

- Modify the permissions for a specific user (Figure 15-17). This is launched from the By User view on the Manage Permissions panel when choosing 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	
Integrated Web Application Server	<input type="checkbox"/> Create server
Servers: ALL	
Role: Developer	
Integrated Web Services Server	<input checked="" type="checkbox"/> Create server
Servers: *ALL	
Role: Operator	
WebSphere Application Server	<input checked="" type="checkbox"/> Create server
Servers: *ALL	
Role: Developer	
IBM HTTP Server	<input checked="" type="checkbox"/> Create server
Servers: *ALL	
Role: Developer	

Back Next Cancel

Figure 15-17 Modify permissions by User view

Removing permissions

The Remove Permissions wizard allows an administrator 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 takes into consideration all aspects of a selected server. If a selected application server is associated with an HTTP Server, the Remove Permissions wizard takes this into account and will also remove the permissions for all servers within that specified web environment. This ensures that the specified user no longer successfully manage 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 15-18). This is launched from the By Server view on the Manage Permissions panel when choosing a specific server.

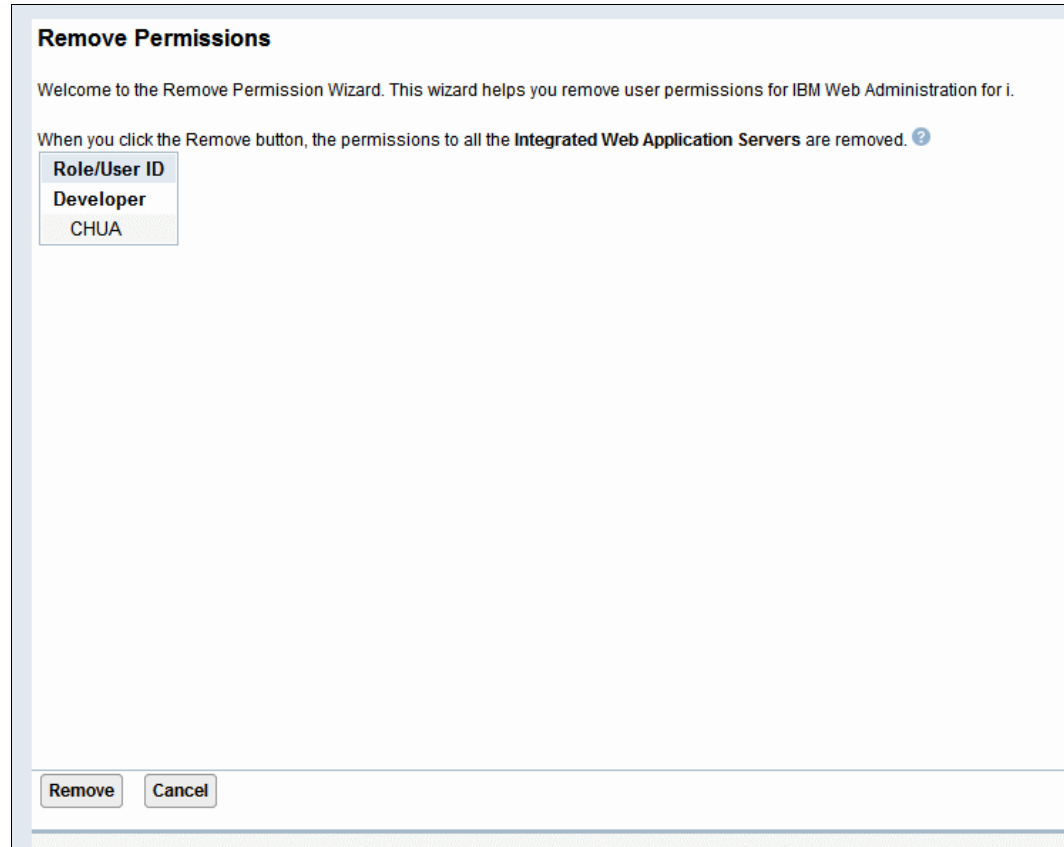


Figure 15-18 Remove permissions by Server view

- Remove all the permissions for a specific user (Figure 15-19). This is launched from the By User view on Manage Permissions panel when choosing a specific user.

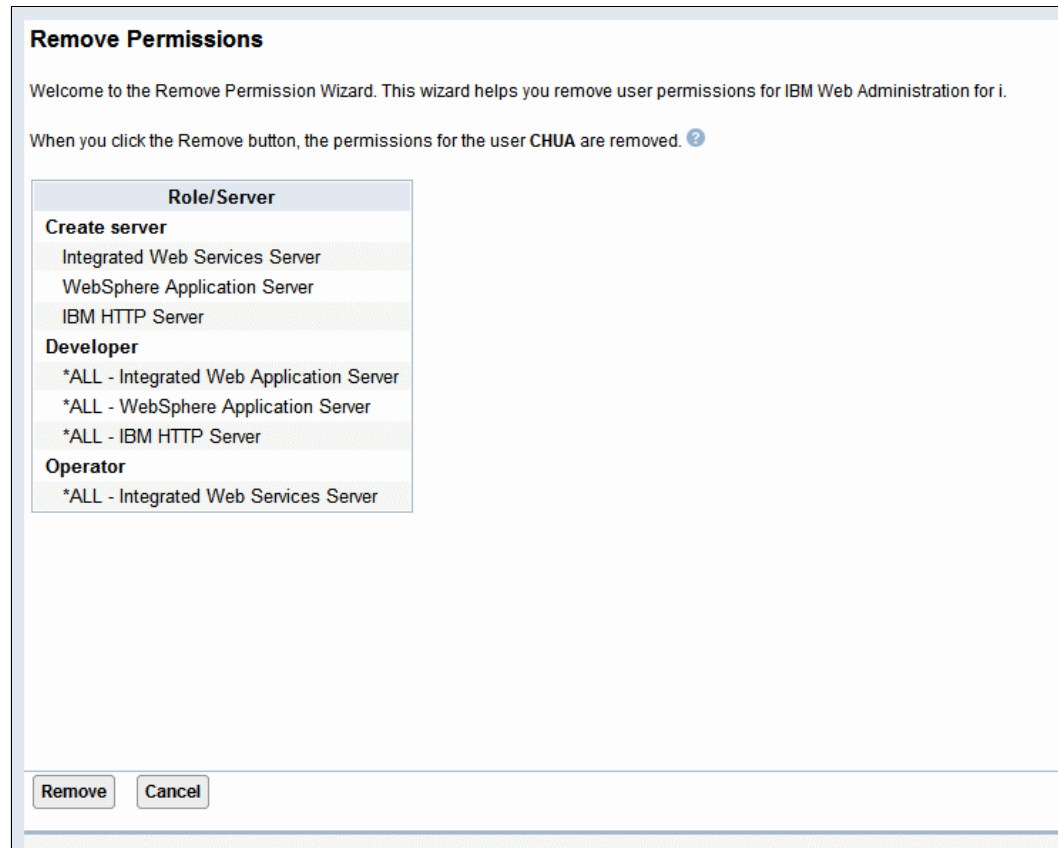


Figure 15-19 Remove Permissions by User view

Click **Remove** to remove the permissions. This operation cannot be undone.

15.3.3 Web Performance Advisor

Web Performance Advisor (Figure 15-20 on page 467) provides a consolidated place to review, evaluate, and modify the performance attributes for your web environments on the IBM i. This tool is a great first stop to ensure you are running at a reasonable set of performance related settings. It has been updated to include the latest performance setting 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 currently using Web Application server V6.0 and upgrading to IBM i 7.1 you need to update the web performance profile. The Classic Java virtual machine is not available for IBM i 7.1. If your Application Server installation is enabled to use Classic, it is not operational until it is enabled to use IBM Technology for Java Virtual Machine. see 15.2.2, “Java for IBM i 7.1” on page 443 for details.

IBM Web Administration for i

Setup

Manage

Advanced

Related Links

All Servers | HTTP Servers | Application Servers

Stopped

Server: IWAIAS - 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 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
<input checked="" type="radio"/>	IWAIAS	V7.1 (int app svr)	<div>Unknown - Web environment is not running</div>
<input type="radio"/>	IWAIAS	Apache-HTTP/Apache/2.2.11 (i5)	<div>Improvements possible</div>

Manage attributes

Export performance profile

Import performance profile

Cancel

Generate

Figure 15-20 Web Performance Advisor

Chapter 15. Web enablers 467



Application development

This chapter discusses the following topics:

- ▶ 16.1, “High-level programming languages” on page 470
- ▶ 16.2, “Control language CL” on page 479
- ▶ 16.3, “PHP” on page 482
- ▶ 16.4, “Lotus products for IBM i” on page 484
- ▶ 16.5, “Native archive and un-archive API support” on page 484
- ▶ 16.6, “IBM Toolbox for Java JDBC” on page 487
- ▶ 16.7, “Application Runtime Expert for i” on page 488

16.1 High-level programming languages

In this chapter, we will highlight changes to the programming languages C/C++, RPG, and COBOL, which are well-established programming languages for a large variety of applications.

16.1.1 C/C++

For C/C++, the following additions and modifications are made in IBM i 7.1:

- ▶ ILE C/C++ predefined macros
- ▶ `do_not_instantiate` and `namemanglingrule` pragma
- ▶ Control Language Command options `DBGENCKEY` and `DECFLTFND`

16.1.2 Predefined macros

These can be grouped either for C or for C++. Most of them are new, others have been modified.

C macros

- ▶ `__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

- ▶ `__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 XL C++ compiler upon which the ILE C++ compiler is based. It returns an integer representing the compiler version. For example, using the IBM i 7.1 compiler with the `TGTRLS(*CURRENT)` compiler option, `__IBMCPP__` returns the integer value 900. 900 means the ILE C++ compiler is based on the XL 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(*RTTIALL)` or `OPTION(*RTTICAST)` compiler option is specified.
- ▶ `__RTTI_TYPE_INFO__` can be defined when the `OPTION(*RTTIALL)` or `OPTION(*RTTITYPE)` compiler option is specified.

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)` or `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. Defined when the `LANGVL(*EXTENDED)` compiler option is in effect.
- ▶ `__DIGRAPHS__` indicates support for digraphs.

16.1.3 Pragmas

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.

Name mangling or name decoration is a technique used to solve various problems caused by the need to resolve unique names for programming entities. It allows you to encode additional metadata information in the name of a function, structure, class, or another datatype 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 as long as 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 with respect to the mangling of `cv`¹-qualifiers in function parameters. It allows you to control whether top-level `cv`-qualifiers are mangled in function parameters or whether intermediate-level `cv`-qualifiers are to be considered when the compiler compares repeated function parameters for equivalence.

16.1.4 Compiler options

With the `DBGENCKEY` compiler option, you can specify the encryption key to be used to encrypt the program source that is embedded in debug views. The length of the key can be between 1 byte and 16 bytes. A key of length 1 byte to 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, it is up to the user to 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, it is recommended that the key is made of characters from the (EBCDIC) invariant character.

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. The possible values for this option are:

- ▶ `*HALFEVEN`,

¹ `cv` stands for *constant* and *volatile*

This is the default setting. It rounds to the nearest value, but, when in a tie, it will choose an even number. For example, 5.22 rounds to 5.2, 5.67 rounds to 5.7, 5.55 and 5.65 round to 5.6.

► ***DOWN**

This value rounds toward zero, or truncate the result. For example, 5.22 rounds to 5.2, 5.67 rounds to 5.6, 5.55 rounds to 5.5, 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, 5.65 rounds to 5.7.

► ***HALFUP**

This value rounds toward to the nearest value and will, in a tie, round away from zero. For example, 5.22 rounds to 5.2, 5.67 rounds to 5.7, 5.55 rounds to 5.6, 5.65 rounds to 5.7.

► ***HALFDOWN**

This value rounds toward to the nearest value and will, in a tie, round toward zero. For example, 5.22 rounds to 5.2, 5.67 rounds to 5.7, 5.55 rounds to 5.5, 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, 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, 5.65 rounds to 5.7.

16.1.5 ILE COBOL

The ILE COBOL language has been enhanced with the following functions:

► **COMPUTATIONAL-5 or COMP-5 now supported**

This is a native binary data type now 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. This support will enhance portability to or from COBOL on other IBM platforms and operating systems. Table 16-1 shows the equivalent SQL data types for the COBOL COMP-5 data type:

Table 16-1 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

► **Ability to specify a non-numeric literal on the VALUE clause for a national data item.**

► XML GENERATE performance improvements and PROCESS options

Performance improvements have been made for XML GENERATE when the APPEND option is specified. Users who have a large number of data records to be appended into a data structure or into a stream file will benefit from these changes. The improvements include the addition of new PROCESS statement parameter XMLGEN with option values:

– NOKEEPFILEOPEN / KEEPFILEROPE

You specify KEEPFILEROPE 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

You specify ASSUMEVALIDCHARS to have XML GENERATE bypass the checking for special characters (less than "<", greater than ">", ampersand "&", and the single and double quote symbols), and for characters not supported by XML that requires being generated as hexadecimal. Otherwise normal checking is done with the default NOASSUMEVALIDCHARS.

► A new CRTBNDCBL/CRTCBLMOD parameter is added to support the encryption of the listing debug view. DBGENCKEY specifies the encryption key to be used to encrypt program source that is embedded in debug views.

► Larger program support is enabled by the CRTBNDCBL / CRTCBLMOD OPTIMIZE parameter, which now supports a new *NEVER option value. This value allows larger programs to compile by not generating optimization code for the program. PROCESS statement option NEVEROPTIMIZE is also added.

► The storage model for a program/module can now be specified using the new CRTBNDCBL/CRTCBLMOD parameter STGMDL with the option values:

– *SNGVL

This value specifies that the program/module is to be created with single-level storage model

– *TERASPACE

This value specifies that the program/module is to be created with teraspace² storage model

– *INHERIT

This value specifies that the program/module is to inherit the storage model of its caller

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 into the QILETS activation group. For all other storage models, the program is activated into the QILE activation group when it is called.

² 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 unallocated areas in between. Teraspace exists no longer than the time between job start and job end. A teraspace is not a space object. This 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.

- ▶ 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

16.1.6 ILE RPG

The ILE RPG language has been enhanced with the following functions:

- ▶ Sort and search data structure arrays: you can sort and search data structure arrays using one of the subfields as a key as shown in Example 16-1:

Example 16-1 Sort and search a data structure array

```
//Sort the custDS array by the amount_owing subfield
SORTA custDS(*).amount_owing;
// Search for an element in the custDs array
// where the account_status subfield is "K"
elem = %LOOKUP("K" : custDs(*).account_status);
```

- ▶ An array can be sorted ascending using SORTA(A) and descending using SORTA(D). The array cannot be a sequenced array (ASCEND or DESCEND keyword). See Example 16-2.

Example 16-2 Sort an array in descending order

```
//Sort the salary in descending order
SORTA(D) salary;
```

- 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 16-3.

Example 16-3 %SCANRPL

```
// Replace NAME with 'Tom'
string1 = 'See NAME. See NAME run. Run NAME run.';
string2 = %ScanRp1('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.

- As shown in Example 16-4, you can use the ALIAS keyword on a Definition specification to indicate that you want to use the alternate 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 alternate names for LIKERECD data structures defined from the records of the file.

Example 16-4 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;

```

- To obtain faster return values, you can now define a procedure with the RTNPARM keyword (see Example on page 477) 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 16-5 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 significantly improved by defining the procedure with this keyword.

The impact on performance due to 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 one 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 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 will require additional automatic storage (even if MYCALLER does not actually call procedure MYPROC at run-time). In certain cases, procedure MYCALLER will not compile due to 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.

Notes:

1. The additional parameter is passed as the first parameter.
2. The %PARMS and %PARMNUM built-in functions include the additional parameter in the parameter count. When the RTNPARM keyword is specified, the value returned by %PARMNUM is one higher than the apparent parameter number.
3. When calling 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.
4. 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 as though there is an additional first parameter passed by reference with the same type as the RPG return value.
5. 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 passed by reference with the same type as the RPG return value.
6. When RTNPARM is specified for the procedure, the maximum number of prototyped parameters is 398.
7. The RTNPARM keyword is not allowed for a Java method call.

- The %PARMNUM(parameter_name) built-in function returns the ordinal number of the parameter within the parameter list (see Example 16-6). The operand for %PARMNUM is the name of a parameter defined as part of a procedure interface. It is especially important to use this built-in function when a procedure is coded with the RTNPARM keyword.

Example 16-6 %PARMNUM built-in function

```
D                                pi
D  name                          100a  const varying
D  id                            10i  0  value
D  errorInfo                      likes(errs_t)
D                                options(*nopass)
/free
// Check if the "errorInfo" parameter was passed
if %parms >= %parmnum(errorInfo);
```

Notes:

1. A parameter defined using a *ENTRY PLIST cannot be specified as the operand for %PARMNUM.
2. The parameter must be specified the same way it appears in the procedure interface parameter list. If the parameter is an array, an index cannot be specified. If the parameter is a data structure, a subfield cannot be specified. If the parameter is a file, a record format cannot be specified.
3. If the RTNPARM keyword is coded for a procedure, the return value is handled as an additional first parameter. The other parameters have a number one higher than the apparent number. For example, if a procedure defined with RTNPARM has two parameters P1 and P2, %PARMNUM(P1) will return 2 and %PARMNUM(P2) will return 3.

- ▶ 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 only intended to be used as an exit program or as the command-processing program for a command
 - A program that is only intended to be called from another programming language
 - A procedure that is not exported from the module
 - A procedure that is exported from the module but only intended to be called from another programming language
- ▶ You can now pass any type of string parameter because an implicit conversion is done for string parameters passed by value or by read-only reference as shown in Example 16-7. For example, a procedure can be prototyped to have a CONST UCS-2 parameter, and character expression can be passed as a parameter on a call to the procedure. This enables you to write a single procedure with the parameters and return value prototyped with the UCS-2 type. To call that procedure, you can pass any type of string parameter, and assign the return value to any type of string variable.

Example 16-7 Passing a string parameter

```
// The makeTitle procedure upper-cases the value
// and centers it within the provided length
alphaTitle = makeTitle(alphaValue : 50);
ucs2Title = makeTitle(ucs2Value : 50);
dbcsTitle = makeTitle(dbcsValue : 50);
```

- ▶ There are two new options (also available through a PTF for IBM i 6.1) for *XML-INTO*:
 - `datasubf`
The `datasubf` option allows you to name a subfield that will receive the text data for an XML element that also has attributes.
 - `countprefix`
The `countprefix` 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 set by the XML-INTO operation.

- ▶ You can create RPG modules and programs to use the teraspace storage model or to inherit the storage model of their caller. With the teraspace storage model, the system limits regarding automatic storage are significantly higher than for the single-level storage model. There are limits for the amount of automatic storage for a single procedure and for the total automatic storage of all the procedures on the call stack.

You can use the storage model (STGMDL) parameter on the CRTRPGMOD or CRTBNDRPG command, or use the STGMDL keyword on the Control specification. If you specify *TERASPACE, the program or module uses the teraspace storage model. If you select *SNGLVL, it uses the single-level storage model. With *INHERIT, the program or module inherits the storage model of its caller.

- ▶ At the same time, the ACTGRP parameter of the CRTBNDRPG command and the ACTGRP keyword on the Control specification have changed: the default value is *STGMDL instead of QILE. ACTGRP(*STGMDL) specifies that the activation group depends on the storage model of the program: when the storage model is *TERASPACE, ACTGRP(*STGMDL) is the same as ACTGRP(QILETS). Otherwise, ACTGRP(*STGMDL) is the same as ACTGRP(QILE).

The change to the ACTGRP parameter and keyword does not affect the default way the activation group is assigned to the program. The default value for the STGMDL parameter and keyword is *SNGLVL, so when the ACTGRP parameter or keyword is not specified, the activation group of the program will default to QILE as it did in prior releases.

- ▶ You can use the ALLOC keyword on the Control specification to specify whether the RPG storage management operations in the module will use teraspace storage or single-level storage. The maximum size of a teraspace storage allocation is significantly larger than the maximum size of a single-level storage allocation.
- ▶ When a module's listing debug view is encrypted, the listing view can only be viewed during a debug session when the person doing the debugging knows the encryption key. This enables you to send debuggable programs to your customers without enabling your customers to see your source code through the listing view. Use the DBGENCKEY parameter on the CRTRPGMOD, CRTBNDRPG, or CRTSQLRPGI command to enable encryption.
- ▶ Rational OpenAccess RPG Edition is supported. Open Access is a full-featured version of the limited support provided by RPG SPECIAL files. It provides a way for RPG programmers to use the simple and well understood RPG I/O model to access devices that are not directly supported by RPG. See 17.5, "Rational Open Access: RPG Edition" on page 519 for a detail explanation.

16.2 Control language CL

This section covers

- ▶ New workload capping commands
- ▶ Retrieve CL source support for ILE CL
- ▶ Longer integer CL variables for ILE CL
- ▶ Showing DO and SELECT nesting levels in 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. Workload capping allows you 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.

The new workload capping commands include:

- ▶ ADDWLCGRP (Add Workload Capping Group) command
- ▶ ADDWLCPRDE (Add Workload Capping Product Entry) command
- ▶ CHGWLCGRP (Change Workload Capping Group) command
- ▶ DSPWLCGRP (Display Workload Capping Group) command
- ▶ RMVWLCGRP (Remove Workload Capping Group) command
- ▶ RMVWLCPRDE (Remove Workload Capping Product Entry) command

For detailed information about the above commands refer to the following web page:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=%2Frbam6%2Frbam6whatsnew.htm>

16.2.2 Retrieve CL source support for ILE CL

Developers can move to Integrated Language Environment® ILE without fear of not being able to retrieve their source statements. IBM i 7.1 has included this enhancement to the ILE module. For completeness, even the DCLPRCOPT (Declare Processing Options) command has been updated to reflect that the ALWRTVSRC keyword now applies to ILE program as well.

The RTVCLSRC (Retrieve CL Source) 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 CRTCLMOD (Create CL Module) command or the CRTBNDCL (Create Bound CL program) command specifying *YES for the ALWRTVSRC parameter. The module which contains the CL source to be retrieved can be a module *MODULE object or a module within a 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 example retrieves the CL source from module MOD1 in 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 located 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 command. This capability is useful when calling API programs and API procedures that define 8 byte integer fields in input or output structures.

Note: LEN(8) can only be specified if the CL source is compiled with the CRTCLMOD (Create CL Module) or the CRTBNDCL (Create Bound CL Program).

16.2.4 Showing DO and SELECT nesting levels in compiler listing.

You can specify a CL compiler option to see the nesting level for all types of DO commands and SELECT commands.

CL source programs contain DO commands 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(*DOSLTVL) and the CRTCLPGM (Create CL Program) command or the CRTCLMOD or the CRTBNDCL.

This compiler option adds a new column to the compiler listing which shows the nesting levels for the following elements:

- ▶ Do DO
- ▶ Do For DOFOR
- ▶ Do Until DUNTIL
- ▶ Do While DOWHILE
- ▶ SELECT

If you do not want to see this nesting level information, you can specify *NODOSLTVL 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 command. 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 being stored with the module *MODULE or ILE program *PGM object. To see the listing view during debug 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 located 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 at a later time 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 running on IBM i.

PHP is widely used for content management, customer relationship management, database access, ecommerce, forums, blogs, wikis, and other Web-based applications.

Zend and IBM have partnered to deliver Zend Solutions for IBM i, a complete PHP development and production environment solution for the IBM i platform.

Zend Solutions for IBM i include:

- ▶ 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 Zend web page:

<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 runtime 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:

- ▶ Preloaded on IBM i 6.1 and IBM i 7.1 starting April 2010
- ▶ Includes extensions and a toolkit that:
 - Enables PHP application to easily access DB2 for i data
 - Take 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, 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 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 commerical-grade web applications in highly reliable production environments.

Zend Server replaces Zend platform. It offers all the features 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 web page shows a comparison between the features 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 life cycle 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 comprehensive set of editing, debugging, analysis, optimization, database tools, and testing
- ▶ Toolkit support for easy integration with IBM i legacy applications and data
- ▶ Customizable and smart context sensitive templates for IBM i Toolkit functions
- ▶ Create/generate PHP Toolkit code quickly
- ▶ Call RPG/CL/COBOL program, execute CL command, retrieve Spooled file entries, access Data Area, Data Queue and User Space.

For more information about Zend products for IBM i refer to the following web pages:

- ▶ Zend and IBM i web page:

<http://www-03.ibm.com/systems/i/software/php/index.html>

- ▶ Zend Products for IBM i page:

<http://www.zend.com/en/solutions/modernize-ibm-i/ibm-i-products>

16.4 Lotus products for IBM i

IBM 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 smarter collaboration. The following Lotus Software products are supported on IBM i 7.1

- ▶ Lotus Domino - Domino 8.5.1 or later
- ▶ Lotus Enterprise Integrator - LEI 8.5.2 or later
- ▶ Lotus Sametime - Sametime 8.5.1 or later
- ▶ Lotus Quickr - Quickr 8.5 or later
- ▶ IBM Forms Server - version 4.0 or later
- ▶ Lotus Forms Server- version 3.5.1 FP2 or later
- ▶ Lotus Workflow - version 7.0 or later
- ▶ IBM Integrated Domino Fax for i5/OS V4R5

Note: Domino Fax for i5/OS V4R5 was withdrawn from worldwide marketing on January 14, 2011 and end of support is April 30, 2012

Only specific releases of these products are supported on IBM i 7.1. Before upgrading to IBM i 7.1 check the most current details about the product releases supported at the following web page:

http://www-03.ibm.com/systems/resources/systems_power_ibmi_lotus_releasesupport.pdf

16.5 Native archive and un-archive 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

16.5.1 QZIPZIP API

Multiple files and directories within IBM i can be compressed and packaged into a single archive file using the QZIPZIP() API.

QZIPZIP syntax is:

```
#include <qziputil.H>
void QzipZip(
    Qlg_Path_Name_T * fileToCompress,
    Qlg_Path_Name_T * compressedFileName,
    char * formatName,
    char * zipOptions,
    char * errorStruct)
```

Parameters

Table 16-2 shows the list of QZIPZIP API parameters.

Table 16-2 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 needs to 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 needs to 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. See section "ZIP00100 Format Description".
zipOptions	Input	This is a pointer that passes in the user's options to the QZIPZIP API in ZIP00100 format.
errorStruct	Input/Output	This 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 preceding that file
- *W data authority to the directory where the compressed file will be written

During the time this API reads a file for compressing it, the file will be locked and shared with reading only mode. The API will release the lock on the file after reading the file completely. If the file that is to be compressed is locked, then an error message will be sent. Further compression will be stopped

ZIP00100 Format Description

Table 16-3 shows the format for passing the user's options to compress files or directories.

Table 16-3 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

16.5.2 QZIPUNZIP API

The contents of the (.zip) archive file can be extracted using QZIPUNZIP to the target IBM i system.

QZIPUNZIP syntax is:

```
#include <qziputil.H>
void QzipUnzip(
    Qlg_Path_Name_T * compressedFileName,
    Qlg_Path_Name_T * dirToPlaceDecompFiles,
    char * formatName,
    char * unzipOptions,
    char * errorStruct)
```

Parameters

Table 16-4 shows the list of QZIPUNZIP API parameters.

Table 16-4 Parameters for QZIPZIP API

Name	Type	Description
CompressedFileName	Input	The name of the archive file that is to be decompressed. The path name needs to be in Qlg_Path_Name_T structure.
dirToPlaceDecompFiles	Input	The directory to place the contents of the archive file. The path name needs to be in Qlg_Path_Name_T structure.
formatName	Input	The format name to pass the user's options for decompressing an archive file. See section "UNZIP100 Format Description" for description of this format.
unzipOptions	Input	This is a pointer that passes in the user's options to the QZIPUNZIP API in UNZIP100 format.
errorStruct	Input/Output	This 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 perceding that file
- *W data authority to the directory where the decompressed file will be written

During the time this API reads a compressed archive file for decompressing it, the file will be locked and shared for reading only. The API will release the lock on the file after reading the file completely. If the file that is to be decompressed is locked, then an error message will be sent. Further decompressing will be stopped.

UNZIP100 Format Description

Table 16-5 shows the format for passing the user's options to compress files or directories.

Table 16-5 Format description

Offset		Type	Field
Dec	Hex		
0	0	CHAR(10)	Verbose option
10	A	CHAR(6)	Subtree option

16.5.3 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.

QZIPUTIL service program has exported QZIPZIP and QZIPUNZIP APIs.

16.6 IBM Toolbox for Java JDBC

JDBC is an application programming interface (API) included in the Java platform that enables Java programs to connect to a wide range of databases.

The IBM Toolbox for Java JDBC driver allows you 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 done to IBM Toolbox for Java JDBC support for IBM i 7.1

16.6.1 JDBC 4.1

Java JDBC interface now supports several features of the latest JDBC 4.1 definitions with DB2 i.

This enhancement allows Java developers to continue to leverage the latest defined options for JDBC.

16.6.2 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.

16.6.3 Database metadata updates

Database metadata is obtained by calling methods of 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.

In order to provide backwards compatibility 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.

16.6.4 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 returned to readers is based on the currently committed version; that is, data prior to the start of the write operation.

This feature also implements a way to direct the database manager to wait for the outcome when encountering data in the process of being updated.

16.6.5 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 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, ResultSets returned from stored procedures or other queries containing 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`.

16.6.6 Long schema name support

IBM i 7.1 DBMS has added support for 128 byte schema names. IBM Toolbox for Java JDBC is also adding support for long schema names.

16.7 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 allows you to build and maintain an application knowledge base, one that can be used to automatically apply its knowledge to verify a system.

ARE provides a GUI that allows you to collect and verify a customized set of information, system settings, and attributes about:

- ▶ Applications
- ▶ IBM i System
- ▶ Runtime Environment

ARE collects the needed information and build 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 will be verified against the information stored in the template, and the results are documented in a set of reports. These reports give a clear and concise view of what has changed, or what is different.

Templates can be updated easily to include new information and settings from ARE GUI interface.

16.7.1 Deployment template

A deployment template represents the expected attributes and state of a deployment and its environment. Typically a deployment consists of application attributes as well as attributes of the environment the application run in. Users can customize different plugins by defining what information to include in the deployment template using the “Deployment Template Editor”. Each plugin is capable of verifying a different aspect of a deployment.

Common examples of a deployment are:

- ▶ IBM i products
- ▶ Application on IBM i
 - ISV application
 - WebSphere Application Server
 - Integrated Web Application Server
 - RPG application
 - Directory tree in IFS
- ▶ Custom selection of software products, system and environment information

The deployment template is used as 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

The Deployment Template Editor allows users to create, edit, import and export deployment templates. The editor enables users to customize their templates. Templates created, or edited can be used to verify systems using the Console.

16.7.2 Customize Plugins For Template

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 most critical to your applications. Once the template customization is completed, you can build it using 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

The File and Directory Attributes plugin verifies attributes such as existence, creation date, last modification date, size, CCSID. Attributes can be verified for files and directories in IFS, and objects in the Library file system. The precise file and directory attributes verified by this plugin are fully customizable.

This plugin can also verify that a symbolic link is truly a symbolic link and not a real file. This is useful to detect cases where a symbolic link has been replaced by an actual copy of the file it is supposed to be a link to.

File and Directory Authorities

The File and Directory Authorities plugin 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 plugin are fully customizable.

Configuration Files

The Configuration Files plugin 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 well as 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 16-1.

Edit Expected Value

Key: newKey

☐ Report only
☒ Verify this value [? Learn more...](#)

Value type: String Comparison: Equals

Expected value:
result1

Severity: Error ☐ Match case

OK Cancel

Description

Equals

The actual value must equal the expected value.

Examples

Actual value	Comparison	Expected value	Result
mystring	Equals	mystring	True
mystring	Equals	somestring	False

Figure 16-1 Edit Expected Key Value

Resource Collector

The Resource Collector plugin allow you the ability to collect and package files and objects from the target system into an archive file for additional review. This feature allow Application Runtime Expert to not only verify specific files and attributes, but can also be used to collect the data needed to completely review and possibly debug an issue on a system.

The plugin allows specific files and directories from IFS to be collected and packaged into a single archive file. It also allow you the ability 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 pre-requisites:

- IBM i Products
- PTFs

IBM i Products

This feature allows you to select specific IBM i products, from the list of all IBM i products. The selected products will be added to the list of products to verify in the template.

CHKPRDOPT (Check Product Option) CL command can be used during product verification for providing useful, product specific information. Select CHKPRODOPT check box to enable this feature, as shown in Figure 16-2.

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	<input checked="" type="checkbox"/>	Error
5770DG1	*BASE	V7R1M0	<input checked="" type="checkbox"/>	Error
5770SS1	3	V7R1M0	<input type="checkbox"/>	Error
5770SS1	12	V7R1M0	<input checked="" type="checkbox"/>	Error
5770SS1	30	V7R1M0	<input type="checkbox"/>	Error
5770SS1	33	V7R1M0	<input type="checkbox"/>	Error

Navigation: > < >> <<

Buttons: Previous Next Filter 7/11

Position to:

Collection name:

Buttons: OK Cancel

Figure 16-2 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 choose the Show all products supported by IBM i option, as shown in Figure 16-3.

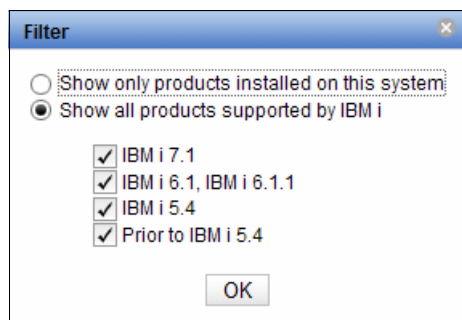


Figure 16-3 Filter IBM i products

PTFs

The PTFs plugin 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 verified is fully customizable. This plugin can verify a group PTF level meets both a minimum and recommended level.

Network

This feature allows you to verify a variety of TCP/IP configuration settings, network characteristics, and ports status:

- ▶ Network Configuration
- ▶ Ports

Network Configuration

The Network Configuration plugin verifies a variety of TCP/IP configuration settings and network characteristics. Much of the TCP/IP configuration that is verified is the configuration settings viewable using the CFGTCP CL command.

The primary purpose of this plugin 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.

Ports

The Port validation plugin verifies if one or more specified ports are active or free on the system. The plugin is also capable of verifying if one or more specified ports are restricted.

Ports on the system can be restricted using CFGTCP and selecting option 4 - Work with TCP/IP port restrictions. If an application requires a port that is already in use or has been restricted, the application may fail to start or not work correctly.

System Environment

This feature allows you to verify various system environments settings including network attributes, environment variables, user IDs:

- ▶ System values and Network Attributes
- ▶ Environment Variables
- ▶ User IDs

- ▶ Scripts and Commands
- ▶ SQL Query Verifier

System values and Network Attributes

The System Values and Network Attributes plugin verifies system values and network attributes are configured the way in which 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

The Environment Variables plugin 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, as well as their values, are an important part of the environment in which any application runs.

User IDs

The User IDs plugin 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 16-4.

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

Previous Next 7/15

Position to:

Collection name:

User IDs to verify [? Learn more...](#)

QDESADM

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 16-4 Verify User IDs

Scripts and Commands

The Scripts and Commands plugin allow you the ability to verify the results of QShell scripts or CL commands on a remote system. This plugin provides a very powerful and simple way to extend the base verification function provided by Application Runtime Expert for i. The results of each verification are customizable, including capability to specify the source of result, expected value, a range of expected values, and more

SQL Query Verifier

The SQL Query plugin 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 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 plugin customization:

- ▶ Custom Plugins
- ▶ Other Resources
- ▶ Plugin Configuration

Custom Plugins

Manage custom plugins that are included in the deployment template. Custom plugins are Java classes that can augment the verification performed by a deployment template. Anyone can write a custom plugin, and once the plugin is created it can be added to any deployment template.

A custom plugin plugs into the IBM Application Runtime Expert for i environment, and is run along with the other plugins selected and customized using the deployment template editor.

Other Resources

Manage other resources, which is anything that is not a custom plugin, that are included in the deployment template. Other resources are commonly needed in conjunction with a custom plugin. For example, a custom plugin may require several other Java classes for it to run. These additional Java classes can be added to the deployment template using the Other Resources page.

Plugin Configuration

There are several advanced plugin features that can be edited via the Plugin Configuration page. Additionally, any plugin that can be utilized by a template, including plugins that are not configured elsewhere in the GUI, can be configured and added to a template using the Plugin Configuration page.

16.7.3 Application Runtime Expert Console

Once 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 describe using the console to verify IBM i systems.

The console is a Web user interface which enables a system, or group of systems, to be verified using a deployment template which was created or imported using the deployment template editor.

Systems are verified by the following sequence of events:

- ▶ The console remotely log in the system
- ▶ The console invoke the ARE to perform the verification, using the specified deployment template
- ▶ The results of the verification goes back to the console so that they can be reviewed

Note: 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 may involve the inspection of many different files, authorities, PTFs, user profiles, and other items.

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 16-5.

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
<input checked="" type="checkbox"/>	9.5.168.119	amonsour	●●●●●●	sample ▾
<input checked="" type="checkbox"/>	9.172.23.77	amansour	●●●●●●	DB2_Extender_Template ▾
<input checked="" type="checkbox"/>	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 16-5 IBM Application Runtime Expert Console

Verifying Systems

Once 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 16-6.

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

Back

Stop all

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

Figure 16-6 System Verification Status

Additional details about the failure can be found by clicking on the ‘View log’ link in 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 16-6.

All three ARE reports (summary, detailed, and XML) are available via the link in the result column. You can also download all report in an archive file using “Download Archive” link.

Summary report

The summary report contains a summary of all problems detected during the verification of a system. Each row in the summary table contains the results for a specific plugin, such as the “Authority Verifier” or “System Value Verifier” plugins. The icon directly before the plugin name indicates the highest severity problem that was found by that plugin. The other items in each row indicate the number of attributes verified by the plugin, and the number of problems 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 16-7 and Figure 16-8 show an example 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 16-7 Summary Report

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 16-8 Summary Report

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.

Note: XML reports includes information about how to fix detected problems. ARE core uses the XML report as a guide for automatically fixing detected problems.

Automatic Fixes

IBM Application Runtime Expert for i offers additional important feature which is the ability to automatically fix problems 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. This is a very convenient feature that allows you to solve certain problems without ever having to go and log into a different system

The following list shows some of the problems that ARE can automatically fix:

- ▶ Authority: This 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 a symbolic link is missing, it has the capability to re-create the link.

Note: It is important to understand that only certain types of detected problems can be fixed directly from the console.

Figure 16-9 shows a summary report in which “Authority Verifier” plugin 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 16-9 Summary report fix action


Clicking on “Fix action” link, opens a new page that summarizes all of the problems detected by that plugin that can be fixed directly from the console. You can select which problems to fix using the check box that precedes each problem description, and then click “Fix” button, as shown in Figure 16-10.

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

Select all Deselect all

▼ Processed problems

Result	Description
	This table is empty.

Fix Back

Figure 16-10 Fix plugin problem

The console begins the process of fixing the selected problems on the target system. Once 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 "areFix.sh" script section in the following document:

http://www-03.ibm.com/systems/resources/systems_i_are_script_interfaces.pdf

16.7.4 Application Runtime Expert Requirements

IBM Application Runtime Expert for i product (5733ARE) requires several IBM i products and PTFs to be installed on the system before installing ARE.

Pre-requisite Products

Following 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 installing ARE. Refer to the following link for up-to-date PTF requirements:

<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 latest PTF numbers refer to the following ARE support link:

<http://www-03.ibm.com/systems/power/software/i/are/support.html>

16.7.5 More Information

For more information on Application Runtime Expert refer to the redpaper:

REDP-4805-00: *Uncovering Application Runtime Expert*.



Rational Products

This chapter discusses the latest Rational Products for Power System Software and IBM i.

- ▶ 17.1, “IBM Rational Developer for Power Systems Software” on page 502
- ▶ 17.2, “Rational Team Concert” on page 508
- ▶ 17.3, “Rational Developer for i for SOA Construction” on page 513
- ▶ 17.4, “Rational Development Studio for i” on page 513
- ▶ 17.5, “Rational Open Access: RPG Edition” on page 519

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 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's 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 life cycle 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
- ▶ Migration to Rational Developer for Power System V8.0

17.1.1 RPG and COBOL Development Tools

Rational Developer for Power System offers the following RPG and COBOL Development Tools:

- ▶ Remote System Explorer
- ▶ iProjects
- ▶ Application Diagram
- ▶ Screen Designer
- ▶ Report Designer
- ▶ Integrated i Debugger
- ▶ IBM i Web Services and Java Tools

Remote System Explorer

The Remote System Explorer, shown in Figure 17-1, is an enhanced and more flexible workstation version of the Programming Development Manager (PDM). It is a workbench perspective that provides access to all development resources of your IBM i server.

Remote System Explorer allows effective management and organization of IBM i by the following features:

- ▶ 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

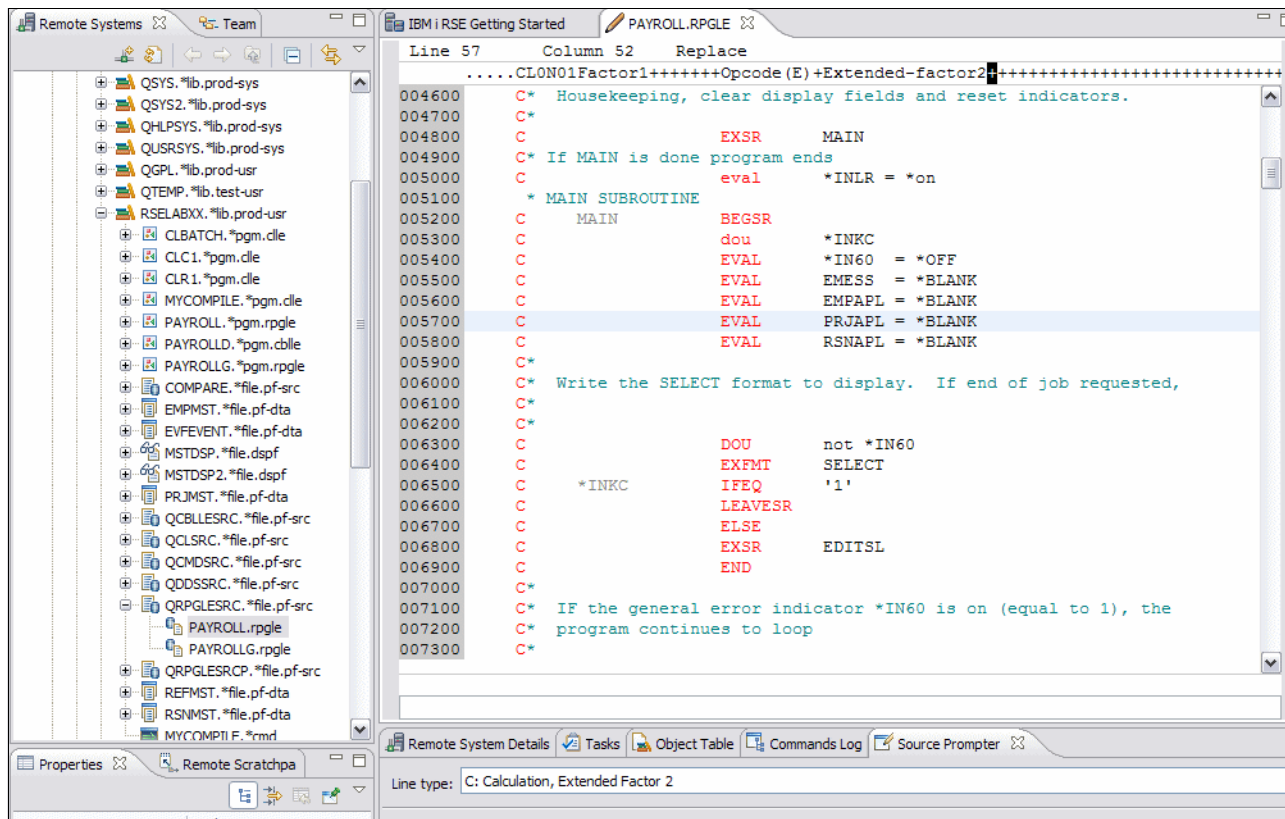


Figure 17-1 Remote System Explorer

iProjects

iProjects, shown in Figure 17-2, allow for disconnected development. A network connection is required only when code updates or build are needed, or when you need to view remote resources for a project.

In disconnected mode you work on files locally, and upload them back to the server once you have finished. While working in the 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.

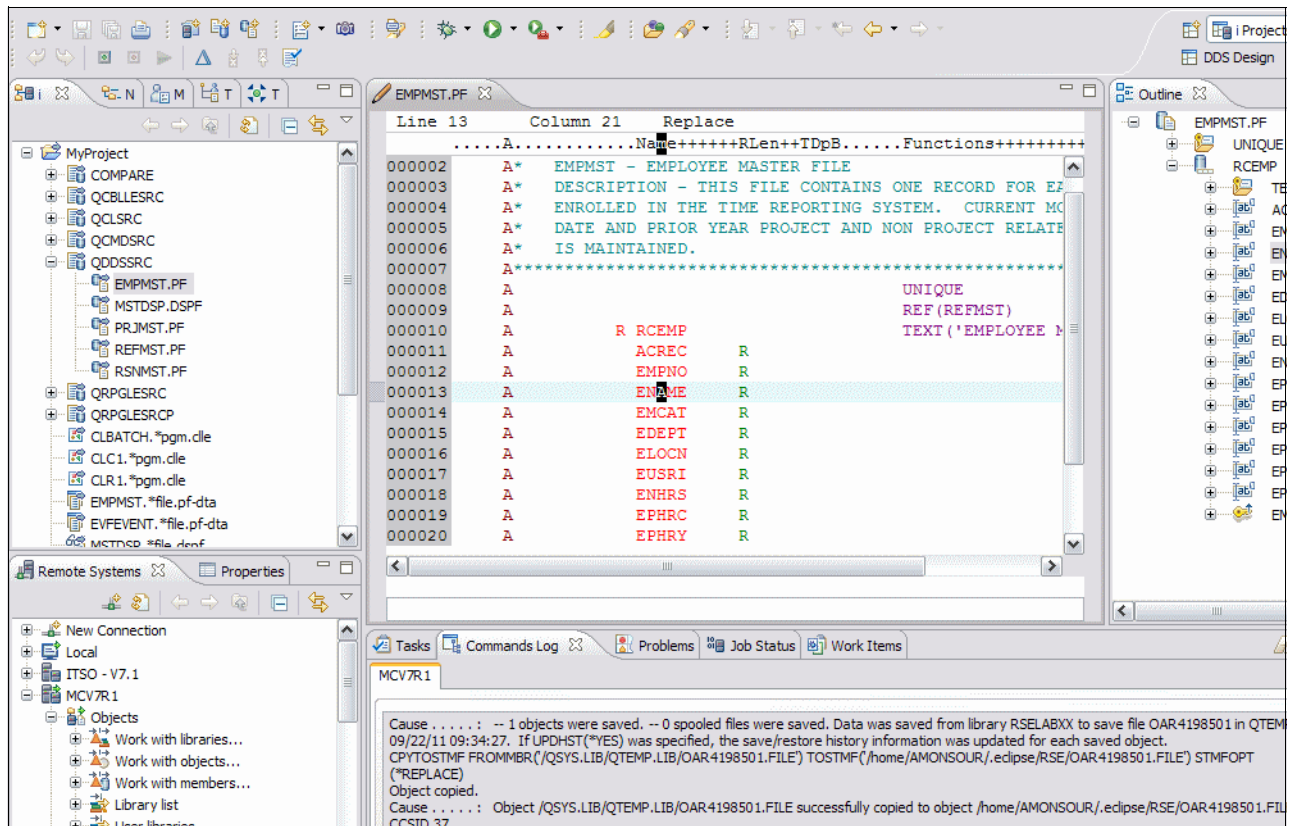


Figure 17-2 i Projects perspective

Application Diagram

The 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 source as input and displays a call graph showing subroutine and procedure call.

- ▶ Program Structure Diagram

This diagram takes a program and service program objects as input and displays the binding relationships between them as well as the modules bounded into 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.

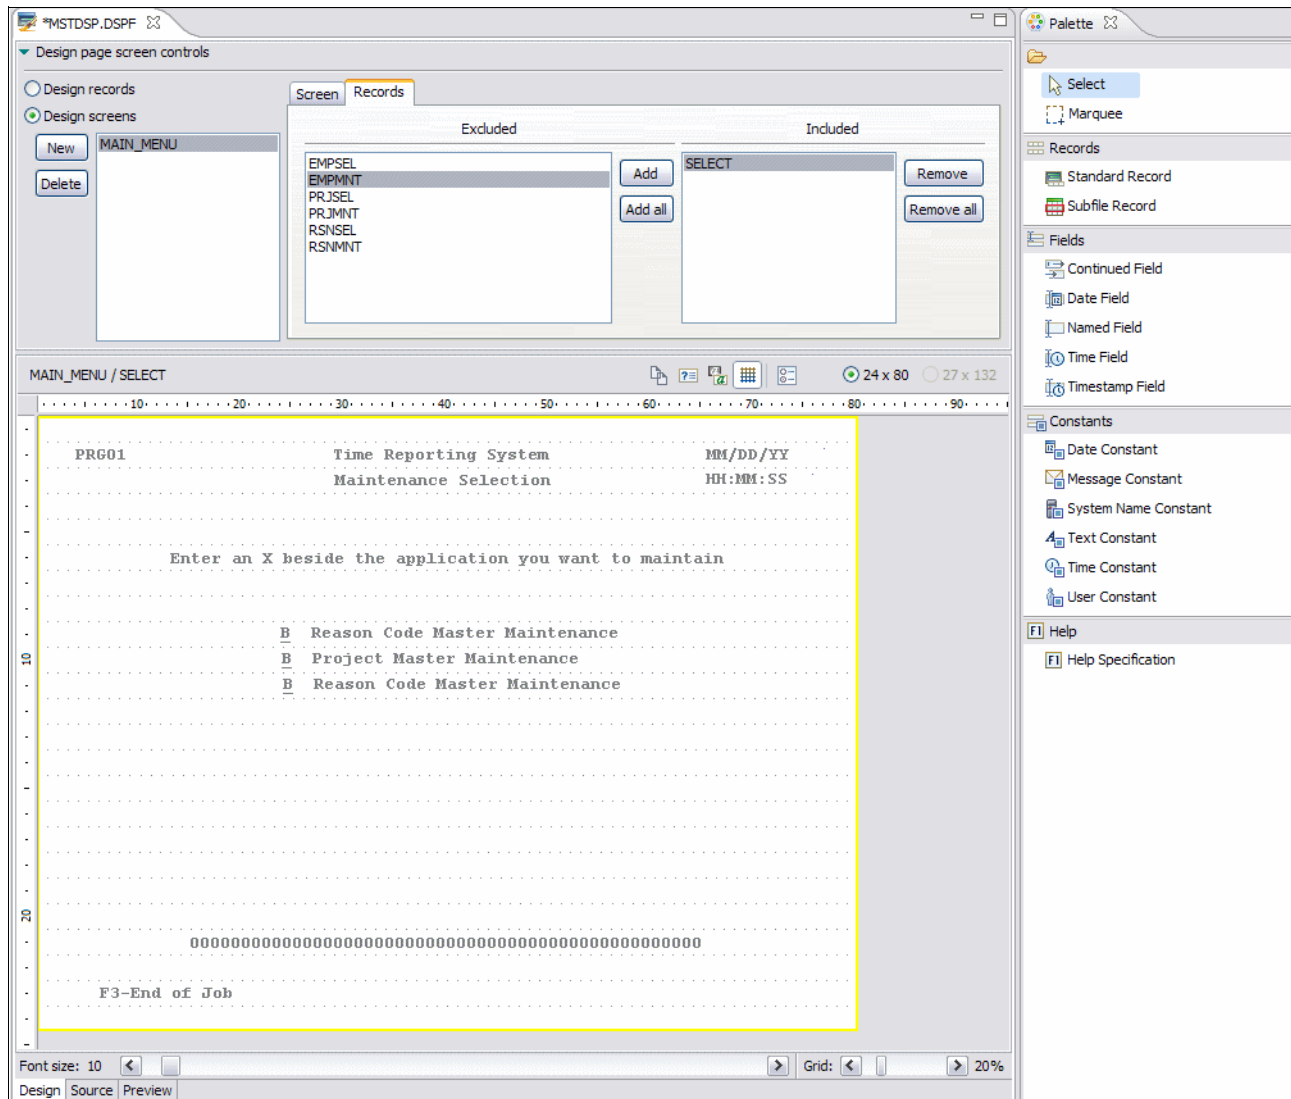


Figure 17-3 Screen designer perspective

Report Designer

The Report Designer provides the capability to graphically design and modify the content of DDS printer files. The Report Design page provides an integrated palette for easy access to design items.

The Report Designer lets you 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

Integrated i Debugger allows you to debug an application that is running on IBM i system. It provides an interactive graphical interface that makes it easy to debug and test you IBM i programs.

Integrated i Debugger provides many program debugging control features including breakpoints setup, watch variables, step through program instructions, examine variables, and examine the call stack.

IBM i Web Services and Java Tools

IBM i Web development tools give you the ability 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 residing on an IBM i server.

The Web Service 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.

Note: 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 System Software provides support for integration with Rational Team Concert client for IBM i.

When used in combination with IBM Power Systems compilers and Rational Team Concert, Rational Developer for Power Systems Software provides a comprehensive application development environment, including compilers, development tools, and collaborative application lifecycle management.

For detailed information about Rational Team Concert, and integration with Rational Developer for Power System Software refer to section 17.2, “Rational Team Concert” on page 508.

Note: The IBM Rational Team Concert 3.0 client product must be installed before installing 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.

Following 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 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 V8.0 (RAD SE).

This new enhancement is ideal for organizations that are developing solutions that are integrating applications written in RPG and COBOL on IBM i with Web Services and Web front-ends that leverage Java, Java EE, Web 2.0, SOA, or Portal.

Power tools allows developers to have all their development tools integrated into one work environment. Power tools also reduces the operation costs by deploying a single development workbench image to all developers.

Enhancements for AIX and Linux

For list of Rational Developer for Power Systems enhancements on AIX and Linux, refer to the following product announcement letter page:

http://www-01.ibm.com/common/ssi/ShowDoc.jsp?docURL=/common/ssi/rep_ca/4/897/ENUS211-344/index.html&lang=en

Version 8.0 Fix Packs

More information about any required Power Systems server PTFs can be found by doing the following steps from IBM Rational Developer for Power Systems Software:

1. Open the Remote System Explorer perspective by selecting select **Window → Open Perspective → Other → Remote System Explorer**
2. Creating a connection to your IBM i by expanding **New Connection → double-click on IBM i**
3. From Remote System Explorer **Expand Your Connection → Right-click Object → Verify Connection**, as shown in Figure 17-4

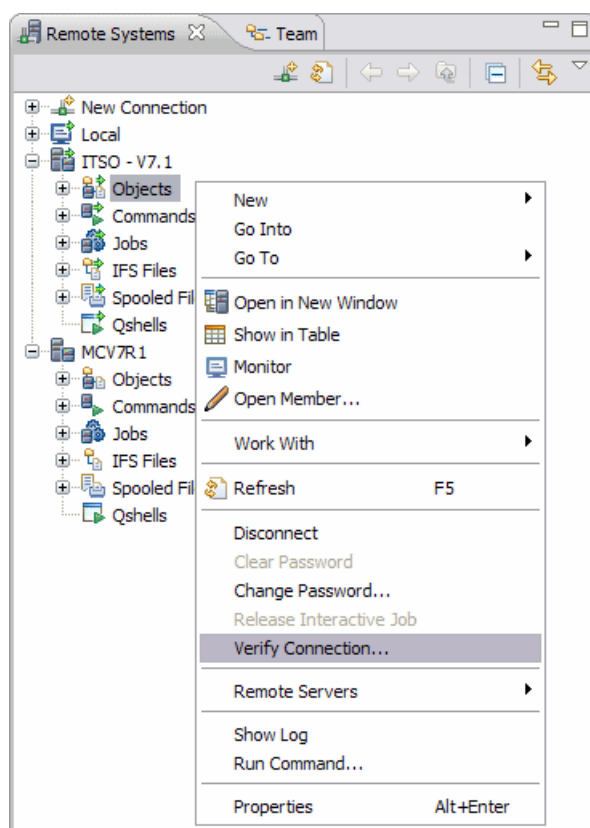


Figure 17-4 Verify connection

4. This displays a dialog that shows which required PTFs are already installed on the system and which ones are missing, as shown in Figure 17-5

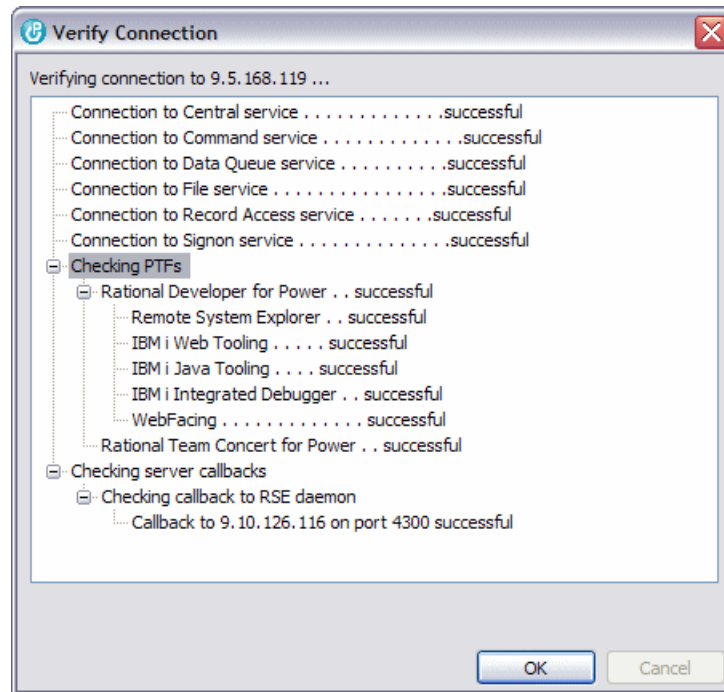


Figure 17-5 Verify PTFs

Refer to Rational Developer for Power Systems Software Support page for information about the latest product Fix Packs.

http://www-947.ibm.com/support/entry/portal/Overview/Software/Rational/Rational_Developer_for_Power_Systems_Software/

17.1.4 Migration to Rational Developer for Power System Software V8.0

You can migrate projects to Rational Developer for Power System V8.0 from previous versions. Migration process is seamless and transparent, you only need to use your existing workspace, and continue working with your projects with the new software version.

Migration to V8.0 can be done from the following products:

- ▶ Rational Developer for Power Systems Software: Migration is supported from V7.5.x and V7.6.x.
- ▶ Rational Developer for i: Migration is supported from V7.5.x

Migration of projects from earlier releases is not supported because of additional features that have been added.

17.2 Rational Team Concert

The Rational solution for Collaborative Life cycle Management (CLM) provides integrations across the Change and Configuration Management, Requirements Management, and the Quality Management Jazz™-based applications, to connect the work of analysts with development and test teams.

Rational Team Concert is the core of Rational CLM solution. Rational Team Concert (RTC) is a team collaboration tool that supports cross-platform development and features native hosting of the Jazz Team Server. RTC includes an integrated set of collaborative software delivery life cycle tools for development, including source control, change management, and build and process management.

RTC has an open, extensible architecture that support a broad range of desktop clients, IDE's, languages, and platforms, as shown in Figure 17-6.

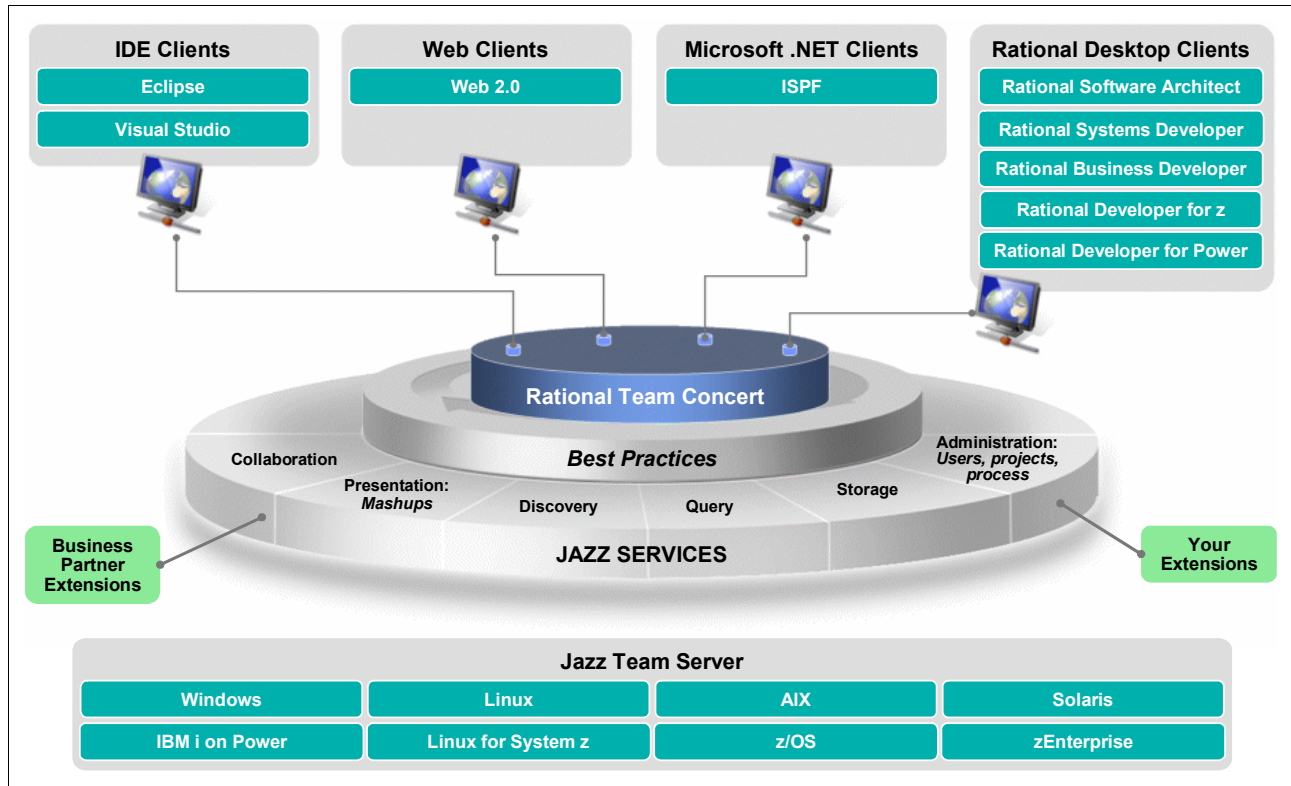


Figure 17-6 Rational Team Concert Architecture

17.2.1 Integration with Rational Developer for Power system

RTC 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 RTC with the individual development environment of Rational Developer for Power Systems Software.

Using RTC and Rational Developer for Power Systems Software together, software development teams can develop IBM i and AIX applications using the tools provided by Rational Developer for Power Systems Software and the planning, team collaboration, build, source control management, defect tracking, and deployment tools provided by RTC.

Using IBM i Projects perspective available with Rational Developer for Power Systems Software, RTC and Rational Developer for Power Systems Software work together so that you can share and modify files managed by Jazz-based source control, in addition to files on the remote IBM i system.

Installing and Configuring the integration

Additional Installation and configuration steps are required to integrate Rational Developer for Power Systems Software and RTC.

There are two ways that you can install the two products 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 System 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

Note: You can not install the RTC client integration for IBM i feature at this point, because the feature is not available if RTC is not present

- b. Install Rational Team Concert
- c. Start IBM Installation Manager and use the **Modify** option to add RTC client integration for IBM i feature to Rational Developer for Power Systems Software product, as shown in Figure 17-7.

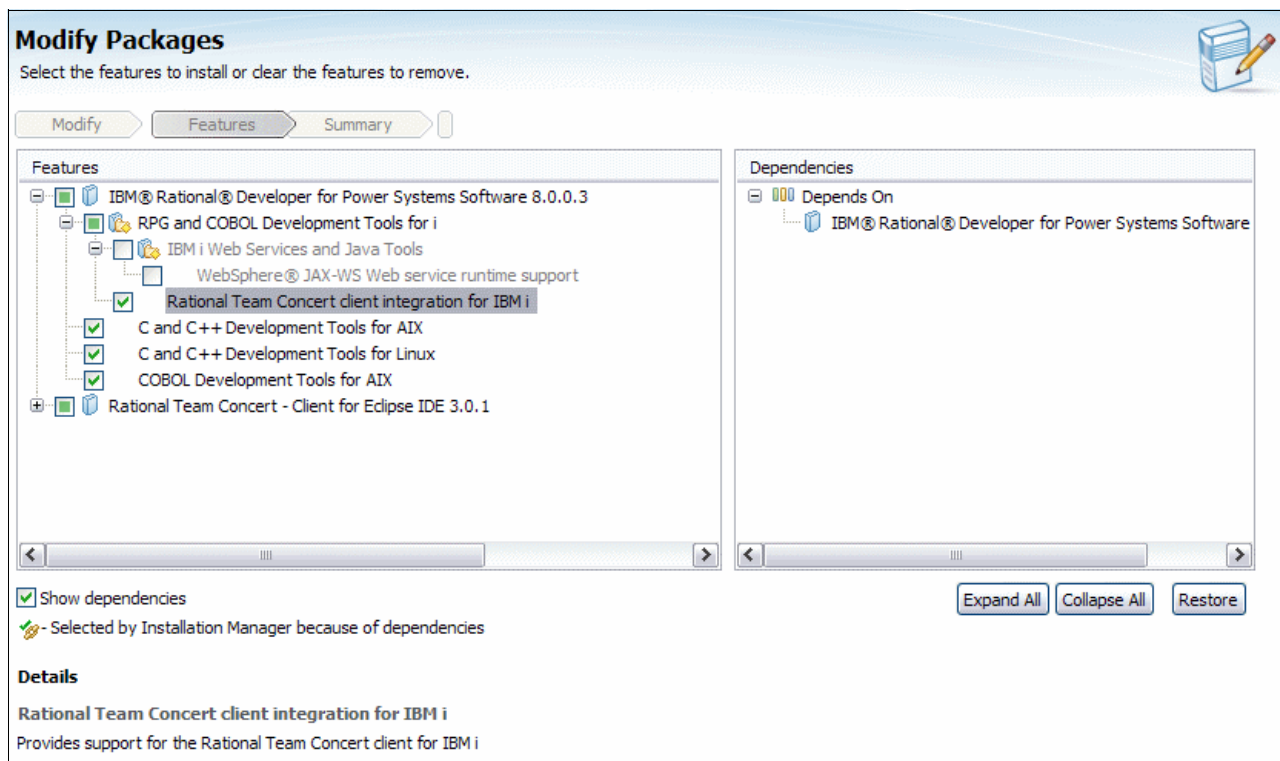


Figure 17-7 Rational Team Concert Client integration

Sharing i Projects in RTC

Sharing IBM i Projects in Rational Team Concert is no different than sharing any other type of Eclipse project. You can manage Power Systems Software source code with the IBM i projects feature.

To make your i project available to other team members:

1. From the context menu for your project, select **Team** → **Share Project** as shown in Figure 17-8

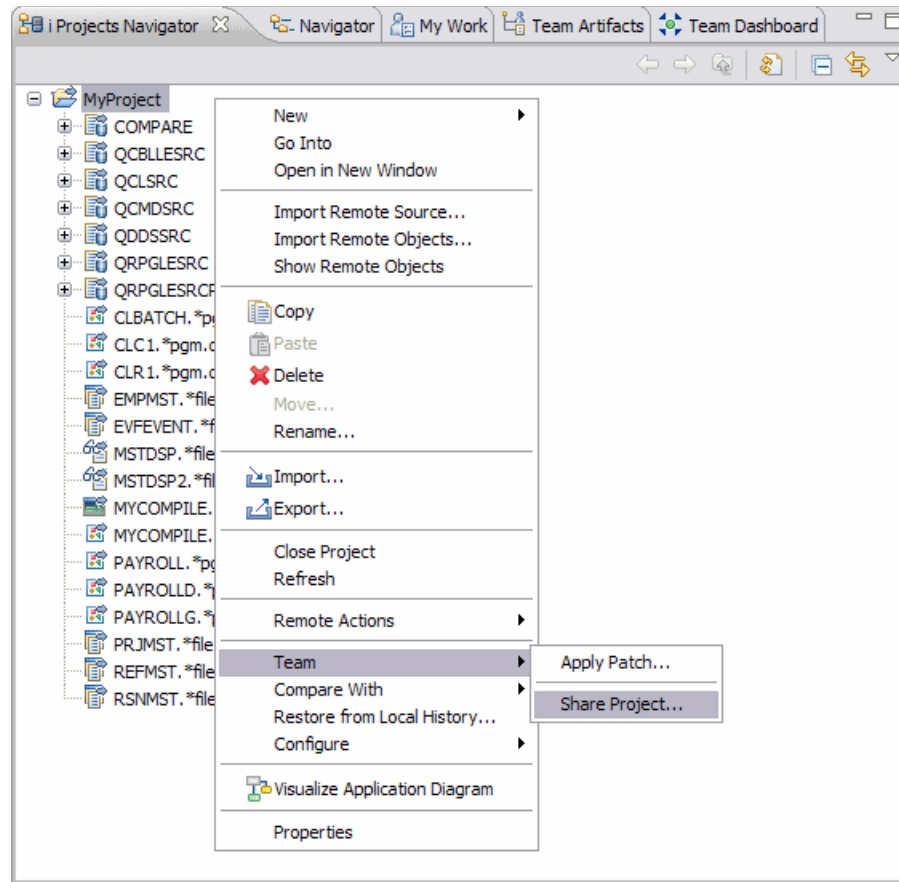


Figure 17-8 Share IBM i project

2. Select **Jazz Source Control**, as shown in Figure 17-9, and click **Next**

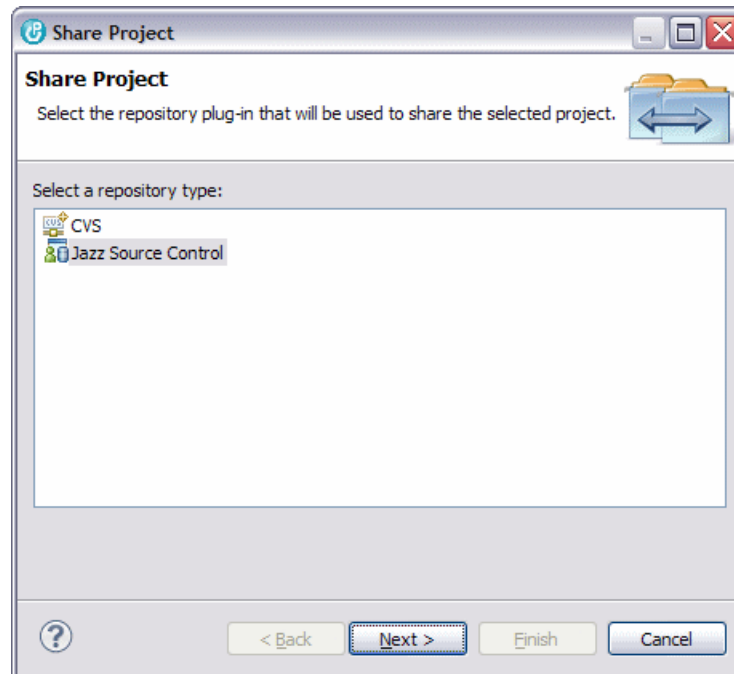


Figure 17-9 Jazz Source Control Repository

3. Specify the repository workspace to use
4. Click **Finish**.

Note: If your project contains any file types or folders that are ignored by Eclipse, the wizard prompts you to review them.

The project is now added to a component in a repository workspace. Changes that you make in the project are copied to the repository workspace when you check them in.

17.2.2 What's New in Rational Team Concert V3.0.1

Here are some highlights of what's new in 3.0.1:

- ▶ 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 System z¹
- ▶ New Lotus Connections Integration
- ▶ Enhanced RTC Client for Visual Studio IDE with CLM support
- ▶ Role based licensing with built in read access across RTC, RRC and RQM
- ▶ Cross product reporting and life cycle queries

¹ Available in RTC Developer for IBM Enterprise Platforms

- Flexible deployment options for CLM applications

17.3 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 developed on Eclipse)..

RDi SOA Version 8.0.3 combines the following products:

- IBM Rational Developer for Power Systems Software: RPG and COBOL Development Tools for i, Version 8.0.3
- IBM Rational Business Developer, Version 8.0.1.2
- IBM Rational Host Access Transformation Services Toolkit, Version 8.0

It accelerates IBM i development teams who need to:

- 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 end users with richer and more productive interface.
- Reuse existing applications and integrate them into automated business processes and new applications

RDi SOA Version 8.0.3 i helps you:

- 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.

Note: Although HATS Toolkit V8.0 is available as a free download, 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.

Refer to Rational Developer for i for SOA Construction web page for more information:

<http://www-01.ibm.com/software/awdtools/developer/rdisoa/>

17.4 Rational Development Studio for i

IBM Rational continues to have a tight collaboration with IBM Systems Technology Group (STG), Software group (SWG), and Research group to provide compilers that exploit the underlying hardware and operating systems for the Power architecture. The RPG, COBOL, C, and C++ ILE compilers for IBM i have been 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.

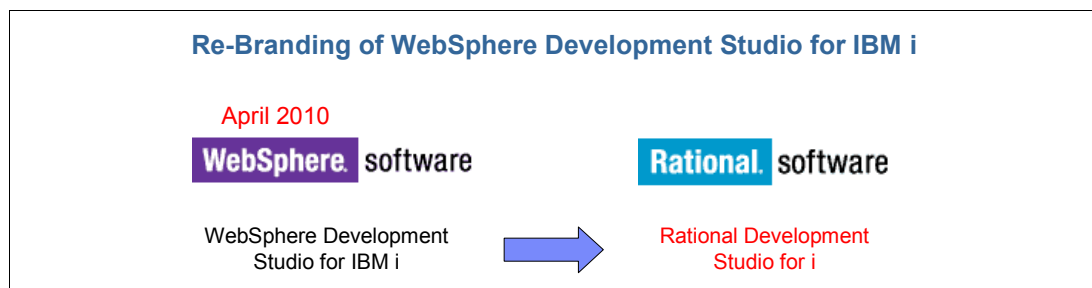


Figure 17-10 Rebranding of WebSphere Development Studio for IBM i to Rational Development Studio for i

Application Development ToolSet (ADTS) and Original Program Model (OPM) Compilers have been previously stabilized, meaning no futures enhancements are made. However, Rational Development Studio for i V7.1 does include enhancements for the ILE RPG and ILE COBOL compilers.

17.4.1 ILE RPG Compiler

The following list details the ILE RPG compiler enhancements available with Rational Development Studio for i V7.1:

- ▶ New options for XML-INTO (see Example 17-1 on page 514)
 - The new datasubf option enables XML-INTO to support XML elements in this form
`<emp type="regular" ID="13573">John Smith</emp>`.

Example 17-1 New options for XML-INTO

```

<emp type="regular" id="13573">John Smith</emp>
RPG coding to get the information with one XML-INTO operation:
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'

```

- ▶ New and updated built-in functions %SCANRPL, %LEN, %PARMNUM (see Example 17-2)
- The %SCANRPL built-in function replaces all occurrences a string with another string.
- %LEN(*MAX) simplifies getting the maximum length of the data part.
- The %PARMNUM built-in function returns a parameter's position in the parameter list.

Example 17-2 New and updated built-in functions %SCANRPL, %LEN, %PARMNUM

Soft-coding the parameter's number makes the code easier to read and maintain.

```
CEEDOD (2 : more parms);           // hard to understand
CEEDOD (%PARMNUM(city) : more parms); // better
```

- Implicit unicode conversion for parameters (see Example 17-3)
 - Implicit CCSID conversion is now supported for prototyped parameters passed by VALUE and by read-only reference (CONST).
 - This reduces the number of code changes that have to be made when a database field is changed from alphanumeric or DBCS to unicode (UCS-2 or UTF-16).

Example 17-3 Implicit unicode conversion for parameters

In this example, there is only a “makeTitle” procedure with a UCS-2 parameter and return value. If the passed parameter is alpha or DBCS, it will be converted to UCS-2 on the call. The procedure will work with the UCS-2 parameter and return a UCS-2 value. This returned value can then be converted on assignment to alpha or DBCS, if necessary.

```
// makeTitle() upper-cases the parameter
// and centers it within the provided length
```

```
alphaTitle = makeTitle(alphaValue : 50);
ucs2Title = makeTitle(ucs2Value : 50);
dbcsTitle = makeTitle(dbcsValue : 50);
```

- Sort and search data structures (see Example 17-4)

Example 17-4 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);
```

- Performance when returning large values (see Example 17-5)
 - RTNPARM keyword greatly improves performance when a procedure returns a large value.

Example 17-5 Performance when returning large values

Performance improvement is especially noticeable when the prototyped return value is a large varying length value

D center	pr	100000a	varying
D			rtnparm
D text		50000a	const varying
D len		10i 0	value

```

D title          s          100a   varying
/free
  title = center ('Chapter 1' : 60);

```

- Optional prototypes (see Example 17-6)
 - If a program or procedure is not called by another RPG module, it is optional to specify the prototype.

Example 17-6 Optional prototypes

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

```

H main(hello)
P hello          b
D               pi          extpgm('HELLO')
D  name          10a   const
/free
  sayHello();
...

P sayHello       b
/free
  dsply ('Hello ' + name);

```

- Support for ALIAS names in externally-described data structures
- When ALIAS is specified, RPG will use the ALIAS name instead of the 10-character standard name.
- Supported on F specs for any file that will not have Input or Output specs generated. Used for LIKERECD data structures.
- Supported on D specs for any externally-described data structure.
- Freeze SEU syntax-checking at V6R1 level
 - Freeze ILE RPG checkers for SEU at the V6R1 level.
 - For customers using SEU to edit ILE RPG source, the syntax checkers will not recognize any features added after V6R1.
- Encrypted debug view (see Example 17-7)
 - Allows programmers to include a debug view with their application that is only visible with an encryption key.

Example 17-7 Encrypted debug view

```

Encrypt the debug view so that the debug view is only visible if the
person knows the encryption key.
CRTBNDRPG MYPGM DBGENCKEY('my secret code')
Then either
STRDBG MYPGM DBGENCKEY('my secret code')
OR
STRDBG MYPGM
and wait to be prompted for the encryption key

```

- Teraspace storage model
- Much higher limits for automatic storage.
- Can compile *CALLER programs with STGMDL(*INHERIT) so 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
- Specify H-spec ALLOC(*TERASPACE) to have teraspace allocations in any storage model

17.4.2 ILE COBOL Compiler

The following are the ILE COBOL compiler enhancements available with Rational Development Studio for i V7.1:

- Improved performance for XML GENERATE (see Example 17-8)
 - If XML is repeatedly generated for the same data structure or record format
 - If the data for the XML never contains special characters, specify process option XMLGEN(ASSUMEVALIDCHARS) to bypass checking XML data for (<,>,&,'"', and characters outside the normal alphanumeric range)
 - If XML is generated to an IFS file, specify process option XMLGEN(KEEPFILEOPEN) to keep the stream file open after an APPEND. Code XML GENERATE FILE-STREAM without the OVERWRITE or APPEND option to close the file.

Example 17-8 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.

```

- OPTIMIZE(*NEVER) supports compiling large programs
 - New COBOL compiler command parameter OPTIMIZE(*NEVER) allows large COBOL programs to be compiled.
 - PROCESS option NEVEROPTIMIZE is added to allow this option to 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.
- Support true integers (COMP-5)
 - COBOL currently supports the COMP-4 binary type which does not support the full range of the binary value. For example, a 2-byte COMP-4 supports values between -9999 and 9999, when a 2-byte integer supports values between -32768 and 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.
- COMP-5 is already supported by System z® COBOL, and customers porting applications from System z have difficulty dealing with their COMP-5 definitions.
- ▶ Freeze SEU syntax-checking at V6R1 level
 - Freeze the ILE COBOL syntax checkers for SEU at the V6R1 level.
 - For customers using SEU to edit ILE COBOL source, the syntax checkers will not recognize any features added after V6R1.
- ▶ Encrypted debug view (see Example 17-7 on page 516)
 - Allows programmers to include a debug view with their application that is only visible with an encryption key.
- ▶ Teraspace storage model
 - Much higher limits for automatic storage.
 - Can compile *CALLER programs with STGMDL(*INHERIT) so 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
 - Specify H-spec ALLOC(*TERASPACE) to have teraspace allocations in any storage model

17.4.3 ILE C

The following are the ILE C compiler enhancements available with Rational Development Studio for i V7.1:

- ▶ Support for encryption of debug data
 - A new debug encryption key compiler option, DBGENCKEY, is supported in the ILE C compiler. This option specifies the key to be used to encrypt source code that is embedded in debug views of a module (for example, the Listing View debug view created by specifying the DBGVIEW(*LIST) compile option). An equivalent option is available on the start debug command, STRDBG, to specify the key for decryption of the encrypted debug views.

17.4.4 ILE C++

The following list details the ILE C++ compiler enhancements available with Rational Development Studio for i V7.1:

- ▶ Decimal floating-point support
 - Allows floating-point computations to be performed using decimal arithmetic (base 10)
 - Avoids potential rounding errors converting binary floating-point data to/from human readable formats
 - Conforms to the decimal formats and arithmetic described in the IEEE 754-2008 Standard for Floating-Point Arithmetic
 - Adds support to ILE C++ compiler, 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, -95/+96 exponent
 - `_Decimal64`, 8 bytes, 16 digits precision, -383/+384 exponent
 - `_Decimal128`, 16 bytes, 34 digits precision, -6143/+6144 exponent
- Conversions to/from C++ built-in data types, such as integers and binary floating-point types
- Includes `DECFLTRND` option on C++ compiler commands (`CRTCPPMOD`, `CRTBNDCPP`) to control compile-time decimal floating-point rounding mode
- Support for C99 features
- Pragma operator
- `__func__` predefined identifier
- hexadecimal floating-point literals
- variable length arrays
- empty arguments for function-like macros
- variable number of arguments for function-like macros
- Support for encryption of debug data

For more information about Rational Development Studio for i visit:

<http://www-01.ibm.com/software/rational/products/devstudio/i/>

17.5 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:

- ▶ Browsers
- ▶ Mobile devices
- ▶ Cloud computing resources
- ▶ Web services
- ▶ External databases
- ▶ XML files
- ▶ Spreadsheets
- ▶ And more

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 a Open Access for RPG solution.

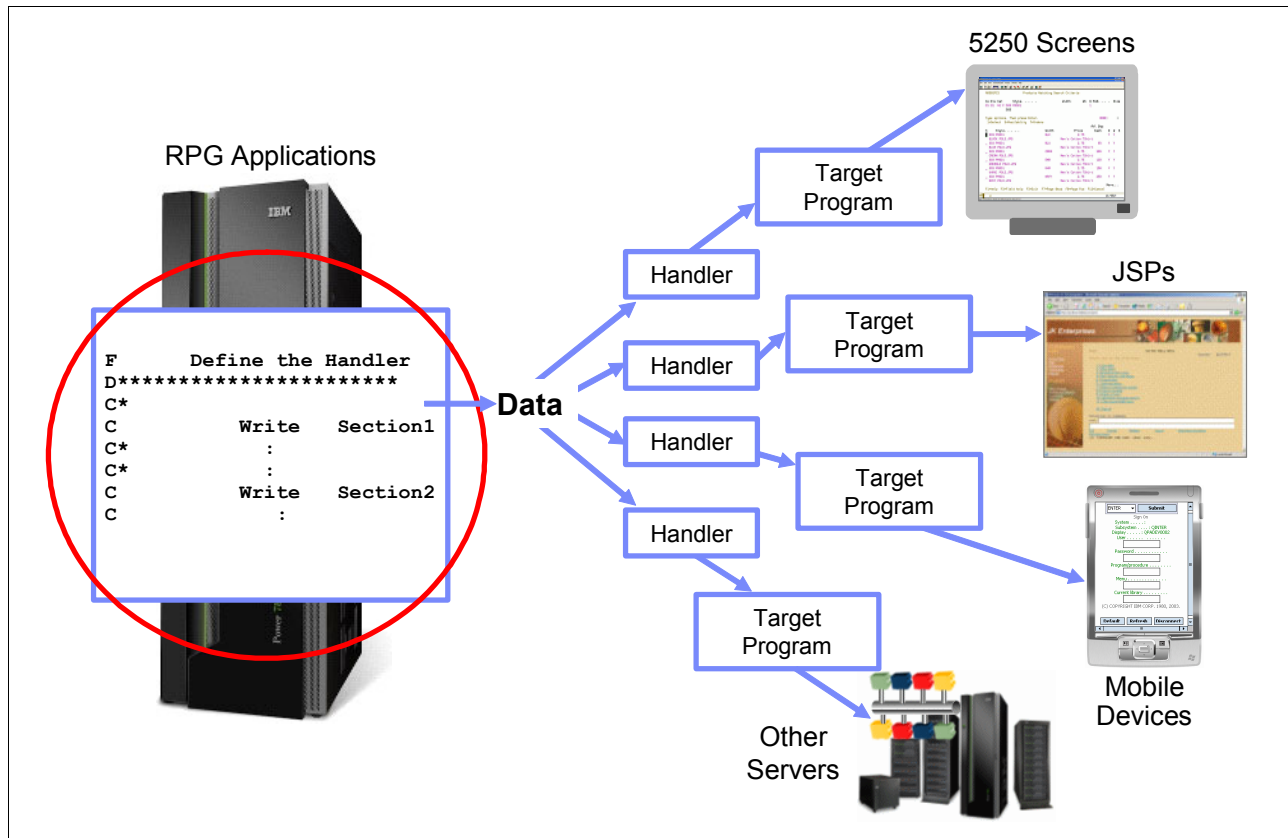


Figure 17-11 Rational Open Access: RPG Edition

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 characteristics of a handler:

- ▶ A handler is a program or a procedure in a service program
- ▶ A handler can be a generic handler able to handle any file of that device type, or it can be a handler 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 approach 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.

Example, an existing application that uses 5250 display files is modified to use Open Access for the WORKSTN files

- ▶ 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 will often be provided by an outside expert such as a software tool vendor or business partner

Handler First

The handler is written before the application is written.

Example: the RPG programmer wants to use a web service that returns information for a specific set of criteria.

- ▶ 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 will support
- ▶ 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 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 17-9 illustrates a web service that provides current weather based on City Name.

- ▶ Think of city name as a key (1)
- ▶ Then an RPG keyed DISK “file” can be used as the interface to the web service (2)
- ▶ Handler Provider has chosen to handle OPEN, CLOSE, and CHAIN (3)
- ▶ Handler provider has put all the required objects into library WEBSVCS: (4)
 - A service program containing 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

Example 17-9 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
```

The RPG Program

- RPG Program using PF CITYWTHR to define records and key information

What Happens

- RPG implicitly opens the "file" and calls handler cityWeather to handle the CHAIN operation
- The handler sets up information needed by the handler-specific data structure so it can handle the CHAIN operation correctly.
 - In this case, perhaps a socket to the web service URL.
- When the RPG program does a CHAIN operation, RPG calls the handler again to handle the CHAIN operation
- After the handler returns, RPG determines the value of %FOUND according to the feedback from the handler
- When the RPG program ends with LR on, RPG implicitly closes file and calls the handler again to handle the CLOSE operation
 - The handler closes the socket to the web service URL.

The Handler

- The handler service program, using PF CITYWTHR to define the records and key information. See Example 17-10

Example 17-10

```
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;
```

Rational Open Access: RPG Edition V1.1.1

IBM Rational Open Access: RPG Edition V1.1.1 provides an evaluation option. It allows you to evaluate and test the product capabilities internally for 70 days without a valid License Key.

Effective on July 29, 2011, IBM Rational Open Access: RPG Edition V1.1.1 is available in the IBM Entitled Software Support (ESS) website for evaluation. You can download the product in ESS for internal evaluation, testing, or demonstration purposes for 70 days without a valid License Key.

After 70 days the product will not operate without a valid license Key. If you want to continue using the product after 70 days, you must then obtain a valid License Key, by purchasing the product.

IBM Rational Open Access: RPG Edition V1.1.1 supports all the hardware models that support IBM i V6.1 and V7.1, and supports IBM i V6.1 and V7.1 operating systems.

Note: The only difference between V1R1M1 and V1R1M0 is the licensing. V1R1M1 can be used up to 70 days without a license, while V1R1M0 always requires a license.

For more information about IBM Rational Open Access: RPG Edition V1.1.1 refer to the product announcement:

http://www-01.ibm.com/common/ssi/rep_ca/3/897/ENUS211-253/ENUS211-253.PDF

17.5.2 Open Access Requirements

IBM Rational Open Access: RPG Edition V1.1 supports IBM i 6.1 and i 7.1 operating systems. These are the requirements for IBM Rational Open Access: RPG Edition:

- ▶ IBM Rational Open Access: RPG Edition (5733-OAR)
- ▶ IBM Rational Development Studio for IBM i 6.1 (5761-WDS) or IBM i 7.1 (5770-WDS) and applicable PTFs.
- ▶ RPG runtime IBM i 6.1 (5761-SS1) or IBM i 7.1 (5770-SS1) and 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, 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, 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 the following web page:

<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 visit:

<http://www-01.ibm.com/software/rational/products/openaccess/>



IBM Systems Director Navigator for IBM i 7.1

This chapter introduces IBM i 7.1 enhancements for the IBM Systems Director Navigator:

- ▶ 18.1, “5250 emulation related enhancements” on page 526
- ▶ 18.2, “Set Target System enhancements” on page 527
- ▶ 18.3, “Database enhancements” on page 529
- ▶ 18.4, “New journal management enhancements” on page 548
- ▶ 18.5, “Integrated Server Administration enhancements” on page 552
- ▶ 18.6, “Printer output enhancements” on page 555
- ▶ 18.7, “File systems enhancements” on page 556
- ▶ 18.8, “Networking enhancements” on page 558
- ▶ 18.9, “Disk management enhancements” on page 559
- ▶ 18.10, “Tape support enhancements” on page 562
- ▶ 18.11, “Performance enhancements” on page 563
- ▶ 18.12, “JS1- Advanced Job Scheduler for i enhancements” on page 618
- ▶ 18.13, “Backup Recovery Media Services: BRMS plugin enhancements” on page 619
- ▶ 18.14, “Additional information” on page 619

18.1 5250 emulation related enhancements

You can now access the same 5250 emulation portlet that is included with IBM i Access for Web from the Systems option within IBM i management in IBM Systems Director Navigator for i, as shown in Figure 18-1. It makes it easier to use one single interface for managing IBM i systems and partitions.

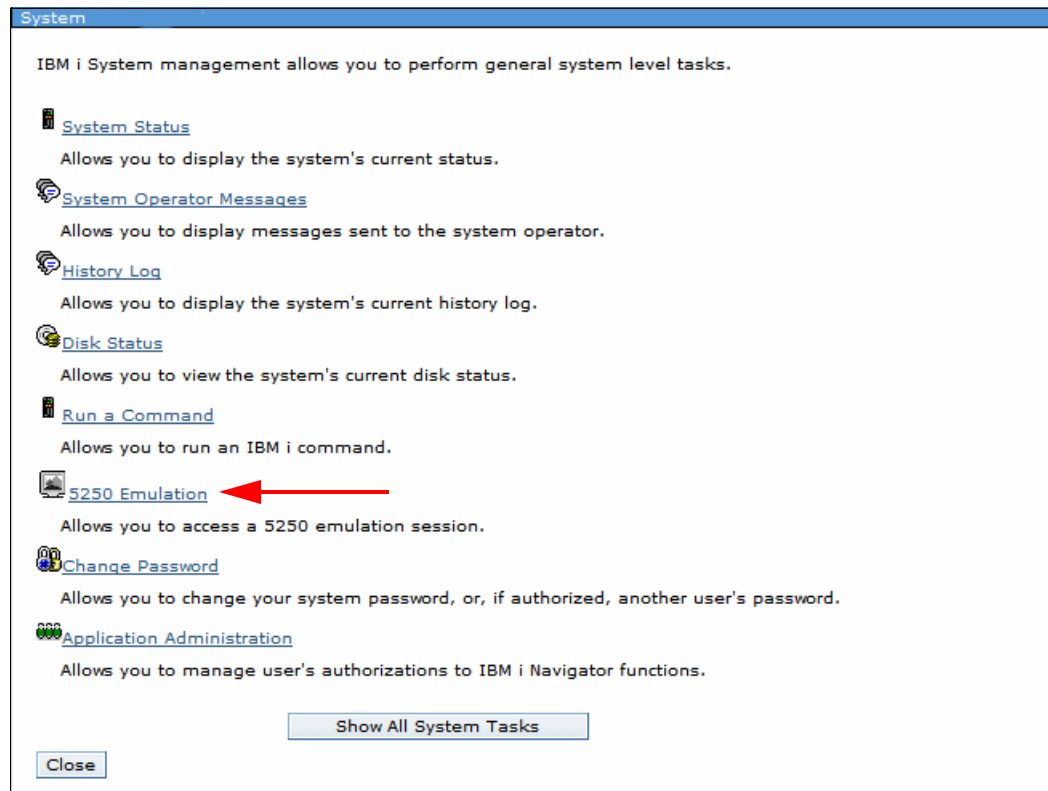


Figure 18-1 5250 Emulation from Systems option in IBM i management

However, as you can see in Figure 18-2, there is no possibility for customization of the 5250 emulation portlet from within the Navigator console.

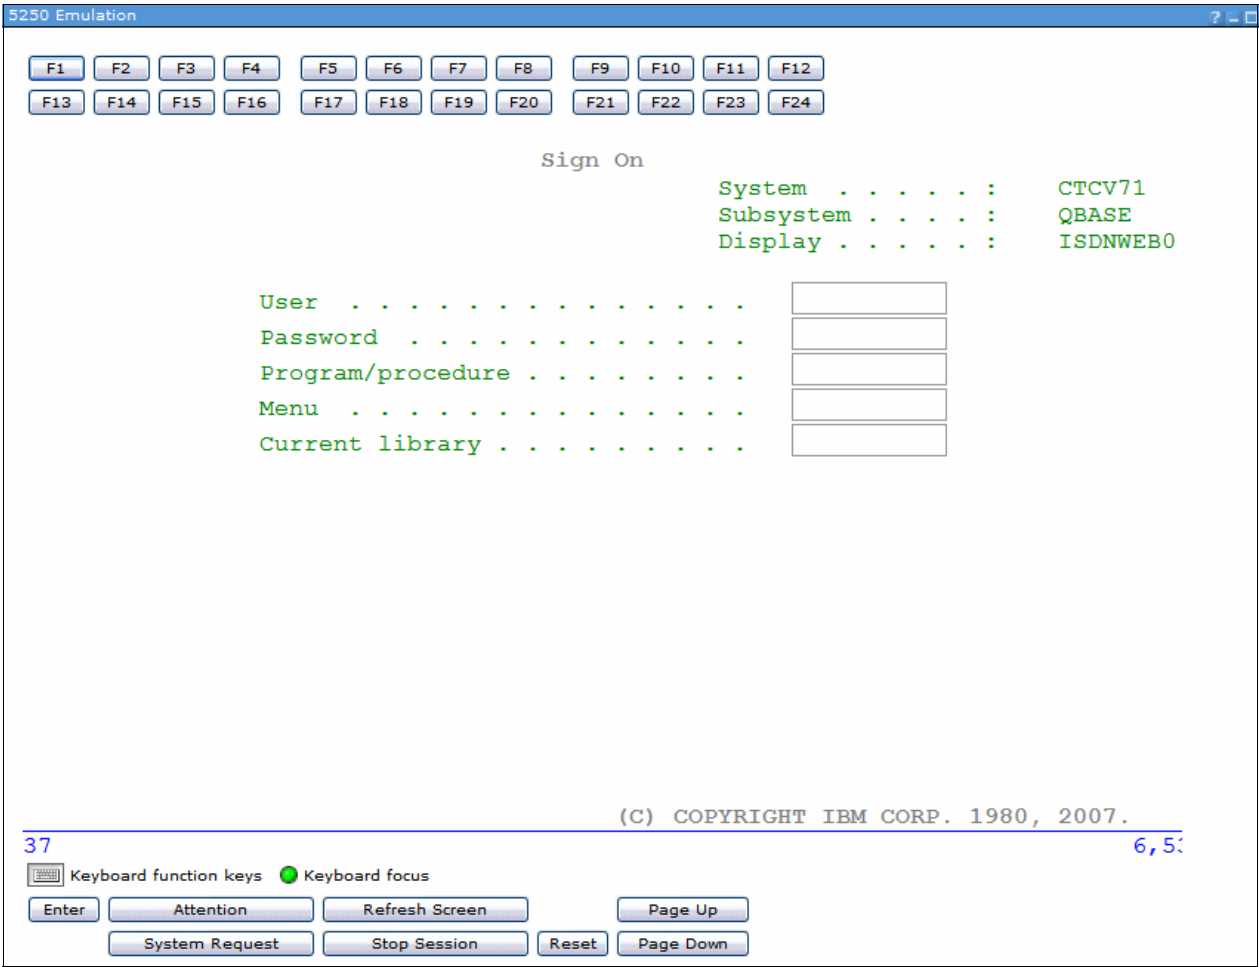


Figure 18-2 5250 mutilation session

18.2 Set Target System enhancements

Systems Director Navigator for i is enhanced to allow clients to work with multiple systems from a single interface using the Set Target System support. A client can now manage an IBM i system or partition by selecting the target system. After this new system or partition is selected, all data returned to the Systems Director Navigator interface is for this new system or partition.

Systems Director Navigator for IBM i 7.1 can now manage a target IBM i 5.4, 6.1 or 7.1 system or partition. The options that are available on that target partition/system can vary and they depend on the running IBM i release on that target system or partition.

As inn Figure 18-3, select the Set Target System task in the task navigation area to specify the system or partition to be managed. Specify the system or partition name, the corresponding user ID and password, and click **OK**.

Set Target System

Target system: ctcv71
Release: v7r1m0

Your target system can be the local system where you are running IBM Systems Director Navigator, or you can specify a different system to manage.

☐ Manage local system
☒ Manage another system

* System name: systemi5.be.ibm.co
* User ID: tomv
* Password:

OK Cancel

Figure 18-3 Setting target system to manage another system or partition

This example is running IBM Systems Director Navigator on a IBM i 7.1 system. Specify to manage another system, which in this case is running IBM i 6.1.

When there, get the IBM i 6.1 functions available within IBM Systems Director Navigator for i.

With the Set Target System feature, the Systems Director Navigator 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 towards the original IBM i 7.1 system, specify that to manage the local system by selecting the **Manage local system** radio button and clicking **OK**, as shown in Figure 18-4.

Set Target System

Target system: systemi5.be.ibm.com
Release: v6r1m0

Your target system can be the local system where you are running IBM Systems Director Navigator, or you can specify a different system to manage.

☒ Manage local system
☐ Manage another system

* System name: systemi5.be.ibm.co
* User ID: tomv
* Password:

OK Cancel

Figure 18-4 Manage local system

From now on, you will be managing the original IBM i running IBM i 7.1 again.

18.3 Database enhancements

This section discusses the following database enhancements:

- ▶ 18.3.1, “On Demand Performance Center enhancements” on page 529
- ▶ 18.3.2, “Health center enhancements” on page 536
- ▶ 18.3.3, “Database management enhancements” on page 537
- ▶ 18.3.4, “Maintenance category enhancements” on page 544
- ▶ 18.3.5, “Long SQL schema names enhancements” on page 547
- ▶ 18.3.6, “Omnifind text search functions enhancements” on page 547

18.3.1 On Demand Performance Center enhancements

This section discusses the following enhancements:

- ▶ New centralized authority model
- ▶ Enhanced filtering support
- ▶ SQL performance monitor enhancement
- ▶ Enhancements to Index Advisor and indexes list

The following sections discuss these enhancements.

New centralized authority model

Support has been added in IBM i 7.1 for a new centralized authority model when working with all the SQL performance tools.

Previously, a system security officer needed to grant *JOBCTL user special authority to enable database analysts and database administrators to use the database tools. Because *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 use of the full set of database tools.

New in IBM i 7.1, the security officer now has additional capability to authorize access to the database analysis tools and the SQL Plan Cache. DB2 for i now takes advantage of the function usage capability available in the operating system.

A new function usage group called *QIBM_DB* has been created. In IBM i 7.1, there are two function IDs in the QIBM_DB group:

- ▶ QIBM_DB_SQLADM (IBM i Database Administrator tasks)
- ▶ QIBM_DB_SYSMON (IBM i Database Information tasks)

The security officer now has flexibility to grant authorities by 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 Systems Director Navigator for i.

Within IBM Systems Director Navigator for i, perform the following as shown in Figure 18-5:

1. Expand **IBM i Management**.
2. Select the **System** category.
3. Select **Application Administration**.

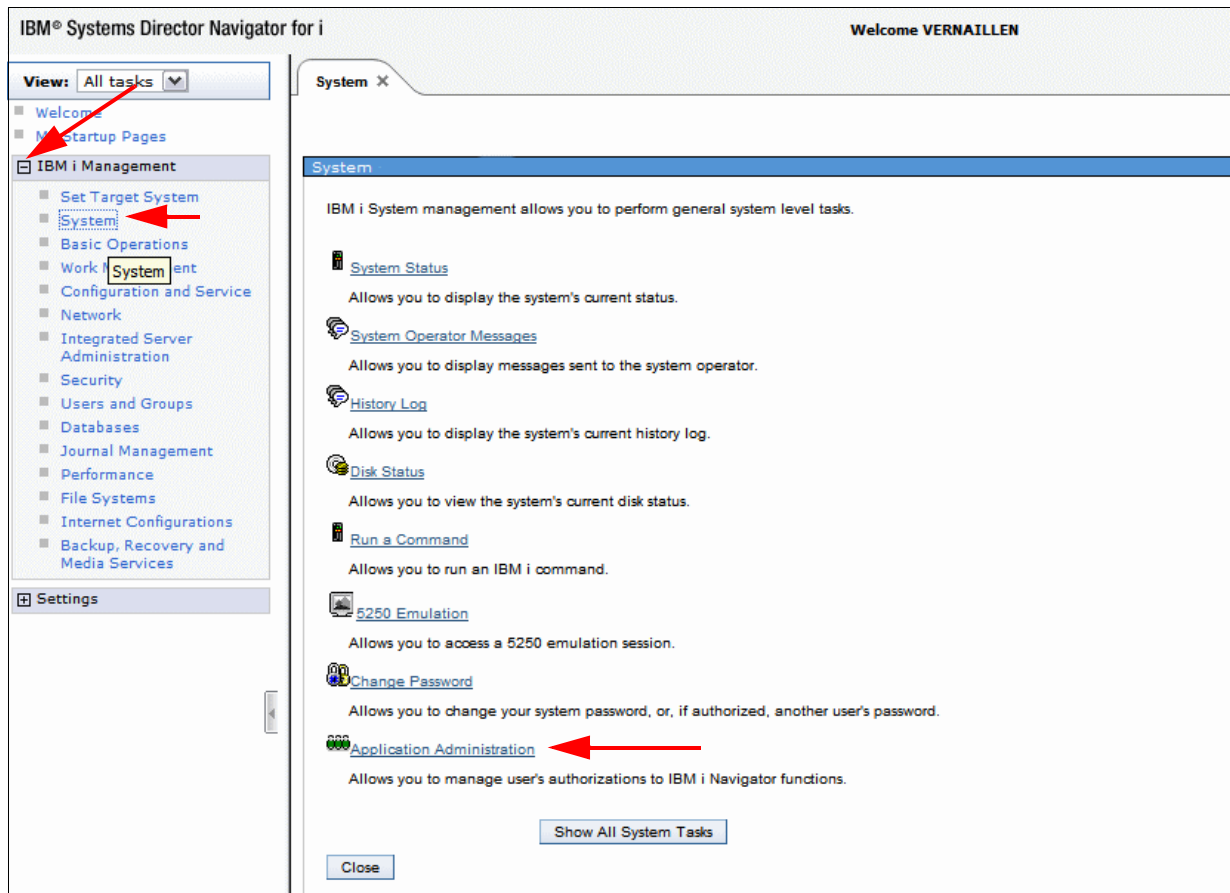


Figure 18-5 Application administration

Now, within Application Administration, perform the following steps, as shown in Figure 18-5 on page 530:

1. Click **Host Applications** on the left side menu.
2. Expand **IBM i**.
3. Expand **Database**.
4. Click **Customize** from the pop-up menu to the right of **Database Administrator**.

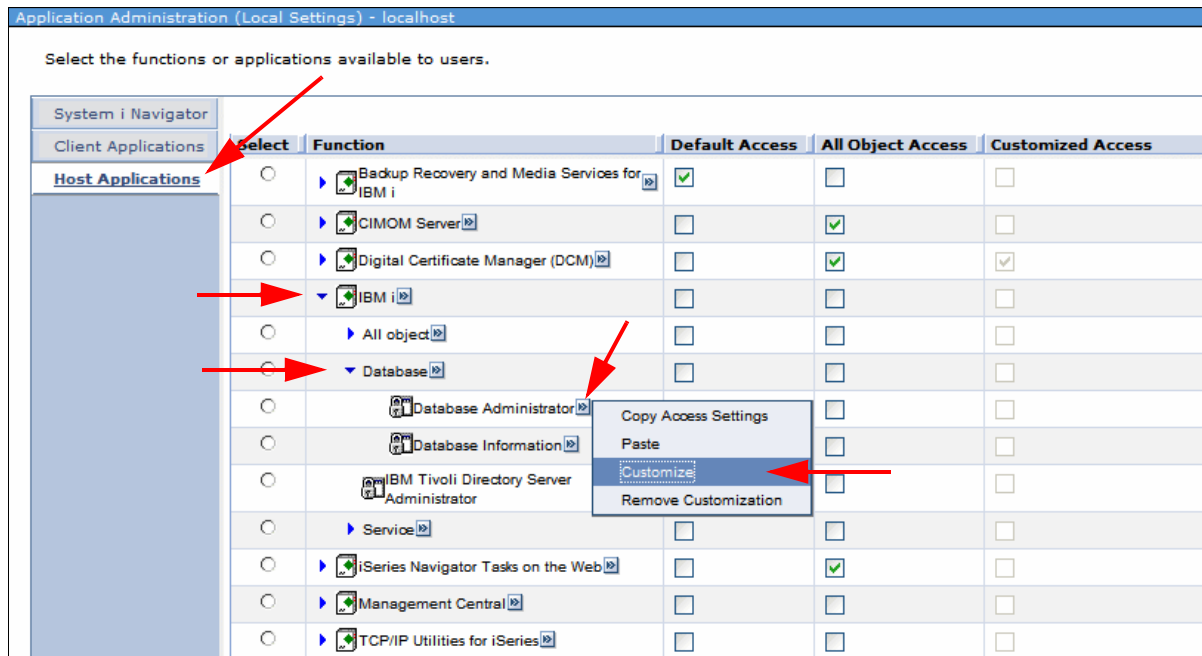


Figure 18-6 Database administrator: Customization

You now have the possibility to further customize the access to the Database Administrator functions, as shown in Figure 18-7.

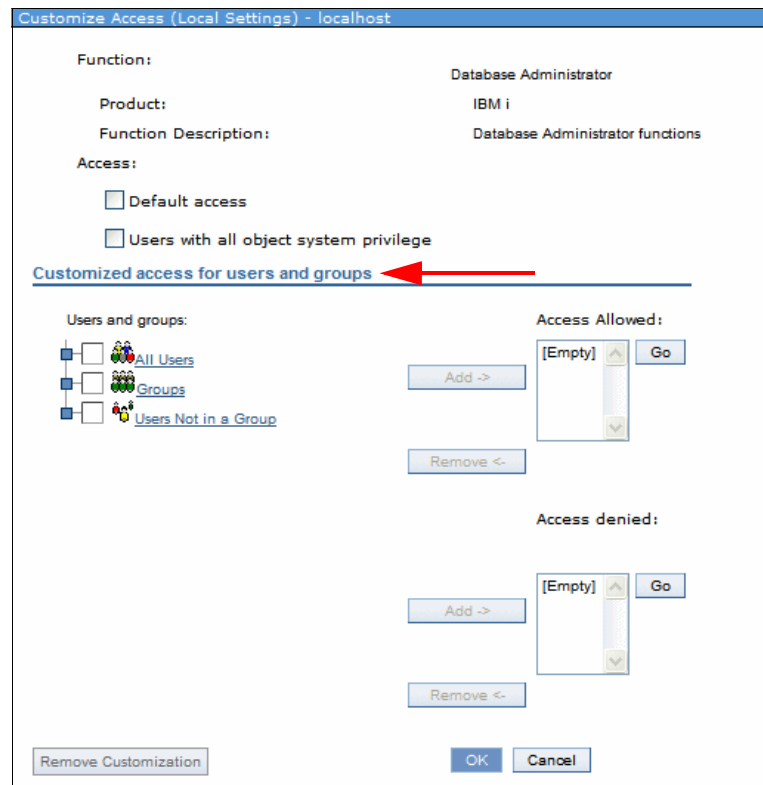


Figure 18-7 Customizing access

Note: The CHGFCNUSG (Change Function Usage) command, with a function ID of QIBM_DB_SQLADM, can also be used to change the list of users that are allowed to perform database administration operations. The function usage controls allow groups or specific users to be specifically allowed or denied authority. The CHGFCNUSG command also provides 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.

Enhanced filtering support

More filters have been added when starting SQL performance monitors. The new filters include the client registers filters shown in Figure 18-8. They allow even more granularity to help reduce the amount of data collected.

SQL Performance Monitor Wizard - Localhost(D1040f40)

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: TOMV
- ☐ Client location:
- ☐ Local port: Use entry from below
- ☐ Query Governor limits: Always collect information when exceeded

Client registers

- ☐ Accounting string:
- ☐ Application name:
- ☐ Program name:
- ☐ Client user:
- ☐ Workstation:
- ☐ Statements that access these objects:

Schema	Name

Activity to monitor

- ☒ Only collect monitor output for user activity
- ☐ Collect monitor output for user and system activity

< Back Next > Finish Cancel

Browse... Remove

Figure 18-8 Filter the database monitor collection

SQL performance monitor enhancement

SQL Details for a Job now has ability to 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 panel, as shown in Figure 18-9.

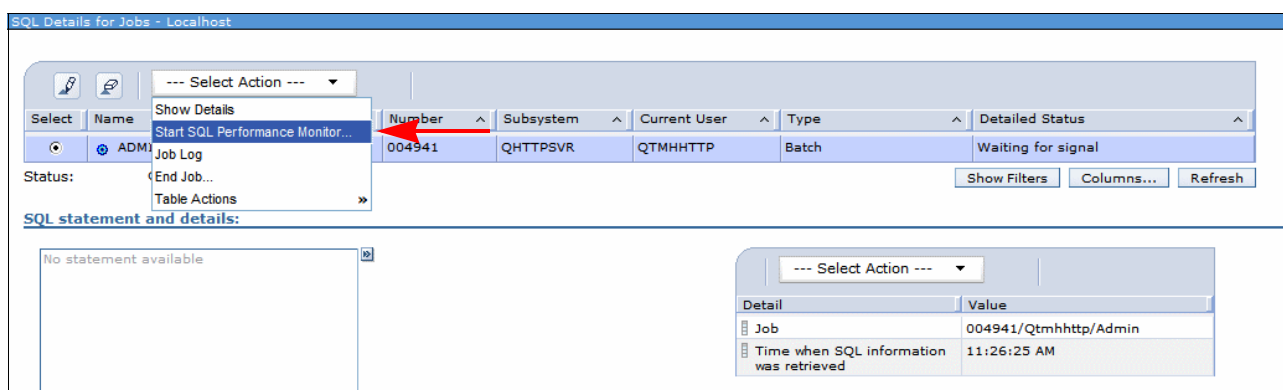


Figure 18-9 Start SQL Performance Monitor from SQL details for a job

Enhancements to Index Advisor and indexes list

Within IBM Systems Director Navigator for IBM i 7.1, there are several enhancements to Index Advisor and indexes list.

Index Advisor enhancements

IBM Systems Director Navigator for IBM i 7.1 now has the following new maintained temporary index (MTI) information within the Index Advisor as shown in Figure 18-10:

- ▶ -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 has been created by the optimizer. MTIs do not persist across system IPLs.
- ▶ -MTI Used
Number of times that this specific MTI has been used by the optimizer.
- ▶ -MTI Last used
Timestamp representing 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 which exactly matches this advice has never been used by the queries which generated this index advice.




Figure 18-10 Index Advisor: MTI information

Indexes list enhancements

The indexes list now includes information for the last time the index was built to allow for planning the amount of time it might take for the next time the index is to be built as shown in Figure 18-11.

Select	Name	Type	Schema	Last Build	Owner	System Name
<input type="checkbox"/>	QASQSPDP	Index	QSYS2	2/24/10 9:21:40 AM	QSYS	QASQSPDP
<input type="checkbox"/>	QASQSPVDP	Index	QSYS2	2/24/10 9:21:39 AM	QSYS	QASQSPVDP
<input type="checkbox"/>	QASQTYPE	Keyed Physical File	QSYS2	10/7/09 3:42:00 PM	QSYS	QASQTYPE
<input type="checkbox"/>	QASQVRLL	Index	QSYS2	2/24/10 9:21:39 AM	QSYS	QASQVRLL
<input type="checkbox"/>	QASQVRLLS	Index	QSYS2	2/24/10 9:21:39 AM	QSYS	QASQVRLLS
<input type="checkbox"/>	QASQVRSL	Index	QSYS2	2/24/10 9:21:39 AM	QSYS	QASQVRSL
<input type="checkbox"/>	QASQVRSS	Index	QSYS2	2/24/10 9:21:38 AM	QSYS	QASQVRSS
<input type="checkbox"/>	QASQXSROSS	Index	QSYS2	2/24/10 9:21:39 AM	QSYS	QASQXSROSS

Figure 18-11 Last built index information

The indexes list now also includes columns for SSD (DB2 Media Preference) and Keep In Memory values, as shown in Figure 18-12.

Select	Name	Type	Schema	Last Build	Keep in Memory	Media Preference	Owner	System Name
<input type="checkbox"/>	QASQSPDP	Index	QSYS2	2/24/10 9:21:40 AM	No	Any	QSYS	QASQSPDP
<input type="checkbox"/>	QASQSPVDP	Index	QSYS2	2/24/10 9:21:39 AM	No	Any	QSYS	QASQSPVDP
<input type="checkbox"/>	QASQTYPE	Keyed Physical File	QSYS2	10/7/09 3:42:00 PM	No	Any	QSYS	QASQTYPE
<input type="checkbox"/>	QASQVRLL	Index	QSYS2	2/24/10 9:21:39 AM	No	Any	QSYS	QASQVRLL
<input type="checkbox"/>	QASQVRLLS	Index	QSYS2	2/24/10 9:21:39 AM	No	Any	QSYS	QASQVRLLS
<input type="checkbox"/>	QASQVRSL	Index	QSYS2	2/24/10 9:21:39 AM	No	Any	QSYS	QASQVRSL
<input type="checkbox"/>	QASQVRSS	Index	QSYS2	2/24/10 9:21:38 AM	No	Any	QSYS	QASQVRSS
<input type="checkbox"/>	QASQXSROSS	Index	QSYS2	2/24/10 9:21:39 AM	No	Any	QSYS	QASQXSROSS

Figure 18-12 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 302.

Note: It might be possible that you need to add the previously described columns to show up, as they are not displayed by default. Select the **Column** option in the Select Action menu to add those columns.

18.3.2 Health center enhancements

In Figure 18-13, an example is shown where a new SQL0901 category has been added to the Environmental Limits tab. This helps to log situations where the database engine has an unknown error and more data is needed to be collected.

Select	Environmental Limit - 3/22/10 3:46:24 PM	Value	Percent of Limit	Status
<input type="radio"/>	Maximum number of activation groups to use SQL per job			
<input type="radio"/>	Maximum number of active descriptors per job			
<input type="radio"/>	Maximum length of SQL statement per job (2 MB)			
<input checked="" type="radio"/>	SQL0901 UNEXPECTED ERROR:1022 MESSAGE: SQL0901 F/QSQROUTS-MODULE/QSQCLNUP-PROCEDURE/SQROUTE_CLEANUP-STMT/4009MESSAGE: SQL0901 F/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/QSQXSRMCMP-MODULE/QSQXSRMCMP-PROCEDURE/CLEANUP-STMT/12547 T/QSQXSRMCMP-MODULE/QSQXSRMCMP-PROCEDURE/CLEANUP-STMT/12547MESSAGE: SQ			
<input type="radio"/>	057826/HOSPERS/QPADEV00F (HOSPERS)	0	No maximum	Normal

Figure 18-13 SQL0901 category added to the Environmental Limits tab

Random and sequential I/O counts are now collected and displayed in the Activity tab, as shown in Figure 18-14.

Select	Activity - 3/22/10 3:18:59 PM	Value	Status
<input type="radio"/>	Insert operations		
<input type="radio"/>	Update operations		
<input type="radio"/>	Delete operations		
<input type="radio"/>	Logical reads		
<input type="radio"/>	Physical reads		
<input type="radio"/>	Full opens		
<input type="radio"/>	Full closes		
<input type="radio"/>	Days used		
<input type="radio"/>	Index query use		
<input type="radio"/>	Index query statistics use		
<input type="radio"/>	Index logical reads		
<input type="radio"/>	Index random reads		
<input checked="" type="radio"/>	Random reads	1,276	Normal
<input checked="" type="radio"/>	Sequential reads	1,276	Normal
<input type="radio"/>	QSYS2.SYSIXADV (SYSIXADV)		

Figure 18-14 Activity tab: Random and sequential reads

18.3.3 Database management enhancements

Support has been added in IBM i 7.1 for the following new database features

- ▶ XML data type columns
- ▶ XML Schema Repository (XSR)
- ▶ Global variables
- ▶ Array types
- ▶ Field procedures
- ▶ Three part aliases

XML Data type Column support

Support is now available for XML Data column within the IBM i 7.1 Systems Director Navigator interface, as shown in Figure 18-15.

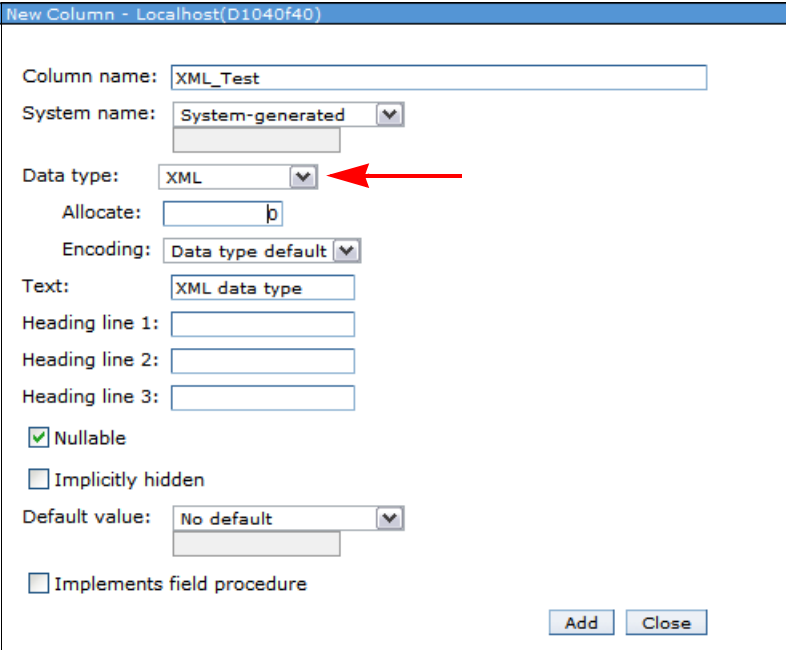


Figure 18-15 XML data type column support

For more information related to this support, see 6.2.1, “XML support” on page 113.

XML Schema Repository (XSR)

An XML Schema Repository is 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 a XML Schema Repository using this IBM Systems Director i Navigator interface.

Global variables

Support is available for global variables within the IBM i 7.1 Systems Director Navigator interface as shown in Figure 18-16.

The screenshot shows a dialog box titled "New Global Variable - Localhost(D1040f40)". It contains the following fields and controls:

- Name:** A text input field containing "Global_one".
- Schema:** A dropdown menu showing "TOMV".
- Data type:** A dropdown menu showing "BIGINT".
- Default Value:** A dropdown menu showing "NULL". To its right are two buttons: "Preview Value" and "Check Syntax".
- Text:** A text input field containing "TestGlobalVariable".
- Show SQL:** A button located below the Text field.
- OK** and **Cancel** buttons are located at the bottom right of the dialog.

Figure 18-16 Global variable support

For more information related to Global Variables, see 6.2.3, "Creating and using global variables" on page 120.

Array support

Figure 18-17 shows the support that is available for array types within the IBM i 7.1 Systems Director Navigator interface.

The screenshot shows a dialog box titled "New Array Type - Localhost(D1040f40)". It contains the following fields and controls:

- Name:** A text input field containing "INTARRAY2".
- Schema:** A dropdown menu showing "TOMV".
- Source data type:** A dropdown menu showing "INTEGER".
- Maximum cardinality:** A text input field containing "200".
- Text:** An empty text input field.
- Show SQL:** A button located below the Text field.
- OK** and **Cancel** buttons are located at the bottom right of the dialog.

Figure 18-17 Array type support

For more information related to array support, see 6.2.4, "Support for arrays in procedures" on page 121.

FIELDPROC support

IBM i 7.1 now has the FIELDPROC support at column level.

This is accomplished by selecting the Implements Field procedure option, when defining the column when you create or alter a table. As shown in Figure 18-18, specify the schema and the corresponding procedure name.

The screenshot shows the 'New Column - Localhost(D1040f40)' dialog box. The 'Column name' is 'column_test'. The 'System name' is 'System-generated'. The 'Data type' is 'CHARACTER' with a 'Length' of '10' and 'Encoding' set to 'Data type default'. There are input fields for 'Text', 'Heading line 1', 'Heading line 2', and 'Heading line 3'. The 'Nullable' checkbox is checked, and 'Implicitly hidden' is unchecked. The 'Default value' is 'No default'. The 'Implements field procedure' checkbox is checked, highlighted by a red arrow. Below this, the 'Field procedure specification' section has 'Schema' set to 'VERMAERE' and 'Program' set to 'PROCE00001'. The 'Parameter list' is an empty text area. 'Add' and 'Close' buttons are at the bottom right.

Figure 18-18 FIELDPROC support

For more information related to FIELDPROC support, see 6.2.6, “FIELDPROC support for encoding and encryption” on page 124.

Three-part names support

Figure 18-19 shows an example of how to define a three-part alias, referencing a table on a remote system. The corresponding SQL statement is also displayed.

New Alias - Localhost(D1040F40)

*Name:

Schema:

Referenced object

Database:

Schema:

Name:

Partition:

Text:

Show SQL

SQL statements:

```
/* Creating alias TOMV.TESTALIAS */  
CREATE ALIAS TOMV.TESTALIAS FOR S10DE18B.VERMAERE.TESTALIAS;  
  
/* Setting label text for TOMV.TESTALIAS */  
LABEL ON ALIAS D1040F40.TOMV.TESTALIAS IS 'test three-name support';
```

Figure 18-19 Three-part name support for an alias

Note: This support is available for a table, a view or an alias.

SQL Privilege statements

Generate SQL now supports the option to include SQL privilege statements. This output option is on by default, as shown in Figure 18-20.

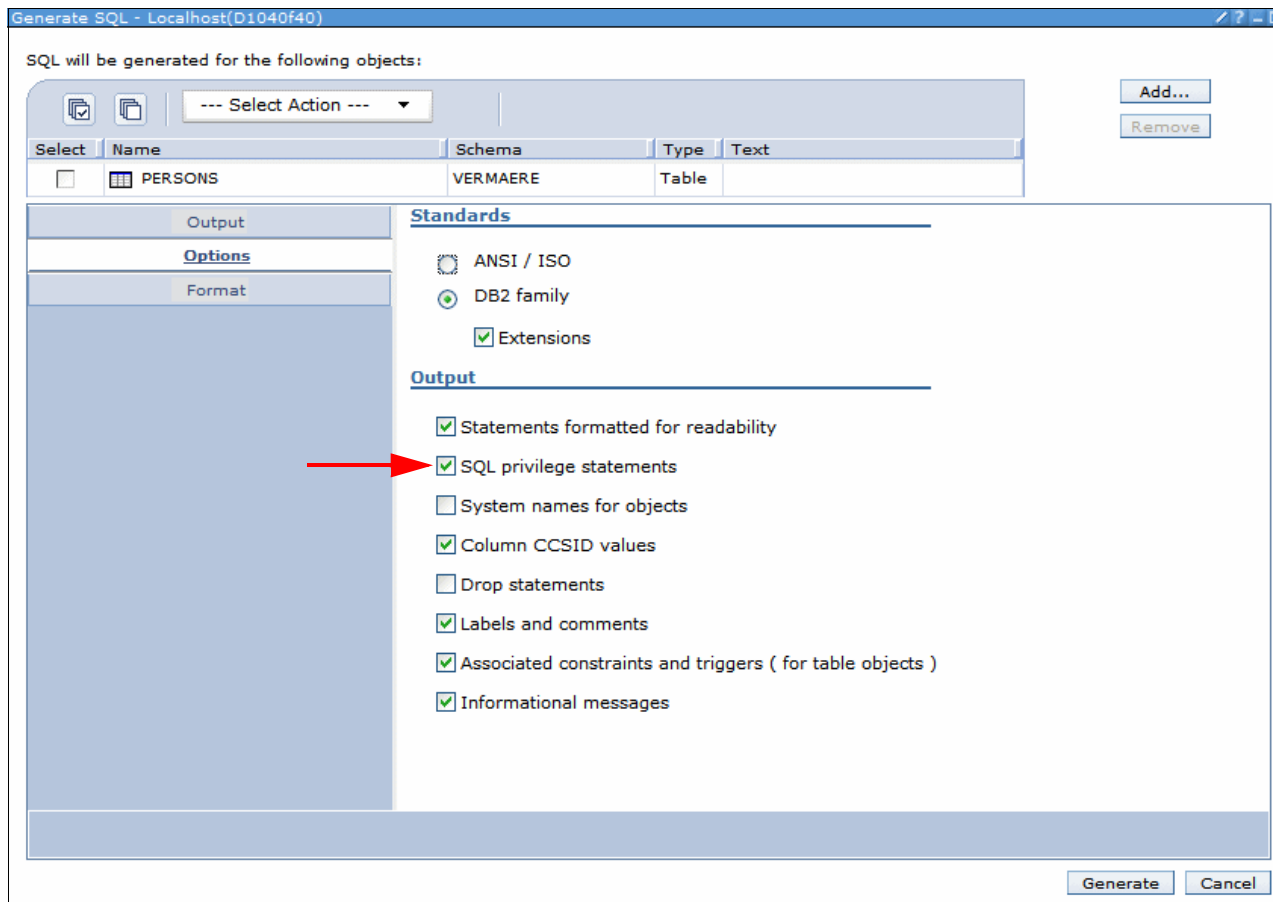


Figure 18-20 Generate SQL: Including SQL privileges

Several usability enhancements

In IBM i 7.1, it is now possible to save list contents to various output formats. In Figure 18-21, the list contents of all schemas within the database is saved into a comma separated value file within the IFS of the system.

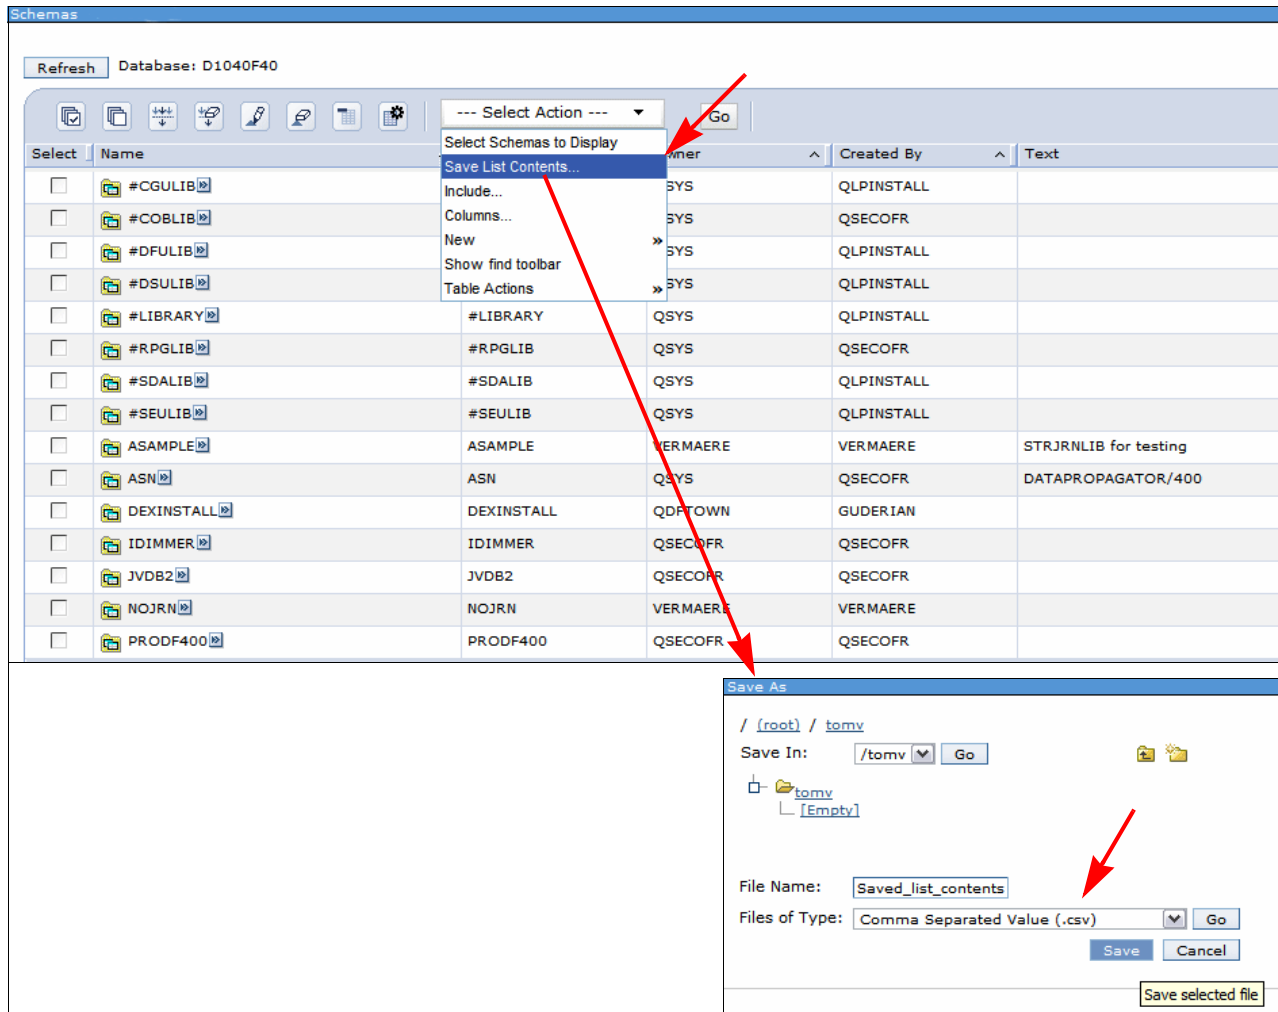


Figure 18-21 Save list contents

Figure 18-22 is an example where it is specified that, from the list of procedures within a given library, only those containing a specific text string within their name is displayed.

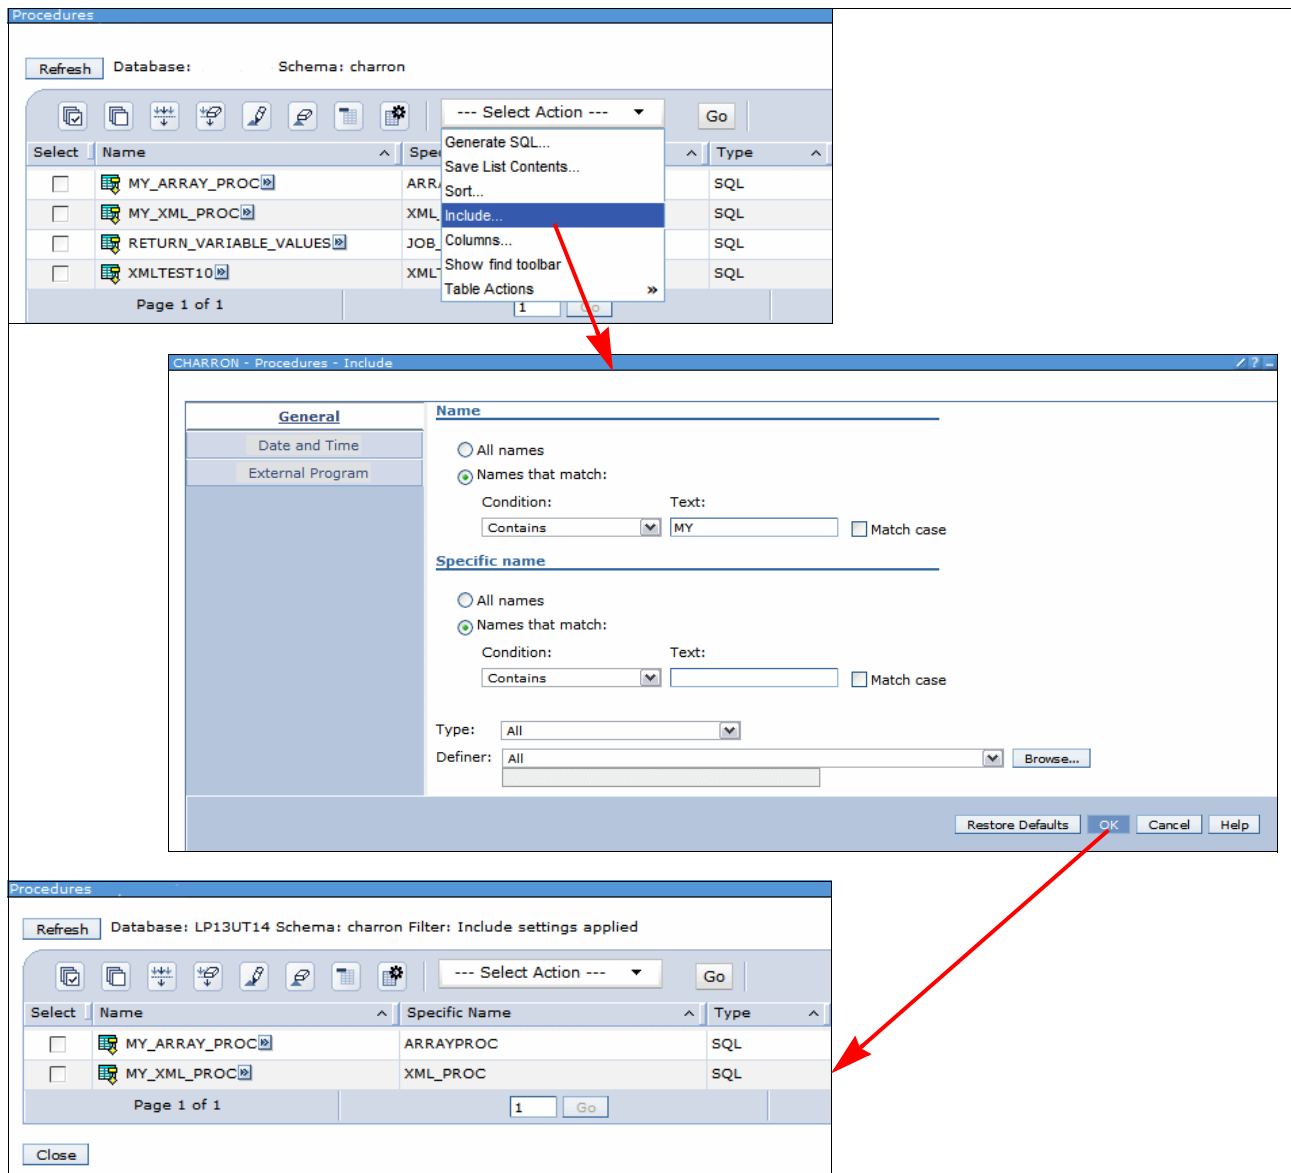


Figure 18-22 Include list support for procedures

Note: This include 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.

18.3.4 Maintenance category enhancements

Figure 18-23 shows the new Maintenance category that helps to organize database maintenance lists for table reorganizations, index builds, text search index builds, table alters, and check pending constraints.

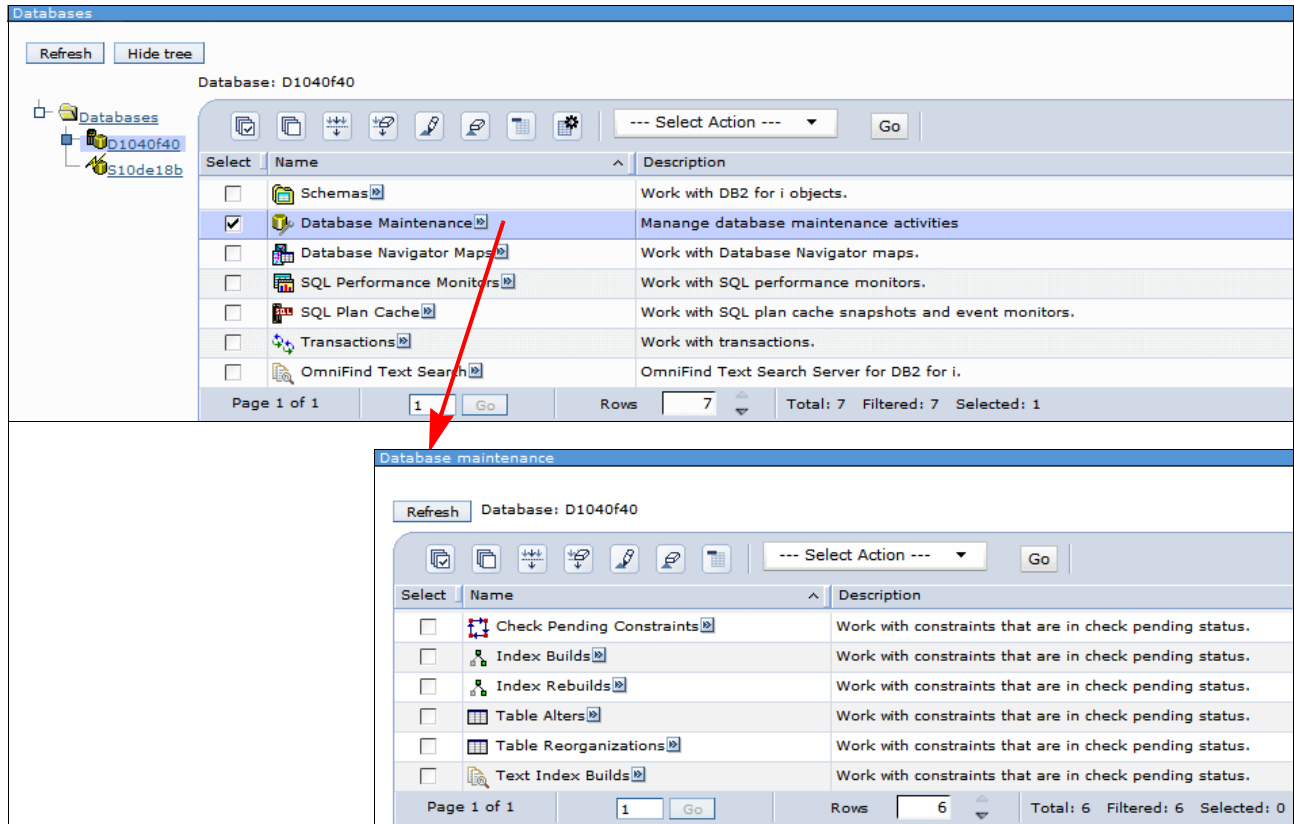


Figure 18-23 New maintenance category for database maintenance

Progress pages have been added for index builds, text search index builds, and table alters. Figure 18-24 shows the progress as the index is built.

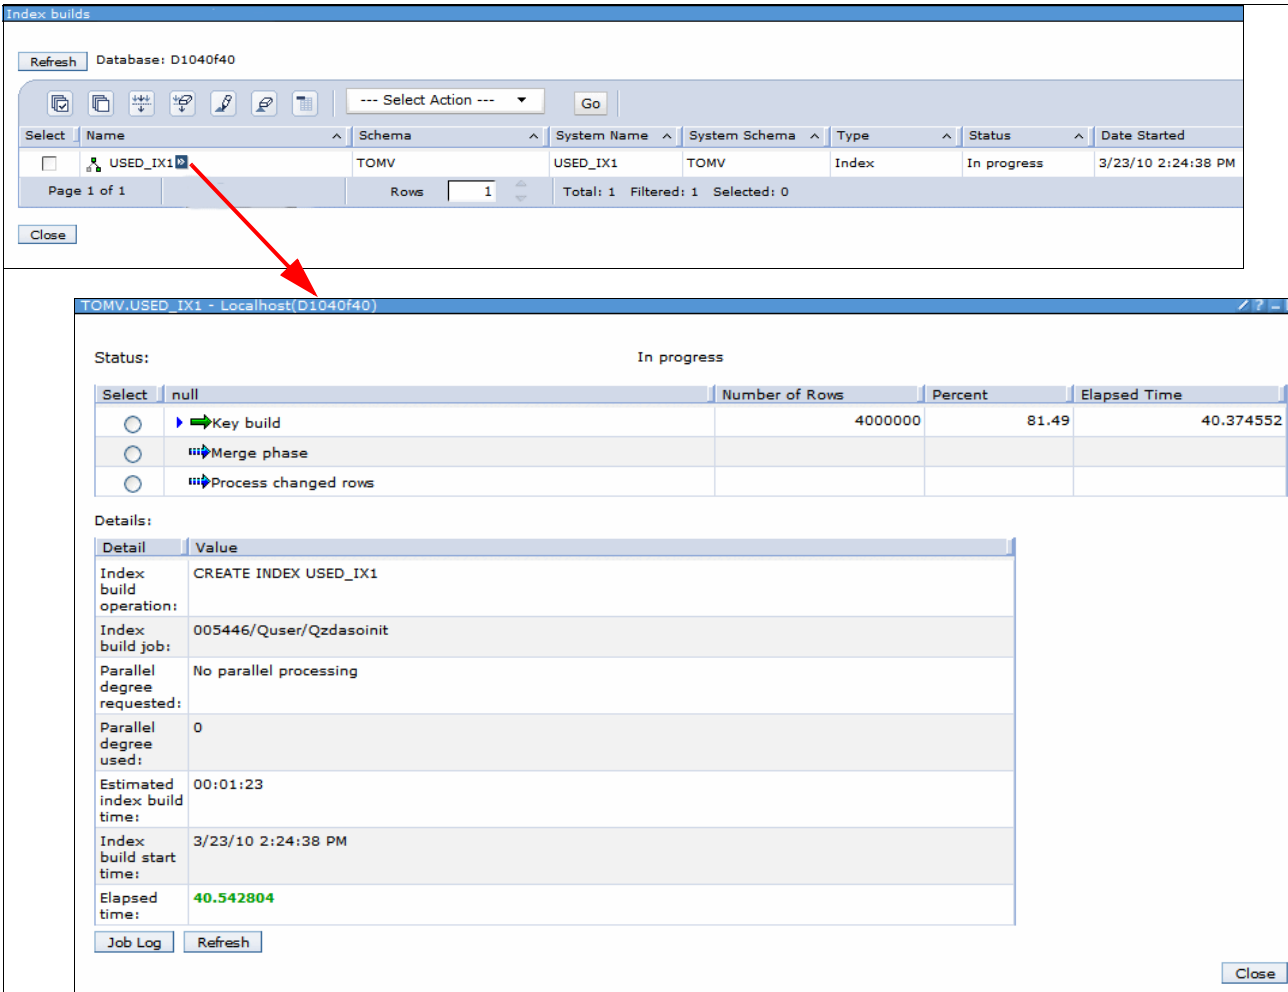


Figure 18-24 Progress pages for index builds

In addition, history information is shown in Figure 18-25 on the table reorganization progress page.

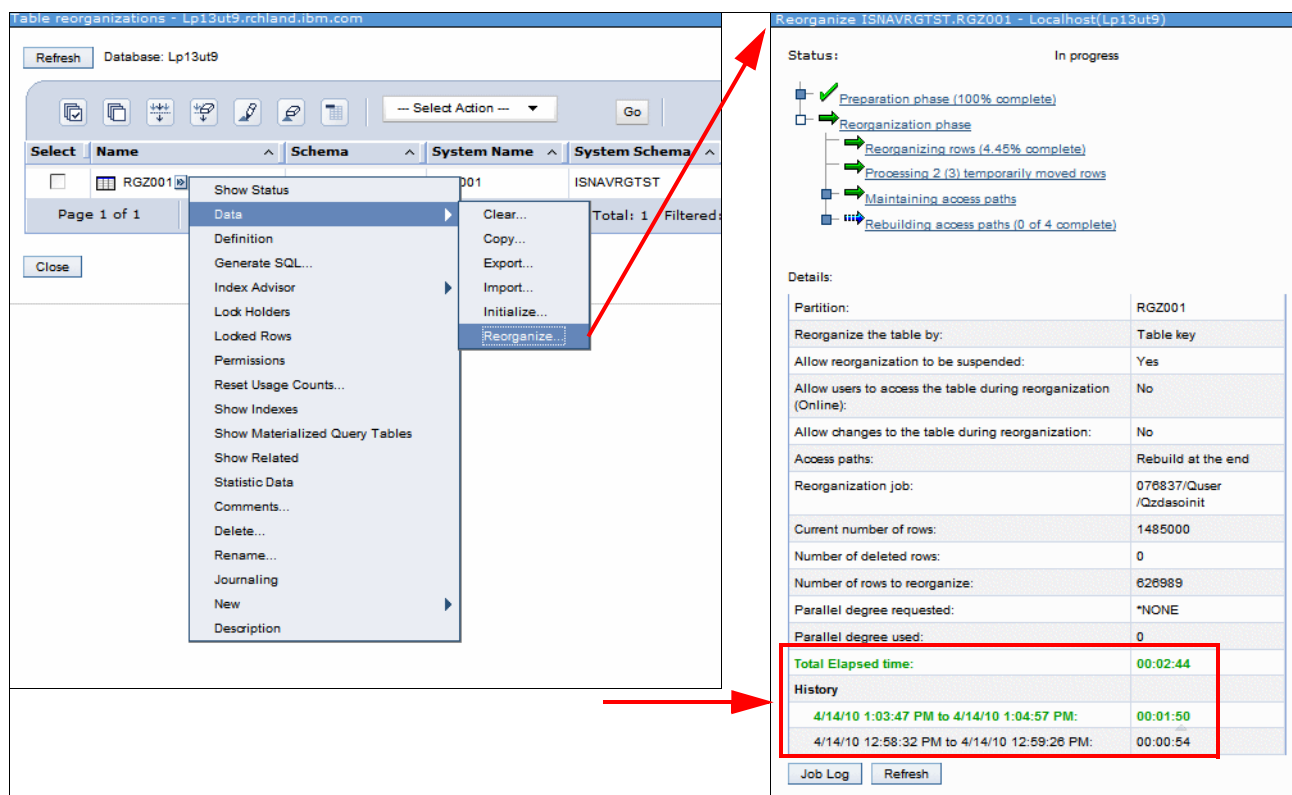


Figure 18-25 Table reorganization: History information

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 reorganize and resume at a later time, you see a row in the history section for the previous instance, and a new row for this current instance.

The Total elapsed time value then includes both the previous instance of the reorganize, plus this current instance of the reorganize. The history applies to only the history of the reorganize for one instance of the reorganize of this table. It does not show prior history of previously completed reorganizes of this table.

Note: The reorganize progress in releases prior to IBM i 7.1 needed to have you find the table and select to reorganize it to see if it is currently being reorganized. This, then, is an easier way to find this, because the Table Reorganizations option in the Database Maintenance category is now available.

18.3.5 Long SQL schema names enhancements

Support for long schema names has been added throughout the database pages to support porting database applications from other databases and more descriptive schema names.

Figure 18-26 defines a new Long Schema Name within our IBM i 7.1 database.

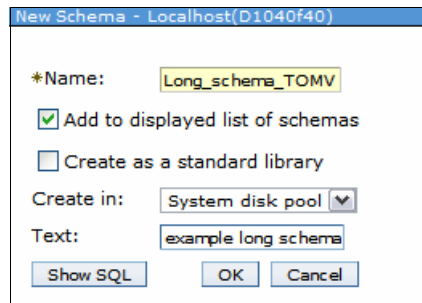
A screenshot of the 'New Schema' dialog box in IBM Systems Director Navigator. The title bar reads 'New Schema - Localhost(D1040f40)'. Inside, the '*Name:' field contains 'Long_schema_TOMV'. Below it, the 'Add to displayed list of schemas' checkbox is checked, and the 'Create as a standard library' checkbox is unchecked. The 'Create in:' dropdown menu is set to 'System disk pool'. The 'Text:' field contains 'example long schema'. At the bottom are three buttons: 'Show SQL', 'OK', and 'Cancel'.

Figure 18-26 Long schema name support

18.3.6 Omnifind text search functions enhancements

Omnifind is a text search product that allows IBM i users to search unstructured text stored in a column of a DB2 for i table. The text 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. This product also allows users to index unstructured data without having to parse it into a structured form such as an SQL table.

In Figure 18-27, the new support is shown that enables managing Omnifind text search servers and indexes.

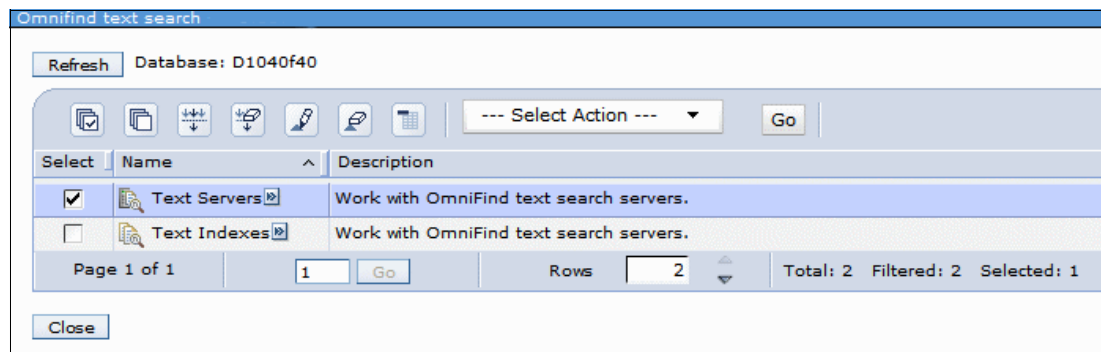
A screenshot of the 'Omnifind text search' window. At the top, there's a 'Refresh' button and 'Database: D1040f40'. Below is a toolbar with icons for various actions and a 'Go' button. A table with two columns, 'Name' and 'Description', lists 'Text Servers' and 'Text Indexes'. 'Text Servers' is selected with a checkmark. The bottom status bar shows 'Page 1 of 1', a 'Go' button, 'Rows: 2', and 'Total: 2 Filtered: 2 Selected: 1'. A 'Close' button is at the bottom left.

Figure 18-27 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 the following web page:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=%2Frzash%2Frzashkickoff.htm>

18.4 New journal management enhancements

You can now manage journal environments through IBM Systems Director Navigator for i. The following are the additional functions added to journal management in i 7.1:

- ▶ Show Journalized Objects
- ▶ Change Receivers
- ▶ Remote Journals
- ▶ Add Remote Journal
- ▶ Properties for a Journal Receiver
- ▶ Activate Remote Journals
- ▶ Deactivate Remote Journals
- ▶ Remove Remote Journals
- ▶ Properties for Remote Journals

The following sections discuss these functions.

For more information, related to the Journal Management capabilities and enhancements, see 6.3, “Availability and consistency” on page 134.

18.4.1 Show Journalized Objects function

The Show Journalized Objects function is now added for a journal as seen in Figure 18-28.

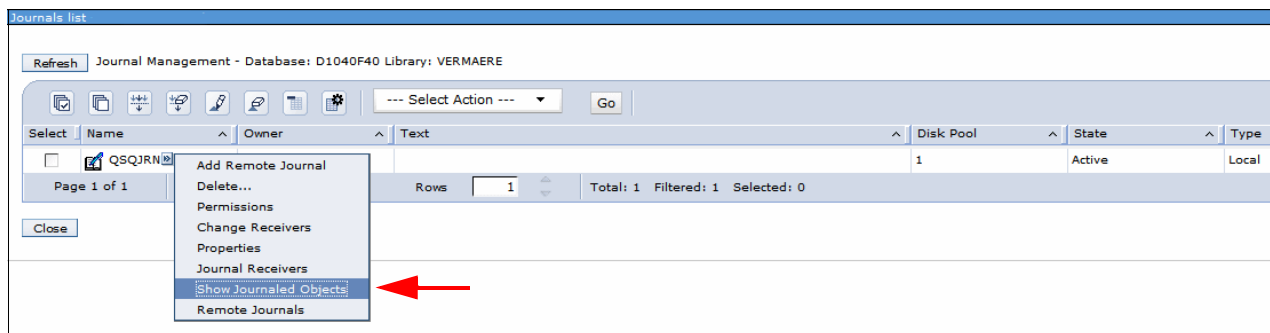


Figure 18-28 Showing journalized objects

This action displays the objects journalized to the journal including files, data areas, data queues, libraries, and integrated file system objects.

18.4.2 Change Receivers function

The Change Receivers action is added for a journal, as shown in Figure 18-29.

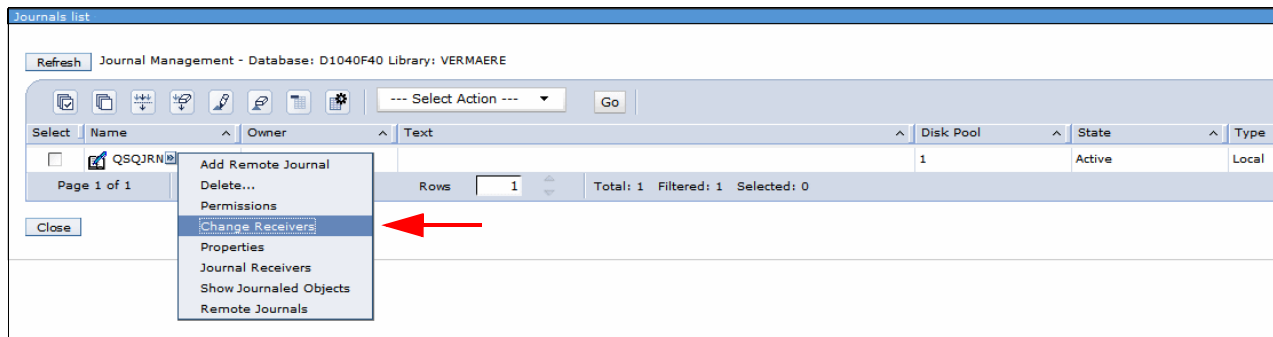


Figure 18-29 Changing receiver

This action allows you 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 more.

18.4.3 Remote Journals function

The Remote Journals action is added for a journal as shown in Figure 18-30.

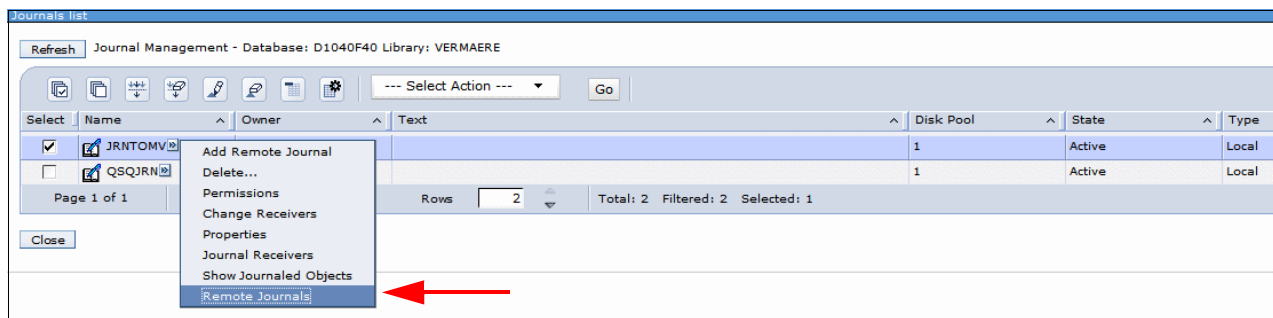


Figure 18-30 Remote journals

This action allows you to see a list of all remote journals associated with a journal. Included in the list are things such as the remote journal state and delivery mode.

18.4.4 Add Remote Journal function

The Add Remote Journal action was added for a journal as shown in Figure 18-31.

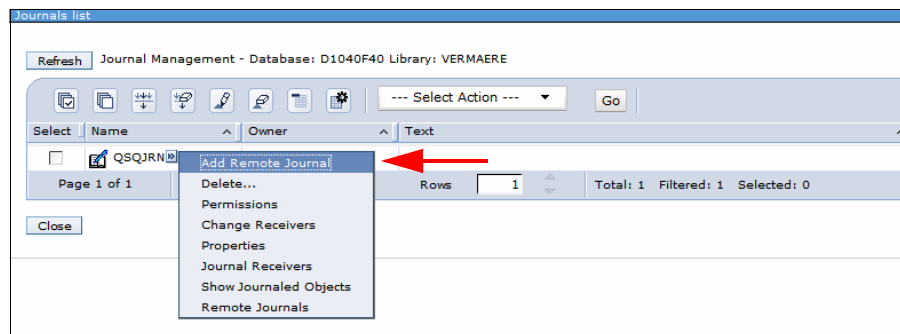


Figure 18-31 Adding a remote journal

This action allows you to define a new remote journal environment

18.4.5 Properties function

The Properties action was added for a journal receiver as shown in Figure 18-32.

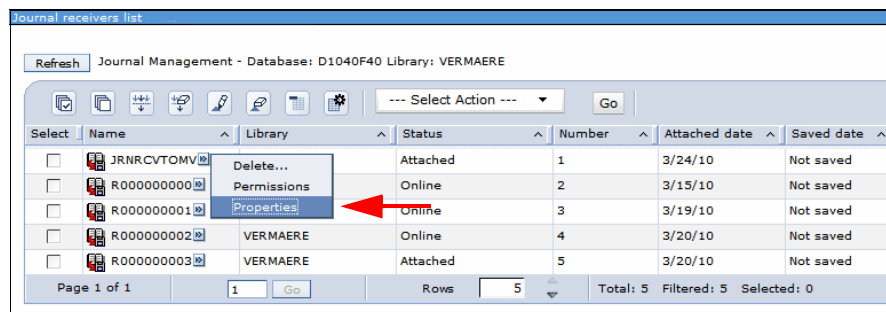


Figure 18-32 Properties action for a journal receiver

This action allows you to view the attributes of a journal receiver. Included in the attributes are things such as the sequence number of the first entry in the journal receiver, the minimize options, the fixed length data options, and more.

18.4.6 Activate function

The Activate action was added for remote journals as shown in Figure 18-33.




Figure 18-33 Activating a remote journal

This action allows you to activate a remote journal and specify the options desired.

18.4.7 Deactivate function

The Deactivate action was added for remote journals as shown in Figure 18-34.




Figure 18-34 Deactivating a remote journal

This action allows you to deactivate a remote journal.

18.4.8 Remove function

The Remove action was added for remote journals as shown in Figure 18-35.




Figure 18-35 Removal of a remote journal

This action allows you to remove a remote journal environment.

18.4.9 Properties function

The Properties action was added for remote journals as shown in Figure 18-36.




Figure 18-36 Properties for a remote journal

This action allows you to view the properties of a remote journal connection. This includes things such as the delivery mode, remote journal type, connection statistics, and more.

18.5 Integrated Server Administration enhancements

The following Integrated Server Administration topics were added to the IBM Systems Director Navigator for i in IBM i 7.1:

- ▶ Create Server
- ▶ Cloning a Server
- ▶ Delete Server
- ▶ Launch Web Console

For more information related to Integrated Server Administration, see Chapter 12, “Integration with BladeCenter and System x” on page 365.

18.5.1 Create Server task

The Create Server task was added for creating the IBM i hosting environment for an iSCSI attached BladeCenter or System x server as shown in Figure 18-37.

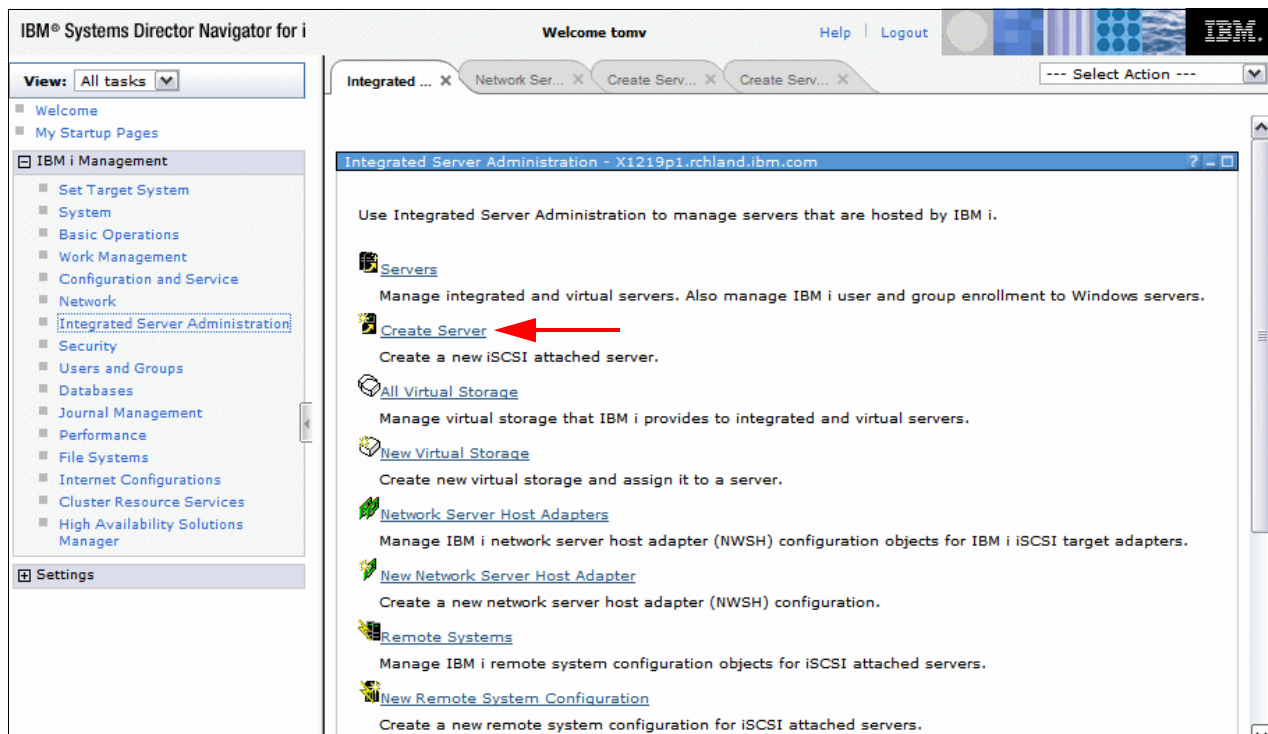


Figure 18-37 Creating a server task

Now the Create Server Wizard starts. This is discussed in more detail in 12.6.1, “Create Server task” on page 373.

18.5.2 Cloning a Server

It is also possible to create a server based on an existing one. This New Based on ... (cloning) a Server task is shown in Figure 18-38 on page 553. Figure 18-38 on page 553. It starts the Create Server Wizard and creates an iSCSI-attached integrated Windows server based on one that was previously installed.

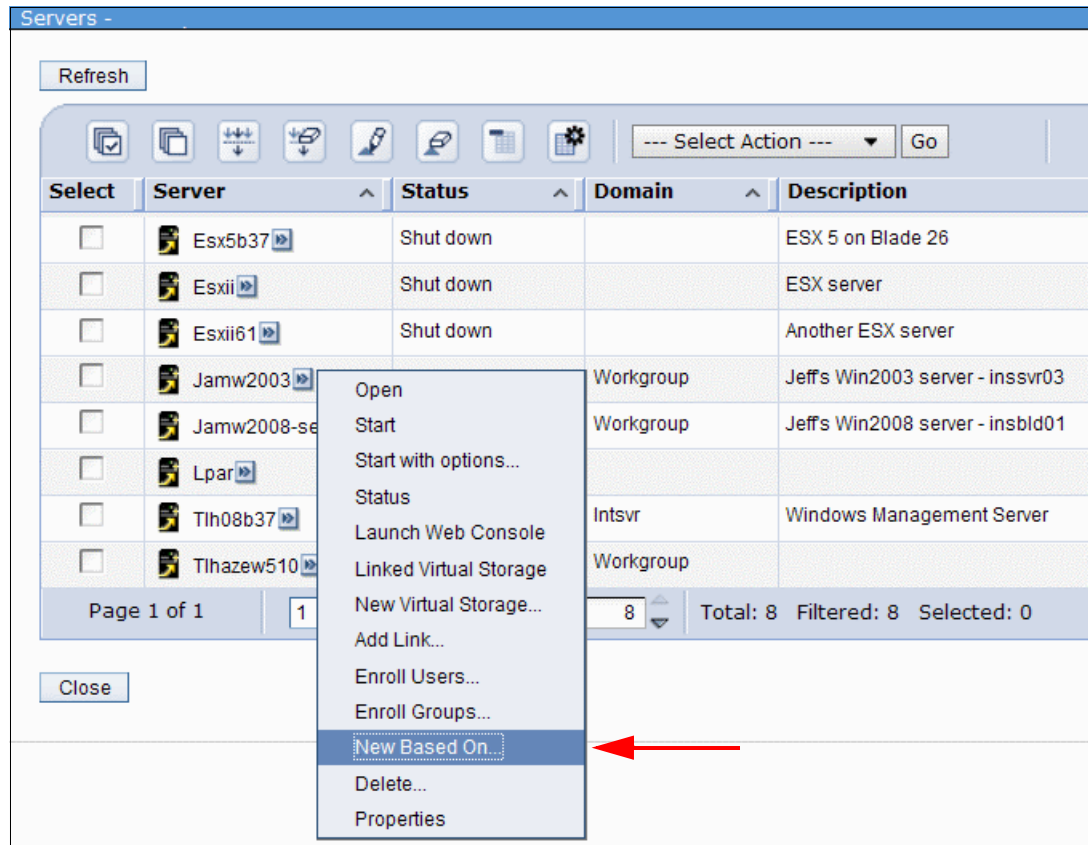


Figure 18-38 New Based On ... (cloning) a Server task

18.5.3 Delete Server task

The Delete Server task was added for deleting the IBM i hosting environment for an integrated server, as shown in Figure 18-39.

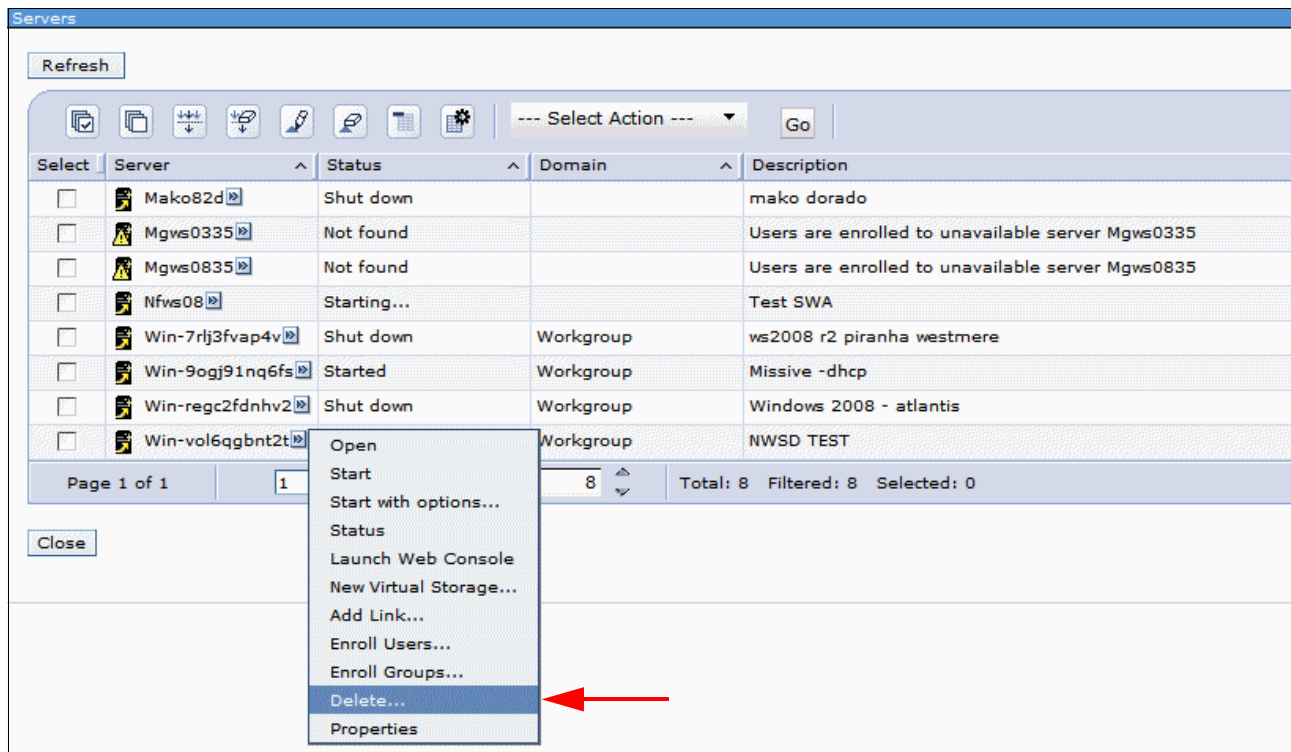


Figure 18-39 Deleting a server

18.5.4 Launch Web Console task

The Launch Web Console task was added for launching the service processor web console for an iSCSI attached integrated server, as shown in Figure 18-40.

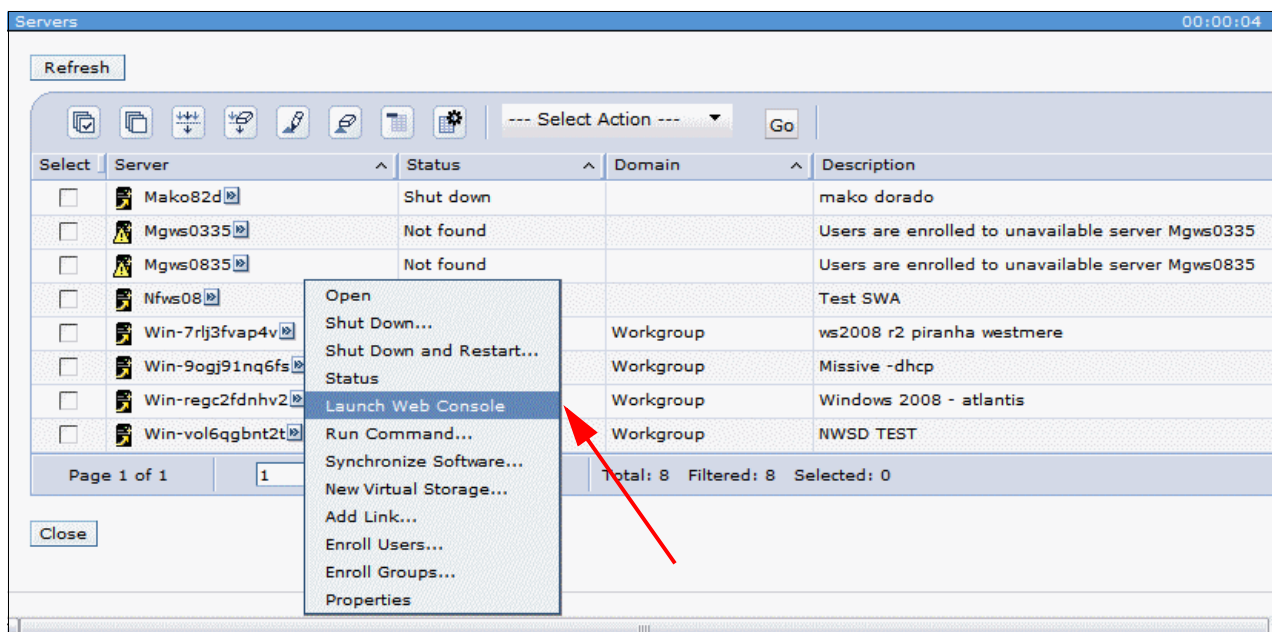


Figure 18-40 Launching the web console

Starting with IBM i 7.1, several iSCSI configuration tasks were now simplified. More information can be found in 12.6, “IBM Systems Director Navigator for i” on page 372.

18.6 Printer output enhancements

In IBM i 7.1, new actions are added to View as PDF and Export as PDF.

18.6.1 View as PDF action

The View as PDF action was added for an AFPDS or SCS printer output file, as shown in Figure 18-41.

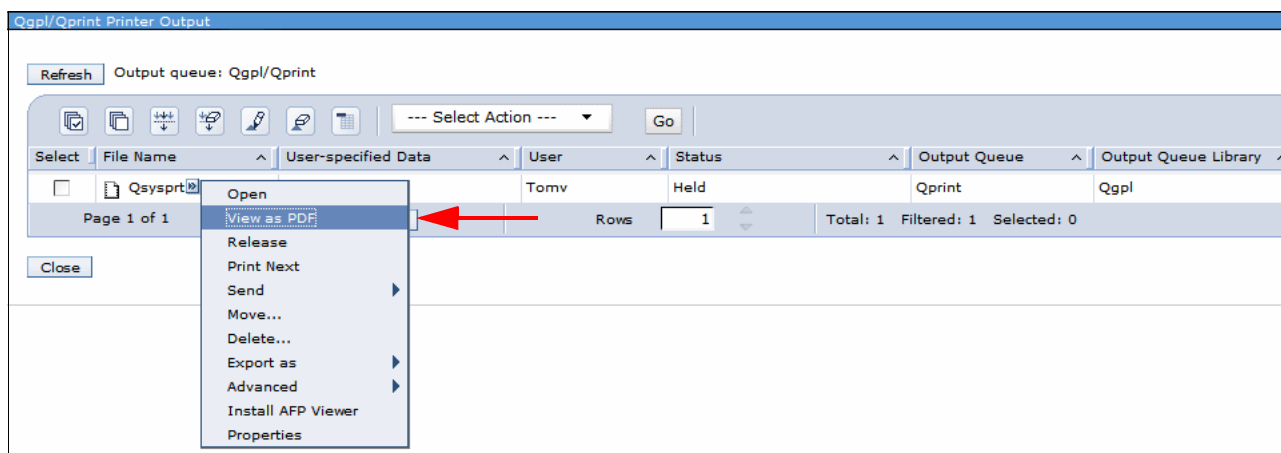


Figure 18-41 Printer output: Viewing as PDF

This function allows you to open and view the contents of a printer output file using Acrobat Reader.

18.6.2 Export as PDF action

The Export as PDF function was added for an AFPDS or SCS printer output file, as shown in Figure 18-42.

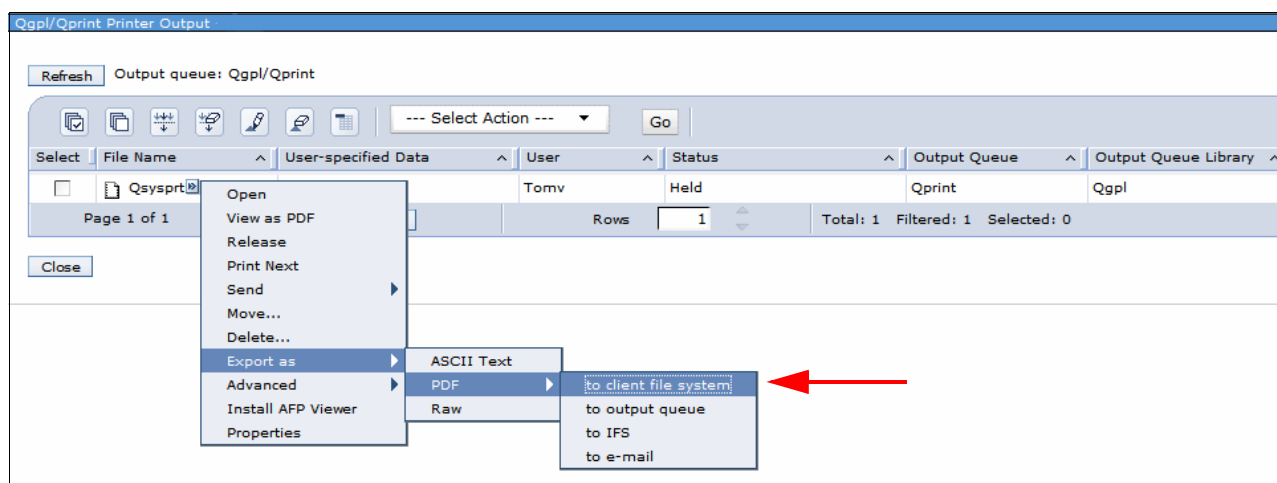


Figure 18-42 Printer output: Exporting as PDF

This function allows you 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

Note: For the latter three options, the InfoPrint Server Licensed Program Product (5722-IP1) is required. Users can use the native IBM Transform Services for i (5770-TS1) Licensed Program Product to export to IFS, but they must map a network drive to the IFS, then select the first option as shown in Figure 18-42 on page 556.

18.7 File systems enhancements

A temporary file system is supported. The temporary file system can be created through 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 18-43.

New User-Defined File System

Auxiliary storage pool: /dev/QASP01

UDFS name: TOMV.TMPUDFS

Description: test with GUI

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 18-43 Temporary UDFS creation

For more information related to temporary user-defined file system support, see 20.3, “Temporary user-defined file systems” on page 633.

18.8 Networking enhancements

Internet Protocol version 6 (IPv6) is enabled for Remote Access Services, as shown in Figure 18-44.

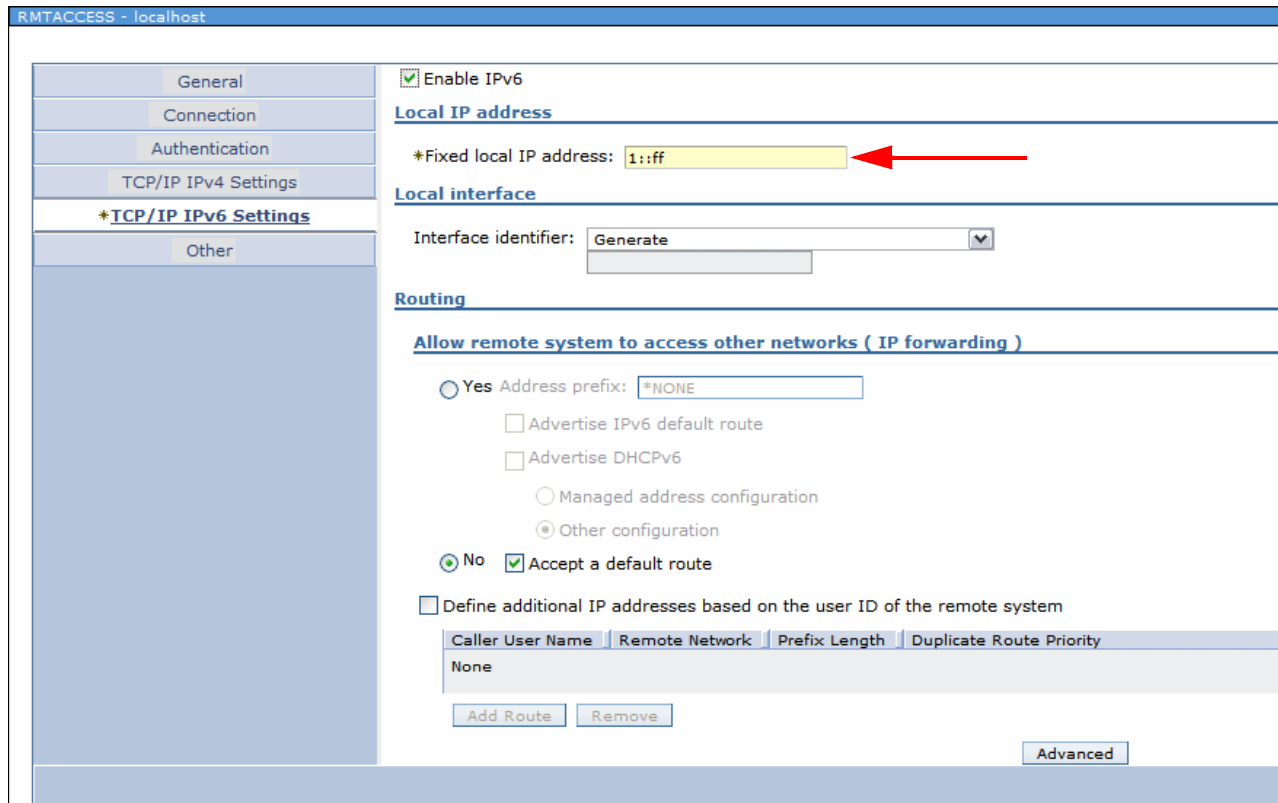


Figure 18-44 IPv6 support for remote access services

IKE version 2 is enabled for VPN, as shown in Figure 18-45.

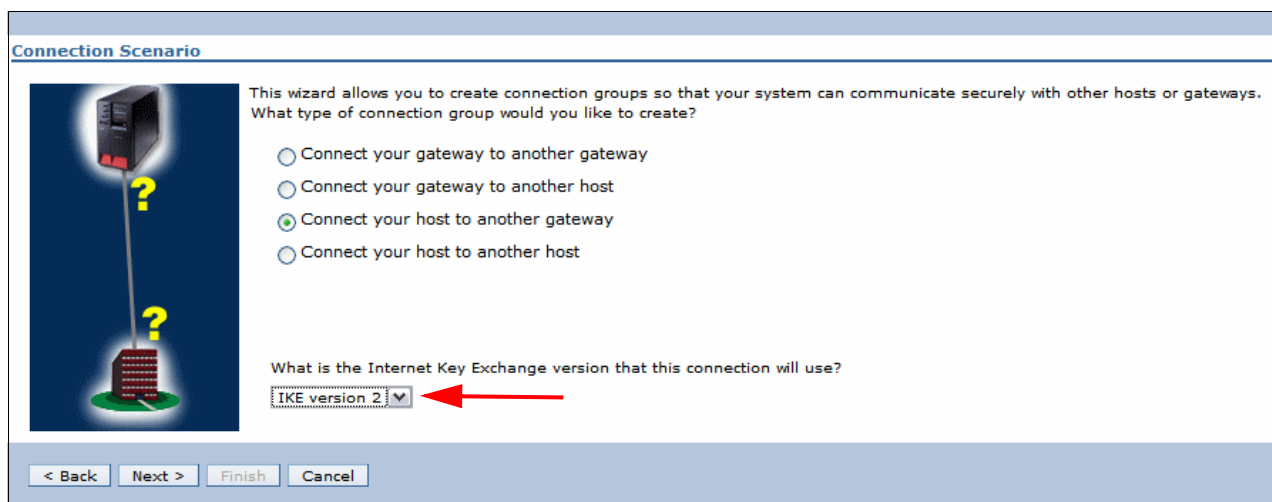


Figure 18-45 Internet key exchange V2 enablement for VPN

For more information related to networking enhancements in IBM i 7.1, see Chapter 10, “Networking enhancements” on page 315.

18.9 Disk management enhancements

There are several enhancements within Systems Director Navigator for IBM i 7.1 related to disk management:

- ▶ Graphical View
- ▶ Start/stop encryption on disk pools
- ▶ Asynchronous transmission delivery for Geographic Mirroring

18.9.1 Graphical view enhancements

To access the graphical view perform the following steps:

1. From the Welcome Page, click **Configuration and Service** → **Disk Units**, as shown in Figure 18-46.

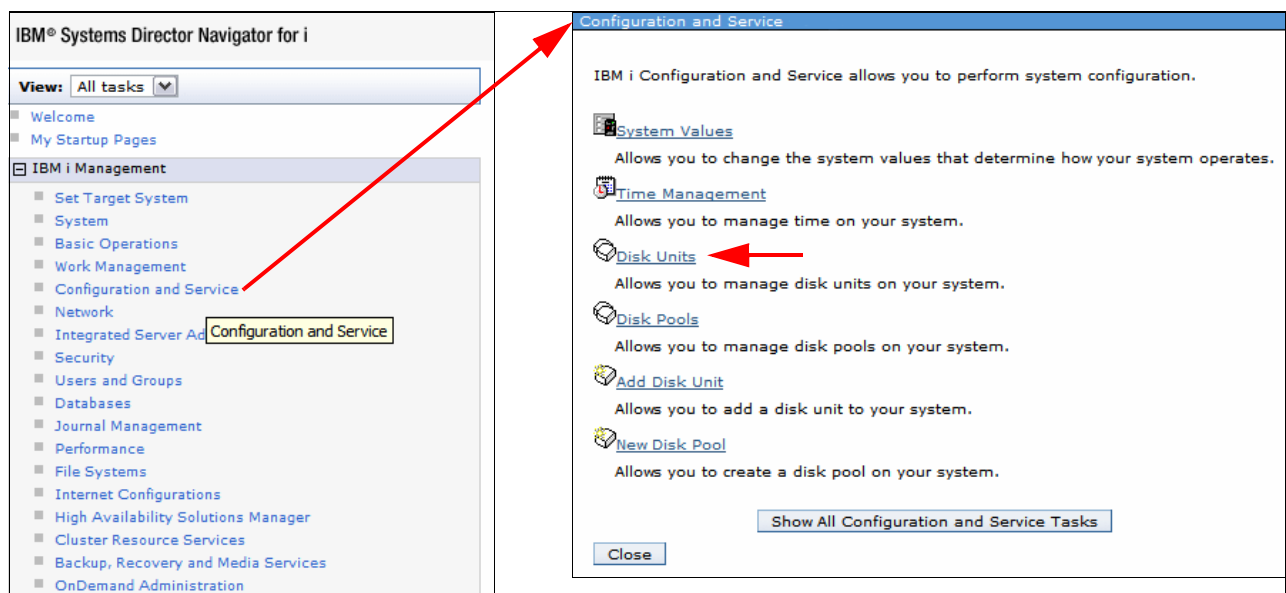


Figure 18-46 Configuration and service: Disk units

- From the Select Action Menu at the top, click **Graphical View** and click **Go**. The graphical view shown in Figure 18-47 is displayed.

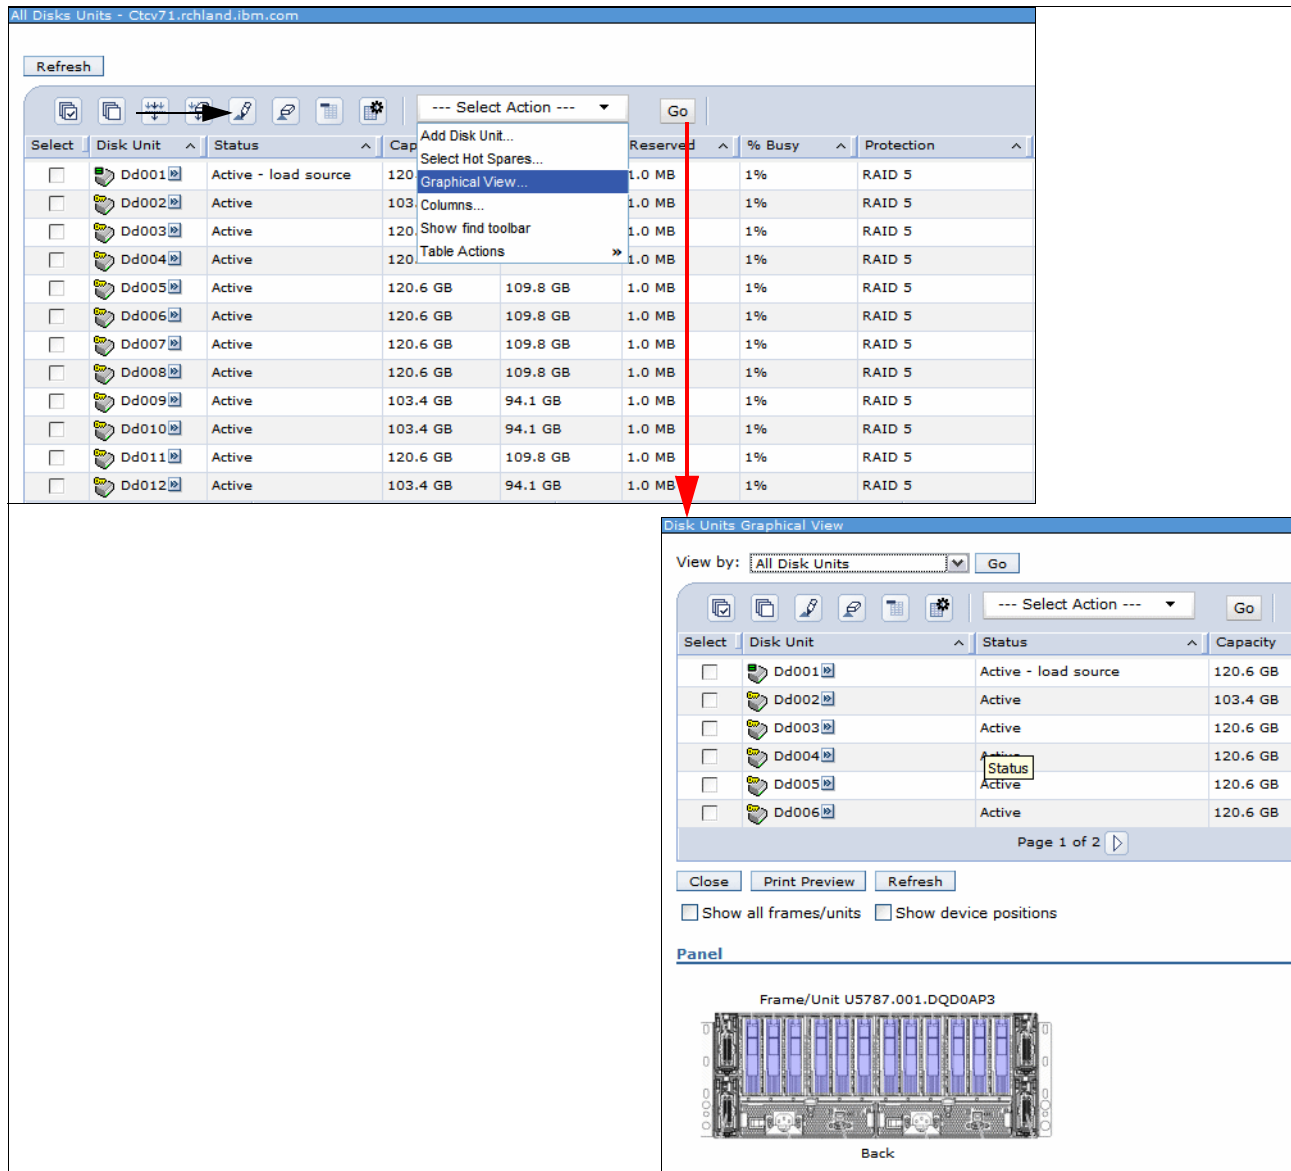


Figure 18-47 Disk units graphical view

18.9.2 Start/stop encryption on disk pools enhancements

The Start/Stop encryption on disk pools is now enabled within IBM Systems Director Navigator for i. Change encryption key is allowed when the encryption is started, as shown in Figure 18-48.

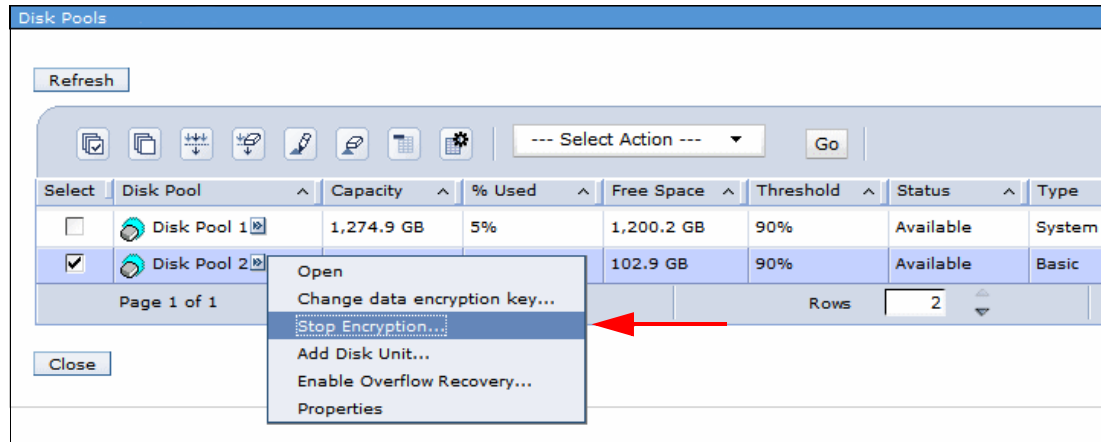


Figure 18-48 Disk management: start/stop encryption

18.9.3 Asynchronous delivery for Geographic Mirroring enhancements

IBM Systems Director Navigator for i also provides asynchronous communications between the source and target servers for geographic mirroring, as shown in Figure 18-49.

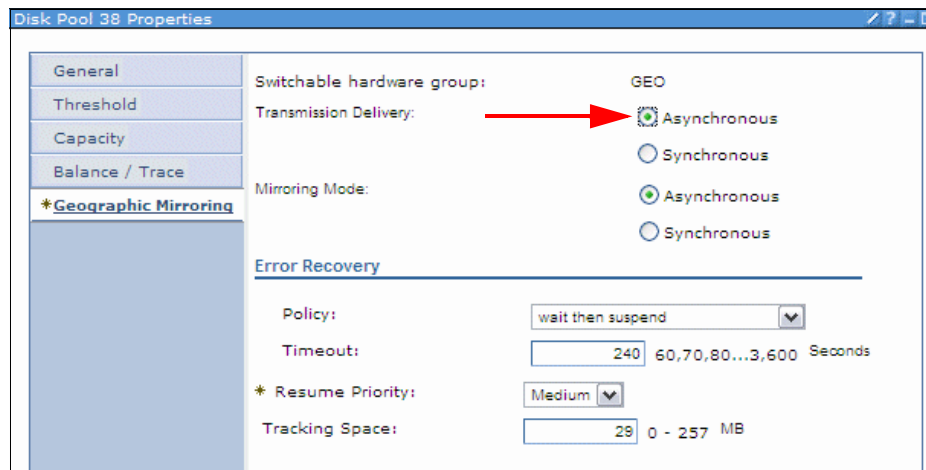


Figure 18-49 Asynchronous delivery mode

For more information related to Asynchronous transmission delivery for Geographic Mirroring, see 5.1.7, “Asynchronous geographic mirroring” on page 100.

18.10 Tape support enhancements

On the Welcome Page, click **Configuration and Service** → **Show All Configuration and Service tasks** to access the new **Tape Devices** option, as shown in Figure 18-50.

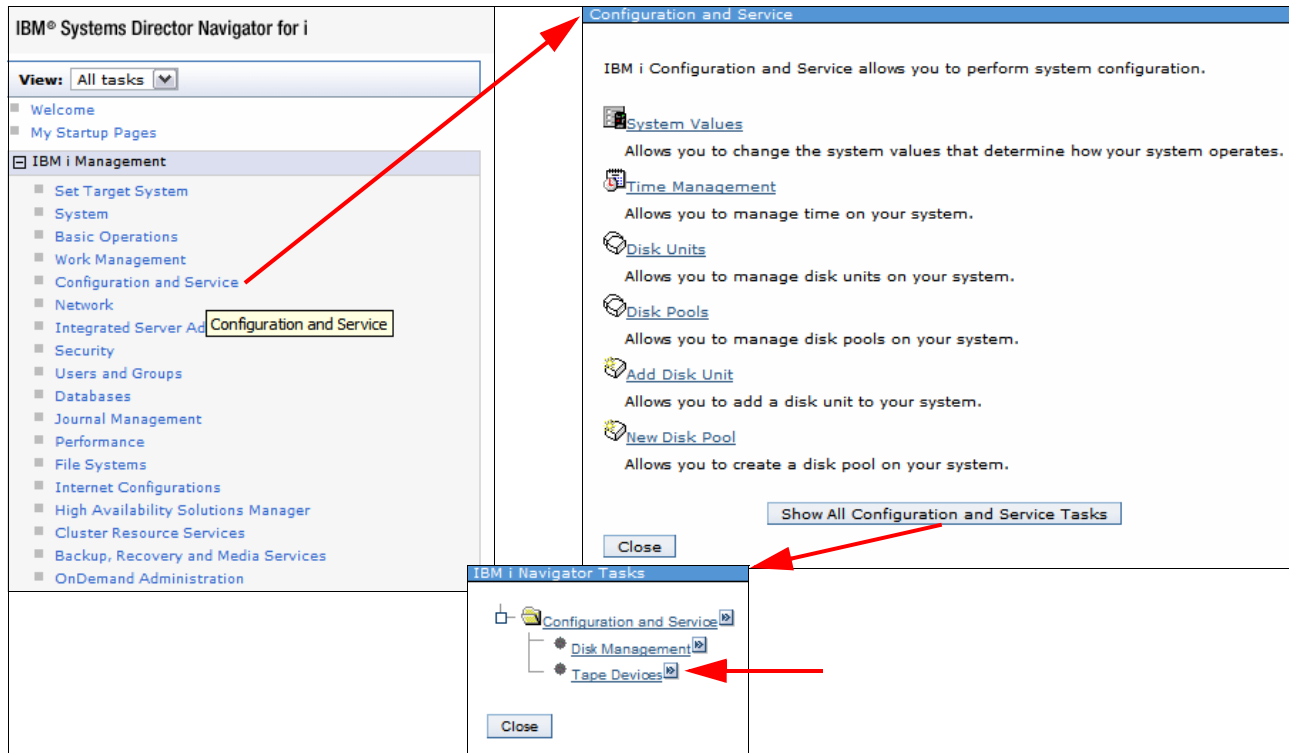


Figure 18-50 Showing configuration and service tasks

From the Tape Devices pop-up menu, you have the following options available, as shown in Figure 18-51 on page 563:

- ▶ Stand-Alone Devices offer the following possibilities:
 - make (un)available
 - look into the properties
 - upgrade the firmware
- ▶ Tape Image Catalogs offer 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

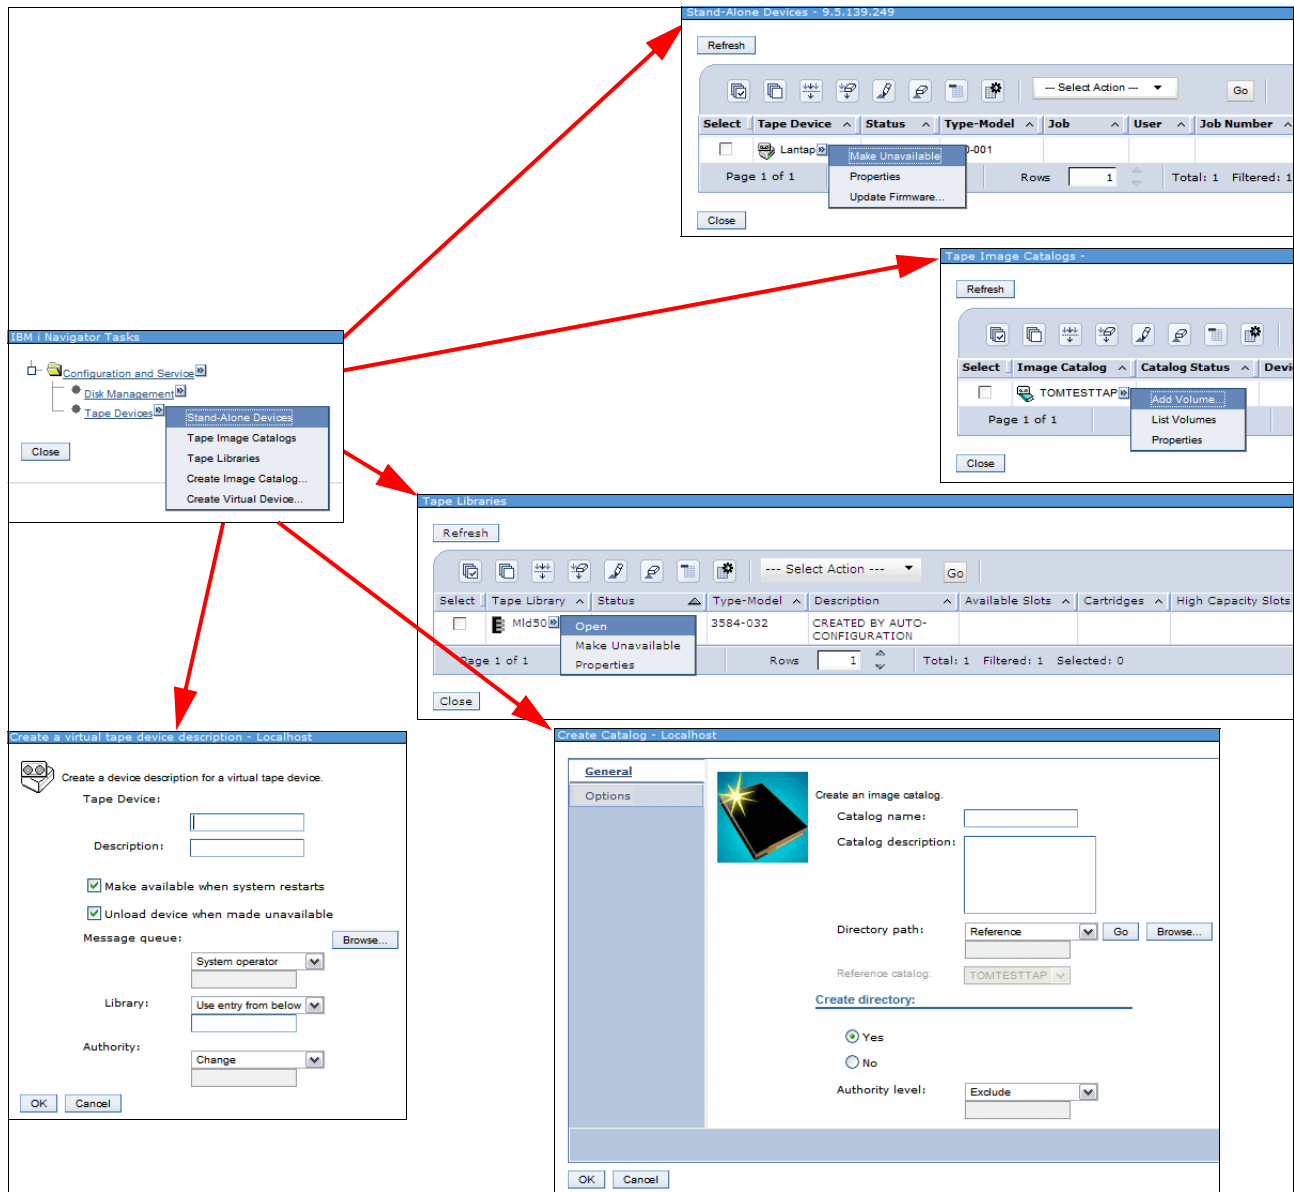


Figure 18-51 New tape support in Systems Director Navigator for i

18.11 Performance enhancements

In order to take advantage of the Performance enhancements within IBM System Director Navigator, we assume the current IBM HTTP Server for i 7.1 Group PTF (SF99368) got installed on the system.

Please refer to the following for more information:

http://www-912.ibm.com/s_dir/sline003.nsf/554c38c4848b77f2862567bd0046e003/56ff79c9506270c98625771400478c7f?OpenDocument

Figure 18-52 shows the interface to work with Performance related within IBM i 7.1.

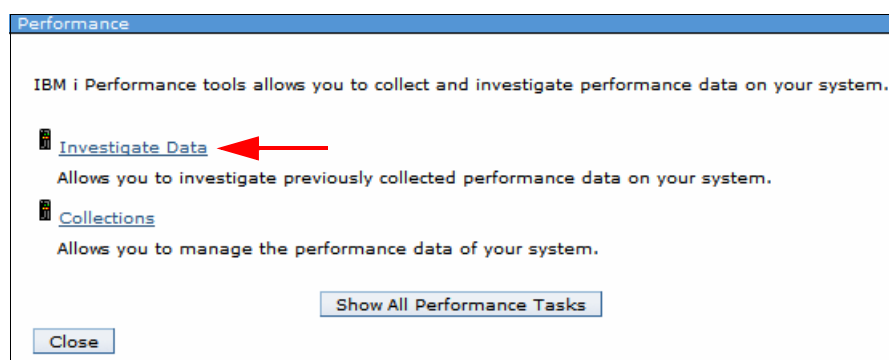


Figure 18-52 Performance

Several enhancements were made to the Performance Data Investigator (PDI), which can be accessed by selecting the **Investigate Data** task as shown in Figure 18-53.

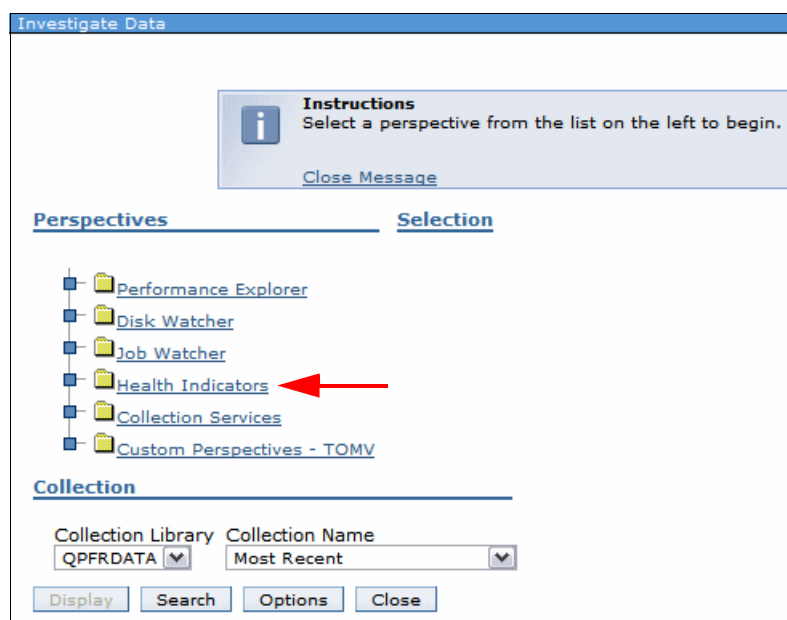


Figure 18-53 Investigate Data option

For more in-depth information to work with Performance related tasks within IBM i 7.1, see Chapter 7, "Performance tools" on page 189.

18.11.1 General health indicators

This new content package helps a system administrator, an IBM i performance specialist, or a general user learn if the partition is performing properly from a general performance management perspective or if there are areas of the system that need to be investigated more closely.

Main system resources and components (such as CPU, DASD, 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 18-54.

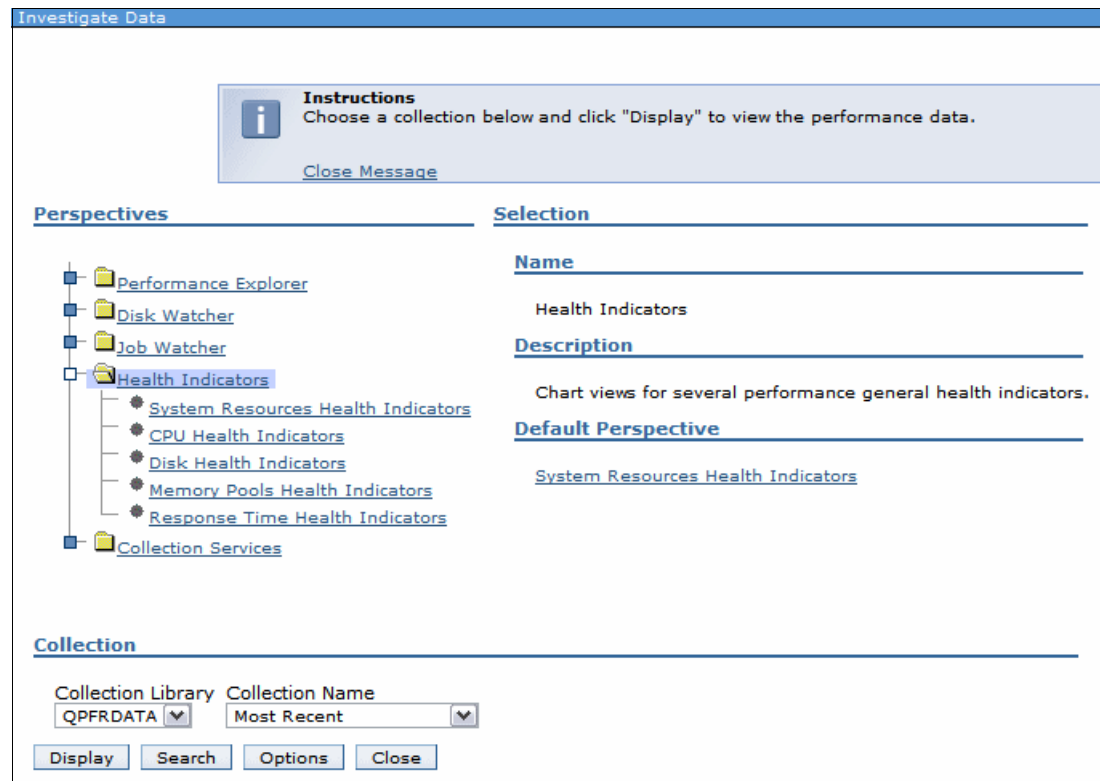


Figure 18-54 Health indicators

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

This new perspective can be used to determine the general health of the main system resources according to the defined thresholds. The chart in Figure 18-55 shows Health Indicators for CPU, Disk, Memory and Communications for the entire selected collection.

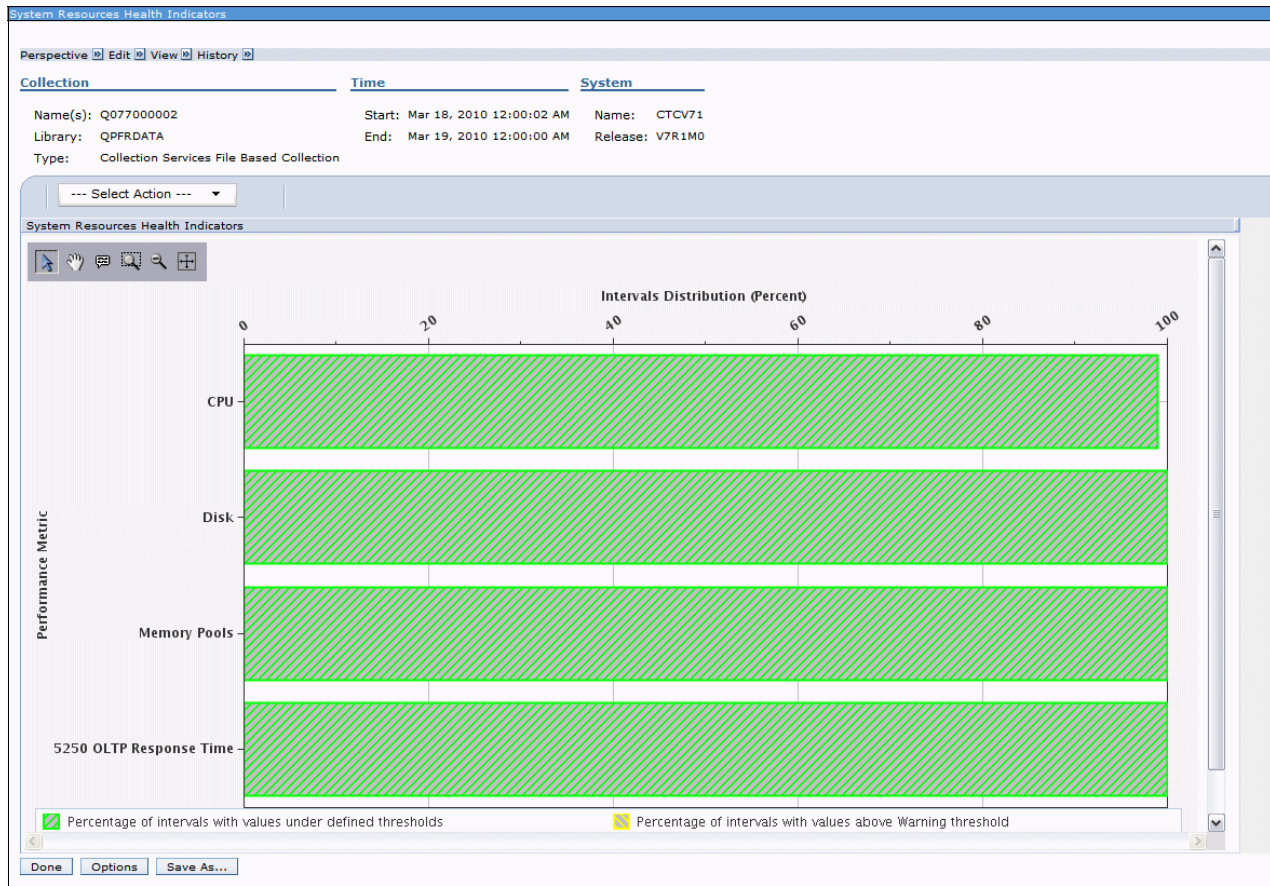


Figure 18-55 System resources health indicators

This perspective can be used 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, the user 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 consist of key metrics associated with the system resource being investigated. The perspectives are available through the Actions menu, as shown in Figure 18-56 and are discussed in more detail in the following sections.

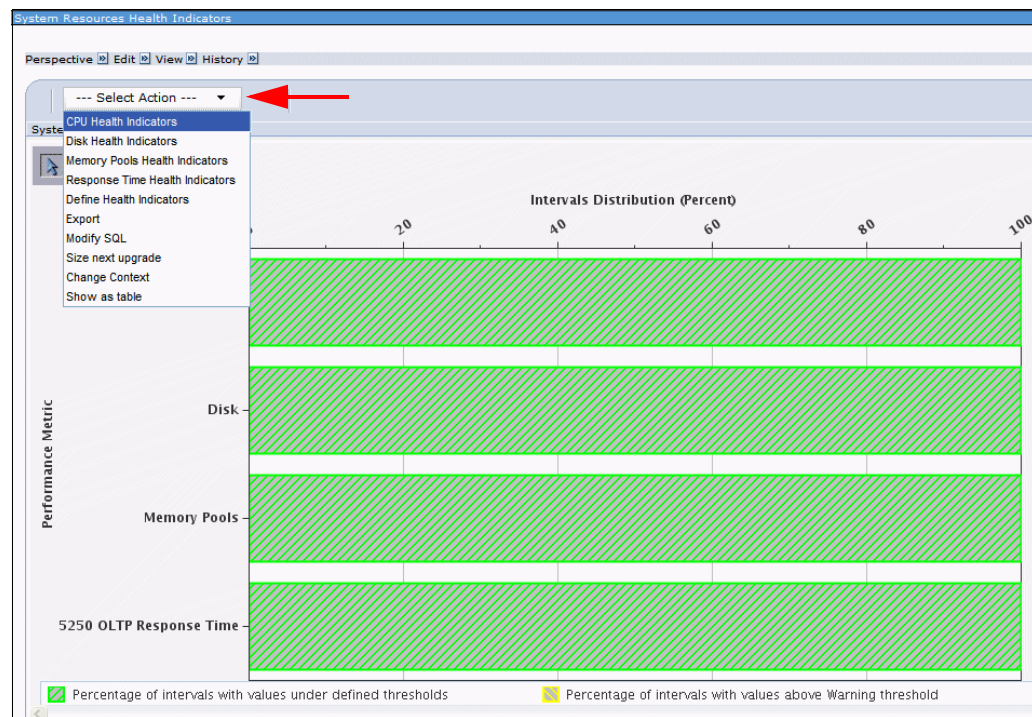


Figure 18-56 System resources health indicators drill-down menu

CPU health indicators

This perspective can be used to determine the proportion of intervals where CPU health indicators exceeded the defined thresholds. The chart in Figure 18-57 shows CPU health indicators by analyzing all collection time intervals according to the defined thresholds for CPU.

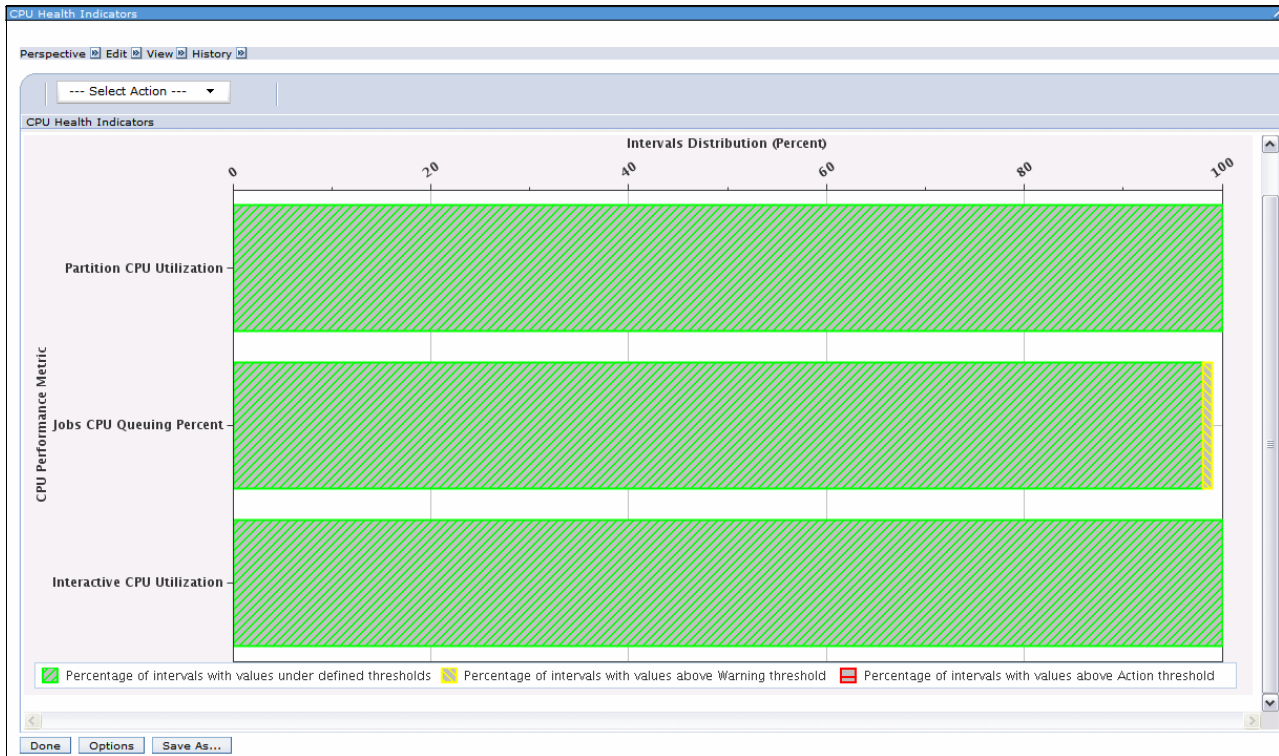


Figure 18-57 CPU health indicators

From the CPU Health Indicators perspective, the user can navigate to the following perspectives:

- ▶ CPU Utilization and Waits Overview
- ▶ CPU Utilization Overview
- ▶ Interactive Capacity CPU Utilization

Disk health indicators

The perspective in Figure 18-58 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 disk health indicators exceeded the defined thresholds.




Figure 18-58 Disk health indicators

From the Disk Health Indicators perspective, the user 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 18-59 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 memory pools health indicators exceeded the defined thresholds.

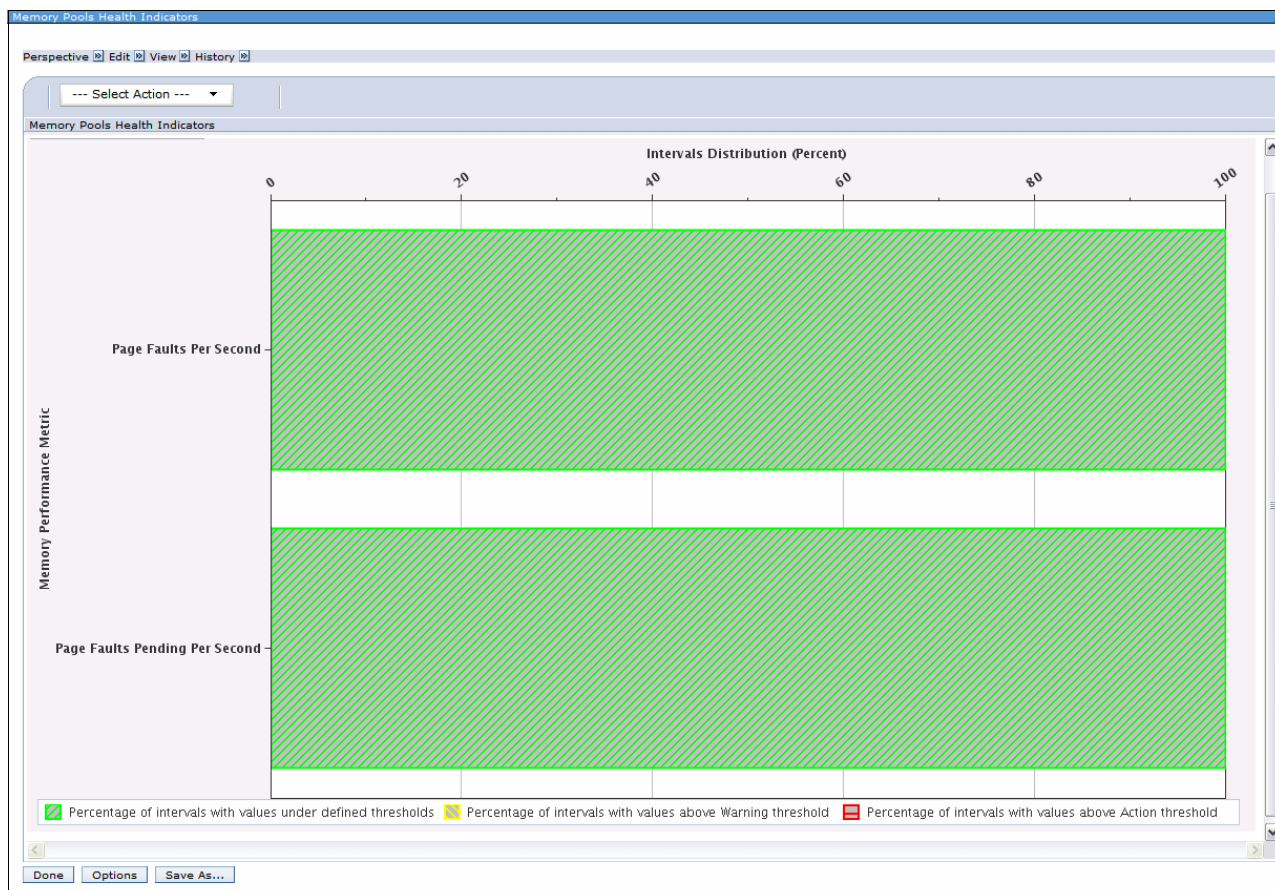


Figure 18-59 Memory pools health indicators

From the Memory Pools Health Indicators perspective, the user can navigate to the following perspectives:

- ▶ Resource Utilization Overview
- ▶ Page Faults Overview

Response time health indicators

The chart in Figure 18-60 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 response time exceeded the defined thresholds.




Figure 18-60 Response time health indicators

From this perspective, the user 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 business requirements. They can be configured by selecting the **Define Health Indicators** action from the Select Action list, as shown in Figure 18-61.

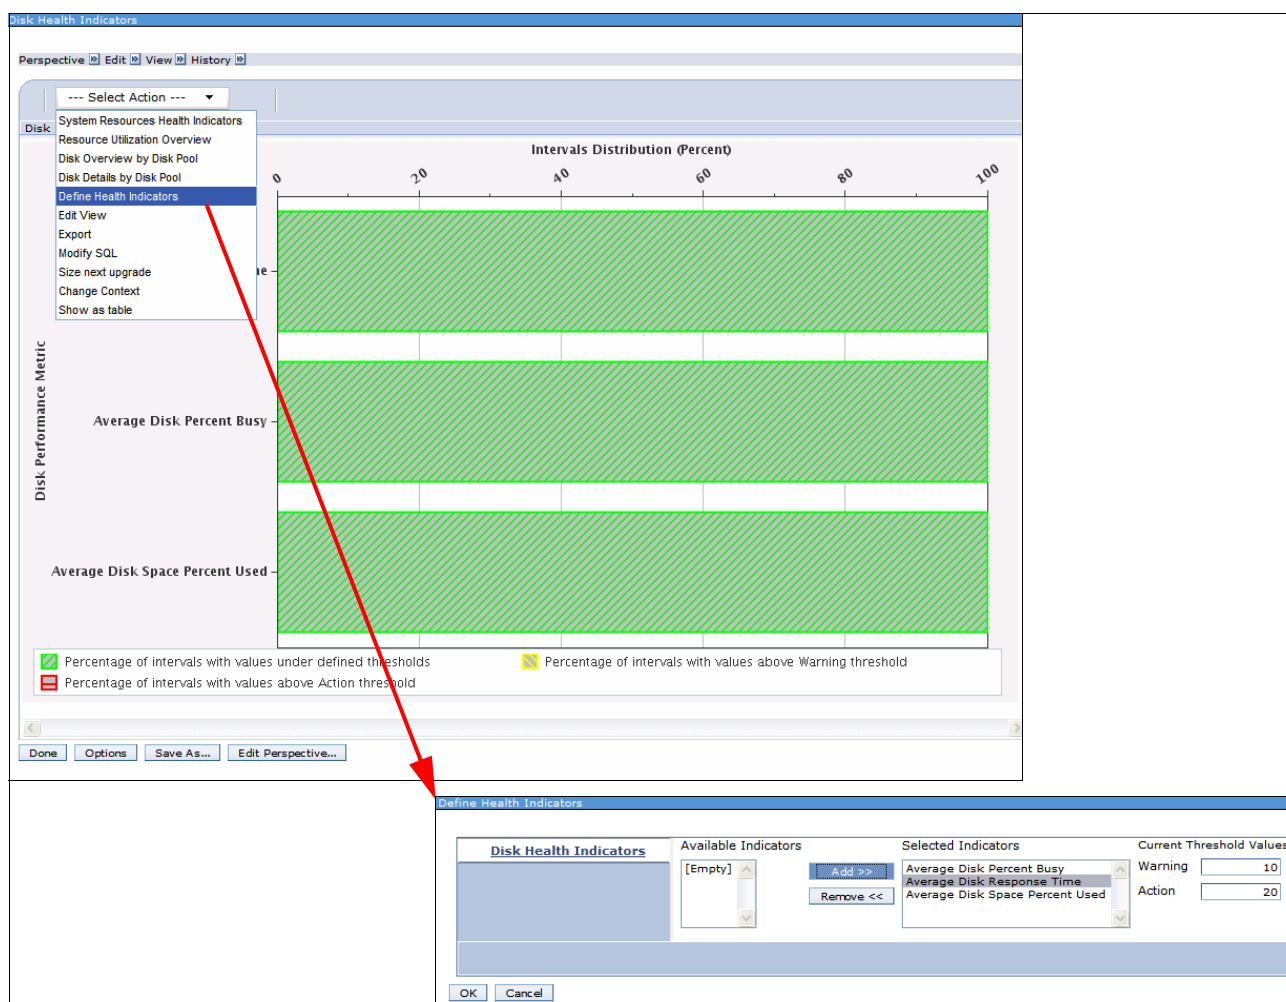


Figure 18-61 Defining the disk health indicators threshold

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 18-61 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 properly defining those threshold values, the user has a point of reference to visually determine if the system is between its defined control limits. Figure 18-62 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.

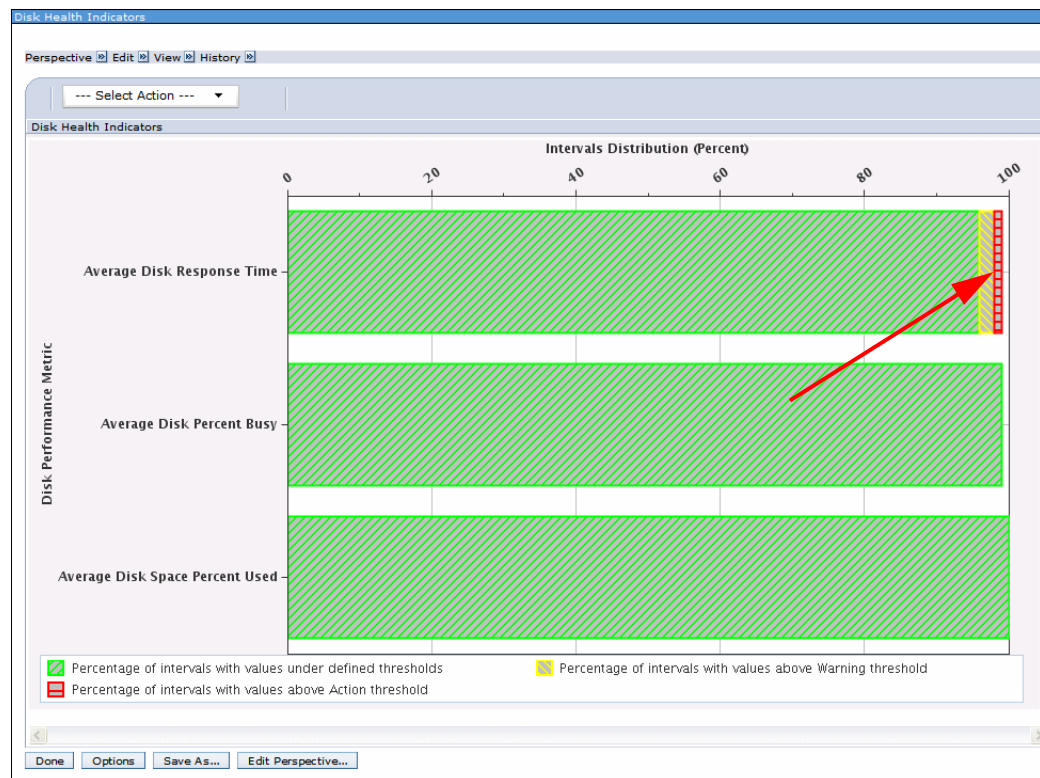


Figure 18-62 Threshold reached for average disk response time

It is possible to define various thresholds for a specific Graphical Health Indicator graph. Figure 18-63 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




Figure 18-63 Multiple thresholds can be defined

All the individual thresholds can be added, removed and tailored to your own specifications.

18.11.2 Viewing cross-partition CPU

Collection Services has the ability to 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.

This is available on Power 6 and later servers, with a minimum firmware level xx340_061.

When this data is available, it can be viewed through several perspectives found under "Physical System". See Figure 18-64.



Figure 18-64 Physical system charts

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 needs to be enabled on the HMC, as shown in Figure 18-65.

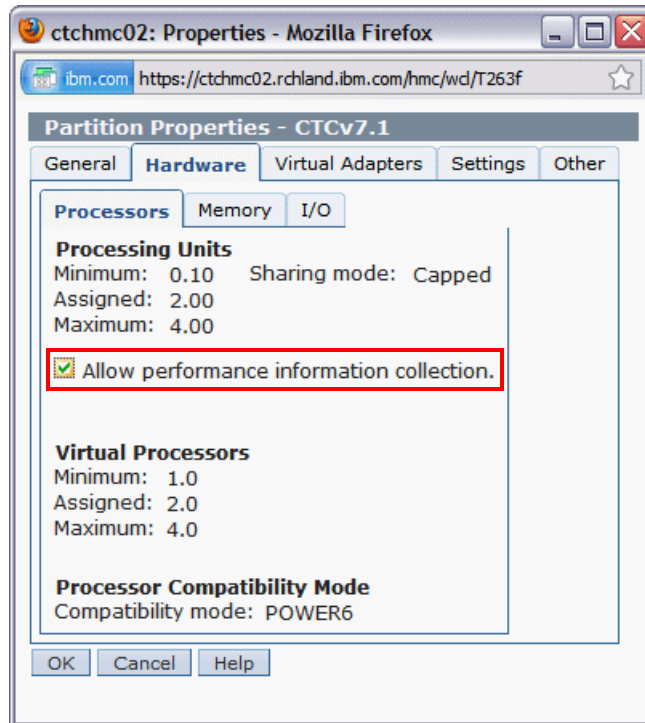


Figure 18-65 Allowing the performance information collection on the HMC

To turn on this option, follow these steps from the HMC:

1. Select **Systems Management** → **Servers**.
2. Click **your IBM i**.
3. Select the partition profile.
4. Click **Properties**.
5. Select the Hardware tab.
6. Click **Processors**.
7. Select **Allow performance information collection**.

Logical partitions overview

Figure 18-66 shows configuration data and CPU use for all logical partitions on the system, including operating system, number of virtual processors, partition memory, donated CPU time, uncapped CPU time used and others.




Figure 18-66 Logical partitions overview

Donated processor time by logical partition

The chart in Figure 18-67 shows the processor time that has been donated by dedicated processor logical partitions that are configured to donate unused processor cycles.




Figure 18-67 Donated processor time by logical partition

Uncapped processor time used by logical partition

The chart in Figure 18-68 shows the uncapped processor time that has been used by logical partitions in excess of their entitled processing capacity configured.




Figure 18-68 Uncapped processor time used by logical partition

Virtual shared processor pool use

The chart in Figure 18-69 shows processing capacity available and processing capacity use for virtual shared processor pools.




Figure 18-69 Virtual shared processor pool use

Physical processors use by physical processor

The chart in Figure 18-70 shows the use percent for each physical processing unit.

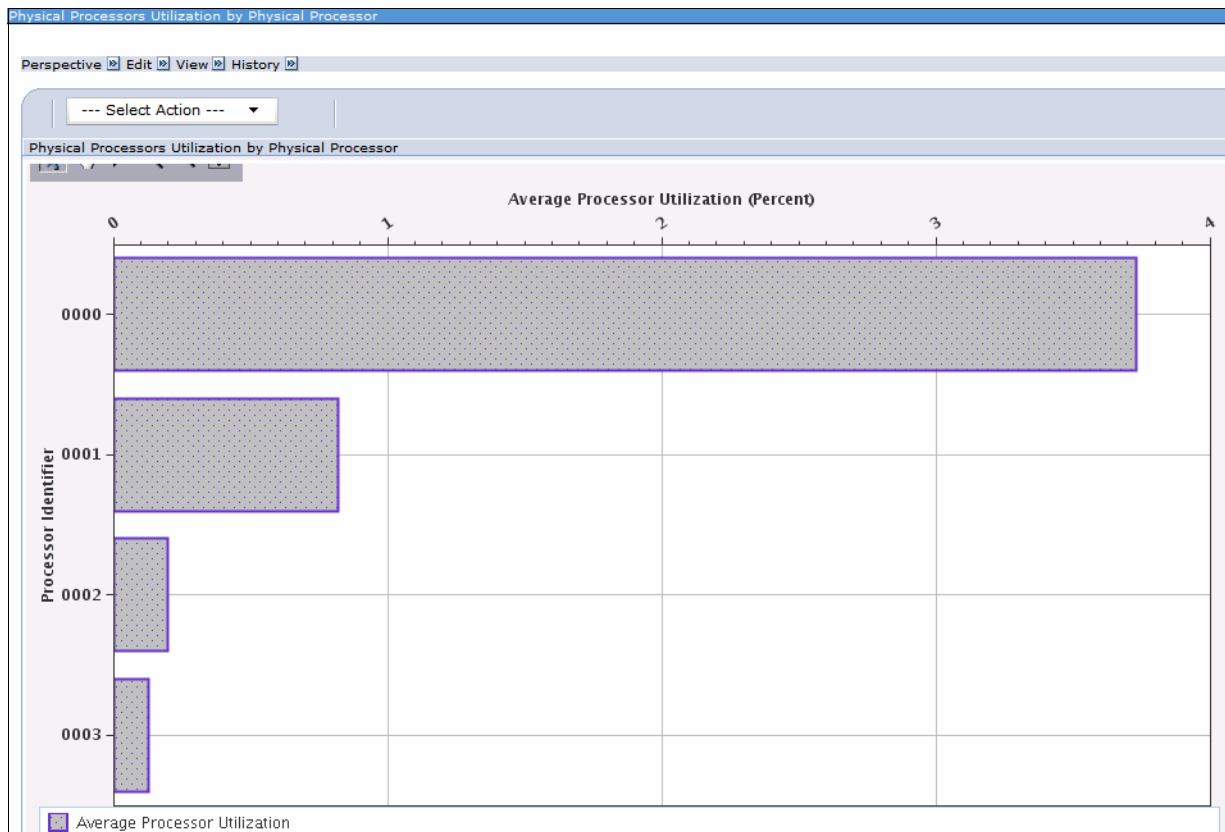


Figure 18-70 Physical processor use by physical processor

Dedicated processors use by logical partition

The chart in Figure 18-71 shows the dedicated processor use by logical partitions.

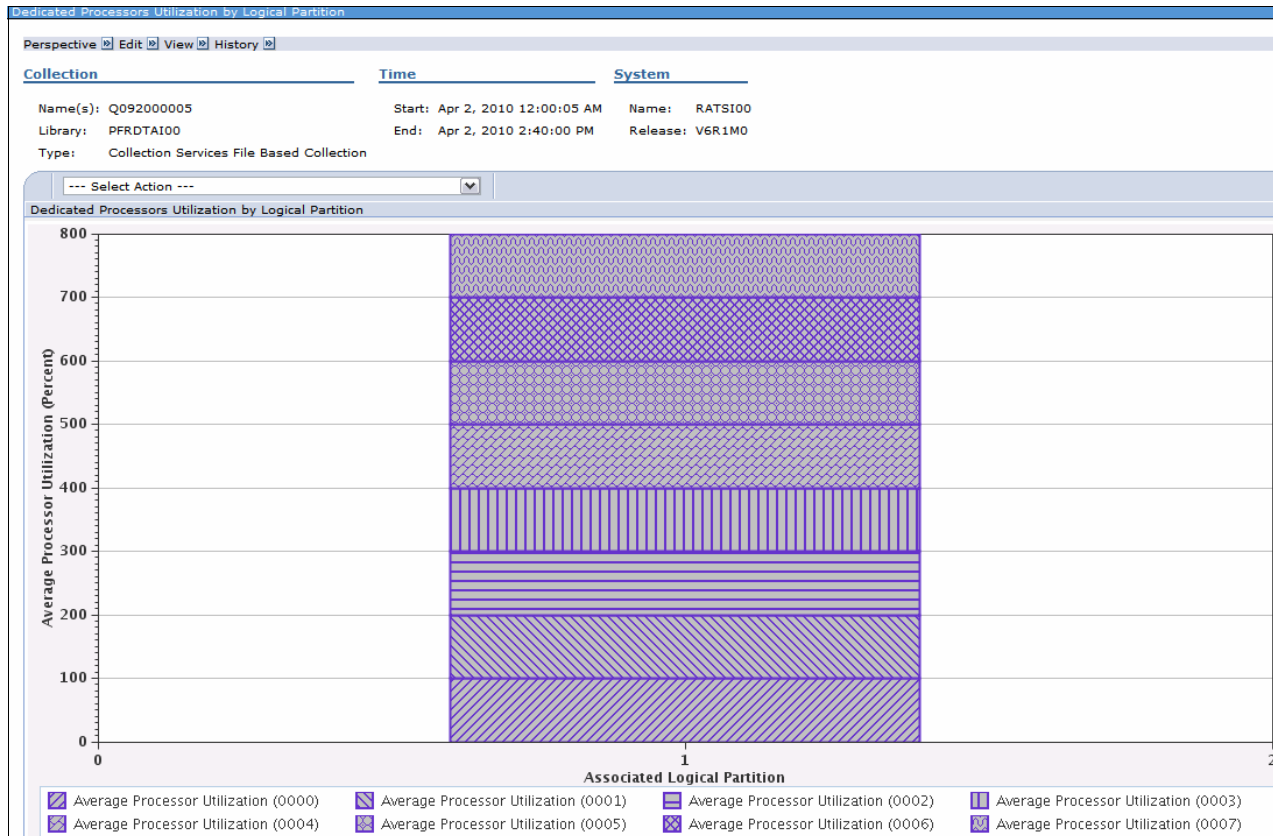


Figure 18-71 Dedicated processors use by logical partition

Physical processors use by processor status overview

The chart in Figure 18-73 on page 583 shows a summary of processors use for the entire collection identifying use for dedicated processors and for shared processors.

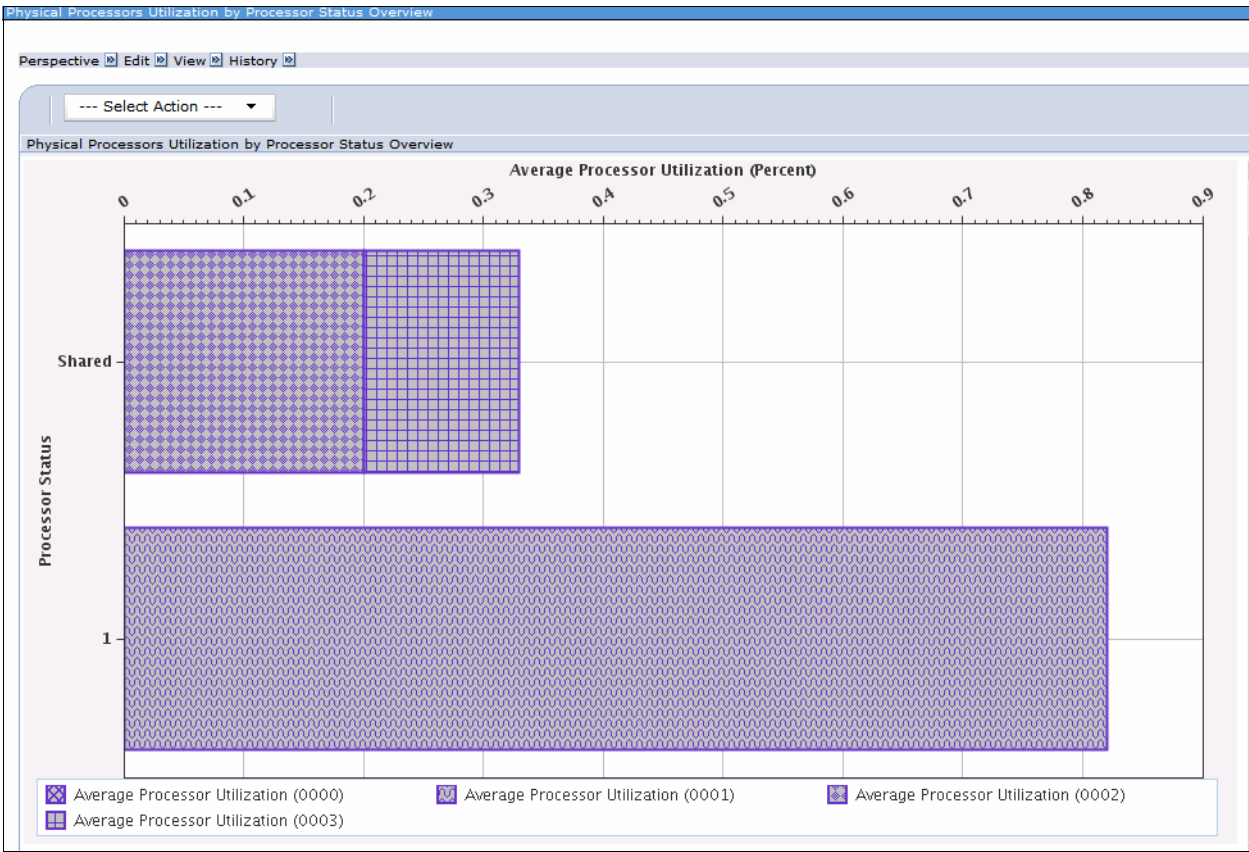


Figure 18-72 Processor use

Physical processors use by processor status detail

The table in Figure 18-73 shows processor use over time. It shows the dedicated processors and the shared processors.

Physical Processors Utilization by Processor Status Detail					
Perspective Edit View History					
<div> <div> <div></div> <div></div> <div></div> <div></div> <div></div> <div></div> </div> <div>--- Select Action ---</div> </div>					
Select	Date - Time	Interval Number	Processor Status	Processor Identifier	Average Processor Utilization (Percent)
<input type="checkbox"/>	Dec 16, 2008 11:06:15 PM	1		0000	14.11
<input type="checkbox"/>	Dec 16, 2008 11:06:15 PM	1		0001	13.94
<input type="checkbox"/>	Dec 16, 2008 11:06:15 PM	1		0002	7.94
<input type="checkbox"/>	Dec 16, 2008 11:06:15 PM	1		0003	3.6
<input type="checkbox"/>	Dec 16, 2008 11:06:30 PM	2		0000	3.29
<input type="checkbox"/>	Dec 16, 2008 11:06:30 PM	2		0001	1.23
<input type="checkbox"/>	Dec 16, 2008 11:06:30 PM	2		0002	0.24
<input type="checkbox"/>	Dec 16, 2008 11:06:30 PM	2		0003	0.13
<input type="checkbox"/>	Dec 16, 2008 11:06:45 PM	3		0000	4.86
<input type="checkbox"/>	Dec 16, 2008 11:06:45 PM	3		0001	2.15
<input type="checkbox"/>	Dec 16, 2008 11:06:45 PM	3	Processor Status	0002	0.56
<input type="checkbox"/>	Dec 16, 2008 11:06:45 PM	3		0003	0.09
Total: 300			Filtered: 300		

Figure 18-73 Physical processors use by processor status detail

Shared memory overview

The chart in Figure 18-74 shows shared memory pool use by the partition on an interval-by-interval basis. This information can be used to find the partition's use of the memory shared pool and metrics that are the sum of activity caused by all partitions using the memory shared pool.




Figure 18-74 Shared memory overview

The performance collection must be turned on for the IBM i partition to collect the data. This option needs to be enabled on the HMC, as shown in Figure 18-75.

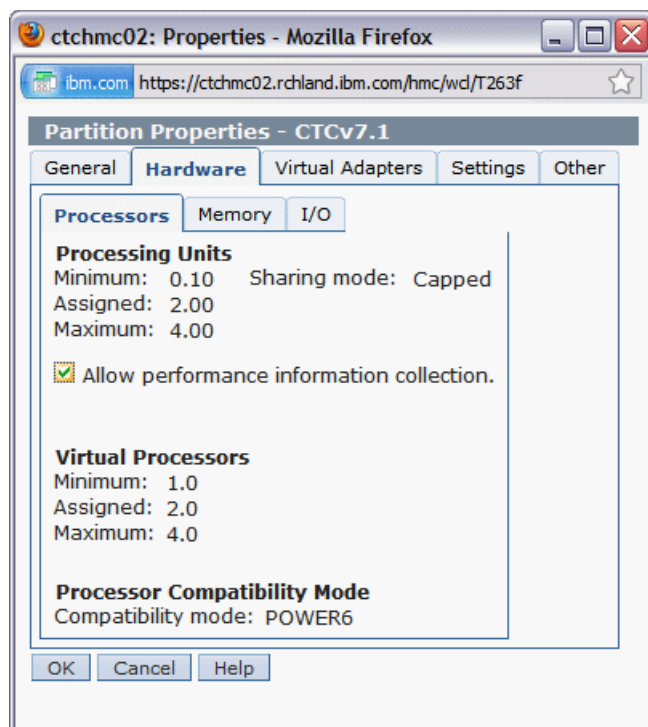


Figure 18-75 Allowing the performance information collection on the HMC

To activate performance collection, perform the following steps from the HMC:

1. Select **Systems Management** → **Servers**.
2. Click **your IBM i**.
3. Select the partition profile.
4. Click **Properties**.
5. Select the Hardware tab.
6. Click **Processors**.
7. Select **Allow performance information collection**.

18.11.3 QAPMCONF Perspective

A new perspective view of the QAPMCONF database file has been added.

It can be found within the Collections Services content package under the Collection Services Database Files folder as seen in Figure 18-76 on page 586.

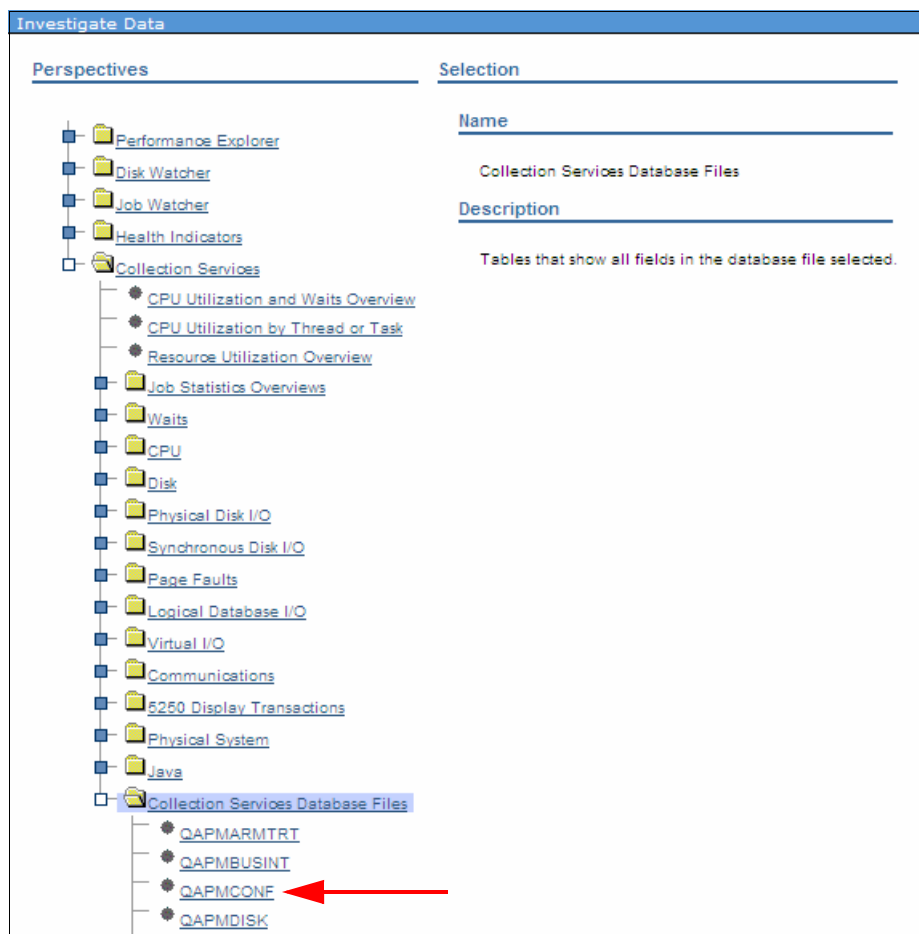



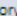


Figure 18-76 QAPMCONF Perspective view

The panel in Figure 18-77 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:	9406	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 18-77 QAPMCONF Perspective

18.11.4 Image and file exportation

Within Performance Data Investigator, you now have an capability to export an image (charts only), a comma delimited file, or a tab delimited file. Figure 18-78 shows an example of exporting CPU Utilization and Waits Overview information to a comma delimited file. You can get this data into a spreadsheet for later manipulation.

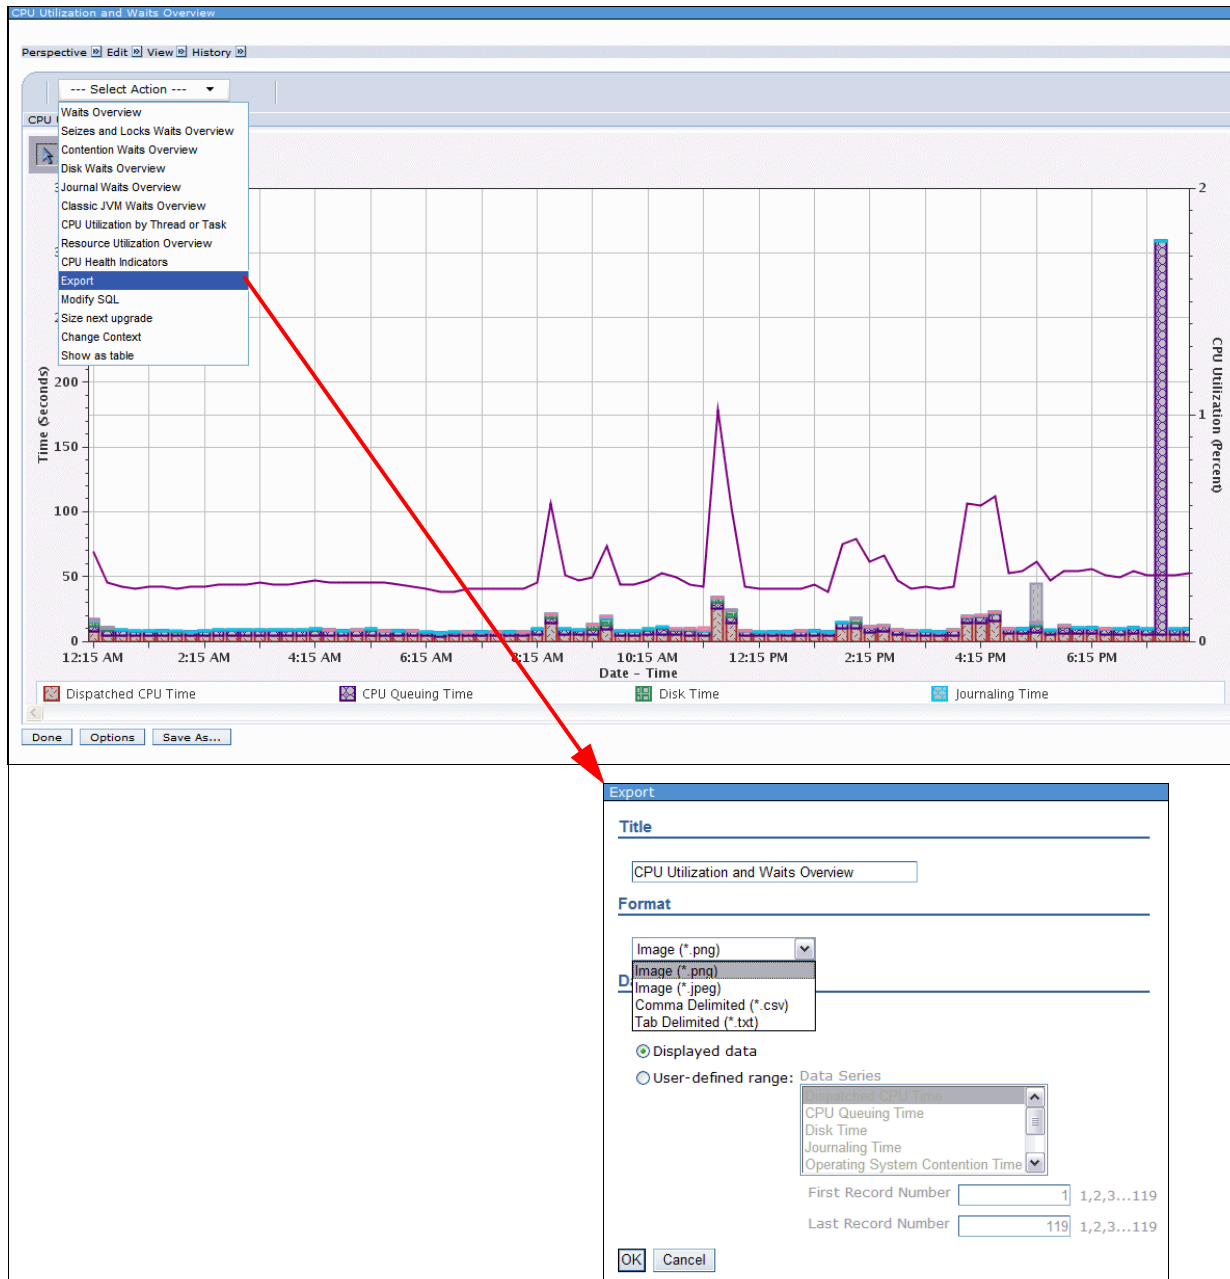


Figure 18-78 Exporting an image, comma delimited file or a tab delimited file

This feature makes it possible to print or export your data to store it for later reference.

The export panel has two main choices:

- ▶ Format which has three options and can be specified by the user
 - Image (*.png or .jpeg) file
 - Comma separated values (CSV) is the default format
 - Tab delimited (*.txt) file
- ▶ Data Range, which has three options and can be specified by the user
 - All data
This option allows the entire data set to be exported to the chosen format
 - Displayed data
This is the default and will produce smaller results but can omit off-panel information
 - User-defined range
This option allows the user to specify exactly which records and which series you want to export

18.11.5 Size the next upgrade

The Investigate Data feature can now send data from your current 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 available at the following web page:

<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 18-79 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 pull-down menu. This brings us to the IBM Systems Workload Estimator web page.

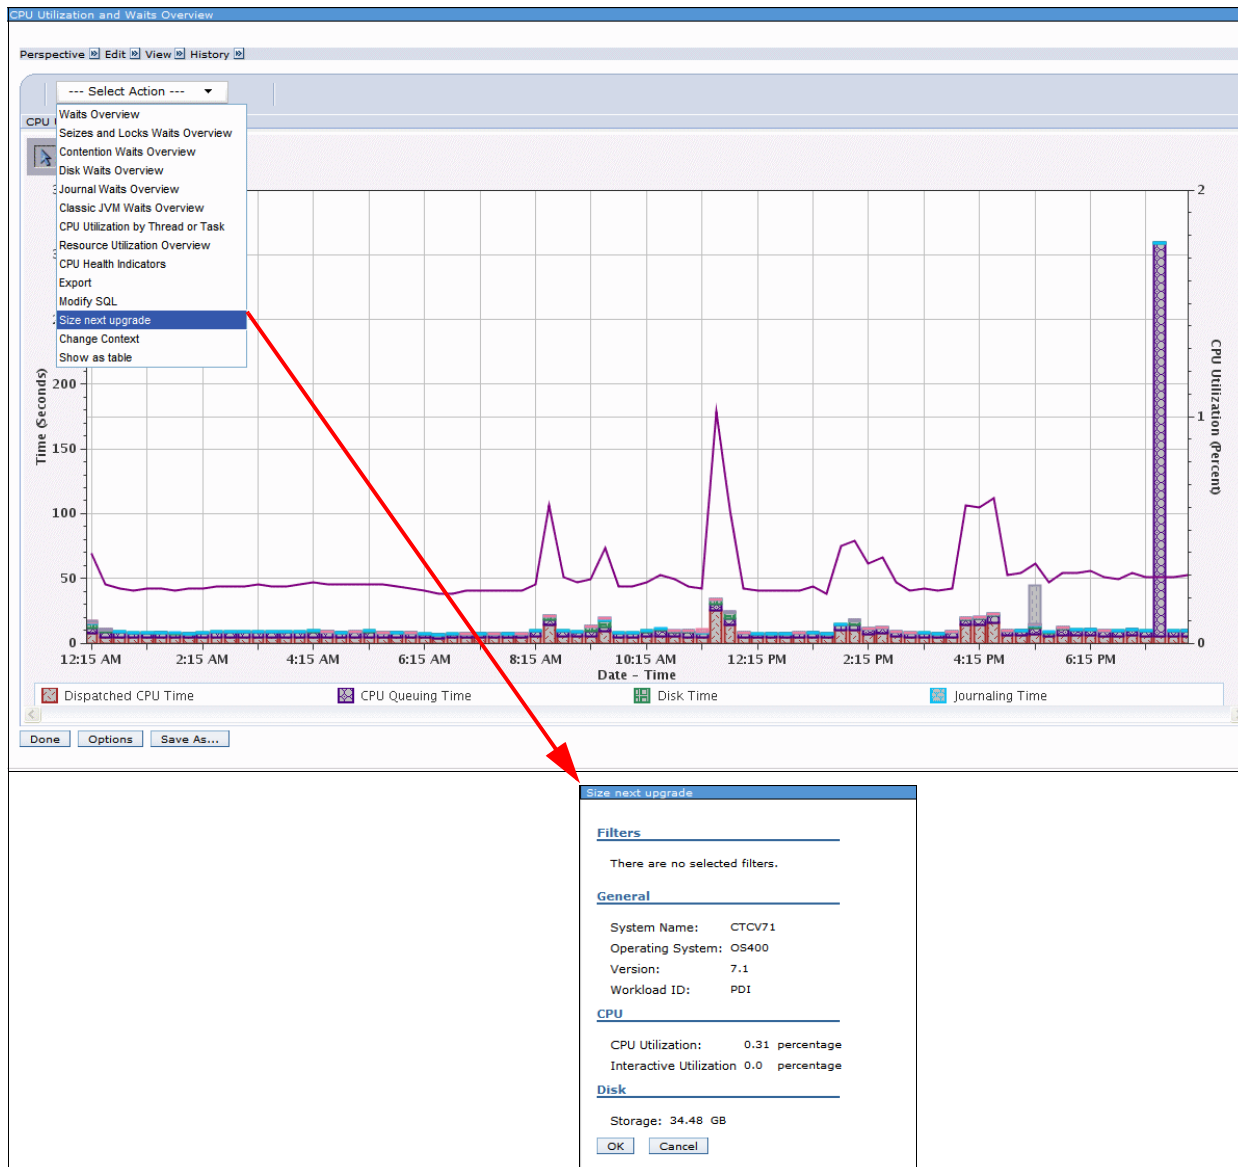


Figure 18-79 Size next upgrade

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. Perform the following steps to do so:

1. From the CPU Utilization Overview perspective shown in Figure 18-80, click the first icon in the top left of the chart to enable the select tool.
2. Select minimum and maximum points in the chart. These are shown like the two red signs that are selected in our case at 6:15 AM and 12:15 PM.

Note: It is important you select bars of the same metric (color/pattern must match) or each selection will clear the previous selections.




Figure 18-80 Time range

3. Select the **Size Next Upgrade** from the Action Menu as shown in Figure 18-81.
A panel where the time range is defined as filtered before is displayed.

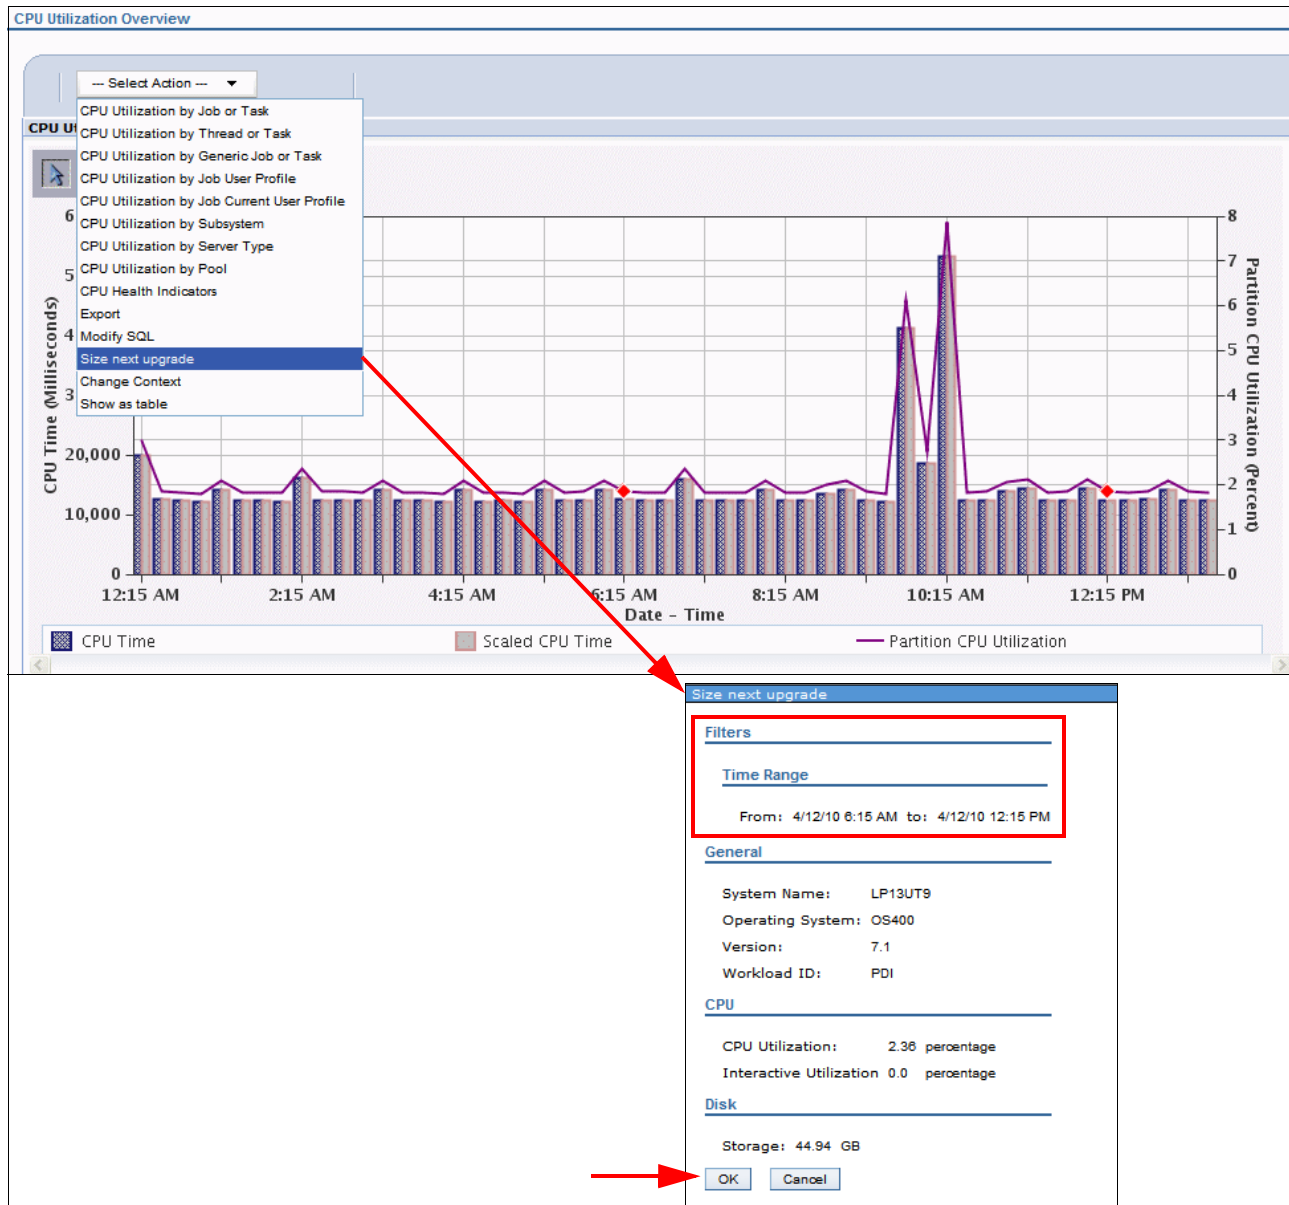


Figure 18-81 Subset of performance related data sent to WLE

4. Click **OK** to send this subset of the data to Workload Estimator.

18.11.6 Java perspectives

The charts contain general information for active Java Virtual Machines (JVM). Figure 18-82 shows the IBM Technology for Java Memory Overview that shows garbage collection information, including heap sizes and other general information related to JVM memory.

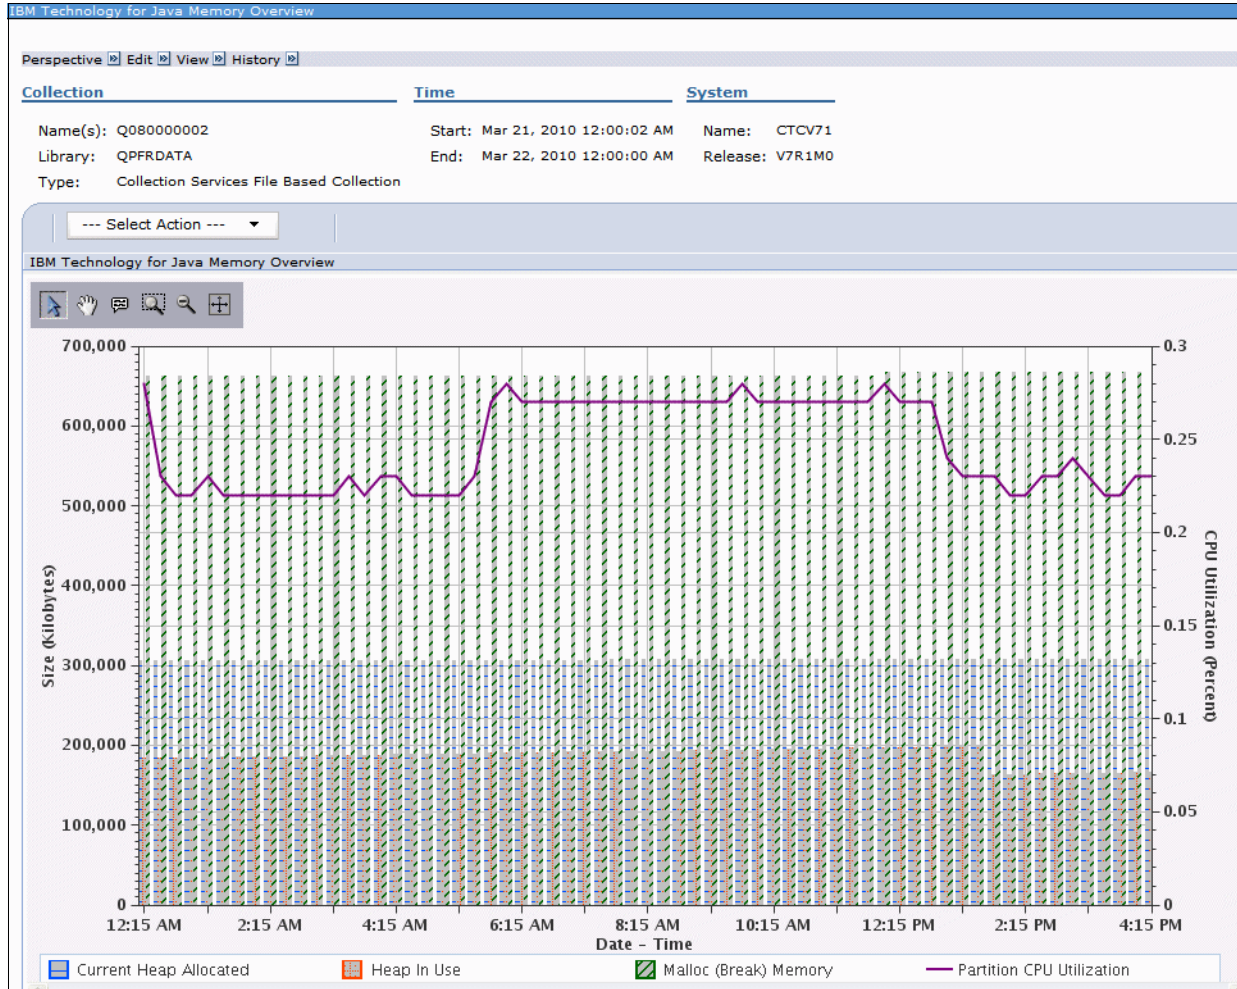


Figure 18-82 Java perspective

18.11.7 Disk response time charts

Within the IBM i 7.1 Collection Services Perspectives, the Disk Response Time Charts are newly available. Navigate to **Disk** → **Disk Response Time** → **Detailed**, as shown in Figure 18-83. The charts show disk I/O metrics segmented by response time and are only available for data collected on IBM i 7.1 or later.

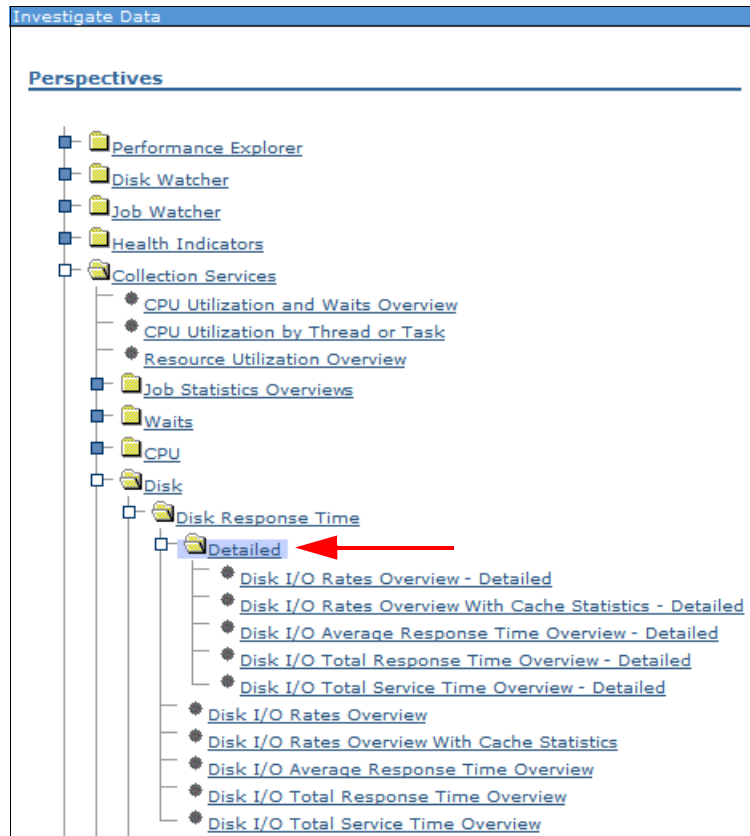


Figure 18-83 Disk response time charts

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

The chart in Figure 18-84 shows disk I/O segmented by the amount of I/O's that occurred when response time was in specific ranges, and the average response time and average service time.




Figure 18-84 Disk I/O rates overview: Detailed

Disk I/O rates overview with cache statistics

The chart in Figure 18-85 shows disk I/O segmented by the amount of I/O's that occurred when response time was in specific ranges, and cache statistics.




Figure 18-85 Disk I/O rates overview with cache statistics: Detailed

Disk I/O average response time overview

The chart in Figure 18-86 shows disk average response time segmented by the amount of I/O's that occurred when response time was in specific ranges, and average response time and average service time.

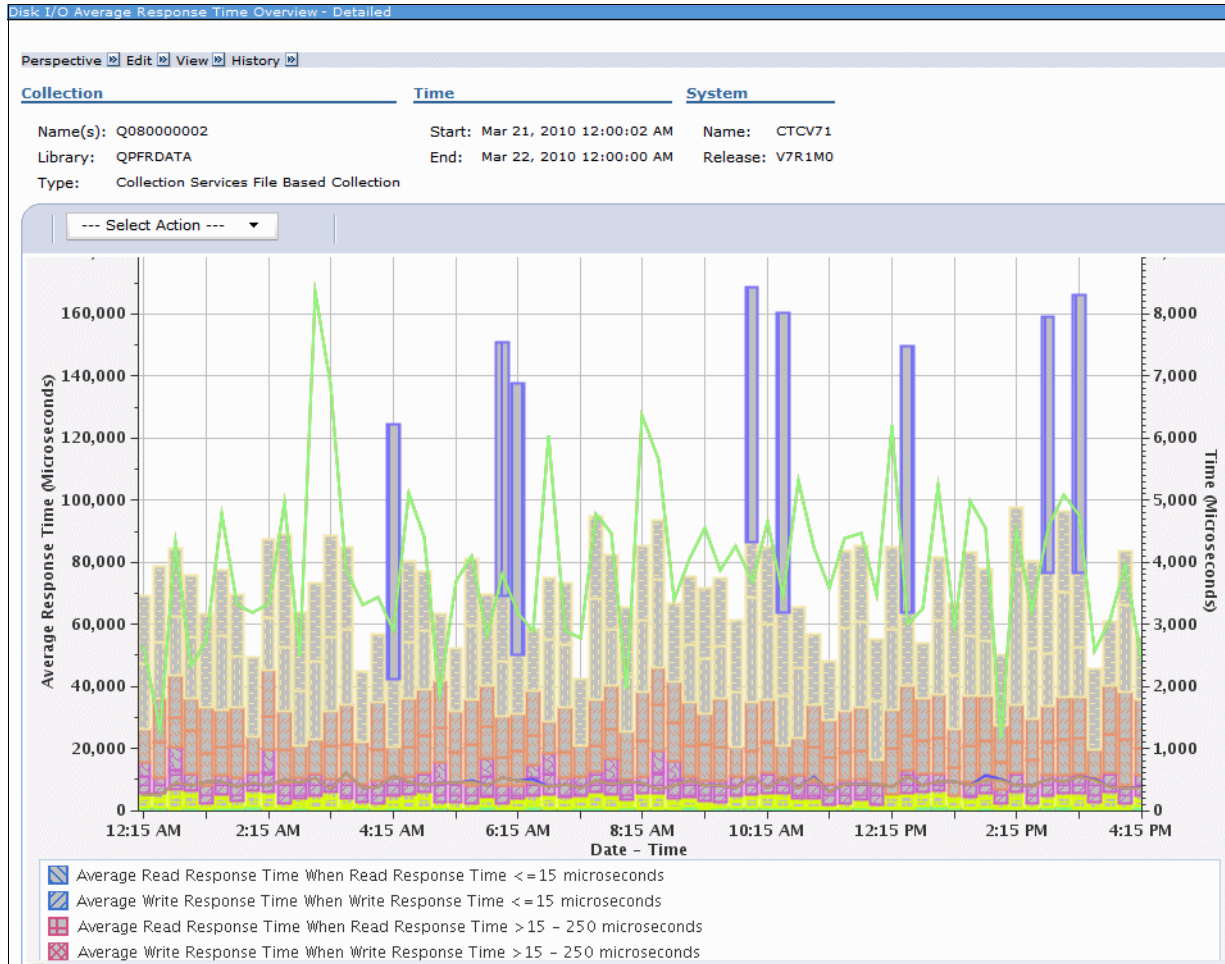


Figure 18-86 Disk I/O average response time overview: Detailed

Disk I/O total response time overview

The chart in Figure 18-87 shows the disk total response time segmented by the amount of I/O's that occurred when response time was in specific ranges, and average response time and average service time.

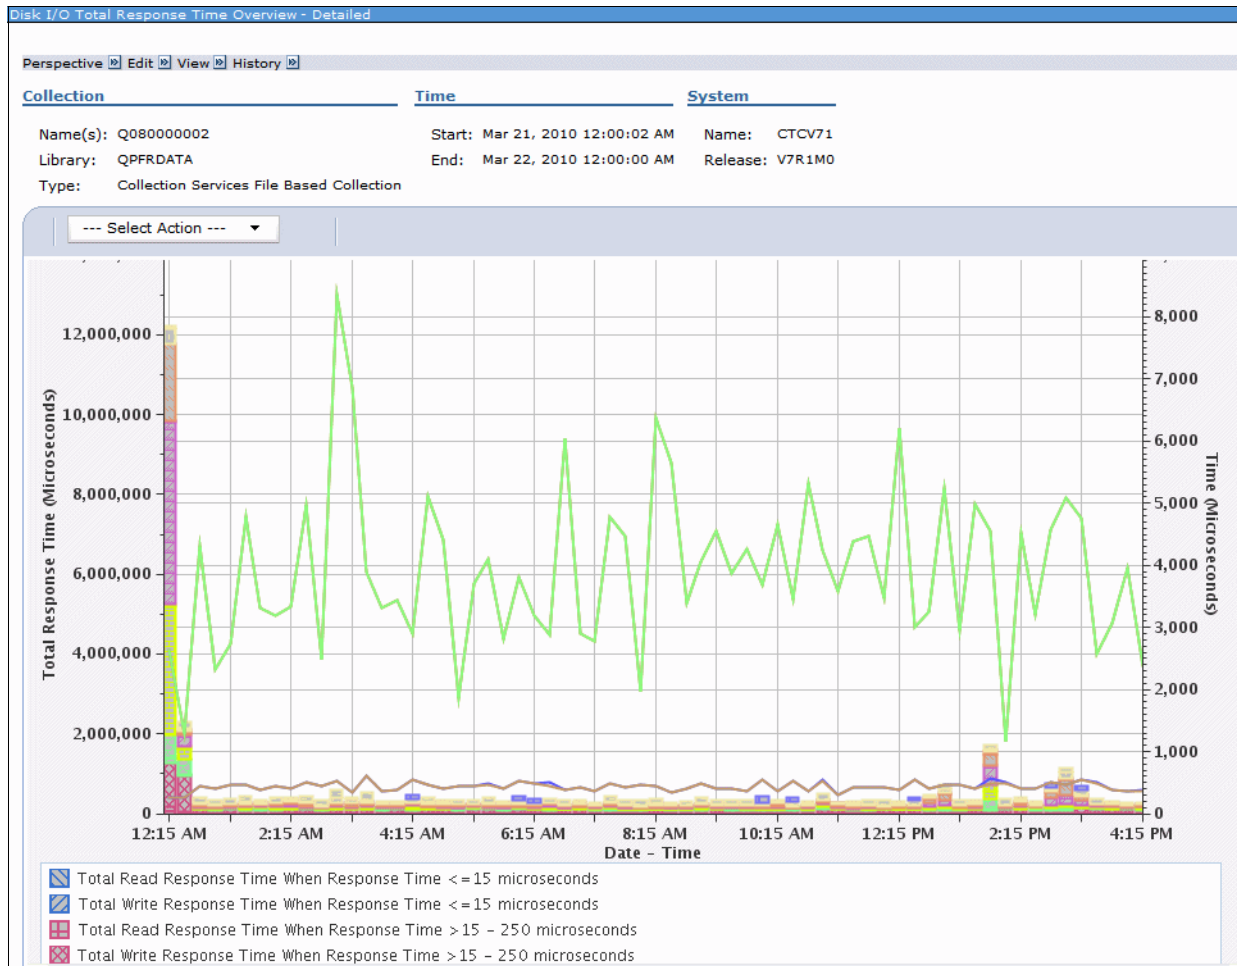


Figure 18-87 Disk I/O total response time overview: Detailed

Disk I/O total service time overview

The chart in Figure 18-88 shows the disk total service time segmented by the amount of I/O's that occurred when response time was in specific ranges, and average response time and average service time.




Figure 18-88 Disk I/O total service time overview: Detailed

For more detailed information related to these enhancements, see “QAPMDISKRB” on page 195.

18.11.8 Interactive perspective development

New charts and tables can be developed from within Investigate Data. Adding views, modifying data series information and modifying SQL statements can all be performed without leaving the page.

Adding views

Perform the following steps to add views:

1. At the bottom of the Investigate Data Perspectives page, click **Options**. This brings you to the page as shown in Figure 18-89

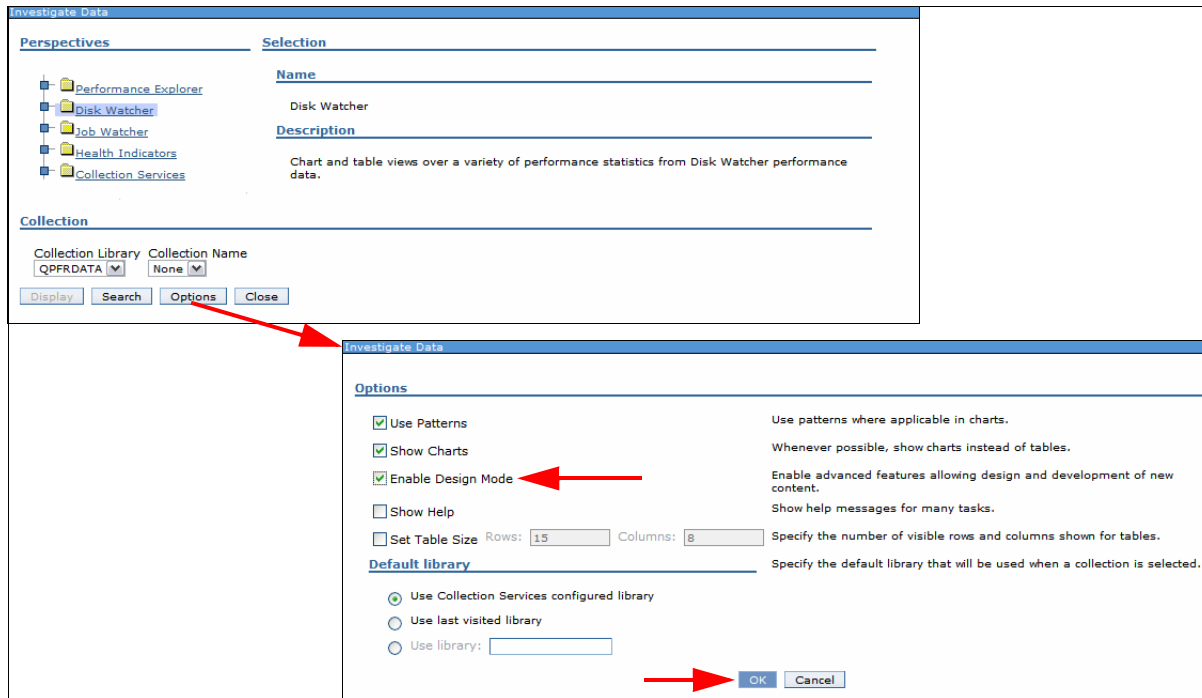


Figure 18-89 Investigate data perspective: Options

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 page shown in Figure 18-90.

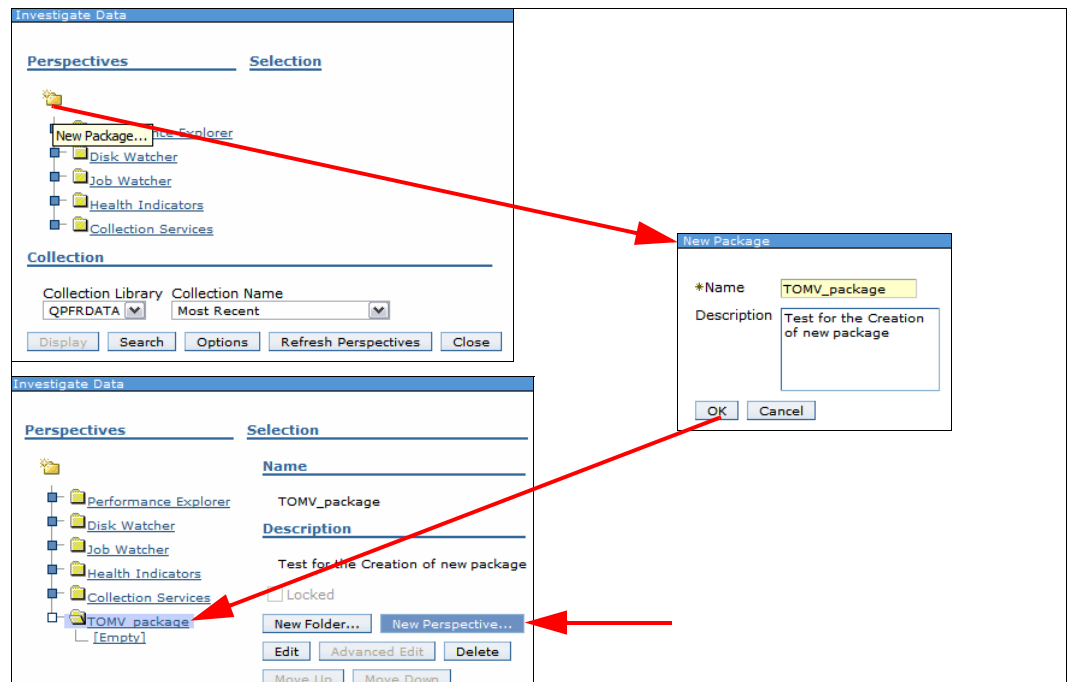


Figure 18-90 Creating a custom content package

4. On the New Package panel, specify a name, and click **OK**. The Package gets created and is displayed.

5. Click **New Perspective**. The **New Perspective** page (Figure 18-91) is displayed.

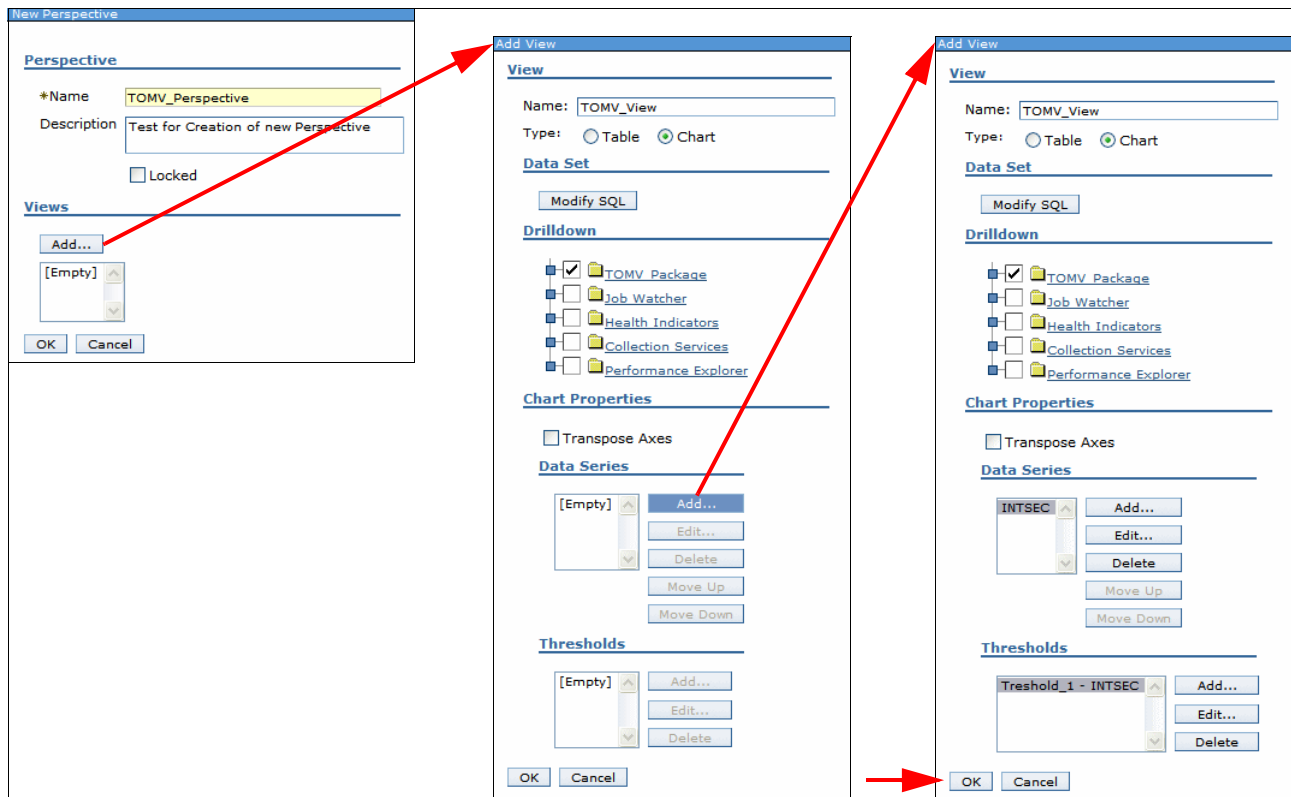


Figure 18-91 Add data series: Thresholds

6. Specify a name and a description for the Perspective and click **Add** to add a view (Table or Chart) with Data series and corresponding Threshold.
7. Click **OK** on the Add View panel and click **OK** on the New Perspective panel. The customized perspective is saved and shows up in the main perspective tree list for further reference, as shown in Figure 18-92.

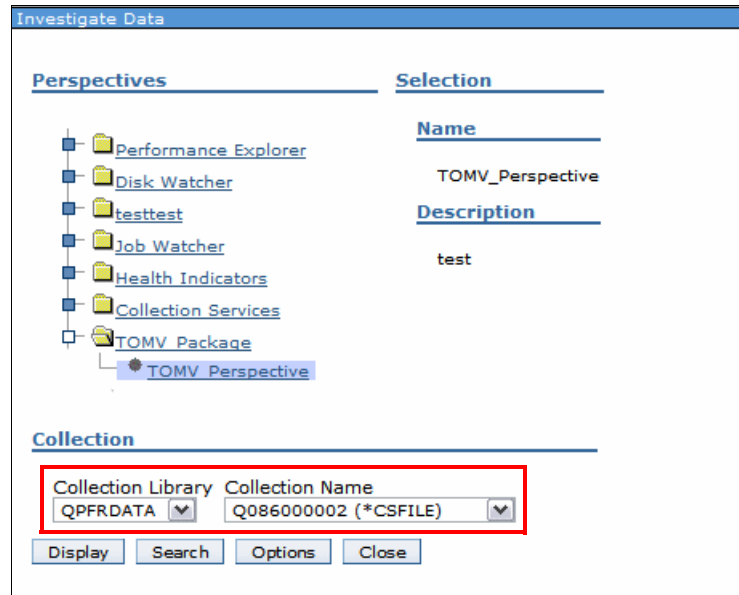


Figure 18-92 Customized package and perspectives added to main perspective tree list

Modifying Data Series information

It is also possible to modify the Data Series information, as shown in Figure 18-93, by clicking **Edit** in the View pane on the Edit Perspective panel.

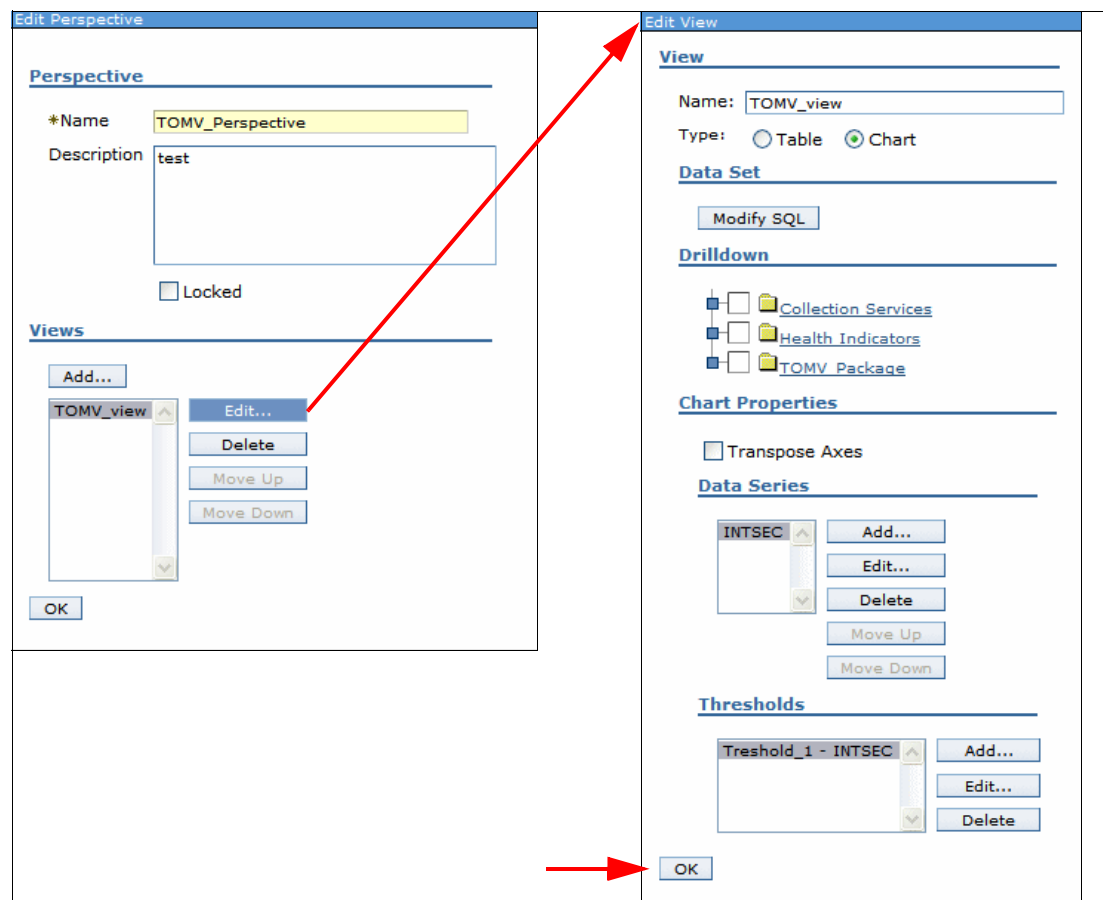


Figure 18-93 Modifying data series

After modifying the View, click **OK** to save the View with the changed information.

Modifying an SQL statement

This feature allows the user not only to view the current SQL, but to modify the SQL and re-render its output to the panel. Figure 18-94 shows an example where **Modify SQL** has been selected from the actions pull-down menu on the CPU Utilization and Waits Overview chart. This displays a panel where the user can modify the SQL statements. Click **OK** to save the modifications.

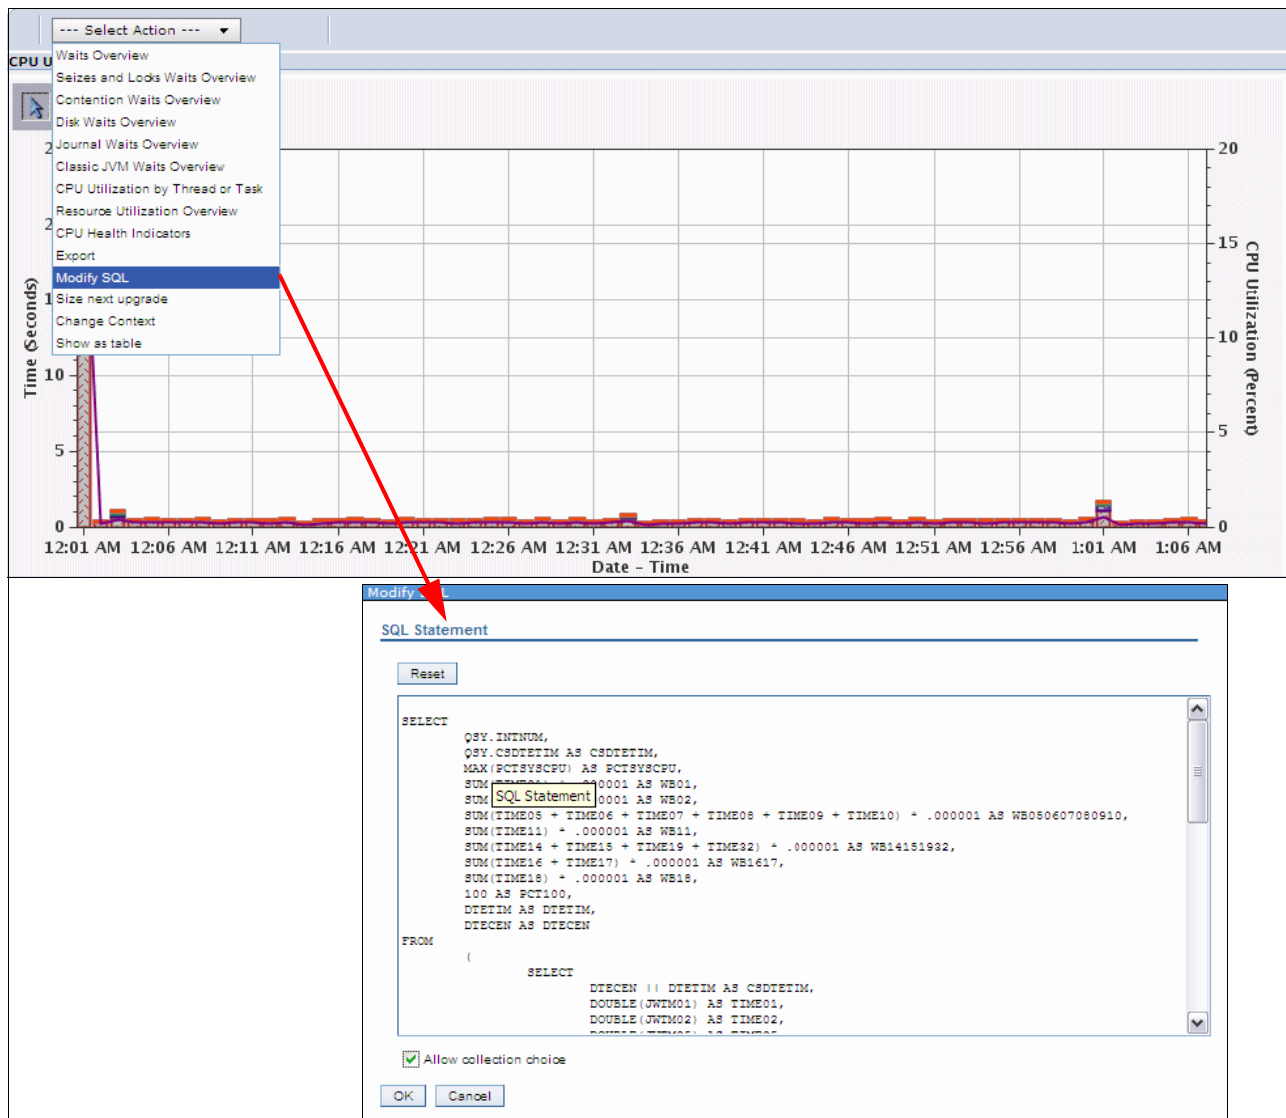


Figure 18-94 Modifying the SQL

Formatting enhancements have been made to the Modify SQL panel, making the text easier to read and understand. The Modify SQL action is available for charts and tables.

18.11.9 Metric finder

With the large number of metrics provided by the Investigate Data feature, sometimes knowing which perspective contains the metric you want to see is no simple task.

The Metric Finder function allows a user to display perspectives based on an interest in a specific metric. This is useful when a user knows what type of information they are looking for but does not know where to start or where certain metrics are included (which perspectives).

To invoke this feature, click the **Search** button when you launch Investigate Data, as shown in Figure 18-95. Then choose your perspective by searching for a metric name.

Note: You can specify a filter to limit the metric names included in the drop-down menu. The filter helps you to search for one specific metric, without knowing its exact name. After entering a filter, click **Apply Filter** to update the metric name list. After a perspective is selected, it is displayed by clicking the **Display** button.

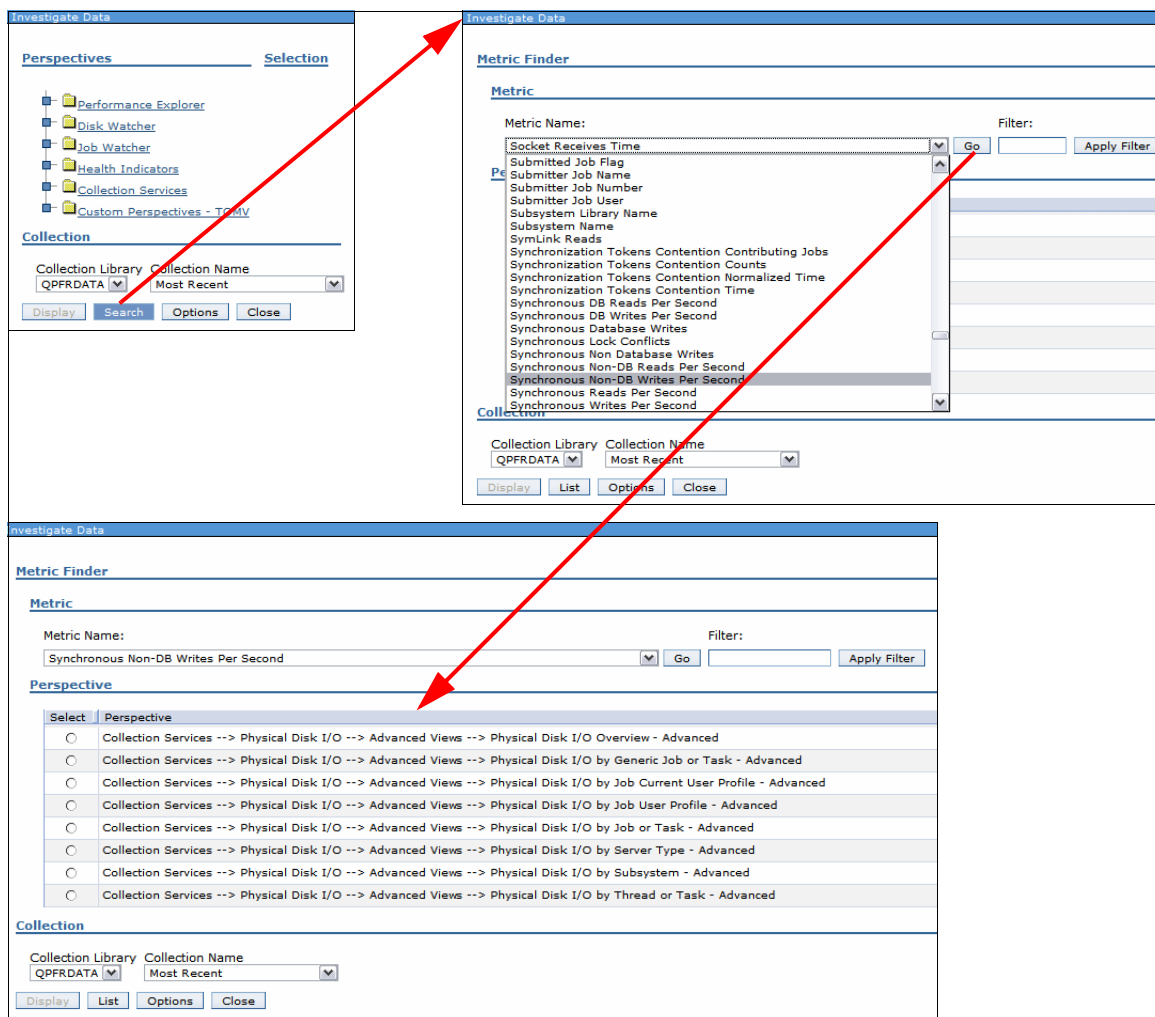


Figure 18-95 Metric finder

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 afterwards.

18.11.10 Performance Explorer content package

A simple Performance Explorer content package is included to start analysis of Performance Explorer data, as shown in Figure 18-96. This can be used to analyze the performance of an application program. The Performance Explorer content package allows you also to view the PEX database files.

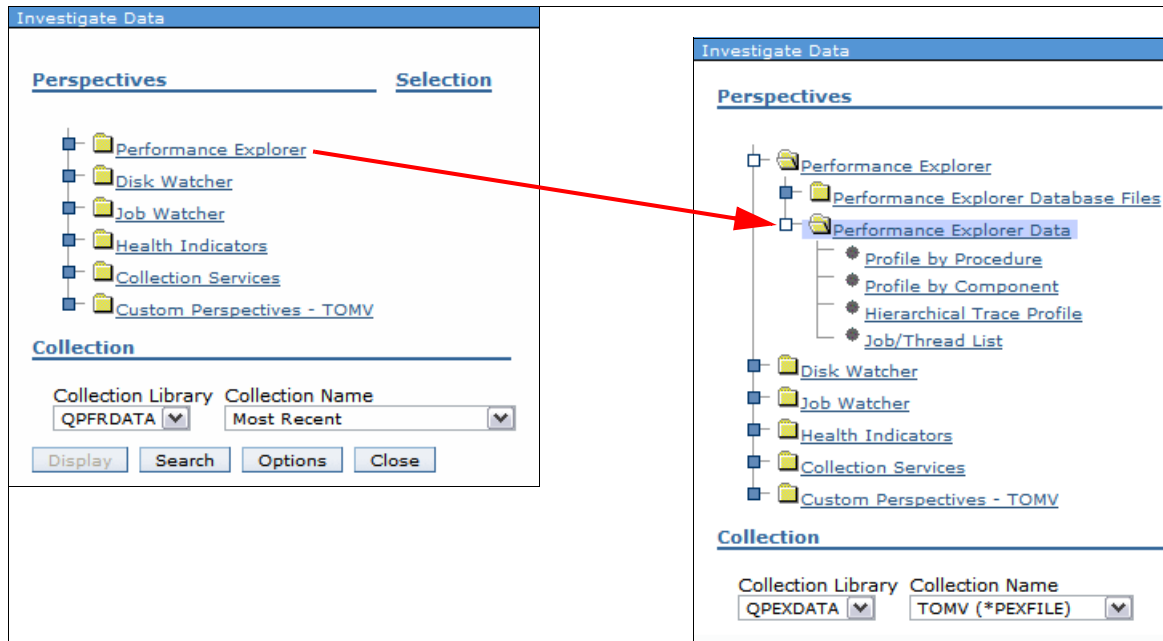


Figure 18-96 Performance Explorer content package

The following PEX Profile perspectives provide function 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 web page:

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

18.11.11 New metrics

Many new perspectives have been added to Investigate Data, and many new metrics have been added to existing perspectives:

- ▶ SAN
- ▶ Virtual I/O
- ▶ Energy Management
- ▶ Communication data

SAN metrics

Within the IBM i 7.1 Collection Services Perspectives, the Disk Response Time Charts are newly available, as shown in Figure 18-97. This is accessed by clicking **Disk** → **Disk Response Time** → **Detailed**. The charts show disk I/O metrics (including SAN attached disks) segmented by response time and are only available for data collected on IBM i 7.1 or later.

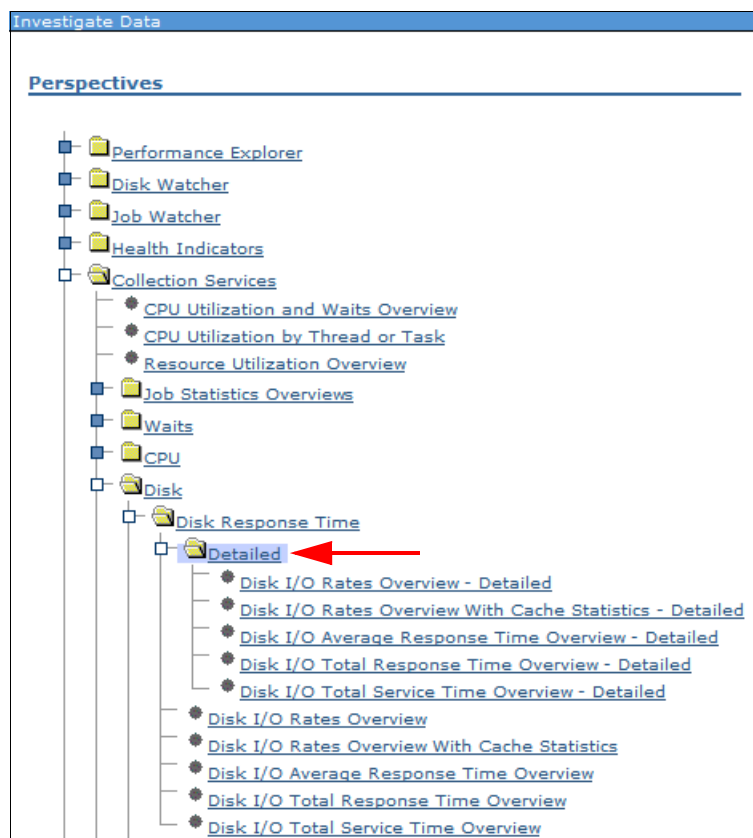


Figure 18-97 SAN metric

For more information about those buckets, see “QAPMDISKRB” on page 195.

Virtual I/O metrics

The charts show I/O operations rates and throughput for virtual I/O devices. Figure 18-98 shows a view of virtual I/O devices categories (Adapter, Disk, Optical, Type).

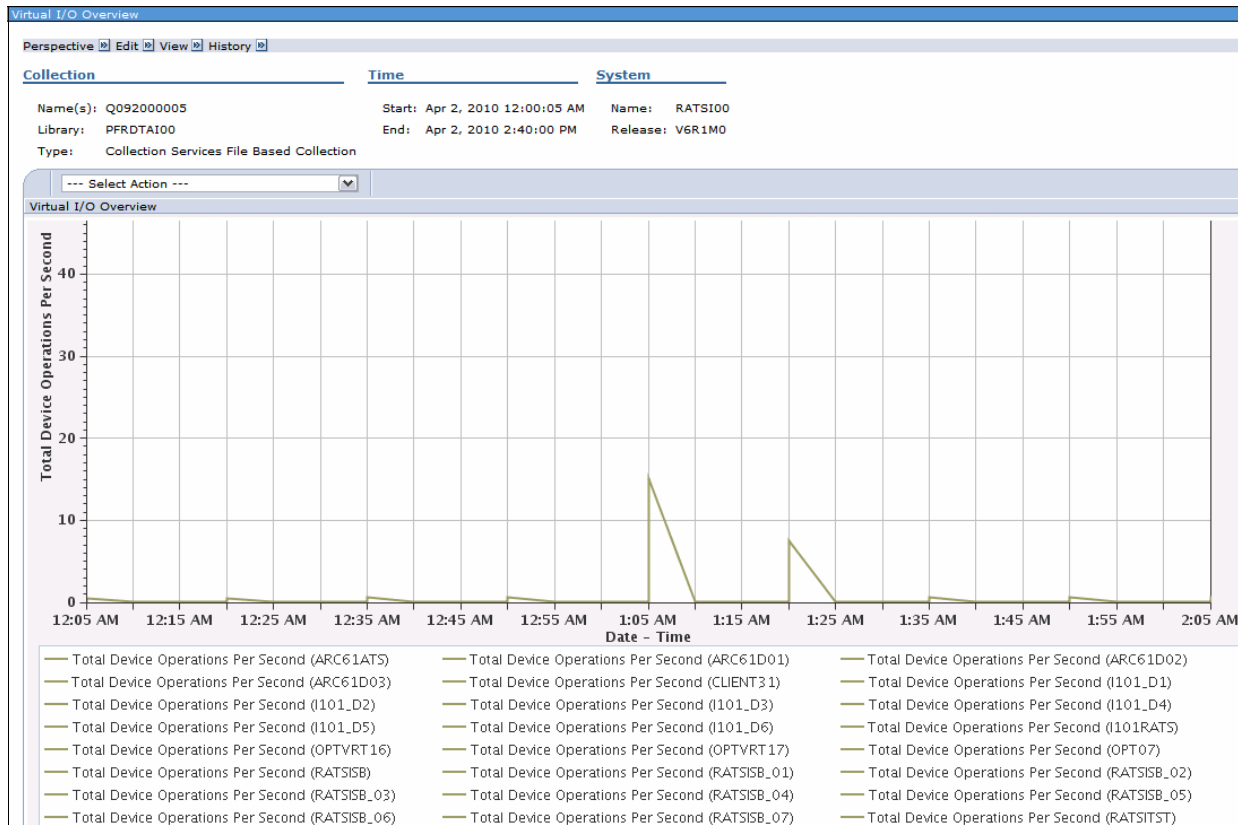


Figure 18-98 Virtual I/O metrics

Figure 18-99 shows a chart with the Virtual I/O Overview displayed by device type.




Figure 18-99 Virtual I/O overview by device type

Energy management metrics

POWER6 systems had so called 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 web page:

<http://www-03.ibm.com/systems/power/hardware/whitepapers/energyscale.html>

IBM recommends using the Active Energy Manager 4.2 (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 system.

For more information related to the AEM, see the following web page:

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 is shown on the CPU Utilization Overview in Figure 18-100.




Figure 18-100 CPU use overview

Figure 18-101 shows the CPU rate (Scaled CPU : Nominal CPU Ratio) for a specific period of time.




Figure 18-101 CPU rate (Scaled CPU: Nominal CPU Ratio)

Communications metrics

The 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 lines information about the partition. Figure 18-102 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 frame retries and MAC errors.

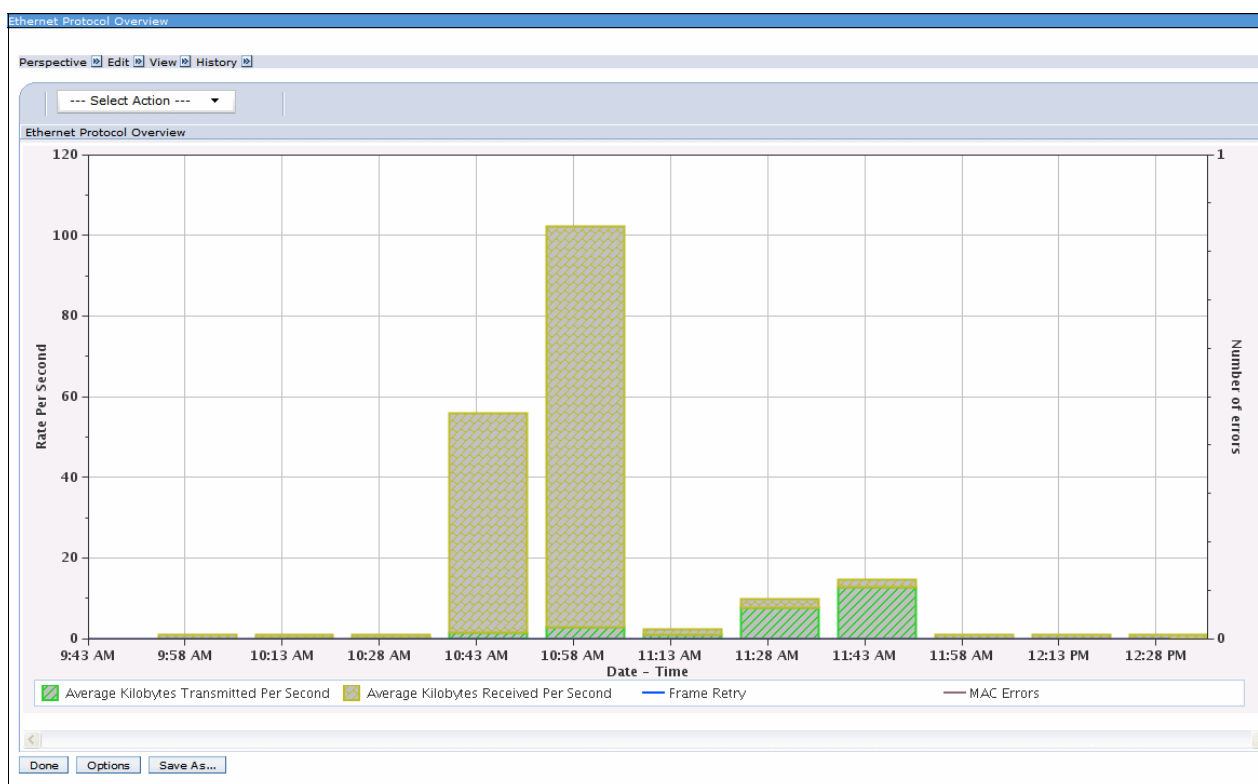


Figure 18-102 Ethernet protocol overview

18.11.12 Miscellaneous enhancements

Miscellaneous changes have been made to improve the overall capabilities of the performance tasks:

- ▶ Improved integration with Active Jobs
- ▶ Improved integration with System and Disk Status
- ▶ New collection information at the top of each perspective
- ▶ New menu bar which allows for quicker navigation
- ▶ More complete History data

Improved integration with Active Jobs

From the Performance category, select **Show All Performance Tasks** and select **Active Jobs** from the pull-down menu, as shown in Figure 18-103.

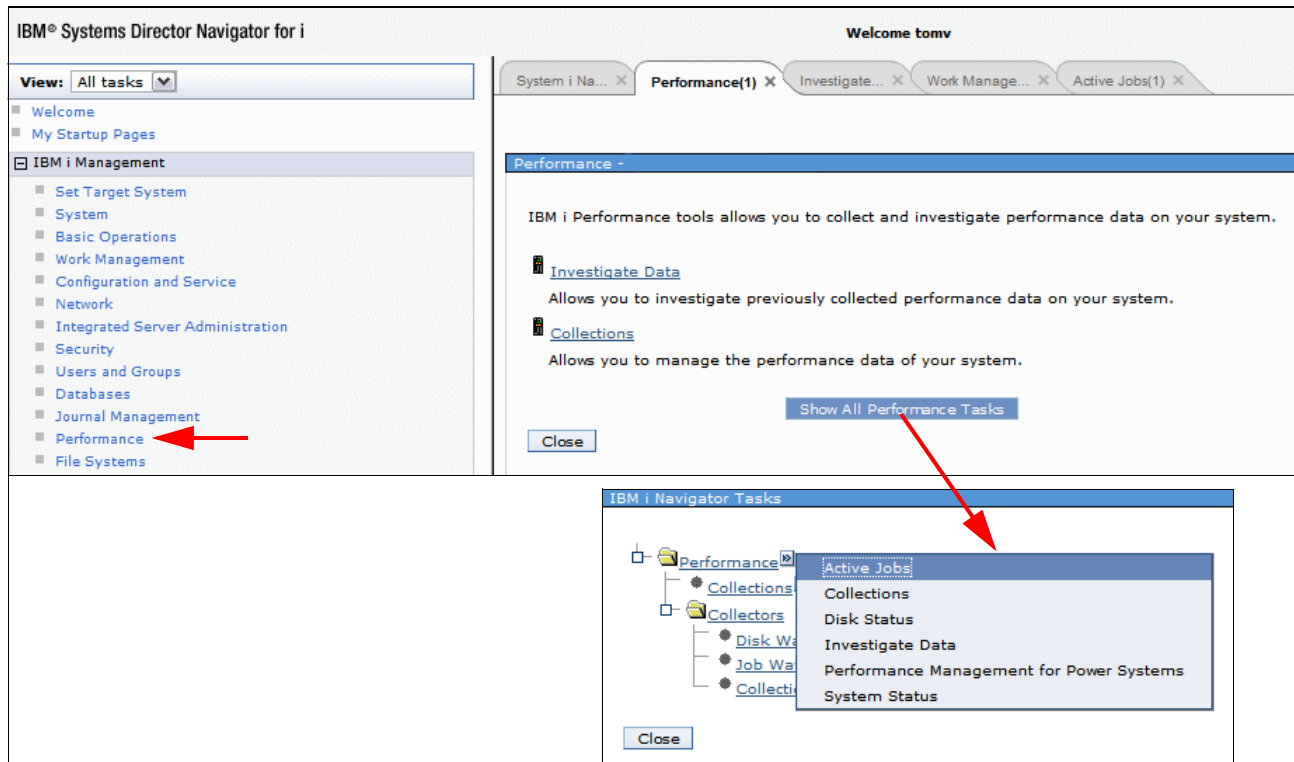


Figure 18-103 Improved integration with active jobs

From within Active Jobs, it is now possible to go directly to performance related information for a selected job as shown in Figure 18-104.

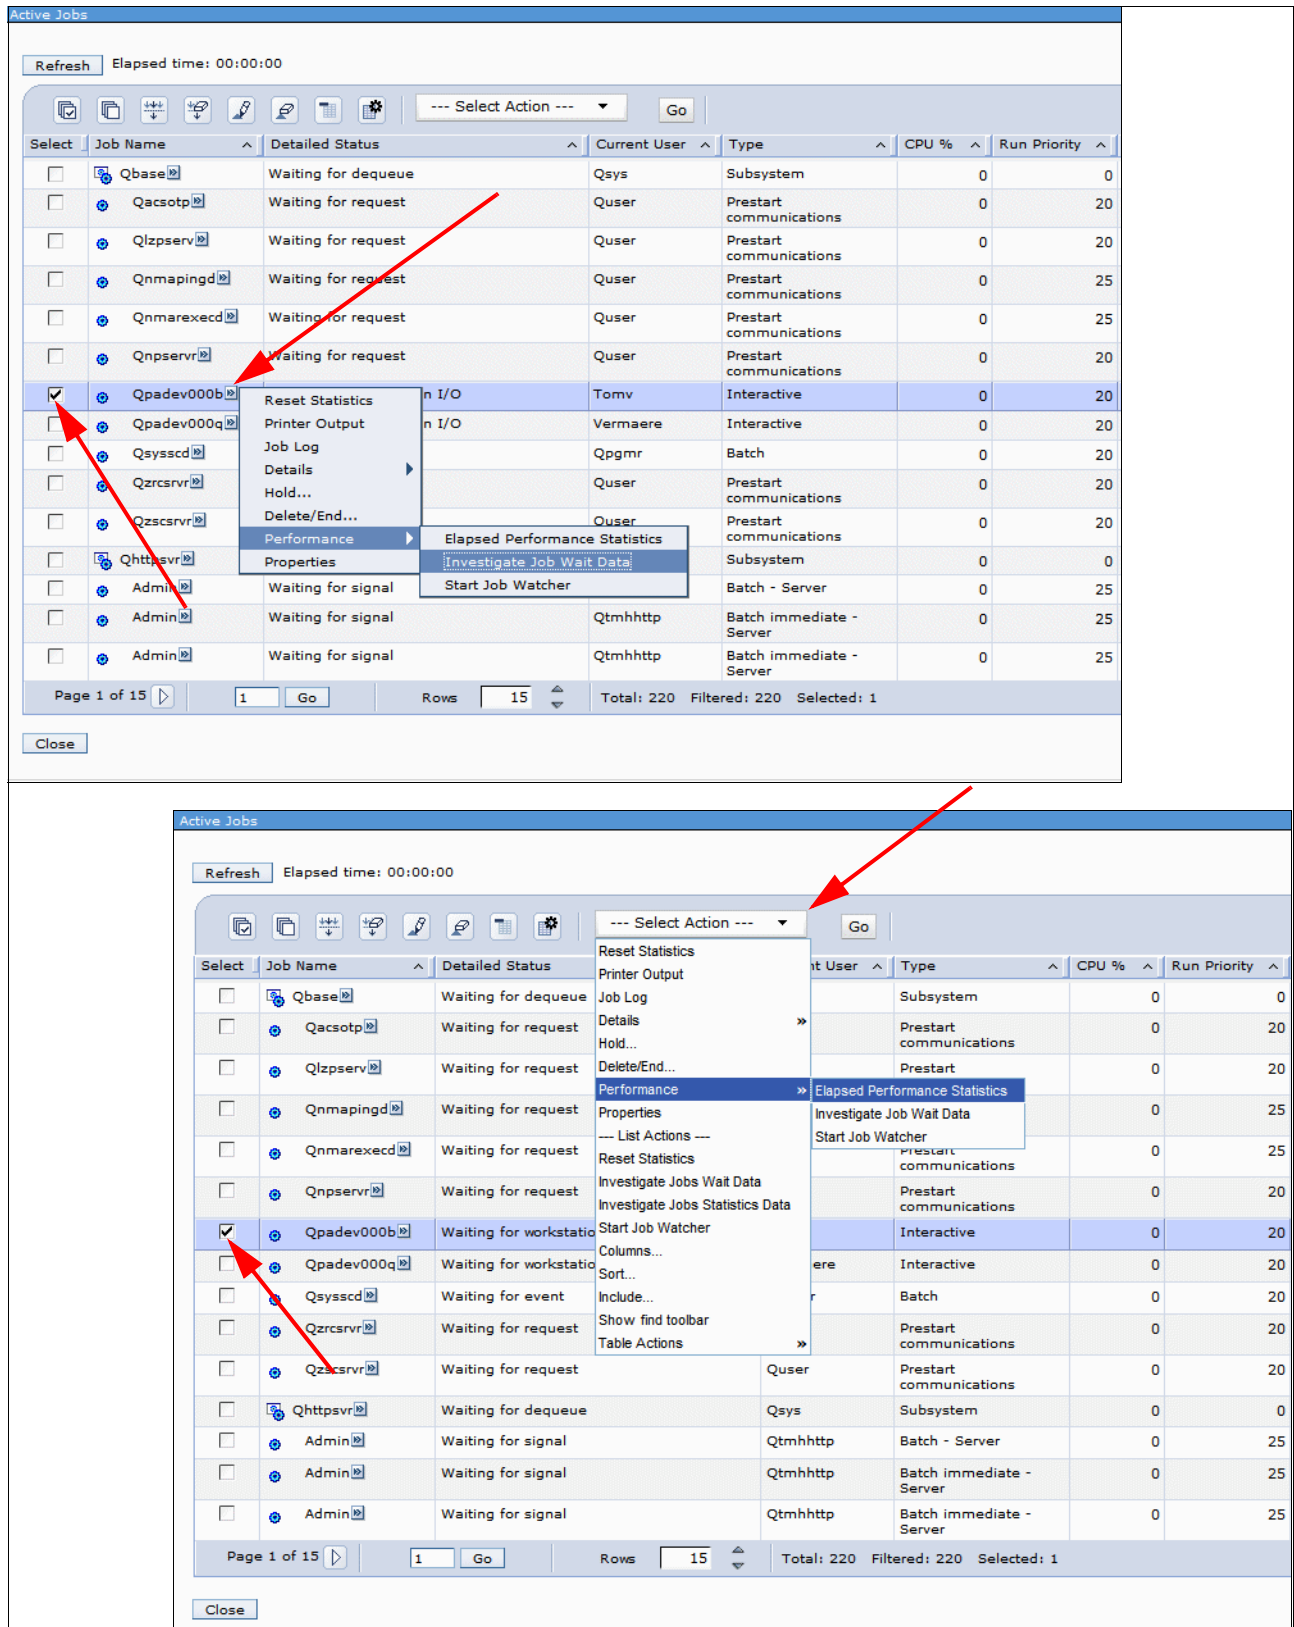


Figure 18-104 Active jobs integration

Note: When selecting the **Investigate Job Wait Data**, the chart rendered only includes data pertaining to that specific job. If you select **Investigate Jobs Wait Data**, the chart rendered includes all jobs that were active during that collection.

The following performance related selections are available either from the pop-up menu at the right side of the selected job or from the Action drop-down menu at the top:

- ▶ Elapsed Performance Statistics
- ▶ Investigate Job Wait Data
- ▶ Start Job Watcher

Improved integration with system status and disk status

From the Performance category, you can select **Show All Performance Tasks** where you can then either select **Disk Status** or **System Status** from the pop-up menu, as shown in Figure 18-105.

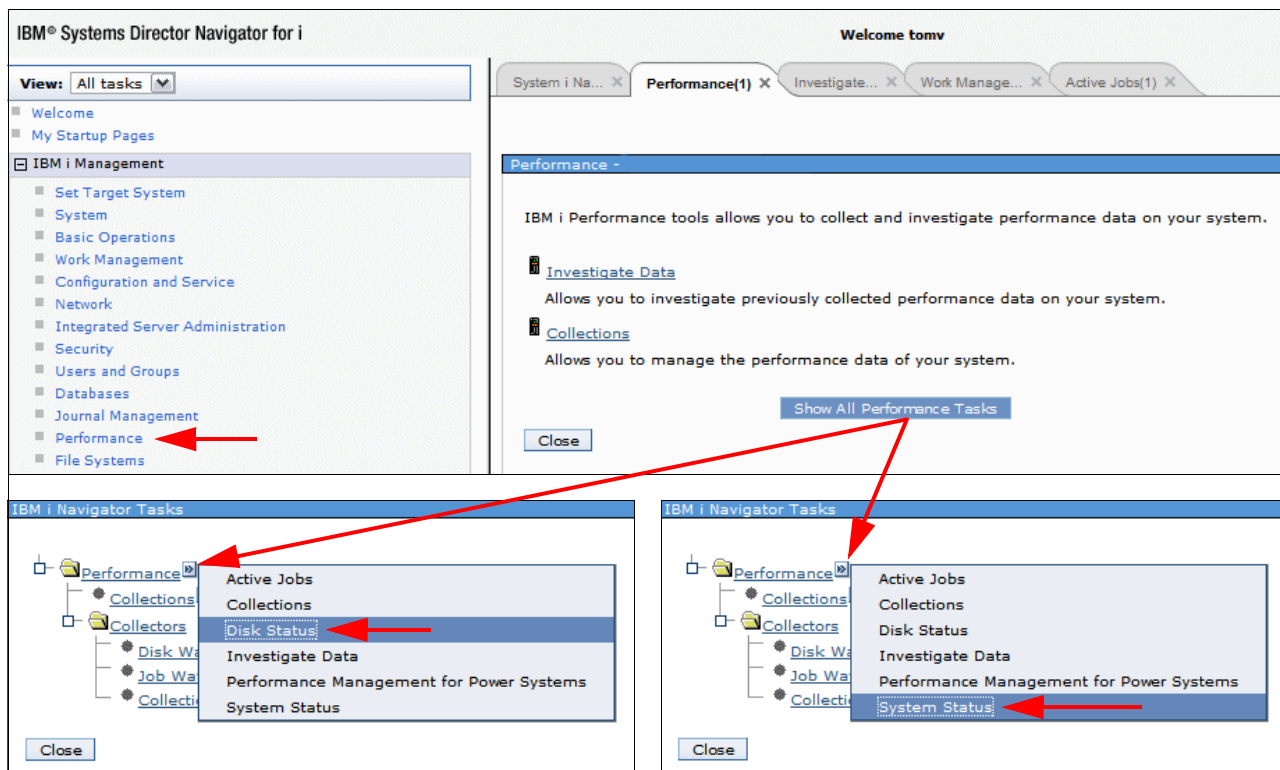


Figure 18-105 Improved integration with disk status and system status

From within System Status and Disk Status, it is now possible to go directly to the performance related information as shown in Figure 18-106.

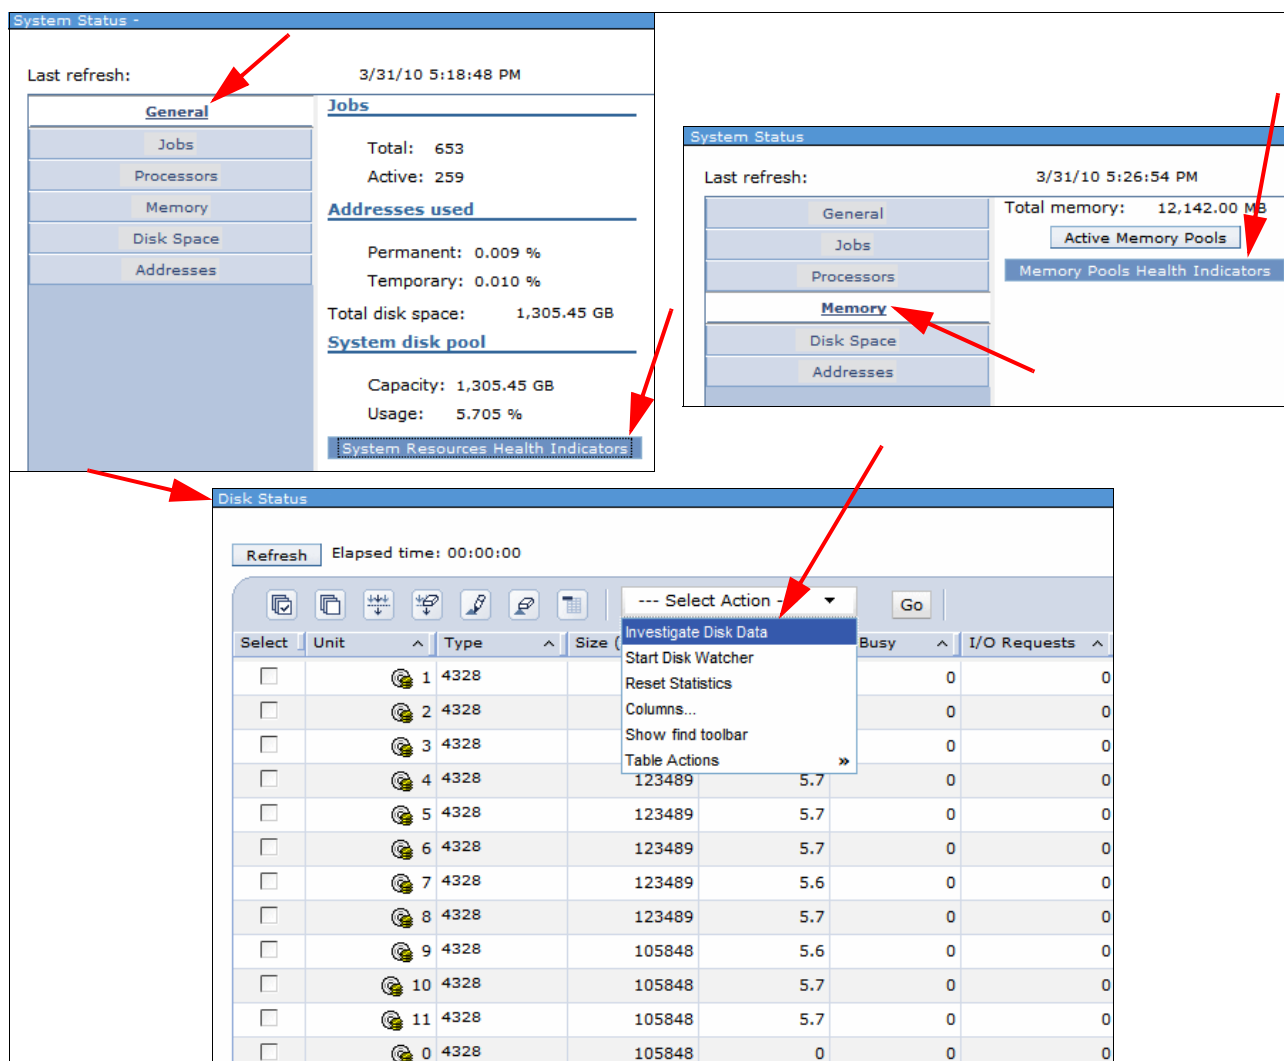


Figure 18-106 System status and disk status integration

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)
- ▶ Investigate Disk Data (from the Select Action list on Disk Status)
- ▶ Start Disk Watcher (from the Select Action 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 viewing a specific perspective.

By default, the name of the currently viewed perspective is now shown as the title of the perspective panel.

Users might find this information not to be useful or take up too much space on the panel. For these users, a new View menu is added with the **Show Context** check box menu, as shown in Figure 18-107. By clearing this check box, this information is hidden. This choice is preserved across sessions.

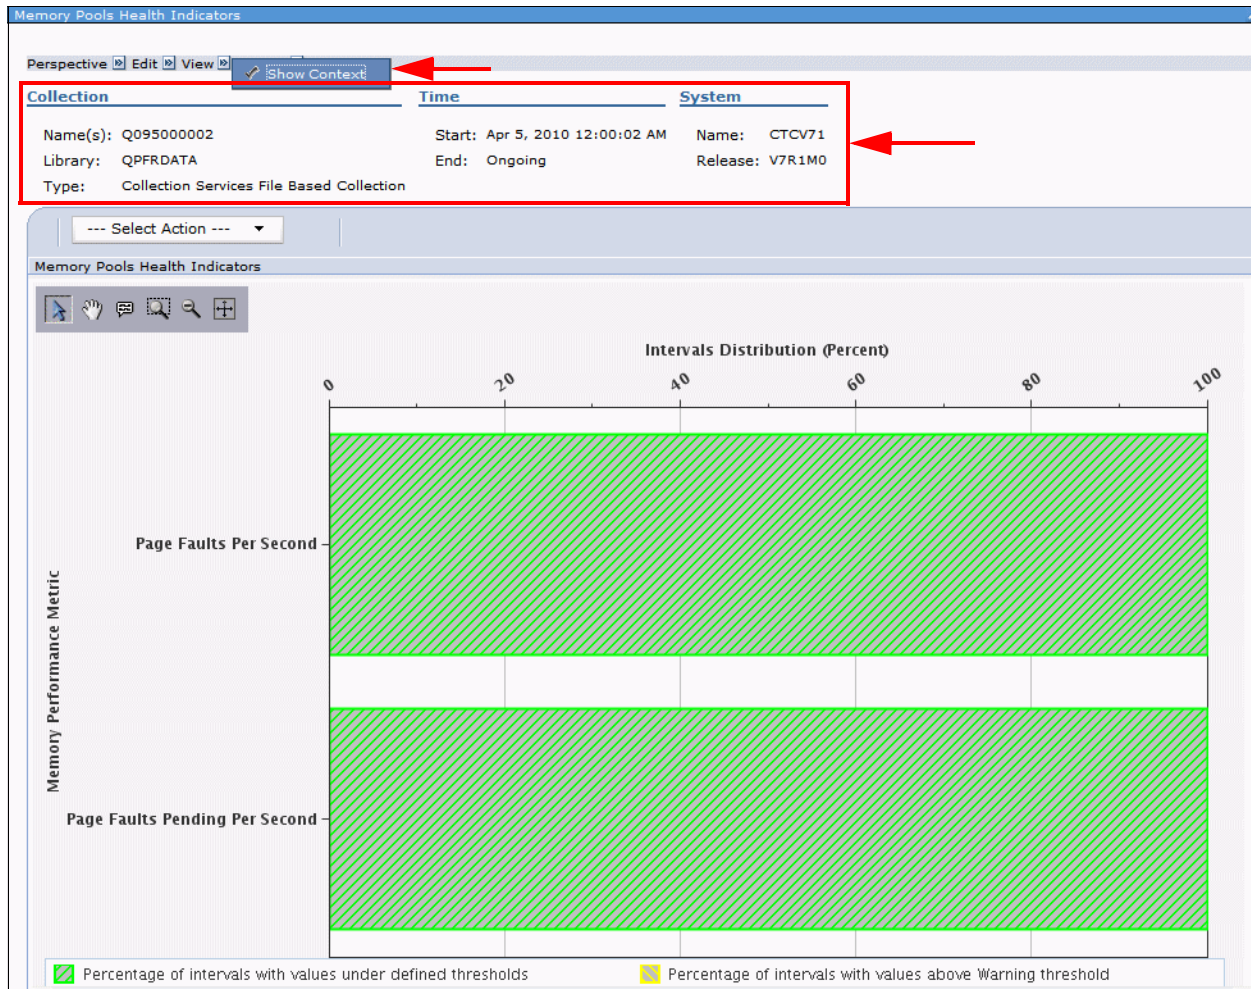


Figure 18-107 Collection information and collection details information

New menu bar

in IBM i 7.1, a new menu system is now added at the top of every perspective, which allows for quicker navigation, as shown in Figure 18-108.

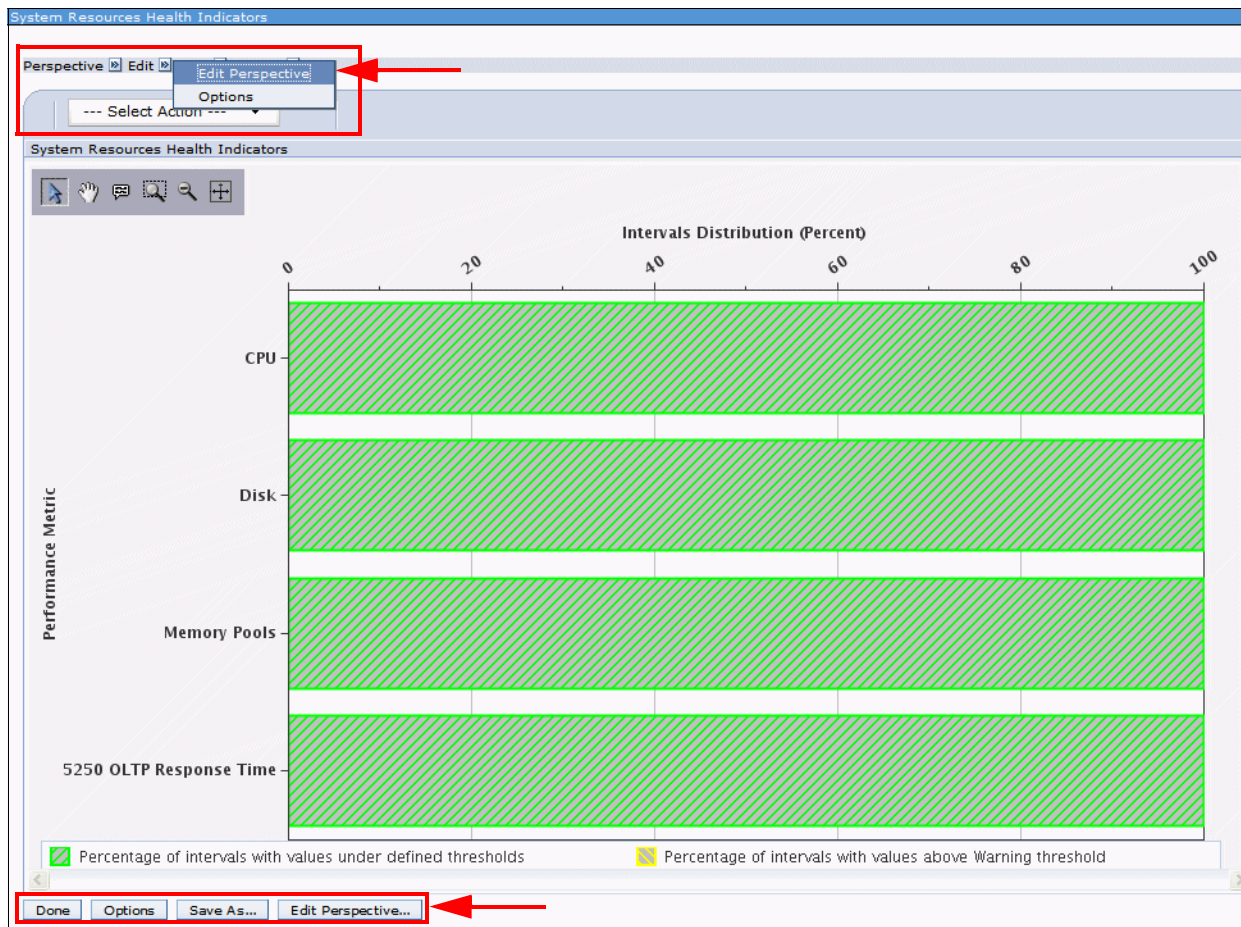


Figure 18-108 New menu bar

The new menu at the top of the perspective has the same actions available as the one at the bottom, but is available without the need to scroll-down, due to the fact that they sometimes might be rendered off-panel. 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 18-109 when clicking **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.

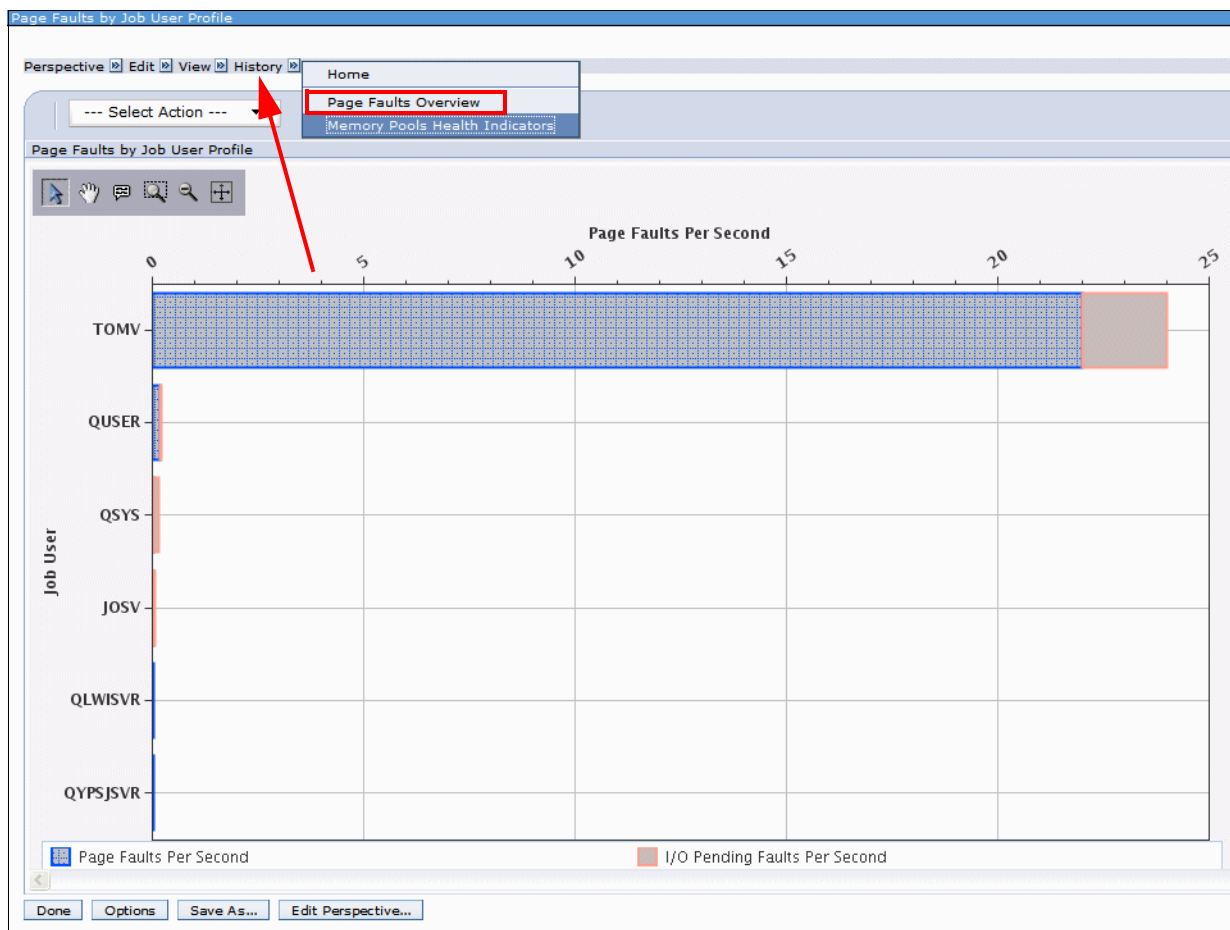


Figure 18-109 History data

From within that menu bar it is then possible to get back to the previous perspectives by clicking the corresponding item in the History Data.

18.12 JS1- Advanced Job Scheduler for i enhancements

With IBM i 7.1, IBM Systems Director Navigator for i has virtually equivalent Advanced Job Scheduler function to the System i Navigator client.

For more information related to Advanced Job Scheduler enhancements, see Chapter 13, “IBM Advanced Job Scheduler for i enhancements” on page 385.

18.13 Backup Recovery Media Services: BRMS plugin enhancements

IBM Systems Director Navigator for i had limited BRMS function in IBM i 6.1. Its capabilities have been greatly 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 27.

18.14 Additional information

Connecting to IBM i IBM Systems Director Navigator for i PDF.

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzatg/rzatgdirector.pdf>

IBM Systems Director Navigator for i, SG24-7789

End to End Performance Management on IBM i, SG24-7808



IBM i Access

This chapter discusses the following IBM i Access topics for IBM i 7.1:

- ▶ 19.1, “IBM i Access for Windows” on page 622
- ▶ 19.2, “IBM i Access for Web enhancements” on page 623
- ▶ 19.3, “System i Navigator tasks on the web” on page 626
- ▶ 19.4, “System i Navigator Database administration functions” on page 628
- ▶ 19.5, “System i Access for Wireless” on page 630

19.1 IBM i Access for Windows

IBM i Access for Windows is a key offering of the IBM i Access family. It provides PC connectivity to IBM i platforms. It extends IBM i applications and resources to the PC desktop through a graphical user interface. This section discusses the IBM i Access for Windows changes for IBM i 7.1.

For more information about IBM i Access for Windows visit the following link:

<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 have been 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
 - Secondary languages can be selected during custom install.
 - Arabic language support has been added.

Note: Additional changes of which you need to be aware:

- ▶ Support for Windows Server 2008 R2 and Windows 7 has been added.
- ▶ Support for Windows 2000 has been removed.
- ▶ Support for 64-bit Itanium processors has been removed.
- ▶ The \QIBM\ProdData\Access\Windows\Image64i install directory has been 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. These 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 now available in 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
- ▶ Connection property 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. These 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
- ▶ Connection property 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 has been 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 PC desktop. The 5250 emulator included with IBM i Access for Windows for IBM i 7.1 provides updated Display and Printer Emulation based on Personal Communications (PCOMM) version 6.0.

19.1.7 Operations Console enhancements

Support for local console direct attached has been removed from the Operations Console for IBM i Access for Windows for IBM i 7.1.

For more information visit the IBM i Planning web page:

<http://www-947.ibm.com/systems/support/i/planning/upgrade/v6r1/planstmts.html>

19.2 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, database and 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 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 Web application server for i instances
 - Support for the WebSphere Application Server V6.0 for i5/OS (Base, Express, and Network Deployment editions) has been 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
 - Requires PTF SI40905 or later for 5770-XH2
 - Support for the IBM WebSphere Portal V6.0 and IBM WebSphere Portal V5.1 environments has been removed

For more information about IBM i Access for Web visit the following link:

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

19.2.1 AFP to PDF Transform

IBM i Access for Web now offers an additional alternative for generating a PDF document from a spooled file using 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:

- ▶ 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 called **Use AFP to PDF Transform** is available to control whether or not 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 go to **Customize** → **Preferences** → **View all preferences** and look under the Print category. The use of this preference is similar to the use of the **Use Infoprint Server if installed** preference prior to 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
Possible values are “Yes” and “No” (the default value is “Yes”)
- ▶ Use AFP to PDF Transform if installed
Possible values are “Yes” and “No” (the default value is “Yes”)

These options control the use of Infoprint Server and the AFP to PDF Transform. They are similar to the policies/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 is 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:
 - V7R1 or later IBM i
 - V7R1 or later 5770TS1 *BASE and option 1
 - V7R1 or later 5770XH2
- ▶ 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, go to **Print** → **Printer output** and select the View PDF icon as shown in Figure 19-1.

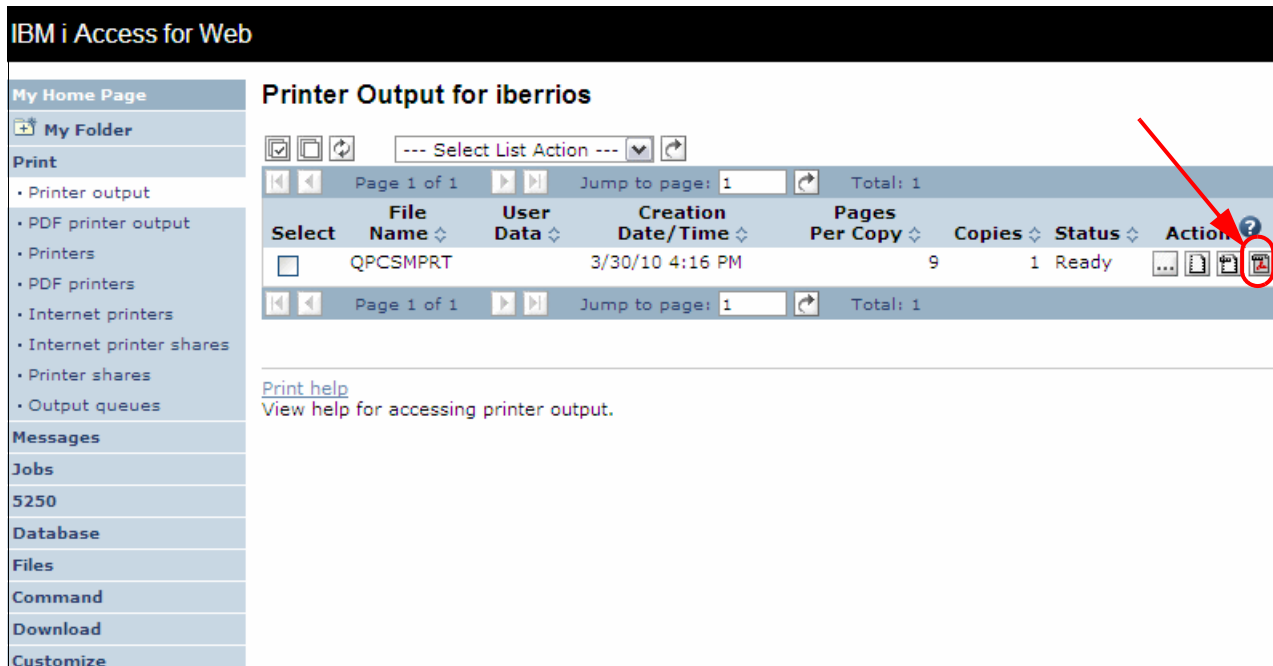


Figure 19-1 AFP to PDF Transform View PDF Icon

The **PDF Output Settings** window is shown in Figure 19-2.

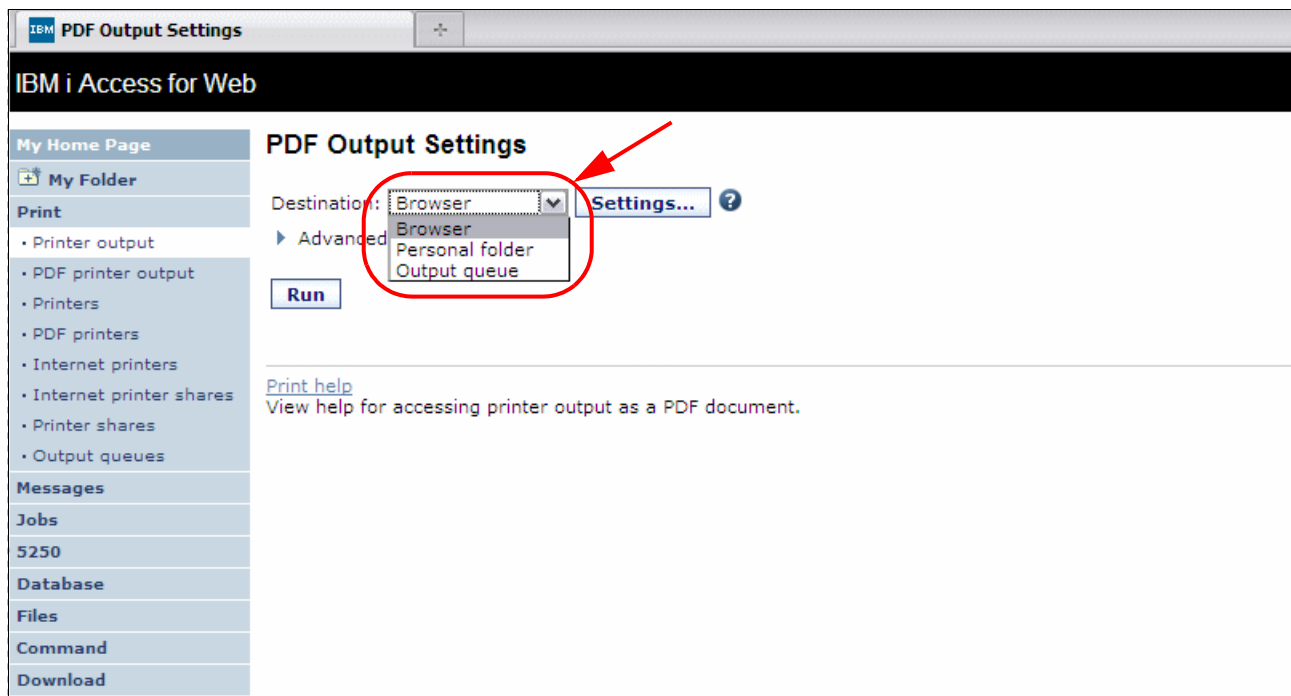


Figure 19-2 PDF Output Settings window

Notice the Destination pull-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 dialog is presented.
- ▶ **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 visit the IBM i information center at the following web page:

<http://publib.boulder.ibm.com/series/>

19.3 System i Navigator tasks on the web

System i Navigator tasks on the web provides a subset of System i Navigator tasks through a Web browser interface. These are the new tasks available for System i Navigator tasks on the web on IBM i 7.1 by category:

- ▶ **System**
 - sysoprmsg - System Operator Messages
- ▶ **Configuration and Service**
 - crtimgcat - Create Image Catalog
 - crtvirtdev - Create Virtual Device
 - graphview - Graphical View
 - mirrorsync - Mirror Synchronization on IPL
 - paritysets - Lists parity sets
 - standalone - Stand-Alone devices
 - imagecatalogs - Tape Image Catalogs
 - tapelibraries - Tape Libraries
- ▶ **Network**
 - stateless - Configure IPv6 Stateless Address Autoconfig
 - tcpipattrIPv6 - Display TCP/IP IPv6 properties
- ▶ **Database**
 - db.crtarray - Create array type
 - db.crtvar - Create global variable
 - db.crtmqt - Create materialized query table
 - db.gblvar - Global variables
 - db.xmlsch - XML schema repository (XSR)
- ▶ **Integrated Server Administration**
 - crtnws - Create server
 - dltnws - Delete server
 - rmtsyswebcon - Launch web console for a remote system configuration
 - srvrpcwebcon - Launch web console for a service processor configuration
 - nwswebcon - Launch web console for an integrated server

- Advanced Job Scheduler
 - actlogprop - Activity log properties
 - esclst - Escalation lists
 - newmail - New email
 - newesclst - New escalation list
 - newjobgrp - New job group
 - newoutqmon - New output queue monitor
 - newrecip - New recipient
 - newrepdstlst - New report distribution list
 - newschjob - New scheduled job
 - notifyprop - Notify properties
 - outqmon - Output queue monitors
 - reciplst - Recipients
 - repdstlst - Report distribution lists
 - resetschjob - Reset scheduled jobs
 - schactprop - Scheduled activity properties
 - schjobprop - Scheduled job properties
 - sentmail - Sent
 - strsch - Start scheduler
 - endsch - Stop scheduler

To execute a task using System i Navigator tasks on the web open a web browser and direct it to the following URL:

`http://hostA:2001/webnav/WnServlet?task=taskid`

In this URL, *hostA* is the IP address or name of the host and *taskid* is the System i Navigator on the web task ID. Figure 19-3 shows an example when using the task ID *sysoprmsg* to display System Operator Messages

Select	From User	Type	Message	Sent
<input type="checkbox"/>	Vermaere	Information	Usage limit of 0 exceeded. Grace period expires in 53 days on 05/17/10.	3/26/10 9:58:58 AM
<input type="checkbox"/>	Vermaere	Information	Grace period expires in 53 days on 05/17/10.	3/26/10 9:58:58 AM
<input type="checkbox"/>	Qsys	Information	IBM i usage limit exceeded - operator action required.	3/26/10 9:49:56 AM
<input type="checkbox"/>	Qsys	Information	IBM i grace period expires in 41 days on 05/05/10.	3/26/10 9:49:00 AM
<input type="checkbox"/>	Qpgmr	Information	User profile MSRJJ disabled for IBM i Support for Windows Network Neighborhood access.	3/26/10 9:46:54 AM
<input type="checkbox"/>	Vermaere	Information	Software problem data for QVPEPRTC has been logged. Refer to help text for additional information.	3/26/10 9:42:23 AM
<input type="checkbox"/>	Vermaere	Information	Usage limit of 0 exceeded. Grace period expires in 57 days on 05/21/10.	3/26/10 9:29:52 AM
<input type="checkbox"/>	Qsys	Information	IBM i usage limit exceeded - operator action required.	3/26/10 8:49:55 AM
<input type="checkbox"/>	Qsys	Information	IBM i usage limit exceeded - operator action required.	3/26/10 7:49:54 AM
<input type="checkbox"/>	Qsys	Information	IBM i usage limit exceeded - operator action required.	3/26/10 6:49:54 AM
<input type="checkbox"/>	Qsys	Information	IBM i usage limit exceeded - operator action required.	3/26/10 5:49:53 AM

Figure 19-3 System i Navigator on the web example for task ID *sysoprmsg*

For more information about System i Navigator tasks on the Web visit the following web page:

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

19.4 System i Navigator Database administration functions

IBM i 7.1 offers the security officer the flexibility to authorize a user or group to Database Administration Functions using the Application Administration function in System i Navigator. The user or users still need the correct object authority.

To work with the Database Administration function in System i Navigator, follow these steps:

1. Launch **Application Administration** as shown in Figure 19-4

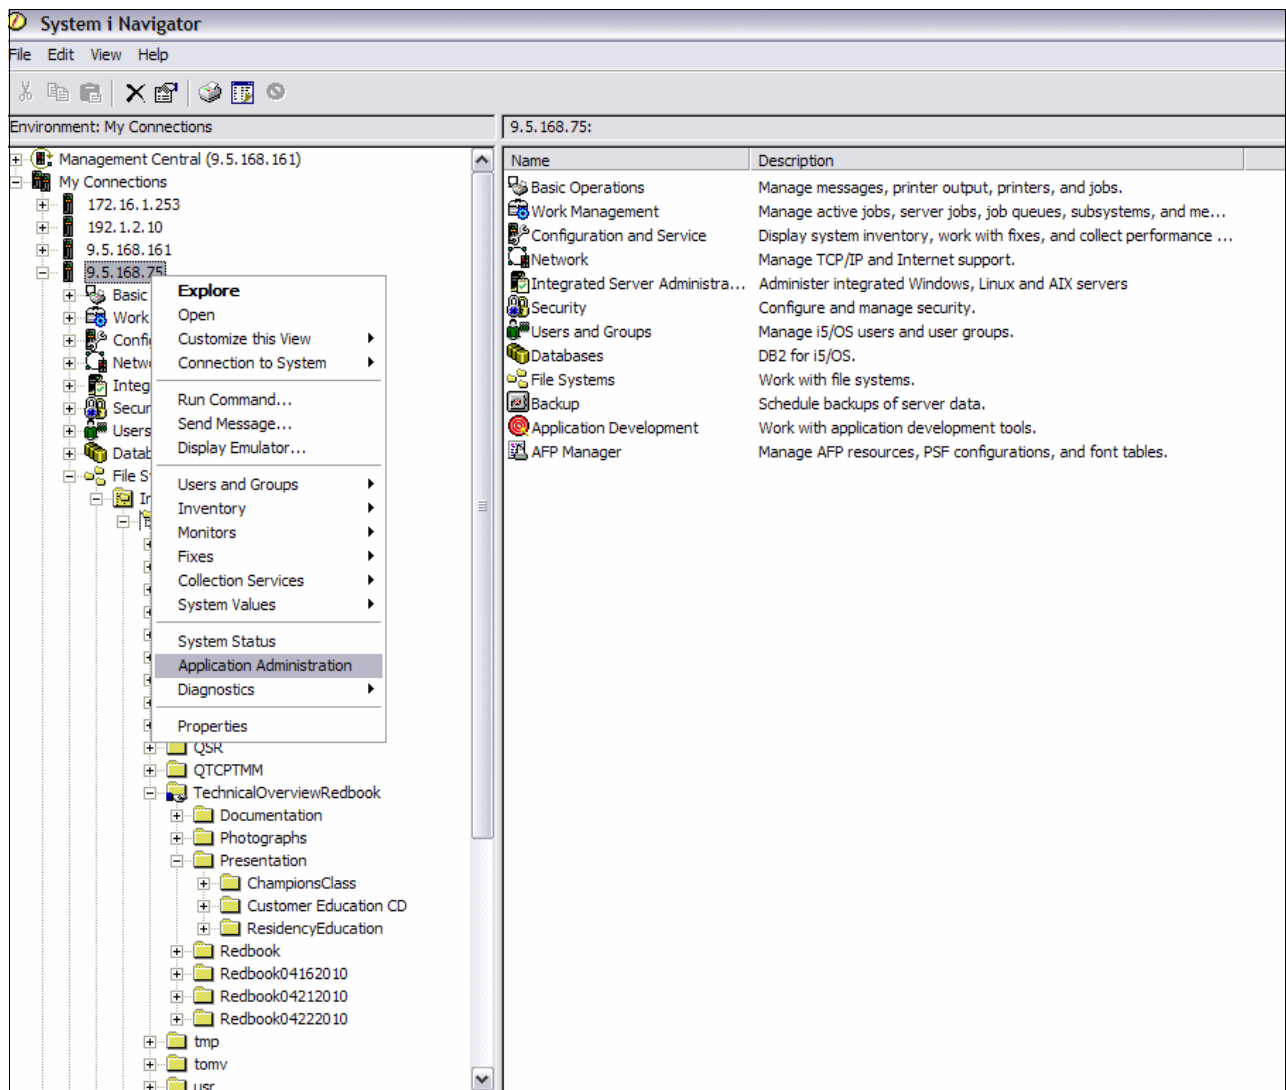


Figure 19-4 Launch Application Administration

- Expand the IBM i and Database folders under the Host Applications tab, as shown in Figure 19-5.

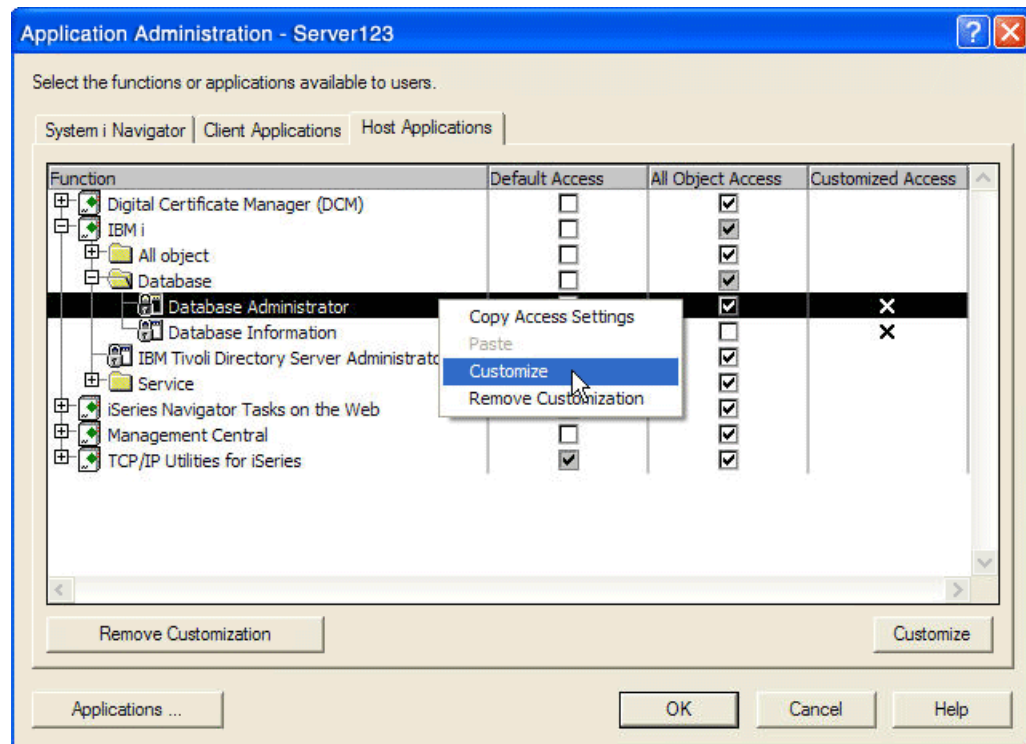


Figure 19-5 Expand the Database group

- Customize the Database Administrator function usage, as shown in Figure 19-6.

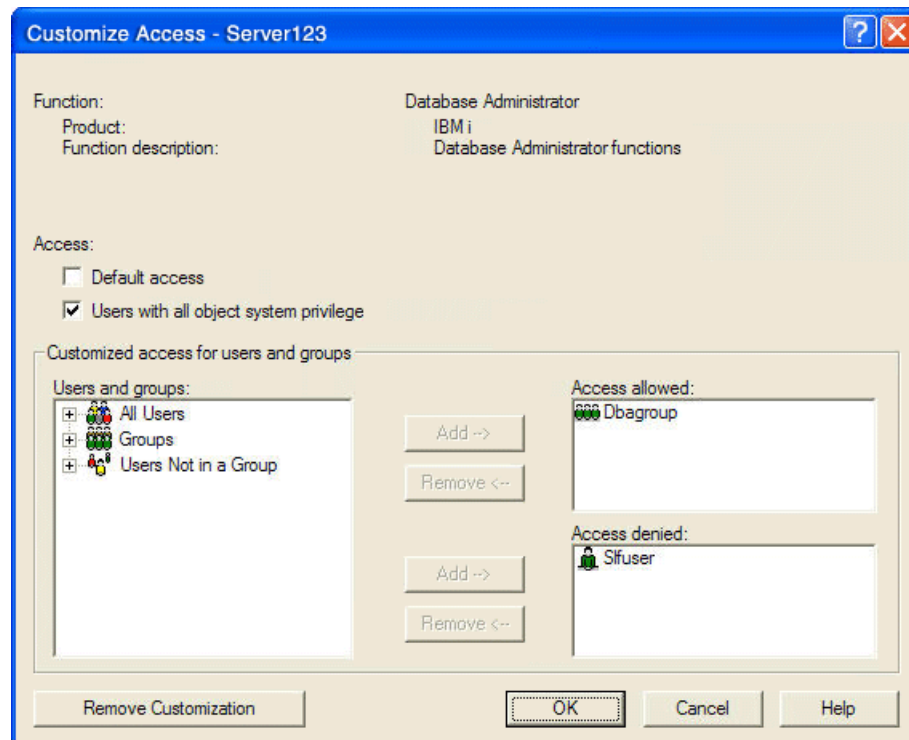


Figure 19-6 Change the QIBM_DB_SQLADM function usage settings

For more information about Database Administration visit the IBM i Information Center at the following web page:

<http://publib.boulder.ibm.com/series/>

19.5 System i Access for Wireless

System i Access for wireless (5722-XP1) is not supported on IBM i 7.1.

For more information visit the IBM i Planning Web Site at the following web page:

<http://www-947.ibm.com/systems/support/i/planning/upgrade/v6r1/planstmts.html>



Miscellaneous enhancements

This chapter introduces the following IBM i 7.1 changes or enhancements not covered in the other chapters of this IBM Redbooks publication:

- ▶ 20.1, “Licensed product program structures and sizes” on page 632
- ▶ 20.2, “Changed or new CL Commands and APIs” on page 633
- ▶ 20.3, “Temporary user-defined file systems” on page 633
- ▶ 20.4, “Watch for Event function (message enhancements)” on page 635
- ▶ 20.5, “IBM Tivoli Directory Server for IBM i enhancements” on page 638
- ▶ 20.6, “Automate Extra IPL for PTF Install” on page 641
- ▶ 20.7, “IBM i workload capping” on page 642

20.1 Licensed product program structures and sizes

IBM i 7.1 operating system and licensed program product (LPP) sizes are documented in the IBM i 7.1 Information Center section “Licensed program releases and sizes” at the following web page:

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzahc/rzahcswslpsze.htm>

The DVD installation media has been consolidated for IBM i 7.1 into three sets of multiple language version media supporting a total of 51 national language versions (NLV).

With IBM i 7.1 there is no offline installation version of the IBM i Information Center shipped on physical media. The IBM i Information Center is available online at the following web page:

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

The following changes to the licensed product program (LPP) structure have been implemented in IBM i 7.1:

- ▶ DHCP has moved from the base OS to option 31 “Domain Name System”
- ▶ Clusters GUI has been withdrawn from option 41 “HA Switchable Resources” and is available with Power HA for i (5770-HAS)
- ▶ IBM HTTP Server i (DG1) option 1 “Triggered Cache Manager” has been removed
- ▶ IBM Toolbox for Java (JC1) has 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) has been withdrawn – IBM Systems Director family provides similar systems management functionality
- ▶ IBM Secure Perspective for System i (5733-PS1, 5724-PS1) has been withdrawn – though continued to be available as a custom service offering only
- ▶ IBM WebSphere Application Server (5733-W61, 5733-W70) minimum required levels for IBM i 7.1 are 6.1.0.29 and 7.0.0.7

Further information about LPP changes is available in the *IBM i Memo to Users* at the following web page:

<http://public.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzaq9/rzaq9.pdf>

Before planning an IBM i release upgrade, refer to the IBM i upgrade planning web site, which provides planning statements about IBM i product changes or replacements at the following web page:

<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 in the IBM i 7.1 Information Center at the following web page:

<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 doing a lot of creation and deletion of temporarily used stream files can most benefit from using temporary UDFSs.

For temporary UDFSs, the system only allocates temporary storage. These temporary files and directories are automatically deleted with 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 it gets cleared at system restarts. The files in a temporary UDFS by their nature should not contain critical data because it is not persistent.

The CRTUDFS command and Systems Director Navigator for i have been enhanced for support of creating temporary UDFSs following a new naming convention. Although names for permanent UDFSs are required to end with .udfs, the names for the new temporary UDFSs follow 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, nor 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
- ▶ Storage used for temporary objects is not accounted to the owning user profile nor to any process

For further information about temporary user-defined file systems see the IBM i 7.1 Information Center at the following web page:

<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 invoked 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 non program 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 on the following trace CL commands:

- ▶ STRTRC (Start Trace)
- ▶ STRCMNTRC (Start Communications Trace)
- ▶ TRCINT (Trace Internal)
- ▶ TRCTCPAPP (Trace TCP/IP Application)
- ▶ TRCCNN (Trace Connection)

Watch for event function was initially available on the trace commands in V5R3M0 of IBM i. Watches were generalized the following release, in V5R4M0, such that the watches were no longer tied only to trace commands. The STRWCH (Start Watch) command and QSCSWCH (Start Watch API) were created for the generalized support. Additionally, the WRKWCH (Work with Watch) command was created to view watches and the ENDWCH (End Watch) command and QSCEWCH (End Watch API) 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 description all the parameters passed to a watch or trace exit program.

20.4.1 Advantages of using watches

Watches can be used as a tool to accomplish the following tasks:

- ▶ Capture data for program debug
- ▶ 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 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 have been expanded to allow generic names including *ALL.

Watch for Message enhancements

Function has been added to expand the message attributes allowed when starting a watch. There are new options for the message ID, message type and severity in 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 watching 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 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*. This allows message CPF1806, CPF1808, CPF1809, and so forth to be handled by starting one watch session
- ▶ Message Type is a new attribute when 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 only want to take an action 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 CRTMSGQ (Create Message Queue) 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 in conjunction 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 only care about messages of severity 99, you can now watch for only those messages.

Figure 19-2 shows the Watch for Message keyword on the STRWCH command for 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 QSCRWCHL (Retrieve Watch List) and QSCRWCHI (Retrieve Watch Information). QSCRWCHL obtains a list of watches on the system. Previously, the list of watches can be viewed using the WRKWCH (Work with Watches) command. QSCRWCHI returns information about a specific watch session. The information returned is similar to the data seen by using 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 ENDWCH (End Watch) command and QSCEWCH (End Watch) API have been 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 most recently started by the same user who is running the end watch function. In 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 all watch sessions with identifiers that begin with the prefix TSTMSG are to be ended.

Beware of 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 greatly affected. That is a watch that 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 caused the CPF1164 and CPF1124 to be generated, hence a never ending loop has begun. 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

In this section we cover the enhancements to the IBM Tivoli® Directory Server for IBM i (LDAP).

20.5.1 Create suffix entries automatically whenever necessary

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 as soon as the first child entry is added.

20.5.2 Administrative roles

IBM Tivoli Directory Server for IBM i now implements a scheme whereby root administrator is able to delegate the tasks at a more granular level, based on the administrative roles of the users defined in the configuration file. These roles are applicable only to the admin group members. 6 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 version 6.2 of Web Administration Tool interface with OS. Web-enablement for LDAP interface on IBM i Navigator: Enable use of LDAP management tool on IBM Systems Director Navigator and IBM i Navigator Tasks for the Web.

20.5.4 Security enhancements

Attribute encryption provides the ability to have arbitrary attributes encrypted 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` 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 but cannot be deleted by a client request. It cannot be modified by a client request either, except for 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 gets 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 his or her own individual password policy. Furthermore, to assist administrators, group password policy is supported to enable effective password management.

20.5.9 Policy enforced for Digest-MD5 binds

The implementation of this feature ensures password policy rules such as account lockout, usage of grace logins, and 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 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 matching 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 made dynamic. For extended operation `readconfig`, addition/deletion/modification of entries having an objectclass of `ibm-slapdReplication/ibm-slapdSupplier` is supported for the scopeValues of `entire/entry/subtree`.

20.5.12 Filtered replication

This allows the directory administrator to control what data is allowed to be replicated to consumer servers by specifying which entries and attributes are to be replicated, based on filters defined by the directory administrator.

20.5.13 Limit number of values 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 returned for an entry and to limit the number of attribute values returned for each attribute in the entry.

20.5.14 Enhanced syntaxes and matching rules

Additional matching rule and syntax support (24 syntaxes, 17 matching rules) has been added for new syntaxes and matching rules from RFCs 2252, 2256, and 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 i 7.1, the Directory Server on IBM i begin to support private IASP.

- ▶ Support database library located in IASP
- ▶ Support change log library located in IASP

20.5.16 Idapcompare utility

The Idapcompare utility compares the attribute value of an entry with a user provided value.

20.5.17 Provide re-entrant LDAP C client library

LDAP C client library have been updated to allow re-entrant.

20.6 Automate Extra IPL for PTF Install

New function is provided in 6.1 and 7.1 to automate any extra IPLs required for a technology refresh PTF or special handling pre-apply PTFs during the PTF installation process. If an extra IPL is required for a technology refresh PTF, we'll save off your PTF install parameters and use them during the next IPL. Instead of seeing "Confirm IPL for Technology Refresh or Special Handling PTFs" panel, you'll see a new English message CPF362E: "IPL required to complete PTF install processing". Actually, if you select Automatic IPL=Y on the "Install Options for PTFs" panel, you won't see any messages or panels, we'll just power down. On the next normal IPL, we'll do your second "GO PTF" during the "PTF Processing" IPL step in the SCPF job, and then automatically do a second IPL of the partition. So when the system IPLs the second time up to sign-on, your PTFs are all activated and ready to go.

Your total PTF install time will be shorter since we won't start any of the system jobs during the first IPL when we restart the partition. However, you'll have a somewhat longer IPL time since we're doing the work you previously did interactively – the second GO PTF to set all PTFs for delayed apply. If you're wondering why the IPL is taking so long and what it's doing, you can always bring up the console. The "Operating System IPL in Progress" panel is shown with the "PTF Processing" IPL step active. Previously, we only displayed "Applying PTFs" for the Activity in this step. But now you will also see "Loading PTFs" or "Setting IPL Action for PTFs" if the previous PTF install was incomplete. When all PTFs have been set for delayed apply, you'll see "IPL requested by PTF processing" status message at the bottom of the screen and then the partition will be restarted again to apply the delayed LIC PTFs. The next time we reach the "PTF Processing" IPL step, we'll display the usual "Applying PTFs" and the IPL will continue.

To take advantage of this new function, you must have the following PTF (PTF management code) temporarily applied BEFORE you perform your PTF install:

- 7.1: SI43585 in HIPER PTF group SF99709 level 30 or higher
- 6.1: SI43939 in HIPER PTF group SF99609 level 94 or higher

For 7.1, if an IPL is required for a technology refresh PTF, the new function only supports installing from a virtual optical device or *SERVICE (PTFs downloaded electronically to save files). If you are installing from a physical optical device, you'll still need to perform the extra IPL and second GO PTF manually. So, if you received your PTFs on physical DVDs, just create yourself an image catalog from the DVDs and use the new support.

20.7 IBM i workload capping

IBM i 7.1 now provides workload capping. Workload capping provides the ability to 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 running on the IBM i system. The user or system administrator can define a workload capping group, assigning a specified number of processing cores to that group. The workload capping group is then assigned to a job or subsystem. Once the assignment has been 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 running (and threads) under the subsystem are not allowed to run on more processing cores than have been designated. The general concept is if a workload is designated to use a single core, the workload will behave as if it is truly running on a single processor core system.

20.7.1 Example of how workload capping works

A user has a batch job that is very CPU intensive. The user needs to run this job during the day but can't afford to impact the performance of their production system. By assigning this batch job to a workload capping group, this job can be put into a "processing container" to help ensure this job is kept to a limited amount of system capacity. If the capping group has a processor core limit of one, then the batch job and any threads running under that job will only be allowed to run 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 will be used at a time. This same concept also applies to jobs running under a subsystem that has been assigned to a workload capping group. All jobs and their associated threads will be limited to the number of processor cores specified in the capping group.

This new capability can help users get better control of the workloads on their systems along with ensuring products are only using a designated number of processor cores. Software vendors can take advantage of the workload capping support to support sub-LPAR licensing. This means that product entitlements can be specified based on 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 tooling that allows you to both register and manage the enforcement of the workload capping. To help users manage and understand the performance of jobs running in a workload capping group, the performance metrics have been updated to include metrics on workload capping.

To learn more about the workload capping support, visit the IBM i Information Center at

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzaks/rzaksworkloadcapping.htm>



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 is also what enables IBM i 5.4 and 6.1 as the OS levels that provide an upgrade path to IBM i 7.1. However, when upgrading from IBM i 5.4 to 7.1 there are special considerations pertaining to object conversions, which are much the same as previously noted when upgrading to IBM i 6.1 from a previous release. This chapter highlights several of these planning details, although it is suggested to review the *i5/OS V6R1 Memo to Users* for additional details and the IBM Redpaper™ publication *IBM i Program Conversion: Getting Ready for 6.1 and Beyond*, REDP-4293, which provides a more complete set of planning considerations. This Redpaper is available at the following web page:

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

This chapter also contains information highlighting 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 License Program Products (LPP) support when upgrading to IBM i 7.1:

- ▶ 21.1, “Installing or upgrading IBM i” on page 644
- ▶ 21.2, “Upgrading from i5/OS 5.4 to IBM i 7.1” on page 647
- ▶ 21.3, “Media delivery changes” on page 649
- ▶ 21.4, “IBM i network upgrade” on page 651
- ▶ 21.5, “Additional considerations for upgrade to IBM i 7.1” on page 653
- ▶ 21.6, “IBM i network installation” on page 659

21.1 Installing or upgrading IBM i

The following sections describe planning considerations when upgrading 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, Power System, or personal computer. You can access the IBM i Information Center at the following web page:

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

Be familiar with the following essential web sites and documents:

- ▶ *Memo to Users*. For IBM i 7.1 this is available online at the following web page:
<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*
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp?topic=/rzahc/rzahc1.htm>

21.1.1 IBM i 7.1 supports POWER5 and later hardware

To determine whether your existing hardware is supported by IBM i 7.1, review Figure 21-1. Notice that IBM i 7.1 supports POWER5 and later hardware. For additional details, see the Upgrade planning web page:

<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

Note: 550* - IBM i 6.1 or later required for POWER6+™ 550

For enterprise clients, IBM i 7.1 is now supported on the 16-core through 256-core Power 795, enabling growth. IBM i supports up to 32 cores in a single partition. IBM Lab Services can be contacted for 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* web page:

<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 5.4 or 6.1 is supported. See Figure 21-2.

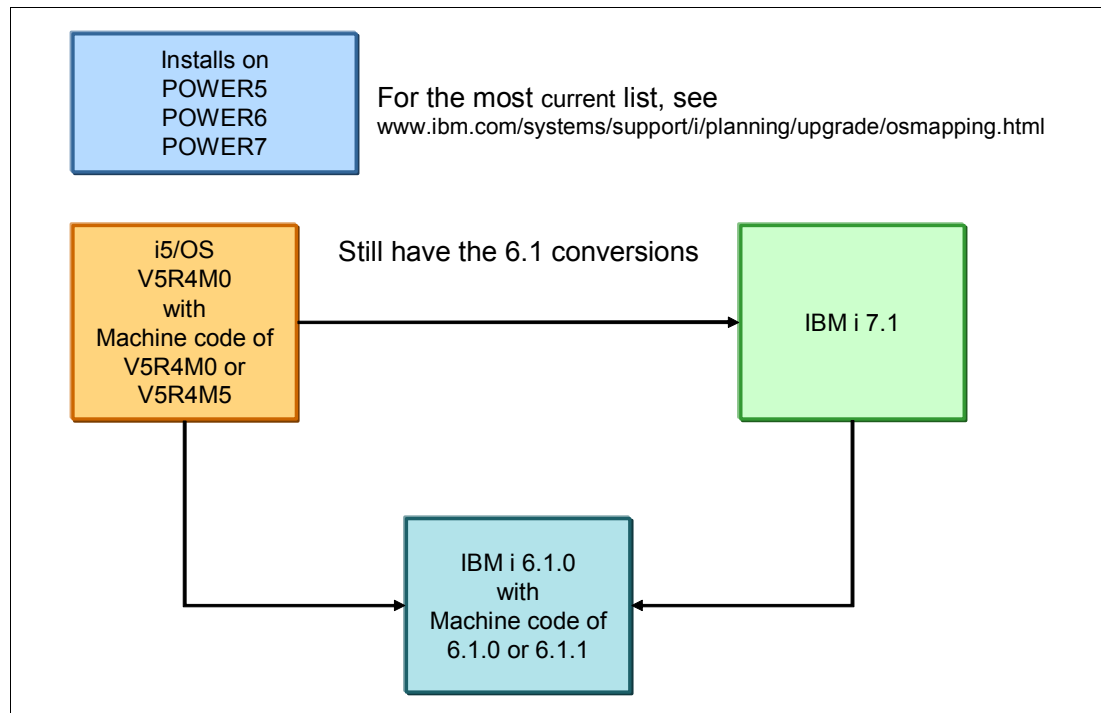


Figure 21-2 Supported upgrade paths to IBM i 7.1

The effective withdrawal from marketing date for IBM i 5.4, and additional selected programs and features, as announced in Software withdrawal 909-003, dated January 27, 2009, has been revised. Withdrawal Announcement 909-285, dated November 10, 2009, has extended marketing availability for IBM i 5.4 from January 5, 2010, to January 7, 2011. This enables systems that might still be running on IBM i 5.2 or 5.3 with the option to still order 5.4 and then perform a two step upgrade. This means you can follow a supported upgrade path to IBM i 5.4, then upgrade to IBM i 6.1 or 7.1.

21.1.4 Program temporary fixes for upgrading to IBM i 7.1

As you plan for your IBM i installation or upgrade, ensure that you review Information APAR II14482 for IBM i 7.1, which is available from the preventive service planning (PSP) section.

To access the PSP, select **Technical Databases** → **Authorized Program Analysis Reports (APARs)** at the following web page:

<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 has been updated for IBM i 7.1. The tool runs on a Windows client and checks the IBM i environment to ensure that all necessary requirements are completed to help ensure a successful upgrade. The tool is available at the following web page:

http://www-912.ibm.com/s_dir/slkbase.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 web page:

<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 POWER5/6/7 server. This must be at a certain minimal level to support IBM i 7.1.

Find the firmware level and IBM i OS levels needed as prerequisite requirements for features that you currently have or plan to add to your system, see the IBM Prerequisite web page:

<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).

Please see Table 21-1 for the mapping between the SLIC Resave level and the IBM i Resave level.

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
10/14/2011	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 with 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 currently on 6.1 there is no additional program conversion requirements for upgrading to IBM i 7.1.

If 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 web page:

<http://publib.boulder.ibm.com/infocenter/iseriess/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 upgrading from IBM i 5.4, allow additional time to analyze your system and adjust your programs for conversions. The length of time required to perform the analysis varies based on the individual system environment. Program conversion can also affect third party 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 5.4 to assist in 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 645

For more complete preparation and planning details see IBM Redpaper *IBM i Program Conversion: Getting Ready for 6.1 and Beyond*, REDP-4293, available at the following web page:

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

The information center has detailed options and instructions:

<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzahc/rzahcswsv5r1con.htm>

21.2.2 Spooled file conversions

When you upgrade from V5R4 to IBM i 7.1, spooled file operations are processed more efficiently than in previous releases due to 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. This can help you determine your best options.

Spooled files 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 5.4 and then restoring them after you have IBM i 7.1 installed.

Additional options are available for managing the spool file conversion after the upgrade. The information center has detailed options and instructions:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzahc/rzahcspooledfileconversions.htm>

Note: Only newly created spooled files or spooled files that have been 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 has occurred.

21.2.3 Integrated file system conversions

As of V6R1, file systems that are not case sensitive in the integrated file system support Unicode Standard 4.0 for names stored in directories.

In 5.4, the integrated file system stores all names in Unicode and supported Unicode Standard 2.0. After upgrading from 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 additional 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(*MONO) on the CRTUDFS (Create User-defined File System) command.

Before you upgrade from V5R4, review Information APAR II14306 and the Redpaper *IBM i Program Conversion: Getting Ready for 6.1 and Beyond*, REDP-4293. These resources help you to analyze your system and help identify objects that are going to be affected by the Unicode conversion. You can then decide if you want to change the names of the affected objects before you upgrade or let the automatic conversion occur. To view the information APAR, see 21.1.4, “Program temporary fixes for upgrading to IBM i 7.1” on page 645

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:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/topic/rzahc/ifsconv.htm>

21.2.4 Backup Recovery and Media Services (BRMS)

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 required to complete software upgrades that include BRMS.

The information center has detailed options and instructions:

<http://publib.boulder.ibm.com/infocenter/iserics/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 desired. 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 has been consolidated. Previously there were 51 sets of unique language version media. This has been reduced to three pre-packaged sets of multiple language version media.

IBM i 7.1 packages 51 NLVs into 3 language groups for keyed media set

This package has the following characteristics. See Table 21-2 on page 650.

- ▶ One language group will deliver 11–21 NLVs, depending on group ordered
- ▶ US English included in each language group
- ▶ Install as many primary and secondary languages as desired from the group
- ▶ For a release upgrade, system will install the same NLV as on current release
- ▶ On a manual install select language / NLV desired and then proceed with installation.
- ▶ Full automatic installation and upgrades from optical media that are changing the primary language need to set the installation language using the QINSTLNG API.
- ▶ Can order more than one group, if desired, 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		
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 complete details see *Media labels and their contents* in the IBM i 7.1 Information Center at the following web page:

<https://publib.boulder.ibm.com/infocenter/iserics/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 applies to all initial and upgrade orders for currently available releases of IBM i, which include 5.4, 6.1, and 7.1. Each additional order for physical media also incurs an additional \$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 web site, see the following web page:

http://www-947.ibm.com/systems/support/software/delivery/en_US/downloadinfo.html

Additionally, a new API called QVOIFIMG (Fill Image Catalog) has been made available for IBM i 7.1, 6.1 and 5.4 through PTFs. This API makes it easier to use image catalogs when working with images that have been 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 645.

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 now, to IBM i 7.1 remotely using install images on a network file server. This enhancement extends the ability to perform not only upgrades, but installations and maintenance activities for remote systems.

Prior to this enhancement you needed to install with physical media or with virtual media located locally on the system being upgraded. To use virtual media required using FTP to manually FTP 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 using network install review Information APAR II14482 (for APAR access see 21.1.4, “Program temporary fixes for upgrading to IBM i 7.1” on page 645) and *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:

- ▶ System must be at IBM i 6.1 or greater.
- ▶ POWER6 or later.
- ▶ The install 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 645

A client partition with a virtual optical device type 632B model 003 can access images located 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

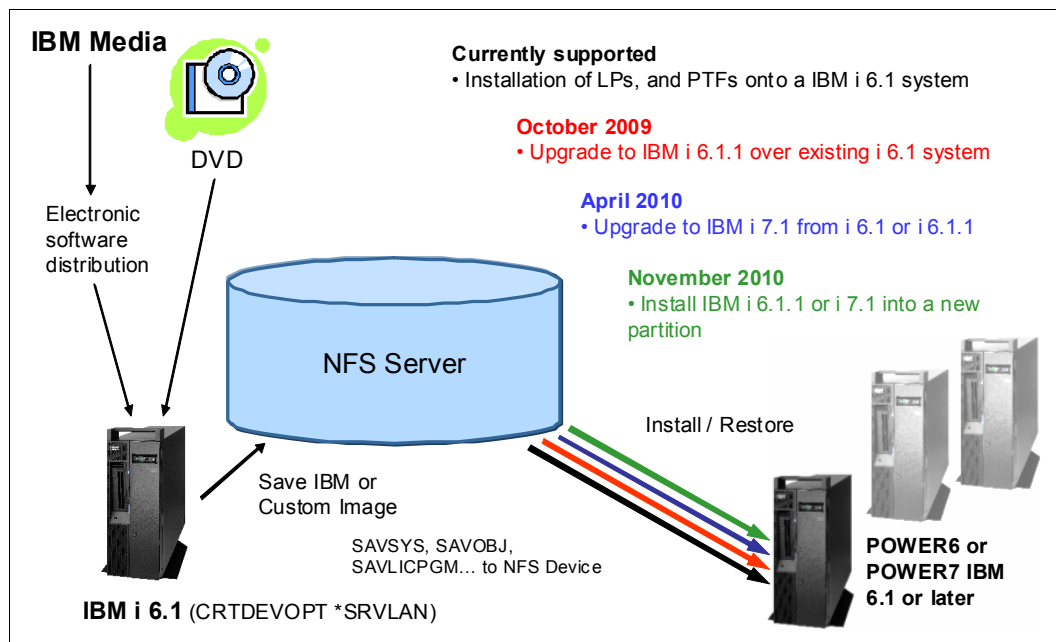


Figure 21-3 Remote install or upgrade using NFS

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 ')
```

Notes:

- (1) The RMTINTNETA is the remote internet address of the Network File system (NFS) server where this virtual optical device will look for virtual image files.
- (2) The NETIMGDIR parameter specifies the network path on the Network File System (NFS) server containing the virtual image files that were prepared for use with this device.

Determine if you need to configure a service tools server. The type of system and configuration determines what type of setup may be required. If LAN console is already configured, no further setup is required. For more information about configuring the Service Tools Server see the following web page:

ftp://ftp.software.ibm.com/systems/support/power/i/nfs_optical_upgrade.pdf

21.5 Additional considerations for upgrade 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 have to ensure that the system is ready to be upgraded before doing so. The following sections detail the changes about which users need to be aware.

21.5.1 IBM i Console changes

Operations Console Direct attached is not supported by IBM i 7.1. If you currently have an Operations Console 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 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 IOP in the Central Electronic 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.

Note: Operations console Direct attached and Twinax console are not supported on any POWER7 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 web page:

<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/iphca/chgconsoles.htm>

The new enhancements on IBM i 7.1 regarding the 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 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, 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
New in IBM i 7.1, you can take over a console type or console device type using the PF key 18. This allows temporary switching 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 preloaded image, and is not HMC managed, then it is required that the console type be set to LAN by

the Using the Console Service functions. For more information about changing consoles, see *Changing consoles* in the IBM systems Hardware Information Center.

Additionally, if LAN console is to use the embedded Ethernet ports, then the Enable Ethernet embedded port (E1) function must be set Using 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 refer to the following URL:

<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/iphbx/functions6521.htm>

Changes in DST/SST Console Security

The changes in the DST/SST console security are highlighted 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	PWLVL 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 planning to change the primary language during the upgrade or installation it is necessary to set the preferred installation language using the QINSTLNG API. This is a new API introduced with IBM i 7.1. For details on using this API see *Set Install National Language Version (NLV) (QINSTLNG) API* at the following web page:

<https://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/apis/qinstlng.htm>

These changes are documented in *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 install actions

Several database files which reside within QSYS2 might not be fully updated during a slip install of Option 1 on an upgrade to IBM i 7.1. See *Memo to Users* for steps to perform before the upgrade:

<https://publib.boulder.ibm.com/infocenter/iserics/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 only supports a one version difference. If all the nodes in the cluster are at the same release, upgrade to the new release, before changing the cluster version. This ensures that all functions associated with the new release are available. See the IBM i Information Center topic *Scenario: Upgrading operating system in an high-availability environment* for detailed actions for an upgrade to a new release. It is available at the following web page:

<http://publib.boulder.ibm.com/infocenter/iserics/v7r1m0/index.jsp?topic=/rzaig/rzaigscenupgrade.htm>

Use the following steps to verify and change the cluster version for a node.

1. In a Web browser, enter `http://mysystem:2001`, where *mysystem* is the host name of the system.
2. Log on to the system with your user profile and password.
3. Select **Cluster Resource Services** from the IBM Systems Director Navigator for i5/OS window.
4. On the Cluster Resource Services page, select the **Display Cluster Properties** task.
5. On the Cluster Properties page, click the General tab.
6. Verify the cluster version setting or change the version to the desired setting.

21.5.5 JAVA considerations

The following sections detail changes that have been made to Java.

IBM Toolbox for Java - 5771-JC1 has moved

With 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 install process automatically removes 5761-JC1 if it is installed on the system.

Java Developer Kit 1.4 - 5761-JV1 Option 6 replaced

The upgrade to IBM i 7.1 automatically removes 5761-JV1 Option 6 and replace with 5761-JV1 Option 13 - J2SE 1.4 64 bit

Java Developer Kit 5.0 - 5761-JV1 Option 7 replaced

The upgrade to IBM i 7.1 automatically removes 5761-JV1 Option 7 and replace 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 replaced

The upgrade to IBM i 7.1 automatically removes 5761-JV1 Option 10 and replace with 5761-JV1 Option 11 - Java SE 6 32 bit and Option 9 - Java SE 6 64 bit

IBM Toolbox for Java - 5771-JC1 has moved

This function has been moved to 5761-SS1 Option 3.

IBM Adopt Authority for Java for i5/OS - 5799-AAJ dropped

This PRPQ provided support for Adopted Authority for Java for the Classic JVM. IBM i 7.1 only supports 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) changes

DHCP in IBM i 7.1 has 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 includes the standard IBM i DHCP server and the new ISC DHCP server. See Chapter 10, “Networking enhancements” on page 315 for more details.

System i Access for Wireless - 5722-XP1 dropped

The recommended product to replace the System i Access for Wireless is a browser based alternative as 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 - changed

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 function 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 that will continue to be supported in IBM i 7.1, however save when active and file level backup are not supported.

IBM Secure Perspective for IBM i - 5733-PS1 dropped

Secure Perspective for IBM i (sold as 5733-PS1 and 5724-PS1) continue to be available by IBM System Lab Services. For more information, see the following web page:

http://www-03.ibm.com/systems/services/labservices/platforms/labservices_power.html

Performance Viewer Open Source Components - 5761-SS1 option 46

This function has been 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 dropped

Replace this product with WebSphere Application Server 7.0

Lotus Domino 7 for i5OS - 5733-LD7 dropped

Upgrade to Domino 8.5.2 (or later) 5733-LD8.

Lotus Domino 8 for System i - 5733-LD8 changed

Upgrade to Domino 8.5.2 (or later) 5733-LD8.

IBM DataMirror iCluster SMB - 5733-ICS change recommended

Use IBM DataMirror iCluster for System i - 5733-ICL. Same or lower prices and more function that SMB

IBM Eclipse Help System removed

The IBMHELP server has been removed in IBM i 7.1. Previously, this TCP server was package with 5761-SS1 option 3 (Extended Base Directory Support). See Figure 21-5.

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

21.6.1 Requirements for IBM i network installation

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 will cause the other PTFs to be ordered and installed)
 - SI39390
 - SI35186
 - SI35189
 - SI35747
 - MF47284
 - MF47285

Client system requirements

- ▶ IBM i 6.1.1 or later
- ▶ Server Firmware
 - POWER 7: 01Ax720 or later; 01Ax710_097 or later (for Physical Adapters)
 - POWER 7: 01Ax730_049 or later (for Virtual Adapters)
 - POWER 6: 01Ex350_085 or later
- ▶ HMC 7.7.2 or later
 - **Note:** Support for setting a Logical Host Ethernet Adapter (LHEA) or a Virtual Ethernet Adapter as the install 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 install source server.

Network configuration requirements

- ▶ The MTU value set on the TCP/IP interface (Server/IBM i) must always match with the MTU value set by the chsysstate command (Client/HMC CLI).
- ▶ Given the fact that chsysstate only supports two different values for MTU (1500 & 9000), on the Server side the Ethernet line description can only be set to MAXFRAMES 1496 or 8996.
- ▶ MAXFRAMES 1496 (line description) translates to MTU 1500 (TCP/IP interface), while MAXFRAMES 8996 (line description) translates to MTU 9000 (TCP/IP interface).
- ▶ Until further notice, only MAXFRAMES 1496 (thus MTU 1500) is supported for Network Installs. If the ethernet line description that you are planning to use on the Server to setup the network install is currently set to MAXFRAMES 8996 (a.k.a. Jumbo frames), you must change it to use MAXFRAMES 1496 before get started.

Refer to APAR II14482 for IBM i 7.1 for PTF and setup information.

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

Also see *IBM Network Installation Using HMC* at the following web page:

http://www-912.ibm.com/s_dir/slkbases.NSF/DocNumber/563261959

Related publications

The publications listed in this section are considered particularly 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 665. Note that a few of the documents referenced here might be available in softcopy only.

- ▶ *IBM Systems Director Navigator for i*, SG24-7789
- ▶ *IBM i 6.1 Technical Overview*, SG24-7713
- ▶ *IBM Power 795 Technical Overview and Introduction*, REDP-4640
- ▶ *IBM Power 770 and 780 Technical Overview and Introduction*, REDP-4639
- ▶ *IBM Power 750 and 755 Technical Overview and Introduction*, REDP-4638
- ▶ *IBM Power 720 and 740 Technical Overview and Introduction*, REDP-4637
- ▶ *IBM Power 710 and 730 Technical Overview and Introduction*, REDP-4636
- ▶ *IBM i Program Conversion: Getting Ready for 6.1 and Beyond*, REDP-4293
- ▶ *Implementing PowerHA for IBM i*, SG24-7405
- ▶ *Getting Started with DB2 Web Query for i*, SG24-7214
- ▶ *PowerVM Virtualization Active Memory Sharing*, REDP-4470
- ▶ *IBM System i5, eServer i5, and iSeries Systems Builder IBM i5/OS Version 5 Release 4 - January 2006*, SG24-2155
- ▶ *PCI, PCI-X, PCI-X DDR, and PCIe Placement Rules for IBM System i Models*, REDP-4011
- ▶ *IBM i 6.1 Independent ASPs: A Guide to Quick Implementation of Independent ASPs*, SG24-7811
- ▶ *IBM Power 550 Technical Overview*, REDP-4404
- ▶ *IBM Power 520 Technical Overview*, REDP-4403
- ▶ *IBM Power 520 and Power 550 (POWER6) System Builder*, SG24-7765
- ▶ *End to End Performance Management on IBM i*, SG24-7808
- ▶ *DS8000 Copy Services for IBM i with VIOS*, REDP-4584
- ▶ *IBM BladeCenter JS23 and JS43 Implementation Guide*, SG24-7740
- ▶ *IBM i and Midrange External Storage*, SG24-7668
- ▶ *IBM System i5 V5R4 Technical Overview Redbook*, SG24-7271
- ▶ *Implementing IBM Systems Director 6.1*, SG24-7694
- ▶ *Security Guide for IBM i V6.1*, SG24-7680
- ▶ *Windows-based Single Signon and the EIM Framework on the IBM eServer iSeries Server*, SG24-6975
- ▶ *IBM System i Security: Protecting i5/OS Data with Encryption*, SG24-7399

- ▶ *PowerVM Virtualization on IBM System p: Introduction and Configuration Fourth Edition*, SG24-7940
- ▶ *VMware VI3 on BladeCenter and System x Integrated with System i*, SG24-7408
- ▶ *Hardware Management Console V7 Handbook*, SG24-7491

Online resources

These web pages are also relevant as further information sources.

- ▶ *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/slkbases/DocNumber/465353483
- ▶ *IBM i Memo to Users 7.1*
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzaq9/rzaq9.pdf>
- ▶ *Installing, upgrading, or deleting IBM i and related software 7.1*
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzahc/sc415120.pdf>
- ▶ *Licensed program releases and sizes*
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp?topic=/rzahc/rzahcswslpsze.htm>

- ▶ *IBM Systems Software Information Center*
<http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp>
- ▶ *IBM Power Systems Hardware Information Center*
<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp>
- ▶ *IBM PowerHA SystemMirror for i*
<http://www-03.ibm.com/systems/power/software/availability/i/index.html>
- ▶ *IBM Systems Management solutions for Power Servers*
<http://www-03.ibm.com/systems/power/software/management/>
- ▶ *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 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>
- ▶ *Lotus Products for IBM i page*
<http://www-03.ibm.com/systems/i/software/domino/index.html>
- ▶ *Lotus Enterprise Integrator for IBM i page*
<http://www-03.ibm.com/systems/i/software/domino/related/lei/index.html>
- ▶ *Lotus Software for IBM i operating system Compatibility Guide*
<http://www-03.ibm.com/systems/resources/releasesupport.pdf>
- ▶ *Connecting to IBM i - IBM Systems Director Navigator for i*
<http://publib.boulder.ibm.com/infocenter/series/v7r1m0/topic/rzatg/rzatgdirector.pdf>
- ▶ *WebSphere Application Server for IBM i*
<http://www-03.ibm.com/systems/i/software/websphere/>
- ▶ *IBM Systems Director page*
<http://www-03.ibm.com/systems/software/director/index.html>
- ▶ *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/>
- ▶ *Backup Recovery & Media Services page*
<http://www-03.ibm.com/systems/i/support/brms/index.html>
- ▶ *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>
- ▶ *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 Advanced Job Scheduler for i* PDF
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzaks/rzaksadjs.pdf>
- ▶ *Print Services Facility for IBM i* page
http://www-03.ibm.com/systems/i/software/print/psfhome_i_ww.html
- ▶ *InfoPrint AFP Utilities for IBM i* page
http://www-03.ibm.com/systems/i/software/print/afputilhome_i_ww.html
- ▶ *AFP Font Collection* page
http://www-03.ibm.com/systems/i/software/print/afpfonthome_m_ww.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>
- ▶ *Profile Data Trace Visualizer (PDTV)* page
<http://www.alphaworks.ibm.com/tech/ptdv>
- ▶ *IBM Systems Workload Estimator* page
<http://www.ibm.com/systems/support/tools/estimator>
- ▶ *Virtualization with IBM i, PowerVM and Power Systems* page
<http://www-03.ibm.com/systems/i/os/index.html>
- ▶ *Power Systems Logical partitioning* document:
<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf>
- ▶ *IBM PowerVM Active Memory Sharing: An Overview*
<http://www-03.ibm.com/systems/power/software/virtualization/whitepapers/ams.html>
- ▶ *IBM i Virtualization and Open Storage Read-me First*
http://www-03.ibm.com/systems/resources/systems_i_Virtualization_Open_Storage.pdf
- ▶ *System Storage Interoperation Center* page
http://www-03.ibm.com/systems/support/storage/config/ssic/displayessearchwithoutjs.wss?start_over=yes
- ▶ *IBM Power Blade Servers - IBM i*
<http://www-03.ibm.com/systems/power/hardware/blades/ibmi.html>
- ▶ *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
- ▶ *Interoperability Guide - IBM BladeCenter*
<http://www-947.ibm.com/systems/support/supportsite.wss/docdisplay?ln docid=MIGR-5073016&brandind=5000020>

- ▶ *IBM i integration with BladeCenter and System x* page
<http://www.ibm.com/systems/i/advantages/integratedserver/library.html>
- ▶ *Linux on integrated servers* page
<http://www.ibm.com/systems/i/advantages/integratedserver/linux/>
- ▶ *Rational Developer for Power Systems Software* page
<http://www-01.ibm.com/software/rational/products/rdp/>
- ▶ *Rational Team Concert for Power Systems Software* page
<http://www-01.ibm.com/software/rational/products/rtcp/>
- ▶ *Rational Team Concert for Power Jazz project* page
<http://jazz.net/projects/rational-team-concert-power/>
- ▶ *Rational Development Studio for IBM i* page
<http://www-01.ibm.com/software/awdtools/wds400/>
- ▶ *IBM Rational RPG Cafe*
<http://www-949.ibm.com/software/rational/cafe/community/rpg>
- ▶ *Rational Host Access Transformation Services* page
<http://www-01.ibm.com/software/awdtools/hats/>
- ▶ *Rational Software Delivery Platform for Power*
<http://www-03.ibm.com/systems/power/software/i/rational/index.html>
- ▶ *Rational Development Studio for i ILE RPG Language Reference*, SC09-2508
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzasd/sc092508.pdf>
- ▶ *Rational Development Studio for i ILE RPG Programmer's Guide*, SC09-2507
<http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/topic/rzasc/sc092507.pdf>
- ▶ *Zend Products* page
<http://www.zend.com/en/products/>
- ▶ *Application Runtime Expert for IBM i* page
<http://www-03.ibm.com/systems/power/software/i/are>

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, access URL <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 from a specific Redbook 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:

1. Locate the book through one of the menu or search options,
2. Select the Redbooks publication. This displays a summary page for the book.
3. Select the “Download PDF” link,.
 - a. A right-click, followed by selecting “Save target as” will download and save the PDF.
 - b. A normal click will display the PDF. Select the ‘Save file’ icon.

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Symbols

SKIP LOCKED DATA 149
.instance.properties 436
.NET Data Provider 622
'SECURE NO 172
"two's complement" representation 132
"Work with Removing Units From Configuration" panel 274
*ADMIN HTTP server 59
*AFPDS 344
*AGG 333
*ALL 138
*ALLOBJ special authority 173
*ALLSYS 34
*BUF128KB 134
*BUF256KB 134
*BUF32KB 134
*BUF64KB 134
*CURAVLCHN 137
*CURCHAIN 137
*CURLIB 137
*CURRENT 137
*DEFAULT 151
*DTAARA 138
*DTAQ 138
*ERR 50
*EXECUTE 132, 347
*FILE 137–138, 159
*FIRST 138
*INZ 50
*IPSCSI 368
*JOBCTL 173
*LIB 138
*LIBL 137
*NO 151
*NONE 138, 344
*OBJEXIST 132
*OPTIMIZE 151
*OUTFILE 37
*PDF 345
*PGM 153
*PURGE option 147
*SAVCFG 34
*SAVSYS 34
*SCS 344
*SECDTA 34
*SECURE 178
*SRVPGM 120, 124, 153
*STGMDL 479
*SYSBAS 139
*TOSTMF 344–345
*USERASCII 344
*YES 151
#5274 290

#5638 290
#5886 288
#5887 288

Numerics

1.5TB/3.0TB LTO-5 SAS Tape Drive 290
10 Gb 366
1200 114
1200 CCSID 114
1208 CCSID 114
128-byte schema names 622–623
13488 CCSID 114
283 GB 288
5.4 112
5250 Emulation 526
571 GB 288
5722-DR1 370, 381
5722-IP1 556
5722-XP1 632
5724-PS1 632
5733-PS1 632
5733-QU2 185
5733-QU3 185, 187
5733-QU4 185, 187
5733-SC1 339
5733-W61 432, 632
5733-W70 432, 632
5761-LSV 381, 632
5770-JS1 Advanced Job Scheduler for IBM i 385
5770-TS1 556
5770TS1 624
5770XH2 624
5799-SD1 312
6.1 112
64-bit Itanium® 622
65535 345
65535 CCSID 114
7.1 112
770TS1 IBM Transform Services for i 347

A

A SELECT list without DISTINCT 113
abnormal end 134
access paths 139
access to multiple remote SQL Server databases 186
Action Threshold 573
activating a remote journal. 136
Active Energy Manager 609
Active Memory Sharing 242
active path 277, 293
Active Technologies 186–187
actlogprop 627
adapter roles 279
adaptive query processing (AQP) 141–143

- Add Server Authentication Entry 179
- ADDASPCPYD 103
- adding views 600
- ADDPFTRG 170
- ADDSVRAUTE 179
- ADDTAPCTG 289
- administering XML schemas 116
- ADTS 514
- advanced analysis 309
- advanced job scheduler capabilities 386
- AEM 609
- AES key 283
- AFP to PDF Transform 623–624
- aggregate function 114, 130, 133
- aggregate functions 112–113
- Aggregate Policy (AGGPCY) parameter 333
- Aggregated resource list parameter 333
- AJS 386
- ALIAS 131
- alias 130
- ALIAS keyword 476
- alias name 130
- ALLOC 479
- ALLOCATE CURSOR 123
- ALLOCATE CURSOR statement 123
- ALLOW_ADAPTIVE_QUERY_PROCESSING 151
- ALLOW_ARRAY_VALUE_CHANGES 151
- ALTER 131
- ALTER TABLE 124–126, 147, 303
- ALTER TABLE/INDEX 141
- ALWCANCEL(*YES) 154
- ALWOBJDIF 53
- ALWOBJDIF(*ALL) 29
- ALWOBJDIF(*COMPATIBLE) 29
- AMS 242
- analytical process 186
- Annotated XML schema decomposition 117
- another alias 130
- another server 130
- AnyNet is no longer supported on IBM i 7.1 332
- AnyNet® 332
- API 170
- API finder 633
- application 131
- application a 121
- Application Administration function 628
- Application Development ToolSet 514
- application encoding 115
- application program 130
- application requestor (AR) 180–181
- application servers (AS) 180–181
- application string 113
- applications 112
- applied permanently 138
- Apply Journaled Change 136
- Apply Journaled Change Extend 136
- appropriate encoding specified to the data 115
- APYJRNCHG 136
- APYJRNCHGX 136
- AQP Handler 143

- AQP Request Support 143
- archive 118
- archive table 118
- archiving data 118
- argument 132
- argument of XMLSERIALIZE 113
- arguments 114, 133
- array 121
- array of another data type 121
- Array Support 538
- array type 121
- ARRAY_AGG 121
- Arrays in procedures 121
- arrays in procedures 112
- ASP 139
- ASP balancer 307–308
- ASP balancer media preference balance type 308
- ASP balancer migration priority 308
- ASP balancer scheduling 309
- ASP copy description 103
- ASP encryption 24, 282
- ASP master key 283
- ASP session 104
- ASSOCIATE LOCATORS 123
- ASSOCIATE LOCATORS statement 123
- associated XML schema 116
- Asynchronous delivery 561
- asynchronous delivery 100
- asynchronous Geographic Mirroring 100
- asynchronous remote journal connections 136
- ata Base file statistics collection (QDBFSTCCOL) 143
- ATOMIC MERGE 119
- attribute-string 131
- audit functions 134
- authentication entry 181
- authorization lists 107
- AUTL 107
- Automatic cancel of QSQSRVR jobs 150
- auxiliary storage 155
- auxiliary storage segments 307
- Availability and consistency 111
- AVG 146

B

- Backup and archive control groups 62
- Backup Recovery and Media Services 10, 27
- barcodes 289
- base module 186
- Beginning date 35
- Beginning time 34
- BETWEEN 114
- BIGINT 122, 132
- binary data types 113
- BIT 133
- BIT function 132
- BIT scalar functions 112, 132
- BITAND 132–133
- BITANDNOT 132–133
- BITNOT 132–133
- BITOR 132

- BITXOR 132–133
- BladeCenter 365
- BLOB 115
- BLOB data type 115
- BM Disk Sanitizer PRPQ 312
- boundaries of a business or logical transaction 138
- BRM expiration (QP1AEP) report 51
- BRM restricted procedure ended 48
- BRM restricted procedure started 48
- BRM0010 52
- BRM1730 53
- BRM6714 47
- BRM6715 48
- BRM6716 48
- BRM6717 47
- BRMS advanced functions panel 62
- BRMS enrolled volume rejected 53
- BRMS file management 62
- BRMS file reorganization 62
- BRMS history 61
- BRMS initial panel 59
- BRMS Log 72
- BRMS log 61
- BRMS Log - Include page 73
- BRMS Online Lotus Server Backup Reference 32
- BRMS recovery policy 37
- BRMS Restore commands 38
- BRMS System Policy 43
- BRMSSYSPCY 43
- Browser 626
- buffer length 134
- building reports 185
- built-in data type 114
- built-in functions to protect the integrity of certain objects 134
- Business Intelligence applications 185
- Business Intelligence data load scenario 118

C

- C 121–122
- C9279610 293
- cache statistics 296
- CALL 122
- CALL statement 112, 123, 130
- cardinality 121
- cartridge ID 289
- CAST a parameter marker 113
- CAST NULL to XML 113
- CAST XML, or NULL to XML 113
- CCSID 114, 129, 155, 344–345
- CCSID 1200 129
- CCSID value 114
- CEC node level mirroring 285
- CecNodeLevelOfProtection 288
- CFGDEVASP 93
- CFGGEOMIR 94
- CFGHTTPSCH 432
- Change DDM TCP/IP Attributes 179
- Change Journal State (QjoChangeJournalState) 136
- Change Journaled Object 135

- Change Logical File Member 162
- Change Network Group panel 44
- Change PF Constraint 154
- Change Physical File Member 162
- Change Recovery Policy panel 42
- Change Remote Journal 136
- Change Remote Journal (CHGRMTJRN) 136
- CHANGE_PLAN_CACHE_SIZE procedure 185
- changeConnection 622
- changeSupport 622
- check constraint 114
- CHECK parameter 154
- CHGASPSSN 102
- CHGCLUNODE 105, 108
- CHGCLUVER 93
- CHGDDMTCPA 179
- CHGDEVMLB 289
- CHGJRNOBJ 135
- CHGLF 148, 303
- CHGLFM 162
- CHGLINETH 339
- CHGNWSCFG 380
- CHGPF 303
- CHGPFCST 154
- CHGPFM 162
- CHGRMTJRN 136
- CHGSRCPF 303
- CHKASPBAL 307
- CHRID attribute 345
- CHRID attribute other than *DEVD 345
- CIM server 106
- CL command finder 633
- CLASS_FUNC 120
- Classic Java 443
- CLI 113, 162, 170
- client 115
- CLIENT ACCTNG 182
- CLIENT APPLNAME 182
- CLIENT PROGRAMID 182
- CLIENT USERID 182
- client workstation name 182
- CLIENT WRKSTNNAME 182
- CLOB 126–128
- cluster administrative domain 107
- cluster device domain 283
- cluster failover 105
- cluster node outage 105
- cluster resource group 103
- clustering GUI 97
- Clusters GUI 632
- COALESCE, 113
- COBOL 122
- CODEDCHARSETID 129
- Collection Duration (MAXSIZE) 230
- collection of rows 114
- Collection Overlap (OVLAP) 230
- Collection Services 156
- collection services 30, 289
- Collection Services Investigator 284
- column description 125

- column name 130
- Column Value Descripto 125
- column's field description 125
- columns with a field procedure 125
- Command to create a virtual line description on another partition to use the bridge 339
- Command to create line description for the physical link. 339
- Command to create the virtual line description for the bridge link. 339
- COMMENT parameter 182
- COMMIT 120
- commit 131
- COMMIT(*ALL) 119, 131
- COMMIT(*CHG) 119, 131
- COMMIT(*CS) 119, 131
- COMMIT(*NONE) 119, 131
- COMMIT(*RR) 119
- commitable operation 115
- Commitment control 138
- commitment control 109, 138–139
- commitment control is 139
- commitment definition 139
- committed status 138
- COMP-5 473
- compatibility 112
- COMPUTATIONAL-5 473
- CONACC 131
- concurrent access resolution 112, 131, 622–623
- concurrent access resolution option 131
- Configuring a shared network connection via Ethernet layer-2 bridging 337
- Configuring a virtual adapter using an Hardware Management Console 337
- Configuring a virtual adapter using IBM i Virtual Partition Manager 338
- Configuring Ethernet link aggregation 333
- CONNECT 130
- CONNECT statement 179–181
- connection 115
- constraint restrictions 112
- constrabt restriction on partitioned tables 112
- CONTAINS 113
- CONTINUE ON EXCEPTION 119
- control decomposition 117
- control group selection parameters 73
- control path drive 295
- Copy from Import File 155
- Copy Services toolkit 294
- Copy Spoolfile 344
- CORRELID 128
- cost-based query optimizer 112
- COUNT 146
- COUNT and COUNT_BIG aggregate functions 113
- countprefix 478
- CPDB8EB 100
- CPU 178, 564
- CPU health indicators 568
- CPYFRMIMPF 155
- CPYFRMIMPF performance 155

- CPYSPLF 344, 348
- CPYSPLF command 346
- CPYSPLF command enhanced to copy to PDF or TIFF stream files 344
- CPYSPLF command prompter 345
- CREATE 131
- create an object 131
- CREATE command 131
- CREATE ENCODED VECTOR INDEX 146
- CREATE FUNCTION 152
- Create Image Catalog 562
- CREATE INDEX 303
- Create Logical File 148
- CREATE OR REPLACE 131
- CREATE PROCEDURE 152–153
- Create Server 373, 552
- Create Server Wizard 374
- CREATE TABLE 124–126, 303
- CREATE TABLE/INDEX 141
- CREATE TRIGGER 152
- CREATE TYPE 121
- CREATE VARIABLE 120
- Create Virtual Device 562
- cross-partition CPU 574
- CRTDEVNWSH 368
- crtimgcat 626
- CRTL 148, 303
- CRTLINETH 339
- crtnws 626
- CRTNWSCFG 380
- CRTPF 303
- CRTSRCPF 303
- CRTUDFS 311, 634
- crtvirtdev 626
- CTLCHAR 345
- CURRENT DATE 125
- CURRENT IMPLICIT XMLPARSE OPTION 115
- current inherit rules associated with a journaled library 135
- CURRENT TIME 125
- CURRENT TIMESTAMP 118, 125
- currently committed version of the data 131
- CVD 125

D

- DAOS 31–32
- DAOS Estimator 32
- DAOS Quick Start Guide 32
- DASD 564
- dashboard 186
- data characteristics of the decoded values 125
- data encryption key 283
- data moving tasks 308
- Data Series 603
- data type 113, 121, 132
- data warehouse 118
- Database Administration 532
- database cross reference 155
- database file 155
- database full opens 156

- Database Information 532
- database manager 122, 131
- Database Monitor 159, 178, 184
- database parallelism 112
- database persistency, 134
- database server 115
- database supplied procedure 148
- DATALINK 125
- DataLink 125
- datasubf 478
- DB Group PTFs 159
- db.crtarray 626
- db.crtvar 626
- db.gblvar 626
- db.xmlsch 626
- DB2 112, 125–126, 129
- DB2 administrators 173
- DB2 Availability and Recovery Enhancements 184
- DB2 component 113
- DB2 database manager 115
- DB2 database server 115
- DB2 enhancements 112
- DB2 Family products 112
- DB2 for Web Query for i 185
- DB2 Group PTF 184
- DB2 Group PTFs 184
- DB2 media preference 302
- DB2 monitoring tasks 174
- DB2 MQ CCSID 129
- DB2 MQ functions 129
- DB2 MQ scalar or table function 129
- DB2 MQ table functions 129
- DB2 MQ tables 129
- DB2 reads statistics 304
- DB2 security enhancements 111, 171
- DB2 Statistics Manager 142
- DB2 Web Query 186–187
- DB2 Web Query Adapter for Microsoft SQL Server 186
- DB2 Web Query BASE product 187
- DB2 Web Query Developer Workbench 186
- DB2 Web Query for i 185–187
- DB2 Web Query for i - 5733-QU2 185
- DB2 Web Query for i Standard Edition 185, 187
- DB2 Web Query Report Broker 187
- DB2 Web Query SDK 187
- DB2 Web Query Software Developer Kit 187
- DB2 Web Query Software Development Kit 187
- DB2 Web Query Spreadsheet Client 186
- DB2 Web Query Standard Edition 187
- DB2 XML schema repository 116
- DB2 XSR 116
- DB2_CURSOR_HOLD 123
- DB2_CURSOR_NAME 122
- DB2_RESULT_SET_LOCATOR 123
- DB2_RESULT_SETS_COUNT 122
- DB2_ROW_COUNT_SECONDARY 119
- DB2.DEFAULT.POLICY 128–129
- DB2.DEFAULT.SERVICE 127–129
- DB2® for IBM® i supplied examples and tools 162
- db2-xdb
 - column 117
 - defaultSQLSchema 117
 - order 118
 - rowSet 117
 - rowSetOperationOrder 118
- DBGENCKEY 474, 479
- DDM 179
- DDM connection over TCP/IP 180
- DDM file I/O operation 180
- DDS BARCODE keyword 353
- DDS-created physical file 125
- DECFLOAT 132
- DECIMAL 132
- DECLARE VARIABLE 114
- decode 124
- decoded data attribute 125
- decomposition annotations 117
- decomposition or shredding 117
- Dedicated Processors Utilization 581
- Dedicated Service Tools 312
- DEFAULT 130
- default value 118
- default values 115
- define one of the new inherit rules 135
- definition of variables 130
- degraded array 276
- DELETE 118
- Delete Server 554
- DESCRIBE CURSOR 123
- DESCRIBE CURSOR statement 124
- DESCRIBE PROCEDURE 122–123
- DESCRIBE PROCEDURE statement 123
- description of the encoded value 125
- descriptor area item 122
- deterministic 134
- DETERMINISTIC_UDF_SCOPE 151
- Developer Workbench 187
- Developer Workbench feature 187
- developers 112
- developerWorks 112, 188
- device domain 103, 283
- DHCP 632
- different data types 132
- dimension tables 118
- dimensional definitions 186
- Disable constraints 170
- Disable or Enable Constraints option 170
- Disk Health Indicators 569
- Disk I/O Average Response Time Overview 597
- Disk I/O Rates Overview 595
- Disk I/O Rates Overview with Cache Statistics 596
- Disk I/O Total Response Time Overview 598
- Disk I/O Total Service Time Overview 599
- disk IO rates 300
- Disk Management 559
- disk response time buckets 284
- disk response time graphs 284
- disk unit path status 280
- Disk Watcher 230
- Display Expiration Schedule 25

- Display File Description (DSPFD) 143
- Display Library Description 135
- Display Recovery for Access Paths (DSPRCYAP) 139
- DISPLAY_JOURNAL 136
- DISTINCT 114
- DISTINCT keyword 114
- Distributed Backup Support 43
- distributed data 130
- distributed system 126
- DLTINTSVR 378–379
- dltnws 626
- document search 113
- Domain Name System 632
- Domino Attachment and Object Service 31
- Donated Processor Time 577
- DOUBLE 132
- DRDA 170, 179
- DRDA access 130
- DRDA connection 179–180
- DRDA® access 130
- drill down 186
- DROP 120
- DROP XSROBJECT 116
- DS5000 293
- DS5000 Copy Services 294
- DS5000 Native Attachment 292
- DS5100 292
- DS5300 292
- DS6000 296
- DS8000 296
- DSPASPSSN 101
- DSPFD 143
- DSPLIBD 135
- DSPRCYAP 139
- dual SAS adapter support 277
- dual SAS adapters 277
- Dual SAS RAID Adapter 277
- DVD/Tape SAS External Storage Unit for Power 795 CEC Rack 290
- Dynamic data movement 303

E

- ead-only queries 131
- EARLIEST_POSSIBLE_RELEASE 159
- easier searches of journal 136
- EBCDIC 114
- ECYELEM 34
- Edit Recovery for Access Paths (EDTRCYAP) 139
- EDTRCYAP 139
- ELSE IGNORE clause 119
- embedded SQL 112
- eMLC 312
- emote journal filtering 134
- emote name 130
- Enable Design Mode 600
- enable the replication of journal entries on a remote system 134
- encode 124
- encoded attribute data type 125
- encoded value 124

- encoded values 124–125
- Encoded vector index 145
- encoded vector index 133
- Encoded Vector Index (EVI) 145
- Encoded vector index symbol table probe 147
- Encoded vector indexes 112
- encoding 124
- encoding and encryption 124
- encoding name 115
- encoding specification 115
- encoding/decoding of data 124
- Encrypted ASP Enablement 282
- Encrypting 481
- Encryption enhancements (FIELDPROCS) 112
- encryption key 561
- ENDHTTPCRL 432
- Ending date 35
- Ending time 35
- endsch 627
- Energy Management 609
- Enhanced audit capability of RUNSQLSTM or STRSQL 182
- Enhanced Save/Restore Limits 29
- enhancements 529, 534
- Enhancements to products related to DB for i 111
- Enterprise Multi-level Cell technology 312
- environment value 170
- environment variable 166
- Escalation 386
- esclst 627
- Ethernet layer-2 bridging 335
- Ethernet link aggregation 332
- EVI 145, 147
- EVI indexes 147
- EVI symbol table scan 146
- exact inverse 125
- Example - Configuring four aggregated network links 334
- Excel 186
- exceptions 186
- EXCLUDING XMLDECLARATION 115
- execution of field procedures 125
- EXP12S 288
- EXP24S SFF Gen2-bay drawer 288
- expiration 62
- Export as PDF 556
- expression 130
- Expressions 112, 130
- expressions 113
- Expressions in a CALL statement 130
- extended attributes 634
- Extended Base Directory Support 632
- Extended Dynamic Remote SQL (EDRS) 167
- extended indicator variable values 130
- extended indicator variables 130
- Extended Integrated Server Support 632
- Extensible Markup Language 113
- External disk storage performance instrumentation 296
- external program 124
- external storage performance data collection 296
- EXTERNAL_NAME 153

externally referenced XML artifacts 116
EXTSTG category 296

F

family of DB2 products 112
Fast restore of a single object 28
feature 5117 136
Fibre Channel link statistics 296
field description 125
field procedure 124–126
field procedure parameter list 125
Field Procedure Parameter Value List (FPPVL 125
Field procedures 24, 124
Field Value Descriptor 126
field-decoded value 125
field-decoding 125
field-encoded value 125
field-encoding 125
field-encoding function 125
FIELDPROC attribute for a column 124
FIELDPROC clause 124
FIELDPROC support 124, 539
FIELDPROC_ENCODED_COMPARISON 151
FILE 159
Filter Messages by Control Groups 72
filter the remote journals 136
find trends 186
firmware 562
FlashCopy 99, 104
flexible service processor 106
for 525
FORCE 155
Force journal entry 155
foreign key 114
FPPVL 125
from 534
FSP 106
FTRINTNETA 184
FTRINTNETA pre-filter 184
FTRSQLCODE parameter 183
FTRUSER 184
FUNCTION 131, 152
function to be performed 125
functions
 %LEN 475
 %PARMNUM 477
 %SCANRPL 475
functions with two arguments 132
FVD 126

G

General Availability Resave 12
General Health Indicators 564
Geographic Mirroring 100
geographic mirroring 282
Global Mirror 99
Global Policy Properties Action 82
Global Policy Properties page 83
Global Statistics Cache 140

Global Statistics Cache (GSC) 142
global variable 120
Global Variables 538
Global variables 112, 120
GRANT 120
Graphical View 559
graphview 626
GROUP BY clause 114
group profile names 184
GSC 142

H

HA Journal Performance 136
HA setups 134
HA Switchable Resources 632
hange Logical File 148
hardware resource name 289
HBA 366
Health Center 536
HH DAT160 drive 290
Hierarchical queries using CONNECT BY clause 160
hierarchical storage management 307
History data 618
host languages 114
hot spare disk 276
Hot Spare for Mirroring 276
hot spare for mirroring 276
How Ethernet layer-2 bridging works on IBM i 336
HSM 307
HTML 115
HTML reports 186
http
 //www-03.ibm.com/systems/i/advantages/integrated-
 server/iscsi/solution_guide.html 30
HTTP Server 632

I

I/O operations 141
IASP 116, 139
IASP. 116
IBM 112
IBM Advanced Job Scheduler for i 386
IBM Developer Kit for Java 632
IBM Disk Sanitizer for i5/OS™ PRPQ 312
IBM HA Assist for i 107
IBM i 112
IBM i 7.1 112–113, 117, 139
IBM i Access for Web 623
IBM i Access for Web AFP to PDF Transform 626
IBM i Access for Windows 386, 622
IBM i Access for Windows Data Transfer 623
IBM i Doctor for IBM i external storage performance anal-
ysis functions 297
IBM i integration with BladeCenter and System x website
30
IBM i multipathing 290, 293
IBM i platform 186
IBM i Prerequisites for Ethernet layer-2 bridging 336
IBM i software 186

- IBM i XML Data Type 622
- IBM i Security Virtual private networking 331
- IBM iDoctor for IBM i 284
- IBM Portable Utilities for i 339
- IBM Portable Utilities for i (5733-SC1) supported versions 339
- IBM POWER 112
- IBM power state change 106
- IBM Software Technical Document 560810071 349
- IBM System Director Navigator for i BRMS initial screen 60
- IBM System Storage DS6000 296
- IBM System Storage DS8000 296
- IBM Systems Director Navigator 140, 284
- IBM Systems Director Navigator for i 47, 140, 285
- IBM Systems Hardware Information Center 281
- IBM Transform Services for i 556
- IBM Web Administration for i 433
- IBM WebSphere Portal V5.1 623
- IBM WebSphere Portal V6.0 623
- IDENTITY 125
- identity column 130
- identity columns 130
- iDoctor for IBM i 297
- IEEE 802.1D standard 336
- ield-decoding function 125
- IFNULL, HEX 113
- IFS 188
- IFS file 155
- IGNORE_DERIVED_INDEX 148
- IGNORE_DERIVED_INDEX(*NO) 148
- IKE version 2 558
- ILE activation 153
- ILE C 149
- ILE C program 153
- ILE C service program 153
- ILE COBOL 473, 517
- ILE compilers 513
- ILE RPG 475, 514
- imagecatalogs 626
- implicit serialization 115
- improve concurrency 131
- Improved Serviceability 30
- improving performance 133
- IN 114
- INCLUDE 133
- INCLUDE aggregates 146
- INCLUDE statement 133
- INCLUDING XMLDECLARATION 115
- incoming remote command (IRC) service 386
- independent ASP 116, 139
- independent ASPs 139
- independent auxiliary storage pool 116
- index advice condenser 148
- index column 114
- indexes 535
- Indexes list 535
- industry-standard SQL 112
- INFINITY 132
- InfoAssist report builder 185

- InfoPrint Server 556
- Information Center 137, 632
- Initialize Media using BRM 50
- inline function 134
- Inlining of SQL scalar function 112
- input argument 132
- input argument as the content 115
- input arguments 132
- input parameters 137
- INSERT 118
- INSERT VALUES clause 113
- INSINTSVR 378–379
- INSLNXSVR 381
- installation enhancements 622
- installation media 632
- INSTEAD OF 118
- INSTEAD OF trigger 177–178
- INSWNTSVR 380
- INTEGER 132
- integer values 132
- integrated file system 188
- Integrated Server Administration 552
- Integrated Server Environment 366
- Integrated Server Support 632
- Integrated Web application server 623
- integration functions 112
- integrity preservation 134
- Interactive Perspective Development 600
- internal representation 113
- Introduction to Ethernet layer-2 bridging 335
- INZMEDBRM 50
- IOP reset 100
- IPv6 108, 558
- IS NULL predicate 113
- iSCSI 366
- iSCSI software targets 366
- ISO 8879 113

J

- JAVA 124
- Java 170, 632
- Java Perspectives 592
- Java™ applications 115
- JDBC 113, 166, 170
- JDBC method 167
- JDK 632
- Job 138
- job CCSID 129
- job default CCSID 129
- Job groups 386
- Job Watcher 230
- journal entry typef X2 171
- Journal Management 548
- Journal management functions 140
- journal management functions 140
- Journal management topic collection 136
- Journal_Codes 137
- Journal_Entry_Types 137
- Journal_Library and Journal_Name 137
- journaling 109, 134, 634

Journals 134
JSR168 Print portlets 624
JVM 593, 632

K

Keeping tables or indexes in memory 147
KEEPINMEM 147
key 114
keyed logical file statistics 304

L

LAND 161
language API 113
LANGUAGE SQL 149
Large Object (LOB) data types 113
large objects 304
large-scale electronic publishing 113
Launch Web Console 378, 554
LENGTH 113
LENR lock 162
level of protection for multi-path 294
LEVELOFPROTECTION 285–286
library control path 295
library, 132
licensed program product 632
LIKE 114
Limit the amount of processing on a RGZPFM cancel 154
link statistics 296
LoadModule 432
LOB 113
LOB type 125
Local receives field 46
locator variable 122
locked data 131
Locked rows 131
logical files 148
logical files over a partition table 148
logical library 295
logical view of the Layer-2 bridging 336
LOGICAL_READS 304
LOR 161
Lotus products 484
LPP 632
LUN level switching 103

M

Maintained Temporary Index 534
Maintenance Enhancements 47
Manage Permissions 459
Managing XML Schema Repositories 116
Mark and Unmark Volumes for Duplication 75
Mark volume(s) for duplication 77
Marked for Duplication column 79
Marker PTF 13–14
master key 283
Materialized Query Tables (MQT) 145
maximum cardinality 121

Maximum Historical Collections (COLNS) 230
maximum length 113
MAXTRIES 230
Media maintenance 62
media preference 162, 302, 308
media preferences for SSD 276
Memo to Users 632
Memory 564
Memory Pools Health Indicators 570
MEMORY_POOL_PREFERENCE 147, 151
menu bar 617
MERGE 113, 118–119
MERGE statement 112, 118–119
message CPDB8EC 139
message data 129
message text area 126
meta data 186
Metric Finder 605
Metrics 606
Metro Mirror 99
MF47377 302
MF47877 304
MF47878 303
MF47879 304
MF47887 303
MF47888 304
MF47889 304
MF47891 303
MF47892 304
MF47893 304
MF48544 307
MF49299 309
MF49371 309
MF49399 307
MF50873 314
MF50875 314
MF52834 314
Microsoft Excel® 186
Microsoft SQL Server 186
Migrating from AnyNet to Enterprise Extender 332
minimize transaction wait time 131
MINRLS 159
mirrorsync 626
Miscellaneous 130
mixed 114
Modifying SQL statement 604
Monitor 178
moves 62
MQ location 127–128
MQ Series 112
MQ Series integration 126
MQREAD 126
MQREADALL 127
MQREADALLCLOB 128
MQREADCLOB 126
MQRECEIVE 127
MQRECEIVEALL 128
MQRECEIVEALLCLOB 128
MQRECEIVECLOB 127
MQSEND 127, 129

- MQSeries 112, 127
- MQSeries® 126
- MQT 145
- MSG result column 129
- MSGCPF5257 166
- MSGCPF9999 166
- msg-data parameter 129
- msg-data return parameter 129
- MSGSQL0438 160
- MSGSQL0443 160
- MTI 534
- multi-data space logical files 148
- multipathing 293
- Multiple e-mail addresse 82
- Multiple scheduling environments 386
- multiple scheduling environments 386
- multi-row UPDATE, DELETE, and MERGE 623
- multi-row UPDATE, DELETE, and MERGE statements 622
- Mware ESX server management 371

N

- N_Port ID Virtualization 247
- namespace declarations 114
- NaN 132
- national language versions 632
- native 124
- Navigate Resources 55
- Navigation to BRMS Functions 54
- nested SQL operation 118
- nested trigger 170
- Network Interface Cards 366
- Network Properties page 84
- Networking 558
- New functionality for DB2 developers 111
- New Package 601
- New Perspective 602
- newesclst 627
- newjobgrp 627
- newmail 627
- newoutqmon 627
- newrecip 627
- newrepdstlst 627
- newschjob 627
- NICs 366
- NLV 632
- no conversion 114
- no persistent storage 178
- Nominal CPU Ratio 610
- non-binary data type 115
- non-Database clients 184
- non-partitioned 130
- non-RDB DDM file connection request 180
- non-RDB DDM files 181
- non-XML data type 115
- NOT ATOMIC MERGE 119
- Notification and report distribution 386
- notifyprop 627
- NPIV 100, 247
- nput XML value 115

- number of elements 121
- number of result sets 122
- NWSD 371
- NWSH 366
- nwswebcon 626

O

- OBJDESC 229
- object references 130
- object signing 634
- Object_Library 137
- Object_Member 137
- Object_Name 137
- Object_ObjType 137
- Object_Statistics table function 159
- OBJTYPE 229
- ODBC 113, 166
- OLAP 187
- OLAP Module 186
- OLE Data Provider 622
- Omnifind 113, 148
- Omnifind 547
- Omnifind™ Text Search Server for DB2 for i 188
- Open Backup Control Groups 64
- OpenSSH 339
- OpenSSL 339
- Operations Console 623
- OPM 514
- OPM *PGM 124
- OPRCTL(*YES) 347
- option 42 of the IBM i Operating System 136
- ORDER BY clause 114
- order of result rows 117
- ordered collection of data elements 121
- Original Program Model 514
- original value 124
- OS Type 57
- ournal receiver 140, 155
- Output file support 37
- Output Queue 626
- outqmon 627
- OVERRIDE 33
- Override recovery element 33
- Overview of Ethernet link Aggregation 332
- OVRDBF 134
- OVRDBF SEQONLY(*YES N). 134

P

- Paging Service Partition 242
- Paging Space Device 242
- Parameter list 125
- Parameter marker enhancements 112
- Parameter marker values for a dynamically prepared statement 113
- Parameter Markers 130
- parameter's value 125
- parent unique constraint 130
- parity optimization method 281
- parity sets 281

- paritysets 626
- partitioned cluster 105
- partitioned key 130
- partitioned SQL table 303
- partitioned table 130
- Partitioned table support 130
- Partitioned tables with referential constraints 130
- passive path 277, 293
- PDF Output Settings 625
- PDFENCRYPT 349
- PDI 564
- Performance 563
- Performance and query optimization 111, 140
- performance data collection 289
- Performance Data Investigator 284, 564
- Performance Explorer 606
- performance optimization 141
- Personal Communications 623
- Personal folder 626
- PHYP 106
- physical file statistics 304
- Physical Processors Utilization 580, 582–583
- plain tex 115
- Plan Cache masking of variable values 178
- Plan Cache properties 185
- Plug-ins 432
- port 112
- Portable Document Format (PDF) 345
- portal environment 623
- POSITION parameter 28
- POWER 112
- POWER Hypervisor 106
- powerful recovery 134
- PowerHA 92
- PowerHA NPIV support 100
- PowerHA SystemMirror for i 92
- PowerVM 242, 290
- PowerVM Virtual I/O Server 290
- pplication requestor (AR) 180
- pre check 185
- Pre check physical file size during restore 185
- predicate 113–114
- preferred controller 293
- preferred media attribute 302
- preferred unit 302
- PREPARE 131
- Prepared Statement Area 154
- pre-requisite to using an OLAP 186
- Prerequisites for Ethernet link aggregation 332
- Preserve SQE Cache Size Across IPLs 184
- PRESERVE WHITESPACE 115
- previous release 130
- primar 114
- primary key 118
- Print Media Movement 49
- printer device descriptions 107
- Printer Output 555
- PROCEDURE 131
- procedure name 123
- Procedure parameters 113

- PROCESS statement 475
- processor tier 186
- Program 138
- PROGRAM NAME clause 152
- PROGRAM NAME on CREATE TRIGGER 161
- PROGRAM TYPE MAIN 153
- PROGRAM TYPE SUB 153
- PROGRAM TYPE SUB clause 153
- programming languages 121
- protect the changes to objects 134
- PRPQ 93
- PRTDEV 107
- PRTMOVBRM 49
- PSA 154
- PSA compression 154
- PSEUDO_OPEN_CHECK_HOST_VARS 151
- PSFTRACE 353
- publishing functions 113

Q

- QALLUSRLNK 32
- QAPMDISK 284
- QAPMDISKRB 284–285
- QAPMJOBOS 156
- QAPMSHRMP 221
- QAPMTAPE 289
- QAPMXSTGD 296
- QAPMXSTGV 296
- QAQQDBMN 174
- QAQQINI 114, 131, 147, 150
- QAQQINI file parameter 114
- QAQQINI properties 150
- QAQQINI value 114
- QAUDJRN 171
- QAUTPROF 348
- QCLUMGT 348
- QCOLSRV 348
- QDB_CHGPF CST 154
- QDBFSTCCOL 143
- QDBRPLAY() API 170
- QDBRTVSN API performance 155
- QDBRTVSN() API 155
- QDBSHR 348
- QDBSHRDO 348
- QDDMDR DASERVER 179–182
- QDDMDR DASERVER server authentication entry special value 179
- QDDMSERVER 180–181
- QDFTOWN 348
- QDIRSRV 348
- QDLFM 348
- QDOC 348
- QDSNX 348
- QFNC 348
- QGATE 348
- QGENOBJNAM 152
- QGENOBJNAM data area 152
- QIBM_DB_SQLADM 169, 173
- QIBM_DB_SQLADM Function usage 169
- QIBM_DB_SQLMON 173

- QIBM_DB_SYSMON 173
- QIBM_DB_SYSMON Function 169
- QIBM_DB_SYSMON Function usage 169
- QIBM_NO_901_LOGGING 170
- QIBM_SP_SECURITY exit point 349
- QIBM_SQL_NO_RLA_CANCEL 166
- QIBMLINK 32
- QILETS 474, 479
- QjoChangeJournalState 136
- QjoRetrieveJournalEntries API 137
- QJOSJRNE API 155
- QL schema statements 114
- QLPAUTO 348
- QLPINSTALL 348
- QM 171
- QMHSNDPM() 160
- QMSF 348
- QNETSPLF 348
- QNFSANON 348
- QNTP 348
- QP0LFLOP (Perform File System Operation) 311
- Qp0lGetAttr()--Get Attributes 311
- QP1ARCY 32–33
- QPEX 348
- QPM400 348
- QRECOVERY 139
- QRECOVERY.QSQ901S 170–171
- QRJE 348
- QSNADS 348
- QSPL 348
- QSPLJOB 348
- QSQBIGPSA 154
- QSQBIGPSA data area 154
- QSQCSRTH data area 154
- QSQSRVR 168–169
- QSQSRVR job 150, 168
- QSQSRVR jobs 150
- QsrCreateMediaDefinition 28
- QsrRestore 28
- QSRRSTO 28
- QSRVAGT 348
- QSVTAP22 432
- QSYS 348
- QSYS.TCPIP_INFO view 163
- QSYS2 137
- QSYS2.CANCEL_SQL 141, 166, 168
- QSYS2.CANCEL_SQL procedure 166
- QSYS2.Display_Journal 137
- QSYS2.DUMP_SQL_CURSORS procedure 165
- QSYS2.FIND_AND_CANCEL_QSQSRVR_SQL 166, 168
- QSYS2.FIND_AND_CANCEL_QSQSRVR_SQL procedure 168
- QSYS2.FIND_QSQSRVR_JOBS and QSYS2.CANCEL_SQL() 168
- QSYS2.FIND_QSQSRVR_JOBS() 168
- QSYS2.FIND_QSQSRVR_JOBS() procedure 168
- QSYS2.GROUP_PTF_INFO view 163
- QSYS2.INDEX_ADVICE 148
- QSYS2.INDEX_ADVICE procedure 148

- QSYS2.Object_Statistics table function 159
- QSYS2.QIBM_SQL_NO_RLA_CANCEL environment variable 166
- QSYS2.SQLQMprofiles 172
- QSYS2.SQLQMProfilesAudit 171
- QSYS2.SYSCOLUMNS 125, 172
- QSYS2.SYSCOLUMNS2 156, 172
- QSYS2.SYSFIELDS 125
- QSYS2.SYSPACKAGESTAT 159
- QSYS2.SYSPARTITIONDISK 157
- QSYS2.SYSPARTITIONINDEXDISK 157
- QSYS2.SYSPROGRAMSTAT 159
- QSYS2.SYSROUTINES 153
- QSYS2.TCPIP_INFO 163
- QSYS2/CANCEL_SQL() procedure 167
- QSYS2/CONDENSEDINDEXADVICE 148
- QSYS2/SYSIXADV 148
- QTCP 348
- QTEMP.CONDENSEDINDEXADVICE 149
- QTEMP/CONDENSEDINDEXADVICE 148
- QTFTP 348
- QTSTRQS 348
- Query Manager 171
- Query Manager profile auditing 171
- query optimizer 112
- Query/400 185–186
- Query/400 import function 186
- queue 126
- QWAS61x 432
- QWAS7x 433
- QZDASOINIT 205
- QZRCSRVS 205

R

- radio 528
- RAID parity sets 281
- RANDOM_READS 141, 304
- rank I/O rates 300
- rank statistics 296
- Rational Development Studio for i 514
- Rational for Power Systems 501
- RCV_ACCEPT_TRUNC_MSG 129
- RCV_CONVERT 129
- RDB DDM file connection request 180
- RDB DDM files 182
- read cache hit% 298
- read only queries 148
- read-only queries 131
- REAL 132
- Receive Media Information column 45
- receive-service 126–128
- reciplst 627
- record lock conflicts 131
- Recovery Defaults 40
- recovery report 36
- Redbooks Web site
 - Contact us xxiv
- reduce the recovery time 134
- Redundant Paging Service partitions 245
- redundant VIOS 290

- redundant Virtual I/O Server 290
- referenced alias 130
- referential constraint 130
- Referential integrity 153
- referential integrity 130
- referential integrity constraints 118
- registration of an XML schema 116
- related 526
- Relational Database Directory 179
- relational database tables 117
- relational database technology 112
- Remote Access Services 558
- Remote command jobs 386
- remote database 130
- Remote Executing Daemon 386
- remote journal 140
- remote journal connection 136
- Remote journal filtering 136
- remote journal filtering 136
- remote journal filtering feature 136
- Remote journaling 136
- remote objects 130
- Remote overrides field 47
- Remote receives field 46
- remote relational database 130
- remote site 130
- remote SQL Server databases 186
- Remove Journalized Changes 136
- Remove of Disk Units 274
- removed permanently 138
- Reorganize Physical File Member 154–155
- repdstlst 627
- replace an object 132
- REPLACE option on CREATE commands 112
- Report and Graph Assistant tools 185
- Report Assistant 186
- Report Assistant tool 186
- reporting functions 186
- reporting server 186
- request XML serialization explicitly. 115
- requirements for solid state drives 302
- Resave 14
- resetschjob 627
- response time buckets 284
- Response Time Health Indicators 571
- Restore (IFS) Object using BRM (RSTBRM) command 38
- Restore all libraries other than the system library 30
- Restore all user data 30
- Restore Command Defaults panel 42
- Restore entire system 30
- Restore Library using BRM (RSTLIBBRM) command 38
- Restore Object using BRM (RSTOBJBRM) command 38
- Restore system data only 30
- RESUBMIT 230
- Result set 112
- result set locator 123
- result set locator values 123
- result set locator variables 123
- result set locators 123

- Result set support 122
- result sets 122–123
- results sets 122
- Retrieve CL source RTVCLSRC 480
- Retrieve Directory Information 311
- retrieve the data from the SQLXML object 115
- RETURN statement 134
- RETURN statement. 134
- REXECD 386
- RGZPFM 154–155
- rmtsyswebcon 626
- RMVJRNCHG 136
- RMVSTG 289
- ROLLBACK 120
- rollback 131
- ROW CHANGE TIMESTAMP 125
- ROW_COUNT 119
- ROWID 125
- RPG 122
- RPG program 124
- RSMHTTPCRL 432
- RST 28
- RSTASP 42
- RSTLIB 28
- RSTLIBBRM 38
- RSTOBJ 28
- RSTOBJBRM 38
- RTNPARM keyword 476
- RTVCLSRC 480
- RTVDIRINF 311
- RTVSYSINF 30
- Run Maintenance Options 67
- Run SQL Scripts 169
- Run Time User Enablement 186
- Run Time User Enablement feature 187
- RUNRMTCMD 386
- RUNSQLSTM 182
- Runtime User Enablement 187

S

- S NULL predicate 113
- SAS 277
- SAS adapter roles 279
- SAS adapter-to-adapter (AA) cable 277
- SAS X cable 277
- SAVCHGOBJ 30
- SAVDOMBRM 31
- Save all libraries other than the system library 30
- Save all user data 30
- Save entire system 30
- Save system data only 30
- save while active 30
- Save While Active Support for Integrated VMware ESX Servers 30
- SAVLIB 30
- SAVPFRCOL/FTPFILE 210
- SAVSYS 283
- scalar functions 113, 132, 161
- schactprop 627
- Schedule Information page 69

Scheduling a job 386
 Scheduling BRMS maintenance 66
 Scheduling Support for BRMS 62
 schema 132, 139
 schema-based decomposition 117
 schjobprop 627
 SCORE 113
 SCPF 348
 SCSI element address 289
 SCSI Log Sense 296
 sector conversion 293
 SECURE Columns 178
 SECURE NO 179
 Secure Perspective for System i 632
 SECURE YES 173, 179
 SELECT 120
 Select Advanced Analysis Command screen 286, 312
 select ASP on RSTLIBBRM and RSTOBJBRM commands 42
 SELECT INTO 131
 SELECT list containing the DISTINCT keyword 114
 Select Recovery Items panel 40
 Selecting the SSDSANITIZE macro. 313
 select-statement 131
 sentmail 627
 SEQUENCE 131
 sequence of row elements 115
 sequence of XML element nodes 114
 SEQUENTIAL_READS 141, 304
 serial-attached SCSI 277
 serialized data 115
 serialized string format 115
 serialized string value 113
 serialized XML value 115
 server authentication entries 179, 181
 server authentication entry 179–180
 service and policy definitions 129
 service processor 554
 Service Processor Manager 370
 service-policy 126–127
 session scope 120
 SET and VALUES INTO 113
 SET CONNECTION 130
 SET CURRENT IMPLICIT XMLPARSE OPTION 115
 Set Object Access 147
 SET OPTION 131, 150
 Set Target System 527
 SETOBJACC 147
 SF99369 30
 SF99707 301
 SGML 113
 Shared Memory 584
 Shared Memory Partitions 242
 Sharing physical Ethernet connections across partitions (Ethernet layer-2 bridging) 335
 SI23390 284
 SI34094 30
 SI34095 30
 SI36892 340
 SI38685 340
 SI39965 340
 SI40092 340
 SI41724 340
 SI43471 347
 SIGNAL SQL statement 160
 SIGNAL statement 160
 SIGNAL support for native triggers 160
 simple SQL scalar functions 134
 simplification of matching rows in tables 118
 single byte 114
 single level store architecture 112
 single processing instruction 115
 single relational database 130
 single SQL statement 130
 single top-level element 114
 single XML value 114
 Size next Upgrade 589
 SKIP LOCKED DATA 131, 149
 Skip locked data 131
 SKIP LOCKED DATA and NC or UR 149
 slice and dice 186
 SMALLINT 132
 smgetstayoffssd 305
 smresetstayoffssd 305
 smsetstayoffssd 305
 smsetstayoffssd macro 306
 software initiator 366
 software vendors 112
 solid state drives 302
 space-efficient FlashCopy 104
 sparse index 145
 Sparse indexes 141, 145
 SPCAUT(*JOBCTL) 347
 SPCAUT(*SPLCTL) 347
 special register 115
 Spool Files 188
 spreadsheet 186
 Spreadsheet Client 186–187
 SQE 142–143, 147–148
 SQL 112, 124–125, 152–153
 SQL Call Level Interface 162
 SQL Call Level Interface (CLI) 167
 SQL CLI applications 166
 SQL data description and data manipulation language 112
 111
 SQL data type 115, 117, 125
 SQL descriptor 123–124
 SQL Diagnostics Area 119
 SQL function 134
 SQL language 112
 SQL Performance Monitor 178, 184
 SQL Plan Cache 178, 184
 SQL plan cache 184
 SQL plan cache size 184
 SQL procedures 121
 SQL pseudo opens 156
 SQL query 114
 SQL Query Engine (SQE) 142
 SQL return code 122

SQL routine performance integer arithmetic 149
 SQL routines 152
 SQL Server 186
 SQL Server Mixed Mode for Batch Processing 169
 SQL Server Mode 150, 168–169
 SQL statement 123, 166
 SQL statement level detail 159
 SQL stored procedures 121
 SQL trigger 170
 SQL trigger program 169
 SQL_CONCURRENT_ACCESS_RESOLUTION 151
 SQL_CONCURRENT_ACCESS_RESOLUTION option 131
 SQL_XML_DATA_CCSID 114, 151
 SQL) 152
 SQL0443 167
 SQL0552 167
 SQL0901 167
 SQL0901 logger education 170
 SQL0952 166–167
 SQLarrays 121
 SQLCA 119, 160
 SQLCODE 122, 167, 183
 SQLDA 122–124
 SQLDAID 123
 SQLDATA 123
 SQLDATA field 124
 SQLDATA fields 123
 SQLERRD(3) 119
 SQLERRD(5) 119
 SQLIND 123
 SQLNAME 123
 SQL-routine-body 134
 SQLSTATE 126, 132, 160
 SQLVAR 123
 SQLXML object methods 115
 srvprcwebcon 626
 SSD 141, 148, 157, 162, 302
 SSD Analyzer Tool for IBM i 302
 SSD versus HDD space usage 306
 SSDSANITIZE 312
 SSDSANITIZE macro help panel. 313
 standalone 626
 Standard Edition 186
 standard labeled tapes 289
 Start Database Monitor 182
 Start Expiration for BRM 51
 Start Journal Library 134–135
 start journaling on new production work files 134
 Start Maintenance for BRM 47, 51
 Start Recovery using BRM 41
 Start/stop encryption 561
 Starting_Receiver_Library 137
 Starting_Receiver_Name 137
 Starting_Sequence 137
 Starting_Timestamp 137
 stateless 626
 statement 114, 123
 statement level 131
 Statements 130
 statvfs()--Get File System Information 311
 STMFOPT 345
 STMFOPT parameter 346
 storage management macros 305
 Storage Pool 147
 storage tiers 308
 stored procedure 121–123
 stored procedure name 123
 stored procedures 121
 Storwize 93
 STRASPBAL 276, 307–309
 STRDBMON 182–184
 STRDBMON Over a View 174
 STRDWMON 230
 Stream File 188
 Stream file option 345
 Stream file option (STMFOPT) parameter 344
 stream file using the Control character 345
 STREXPBRM 51
 STRHTTPCRL 432
 string value 113
 STRIP WHITESPACE 115
 STRJRNLIB 134–135
 STRJWMON 230
 STRMNTBRM 47, 51
 STRPAMON 230
 STRRCYBRM 33, 41
 strsch 627
 STRSQL 169, 182
 subselect of a fullselect that is not UNION ALL 114
 SUM 146
 support 527
 support of more than 999,999 objects 32
 SVC 93
 SWA 30
 sweeper function 308
 switched disks 103
 synchronization 107
 SYSBAS 274
 SYSCOLUMNS 156
 SYSCOLUMNS2 156
 SYSIBM.MQPOLICY_TABLE 129
 SYSIBM.MQSERVICE_TABLE 129
 SYSINDEXSTAT 141, 143, 304
 sysoprmsg 626
 SYSPARTITIONDISK 306
 SYSPARTITIONSTAT 304
 SYSPROC.SET_COLUMN_ATTRIBUTE 172–173, 178
 SYSPROC.SET_COLUMN_ATTRIBUTE() 172
 SYSPROC.XDBDECOMPXML 117
 SYSPROC.XSR_ADDSCHEMADOC 116
 SYSPROC.XSR_COMPLETE 116
 SYSPROC.XSR_REGISTER 116
 SYSPROC.XSR_REMOVE 116
 SYSTABLESTAT 141, 143
 system disk pool 139
 System i Access for Wireless 632
 System i Access for wireless 630
 System i Navigator 169, 184
 System i Navigator tasks on the web 626

- System i Navigator's Run SQL Scripts 167
- System Managed Access Path Protection (SMAPP) 139
- system name 30, 120
- system performance 139
- System Resources Health Indicators 566
- System Service Tools 312
- System-managed access-path protection 139
- Systems Director - IBM i Work Management all tasks menu 394
- Systems Director Navigator for i 135
- SYSTOOLS 162
- SYSTOOLS procedures 162

T

- t 527
- table 113
- table column 121
- table function 137
- Table functions 127
- table or the result of a query 114–115
- Tape Image Catalogs 562
- Tape Libraries 562
- Tape library 295
- tape library 289
- Tape Support 562
- tapelibraries 626
- TAPMLBxx 295
- TCP/IP network 179
- tcpipattrIPv6 626
- Technology Refresh 12, 17
- Technology Refresh PTF 12
- Technology Refresh PTF Group 12
- TELNET 184
- temporary file systems 634
- temporary table 122
- temporary UDFS 633
- TERASPACE 474, 479
- text search collection 188
- TEXT_SEARCH_DEFAULT_TIMEZONE 151
- TGTRLS 30, 150
- The DB2 Web Query Adapter for Oracle JD Edwards 186
- The DB2_RESULT_SET_ROWS 123
- The thread has an uncommitted transaction 139
- Thin Provisioning for DS8700 and DS8800 storage servers 301
- Thin Provisioning for DS8700, DS8800 storage servers and for VIOS shared storage pools 301
- three-part name aliases 130
- three-part names 130
- Three-part names support 130
- thresholds 572
- TINYINT 162
- TMPUDFS 556
- To stream file 344
- To stream file (TOSTMF) parameter 344
- Toolbox for Java 632
- tools 112
- TOSTMF 344
- TOSTMF parameter 346
- Total DB Opens job level instrumentation 156

- Traditional Database Monitor 178
- transform 124
- transformation 125
- Transforms – AFP to PDF Transform 347
- TRANSLATE 161
- transparent encryption 124
- transparent encryption/decryption 124
- TRCASPBAL 307, 309–310
- TRIGGER 131–132, 152
- trigger 118
- Trigger correlation variables 113
- trigger performance 153
- Triggered Cache Manager 632
- truncation 129
- type annotations 115

U

- UDFS 556, 633
- UDFS media preference 302, 310
- UDTF 137
- UNASSIGNED 130
- Uncapped Processor Time 578
- Unicode CCSID 114
- Unicode specific version 114
- Uniform Resource Identifier 116
- unique 114
- unique index 130
- UNIT parameter 287
- UNIT parameter 302
- UNIT(*SSD) 148
- Unmark volume(s) for duplication option 78
- UNNEST 121
- UPDATE 118, 131
- UPDATE SET clause 113
- Update System Information 36
- upgrade 186
- upsert 118
- URI 116
- USE CURRENTLY COMMITTED 131
- Use currently committed 131
- USER 125
- User 138
- user class 120
- User Defined Function 170
- user defined functions 113
- user defined types 113
- User expiration interval 25
- USER_CLASS 120
- user-defined data type 125
- user-defined file systems 633
- User-defined function arguments and result 113
- USPS barcodes 353
- USREXPDATE 24
- UTF-16 114
- UTF-8 114

V

- Validate constraints without checking 154
- value 114

- VALUE clause 473
- VARCHAR 126–127
- VARIABLE 132
- variable length fields 304
- variable name 120
- variable number of result sets 122
- VARIANCE 146
- Verify System Information 36
- VIEW 132
- View as PDF 555
- Viewing scheduled tasks 70
- VIOS 242
- VIOS Shared Storage Pools 301
- Virtual Fibre Channel 248
- Virtual I/O 290, 608
- Virtual I/O Server 273
- Virtual I/O Server (VIOS) 242
- Virtual SCSI 247
- Virtual Shared Processor Pool 579
- Visual Studio 2008 622
- VMware ESX servers 30
- VMware support 370
- VPN 558

W

- Wait for outcome 131
- Warning Threshold 573
- Web Application Server environment 623
- web based authoring tools 186
- web based tools 185
- Web crawler support 432
- web enable 185
- Web log monitor 447
- Web Performance Advisor 466
- Web Services 187
- WebSphere 129
- WebSphere Application Server 442, 632
- WebSphere Application Server V6.0 623
- WebSphere Application Server V7.0 623
- WebSphere Application Servers 432
- WebSphere Development Studio for IBM i 514
- WebSphere MQ 126, 129
- weighted ASP balancer 307
- well-formed XML document 113
- WHEN clause 119
- Windows 2000 622
- WITH HOLD, 123
- WITH ORDINALITY 122
- Work with Hardware Resources 339
- Work with Journal Attributes 136
- Work with Media Information 39
- Work with Media Library Media 50
- Work with Media using BRM 50
- Work with Saved Members 52
- Work with Saved Objects 52
- Workload Estimator 589, 664
- Workstation customizing object 344
- Workstation customizing object (WSCST) parameter 344
- Workstation customizing object name 345
- write cache hit% 298

- WRKHDWRSC 339
- WRKJRNA 136
- WRKMEDBRM 50
- WRKMEDIBRM 32, 39, 52
- WRKMLMBRM 50
- WSCST 344
- WSCST parameter 346

X

- xmforest 114
- XML 113, 115
- XML AS CLOB 114
- XML AS DBCLOB 114
- XML attributes from columns, 114
- XML content 113
- XML data 113–115
- XML data can 114
- XML Data Type 113, 623
- XML data type 113–114
- XML Data type Column support 537
- XML data types 113–114
- XML declaration 115
- XML decomposition enhancements (order of result rows) 117
- XML document 113–114, 117
- XML document and annotation 113
- XML element 114
- XML file reference variables 113
- XML GENERATE 474
- XML host variables 114
- XML instance documents 116
- XML locators 113
- XML Publishing functions 114
- XML scalar functions 113
- XML schema 116–117
- XML schema repository 116
- XML schema repository 116
- XML Schema Repository (XSR) 537
- XML schema validation 115
- XML schemas 116
- XML sequence 113–114
- XML Serialization 115
- XML serialization 115
- XML suppor 112
- XML Support 113
- XML value 113–115
- XML values 113
- xmlagg 114
- xmlattributes 114
- XMLCAST a parameter marker 113
- xmlcomment 114
- xmlconcat 114
- xmldocument 114
- xmlelement 114
- xmlgroup 114
- xmlnamespaces 114
- XMLPARSE 113
- xmlparse 114
- XMLPARSE function 115
- XMLPARSE functions 114

- xmlpi 115
- xmlrow 115
- XMLSERIALIZE 113, 116
- xmlserialize 115
- XMLSERIALIZE function 115
- XMLSERIALIZE invocation 115
- XMLSERIALIZE scalar function 115
- xmltext 115
- xmlvalidate 115
- XOR 161
- XSLT processor 115
- xmltransform 115
- XSM 100
- XSR 116–117

Z

- zlib 339



Redbooks

IBM i 7.1 Technical Overview

(1.0" spine)

0.875" x 1.498"

460 <-> 788 pages



IBM i 7.1

Technical Overview



**Boost efficiency
with enriched
virtualization and
storage integration**

**Exploit enhanced
database functions
and new application
development tools**

**Benefit from
centralized system
management and
performance tooling**

This IBM Redbooks publication introduces a technical overview of the main new features, functions, and enhancements available in IBM i 7.1. It gives a summary and brief explanation of new capabilities and what has changed in the operating system, and also discusses many of the licensed programs and application development tools associated with IBM i.

Many of the new and enhanced functions are described:

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

The information provided in this book is useful for clients, business partners and IBM service professionals 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